

Chapter 7

Day 4: Linear Systems of Algebraic Equations

7.1 Schedule

- 0900-0915: Debrief
- 0915-1000: Synthesis
- 1000-1030: Applications of LSAE
- 1030-1045: Coffee
- 1045-1115: Applications of LSAE
- 1115-1200: Concept Map for Eigenfaces

7.2 Debrief

- Please discuss your overnight work with your table-mates, create a set of key concepts, and a set of ideas that you are still confused by.

7.3 Synthesis

We will increasingly use a computational tool like MATLAB to compute determinants, matrix inverses, and the solutions to linear systems of algebraic equations. In this synthesis section we will explore the theoretical foundation of these algorithms - the so-called LU decomposition.

Gaussian Elimination

The basic process of *elimination of variables* can be formalized and is known as Gaussian Elimination. Here we will briefly introduce it but you can consult other sources, such as the *Gaussian Elimination* page at WolframMathWorld for more details.

Rather than writing equations, we can cast a LSAE in matrix form and perform *Gaussian Elimination* on the augmented matrix $[\mathbf{A} \ \mathbf{b}]$.

For example, the linear systems of algebraic equations

$$\begin{aligned}2x_1 + 3x_2 &= 6 \\4x_1 + 9x_2 &= 15\end{aligned}$$

can be written as the following augmented matrix

$$\begin{bmatrix} 2 & 3 & 6 \\ 4 & 9 & 15 \end{bmatrix}$$

Thinking now in terms of rows, we replace the second row with row 2 - 2 row 1 to give

$$\begin{bmatrix} 2 & 3 & 6 \\ 0 & 3 & 3 \end{bmatrix}$$

This matrix is now in so-called *echelon* form: we can find the solution to the original LSAE by first solving the equation implied by the last row and then back-substituting into the equation implied by the previous row.

Exercise 7.1

1. Set up the augmented matrix for the following example (you will recognise this from the last assignment)

$$\begin{aligned} 2x_1 + x_2 &= 13 \\ 4x_1 + 3x_2 &= 33 \end{aligned}$$

and perform *Gaussian Elimination* to reduce the augmented matrix to *echelon form*. Interpret the resulting system and determine the solution(s).

LU Decomposition

The steps used to solve a LSAE using Gaussian Elimination can also be used to *decompose* a matrix into a product of two matrices: a *lower-triangular* matrix \mathbf{L} and an *upper-triangular* matrix \mathbf{U} . Here we will briefly introduce it but you could consult other sources, such as the *LU Decomposition* page at WolframMathWorld for more details.

In Gaussian Elimination we execute a set of row operations. In our ongoing example, we replaced row 2 with the result of row 2 - 2 row 1. This action can be neatly represented in terms of a matrix operation. Let's multiply the original matrix equation $\mathbf{Ax} = \mathbf{b}$ with the transformation matrix

$$\mathbf{M} = \begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix}$$

to form $\mathbf{MAx} = \mathbf{Mb}$. Note that this transformation leaves row 1 of \mathbf{A} unchanged, and it replaces the row 2 with row 2 - 2 row 1. The product \mathbf{MA} is therefore an *upper-triangular* matrix \mathbf{U}

$$\mathbf{U} = \begin{bmatrix} 2 & 3 \\ 0 & 3 \end{bmatrix}$$

and the LSAE is now expressed as $\mathbf{Ux} = \mathbf{Mb}$. If we now multiply this expression by \mathbf{M}^{-1} we obtain

$$\mathbf{M}^{-1}\mathbf{Ux} = \mathbf{b}$$

The inverse of \mathbf{M} is straight-forward to write down because it "undoes" the row operations

$$\mathbf{M}^{-1} = \begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix}$$

Notice that this matrix is just a *lower-triangular* matrix \mathbf{L} . The LSAE now reads

$$\mathbf{LUx} = \mathbf{b}$$

We have therefore *decomposed* the original matrix \mathbf{A} into the product of \mathbf{L} and \mathbf{U} ,

$$\mathbf{A} = \mathbf{LU}$$

How does this help, you might be asking? First of all, knowing the decomposition of \mathbf{A} into \mathbf{LU} allows us to solve the original LSAE $\mathbf{Ax} = \mathbf{b}$. Here is how.

Let's define a new vector $\mathbf{y} = \mathbf{Ux}$. Then the original LSAE can be expressed as

$$\mathbf{Ly} = \mathbf{b}$$

which is straight-forward to solve by *forward-substitution* because \mathbf{L} is *lower-triangular*,

$$\begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 6 \\ 15 \end{bmatrix}$$

and the solution for \mathbf{y} is $y_1 = 6$, $y_2 = 3$. We can now solve $\mathbf{Ux} = \mathbf{y}$ for \mathbf{x} using *forward-substitution* because \mathbf{U} is *upper-triangular*,

$$\begin{bmatrix} 2 & 3 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 6 \\ 3 \end{bmatrix}$$

and the solution for \mathbf{x} is $x_1 = 1$, $x_2 = 3/2$.

Second of all, and more importantly, knowing the decomposition of \mathbf{A} into \mathbf{LU} allows us to solve any LSAE involving \mathbf{A} . Need to solve the LSAE with a different \mathbf{b} ? No problem, just use the \mathbf{LU} decomposition that you already computed and away you go. No need to redo all the steps of *Gaussian Elimination* just because \mathbf{b} changed. Need to solve a LSAE for lots of different \mathbf{b} 's? No problem, just use the \mathbf{LU} decomposition that you already computed and away you go. Finally, if you want to compute the inverse or determinant of a matrix this is easy too using LU decomposition as we show next.

There is an algorithm in MATLAB, *lu*, which does LU decomposition for you, but you should not necessarily expect to get the same \mathbf{L} and \mathbf{U} , even for this example. (There are a variety of ways to define the \mathbf{L} and \mathbf{U} matrices, but this is beyond the scope of this section.)

Exercise 7.2

1. Consider the appropriate matrix from the last exercise and perform *LU Decomposition*. Check your answer by confirming that $\mathbf{A} = \mathbf{LU}$. (Please note that you perform LU decomposition on the original matrix \mathbf{A} , not the augmented matrix.)

Determinant

The basic algorithm for computing a determinant of \mathbf{A} is to first perform LU decomposition, and make use of the following property:

The determinant of an upper-triangular or lower-triangular matrix is just the product of the diagonal entries.

We already met another property of determinants, namely that the determinant of a product is just the product of the determinants. Therefore, $\det(\mathbf{A}) = \det(\mathbf{L})\det(\mathbf{U})$, each of which is just the product of the diagonal entries.

Exercise 7.3

1. Consider the appropriate matrix from the last exercise and find the determinant using the LU decomposition previously determined. Check your answer using `det` in MATLAB.

Inverse

The basic algorithm for computing the inverse of \mathbf{A} is to first perform LU decomposition, and make use of the following idea. \mathbf{B} is the inverse of \mathbf{A} if it satisfies the following property

$$\mathbf{AB} = \mathbf{I}$$

The columns of \mathbf{B} are just the solutions of a LSAE with a different \mathbf{b} . For example, in the two by two case we can solve

$$\mathbf{Ax} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

and then

$$\mathbf{Ax} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

and if we fill the columns of \mathbf{B} with the solution to these LSAE we will have constructed the inverse. Since we already have the LU decomposition of \mathbf{A} we simply solve each case using the technique already presented.

For example, the first column of \mathbf{B} is determined as follows: First we solve $\mathbf{Ly} = \mathbf{b}$

$$\begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

to give $y_1 = 1$ and $y_2 = -2$. Now we solve $\mathbf{Ux} = \mathbf{y}$

$$\begin{bmatrix} 2 & 3 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ -2 \end{bmatrix}$$

and the solution for \mathbf{x} is $x_1 = 3/2$, $x_2 = -2/3$. This is the entries in the first column of the inverse.

Repeating this process for $\mathbf{b} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ will give the second column of the inverse which now reads

$$\mathbf{A}^{-1} = \begin{bmatrix} 3/2 & -1/2 \\ -2/3 & 1/3 \end{bmatrix}$$

Exercise 7.4

1. Consider the appropriate matrix from the previous exercise and find the inverse using the LU

decomposition previously determined. Check your answer using *inv* in MATLAB.

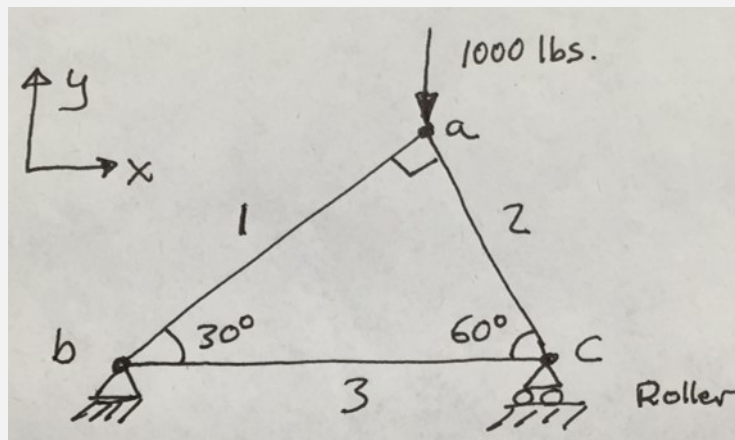
7.4 Applications of LSAE

Choose at least one of the following problems involving linear systems of algebraic equations.

Truss Analysis

Exercise 7.5

Systems of linear equations often come up in engineering when evaluating the strength and stability of structures under load. A *truss* is a simplified model of a structure. It consists of a collection of straight, rigid elements or sections that are long compared to the dimensions of their cross-section. Sections are connected only at their ends through frictionless, pin joints (Remember them? They can only constrain translation but not rotation, i.e., they can only apply force but not moments to a section). This means that sections of a truss are either in tension or compression (axial forces along its length). The roller can be assumed to be frictionless, and thus only exerts normal force. In analyzing trusses it is often assumed that the weight of the sections (dead load) is relatively small, and can therefore be neglected. The method of joints is a classic technique for determining the forces acting on all of the sections of a truss that is in static equilibrium. Here are the steps:




1. Draw a free body diagram for every pin in the truss. Note that the forces acting at the pin have to be in the directions implied by the things the pin is attached to!
2. Write out the equations of static equilibrium, $\sum \vec{F} = 0$, for every one of the pins. Note that some of your forces will be known forces (e.g., external loads), and some will be unknown reaction forces.

3. Express these equations in the matrix form $\mathbf{Ax} = \mathbf{b}$.
4. Evaluate whether the system is statically determinate or not. Note the connection to types of solutions to linear equations here: if you look at the form of \mathbf{A} , you should be able to tell whether the system is statically determinate!
5. Find the solution.

Circuit Analysis

Exercise 7.6

Systems of linear equations naturally arise in circuit analysis, although very few courses on circuits use these anymore. They do, however, form the backbone of circuit design software tools. You've met the relevant physical ideas/models before which are based on Kirchoff's circuit laws: the sum of currents into any node must be zero, and the sum of voltages around any loop must be zero. For the circuit shown in the figure:




1. Set up the linear system of algebraic equations required to solve for the three unknown currents (assuming that the resistors and the voltage sources are known.)
2. Find the solution if all of the resistors are 1 ohm, v_1 is 5 volts, and v_2 is -6 volts.

Chemical Analysis

Exercise 7.7

The complete combustion of propane, C_3H_8 , with oxygen, O_2 yields carbon dioxide, CO_2 , and water, H_2O . Based on conservation of mass, this reaction can be written as



Determine the coefficients in the combustion equation. Note that you will need to learn how to "balance" a chemical reaction.

7.5 Concept Map for Eigenfaces

For the facial recognition project we will be primarily focusing on an early facial recognition software algorithm, Eigenfaces, which is still used for face detection, and introduces some other concepts that are extremely important in both facial recognition and other tasks.

We would like you to spend some time developing an understanding of what you know, and what you don't know about facial recognition using Eigenfaces. A good way to do this is to break down the concept until you get to the point that you have terms that you *do* know:

1. Write the key term at the top or in the center. Circle it, since you don't know it.
2. Research it, and identify terms that are immediately associated with it. Write them down and connect them.
3. Circle new terms you don't understand, and break these down too.

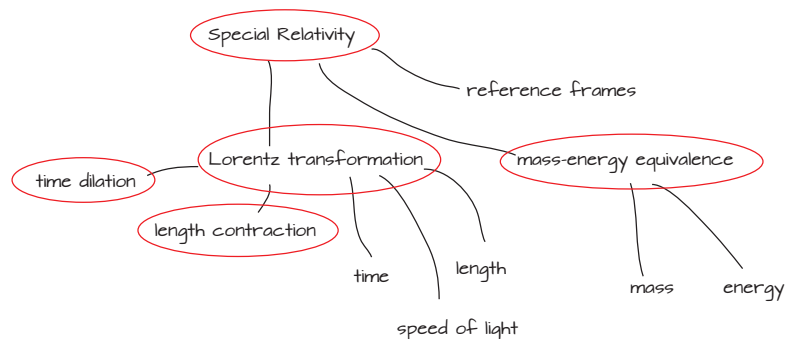


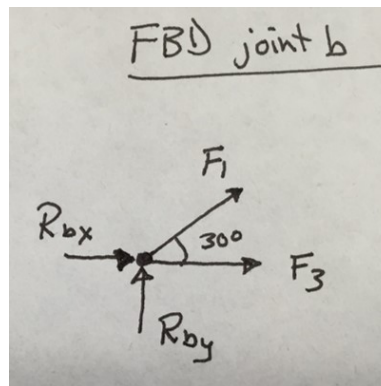
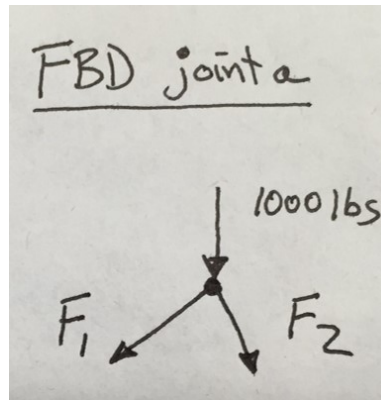
Figure 7.1: If you were trying to break down special relativity, a *portion* of your breakdown might look like this...

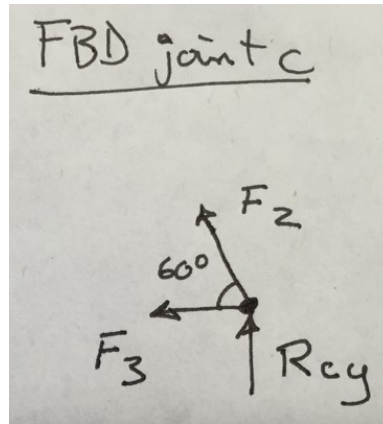
Once you've done your breakdown, try to make the following lists individually:

1. Relevant fundamental mathematical terms that I don't know
2. Relevant fundamental mathematical terms that I do know
3. Ideas specific to facial or image recognition that I don't know
4. Ideas specific to facial or image recognition that I do know

Solution 7.5

1. Note that the pin joint, b, cannot move in space so the ground must apply unknown reaction forces in both the x and y directions. The pin joint, c, is attached to a roller so it is free to move in the x direction but cannot move in the y direction. Therefore, the ground only applies a reaction force (unknown) in the y direction. For the entire truss, there is one known applied external force (1000 lbs) and six unknown axial and reaction forces (F_1 , F_2 , F_3 , R_{bx} , R_{by} , and R_{cy}).





2. For each joint, $\sum F_x = 0$ and $\sum F_y = 0$. Thus we have the following six equations:

$$-F_1 \cos 30 + F_2 \cos 60 = 0$$

$$-F_1 \sin 30 - F_2 \sin 60 = 1000$$

$$R_{bx} + F_3 + F_1 \cos 30 = 0$$

$$R_{by} + F_1 \sin 30 = 0$$

$$-F_3 - F_2 \cos 60 = 0$$

$$R_{cy} + F_2 \sin 60 = 0$$

3. These equations can be written as $\mathbf{Ax} = \mathbf{b}$ where

$$\mathbf{A} = \begin{bmatrix} -\cos 30 & \cos 60 & 0 & 0 & 0 & 0 \\ -\sin 30 & -\sin 60 & 0 & 0 & 0 & 0 \\ \cos 30 & 0 & 1 & 1 & 0 & 0 \\ \sin 30 & 0 & 0 & 0 & 1 & 0 \\ 0 & -\cos 60 & -1 & 0 & 0 & 0 \\ 0 & \sin 60 & 0 & 0 & 0 & 1 \end{bmatrix},$$

$$\mathbf{x} = \begin{bmatrix} F_1 \\ F_2 \\ F_3 \\ R_{bx} \\ R_{by} \\ R_{cy} \end{bmatrix},$$

$$\mathbf{b} = \begin{bmatrix} 0 \\ 1000 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix},$$

4. We have 6 unknowns, and 6 equations, so this is a determinate situation.

5.

$$\mathbf{x} = \begin{bmatrix} -500 \\ -866.03 \\ 433.01 \\ -5.6843e^{-14} \\ 250 \\ 750 \end{bmatrix}$$

Section 3 is the only one in tension (F_3 is positive). 1 and 2 should be in compression, and it makes sense that 2 has more compression than 1 because they support the same load, but 2 is more vertical. The horizontal reaction force is 0 because there is no net horizontal force on the system, and the sum of the vertical reaction forces is 1000 *lbf* as we expect.

Solution 7.6

1.

$$\mathbf{A} = \begin{bmatrix} -R_1 & -R_2 & 0 & 0 & 0 \\ 0 & R_2 & -R_3 & -R_4 & 0 \\ 0 & 0 & 0 & R_4 & R_5 \\ 1 & -1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 & 1 \end{bmatrix},$$

$$\mathbf{x} = \begin{bmatrix} i_1 \\ i_2 \\ i_3 \\ i_4 \\ i_5 \end{bmatrix},$$

$$\mathbf{b} = \begin{bmatrix} -V_1 \\ 0 \\ V_2 \\ 0 \\ 0 \end{bmatrix}$$

2.

$$\mathbf{x} = \begin{bmatrix} 3.8750 \\ 1.1250 \\ 2.7500 \\ -1.6250 \\ -4.3750 \end{bmatrix}$$

Solution 7.7

