21-241 MATRICES AND LINEAR TRANSFORMATIONS SUMMER 1 2012 COURSE NOTES DAY 1

PAUL MCKENNEY

1. Sets, logic, and proofs

21-241 is a mathematics course, so inevitably we will make use of mathematical concepts introduced in the course 21-127, Concepts of Mathematics. Since this is not an official prerequisite for the course, I will try to keep this to a minimum. If you're confused about one of these topics, you can come to me during office hours to talk about it; I also invite you to consult [?]. I've included here a table defining some of the logical notation that we'll be using. Below, S is a set, and P and Q are logical statements.

Notation	Definition	Notes
$x \in S$	"x is a member of S"	Here x is a *free variable*.
$\{x \in S \mid P(x)\}$	"the set of all x in S such that $P(x)$"	
$\mathbb{N}$	$\{0, 1, 2, \ldots\}$	The *natural numbers*. Note that $0 \in \mathbb{N}$!
$\mathbb{Z}$	$\{\ldots, -1, 0, 1, 2, \ldots\}$	The *integers*.
$\mathbb{Q}$	$\{p/q \mid p, q \in \mathbb{Z} \wedge q \neq 0\}$	The *rationals*.
$\mathbb{R}$	the set of all real numbers	We can't provide a true definition here..
$\mathbb{C}$	the set of all complex numbers	
$\neg P$	"not P"	
$P \wedge Q$	"P and Q"	
$P \vee Q$	"P or Q"	This is an *inclusive* or, ie, if *both* P and Q are true, then we still consider $P \vee Q$ to be true.
$P \implies Q$	"P implies Q"	$P \implies Q$ is considered a statement of its own, with its own truth value, which I invite you to deduce from the truth values of P and Q. The case when P is false is particularly interesting.
$P \iff Q$	"P if and only if Q"	This is exactly the same as $(P \implies Q) \wedge (Q \implies P)$.
$\forall x\ P(x)$	"for all x, $P(x)$"	Here and below x is a *bound variable*.
$\exists x\ P(x)$	"there exists an x such that $P(x)$"	

2. Solving systems of linear equations

Consider the following chemistry problem (which I have lifted directly from Hefferon's book [**?**]). I've got x molecules of toluene, C_7H_8, and y molecules of nitric oxide, HNO_3. Putting them together, I can produce trinitrotoluene (TNT), which has the form $C_7H_5O_6N_3$, with some water (H_2O) byproduct. Say I produce z molecules of the first and w of the second. I'd like to balance this equation;

$$xC_7H_8 + yHNO_3 \to zC_7H_5O_6N_3 + wH_2O$$

That is, I'd like to find values of $x, y, z,$ and w so that the number of atoms of each type is the same before and after the reaction. Counting them each individually, this leads to the following system of equations;

$$7x = 7z \tag{1}$$
$$8x + 1y = 5z + 2w \tag{2}$$
$$1y = 3z \tag{3}$$
$$3y = 6z + 1w \tag{4}$$

Of course, we're only interested in nonnegative integer values of x, y, z, w in this situation (since these represent the number of molecules we're dealing with), so we'll have to watch out for that when we're searching for solutions. Equation (1) clearly implies $x = z$, and so subtracting equation (3) from equation (2), we get

$$8z = 2z + 2w \implies 6z = 2w \implies w = 3z$$

So if (x, y, z, w) is any solution, then $x = z$, $y = 3z$, and $w = 3z$. Moreover, it's easy to check that no matter what choice I make for the value of z, then letting $x = z$ and $y, w = 3z$, I get a solution to the above system, and if z is a nonnegative integer then so are x, y, and w. So I've found a solution; in fact, infinitely many!

In general, there is a straightforward process which will tell us all of the solutions to a given system of linear equations, if there are any; and that there aren't any, if there aren't any. This is called *Gaussian elimination*. To describe it I'll need to formally define some of the concepts we've been using informally up to now.

Definition 1. Let $x_1, \ldots, x_n$ be variables. A *linear combination* of $x_1, \ldots, x_n$ is an expression of the form $a_1x_1 + \cdots + a_nx_n$, where $a_1, \ldots, a_n \in \mathbb{R}$. $a_1, \ldots, a_n$ are called the *coefficients* of the combination. A *linear equation* in variables $x_1, \ldots, x_n$ is an equation of the form $a_1x_1 + \cdots a_nx_n = b$, where $b \in \mathbb{R}$. b is called the *constant* of the equation.

A *system of linear equations* is simply a finite set of linear equations in the same variables, e.g.

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2 \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m \end{aligned}$$

A tuple $(s_1, \ldots, s_n) \in \mathbb{R}^n$ is a *solution* to this system of equations if the equations obtained by replacing x_i by s_i for each i, in each equation, results in unanimously true statements. If S is a system of linear equations then we write $\mathrm{sol}(S)$ for its set of solutions. So $\mathrm{sol}(S) \subseteq \mathbb{R}^n$.

When a system of equations is written like the above, we often refer to the ith equation down as *row* i, or ρ_i. Of course, the order of the equations does not matter when considering solutions; it only matters to our written system. The first term $a_{ij}x_j$ in row i such that $a_{ij} \neq 0$ is called the *leading term* of that row, and a_{ij} the *leading coefficient.* Note that a row may not have a leading term, e.g., if all of the coefficients are zero;

$$0x_1 + 0x_2 + \cdots + 0x_n = b$$

A system is in *echelon form* if the leading term of each row (except the first) is strictly to the right of all the leading terms of the rows above it, and all of the rows without a leading term are below the ones that are. Visually;

$$\begin{aligned} a_{1j_1}x_{j_1} + \cdots + a_{1j_2}x_{j_2} + \cdots + a_{1j_i}x_{j_i} + \cdots + a_{1n}x_n &= \quad b_1 \\ a_{2j_2}x_{j_2} + \cdots + a_{2j_i}x_{j_i} + \cdots + a_{2n}x_n &= \quad b_2 \\ &\quad\ \vdots \\ a_{ij_i}x_{j_i} + \cdots + a_{in}x_n &= \quad b_i \\ 0 &= b_{i+1} \\ &\quad\ \vdots \\ 0 &= \quad b_m \end{aligned}$$

Here $j_1, j_2, \ldots$, etc are the column indices of the leading terms in rows $1, 2, \ldots$, respectively, and row i is the last with a leading term.

For instance, the following is in echelon form;

$$\begin{aligned} x - 2y + z &= 0 \\ z &= 2 \end{aligned}$$

However, the following system is not in echelon form.

$$\begin{aligned} y + 3z &= 1 \\ -x - \quad y - \quad z &= 0 \\ 5y - 2z &= -1 \end{aligned}$$

The object of Gaussian elimination is to produce, through a series of operations on a given system of linear equations, a final system which is in echelon form. There are three different operations that we use;

- swap: Swapping two of the equations, written $\rho_i \leftrightarrow \rho_j$ (where $i \neq j$).
- scaling: Multiplying some row by a nonzero number; $\rho_i \to \lambda\rho_i$, where $\lambda \neq 0$.
- row combination: Replacing some row by the sum of itself with a multiple of another row; $\rho_i \to \rho_i + \lambda\rho_j$. Here λ may be zero but i and j must be different.

There is a systematic way of applying these operations to get a system in echelon form, but first we need to know that they won't change the solution set we're after. This is proven in the following theorem.

Theorem 1. *If S is a system of linear equations and T is the result of applying one of the above operations to S, then S and T have exactly the same set of solutions, ie,* $\mathrm{sol}(S) = \mathrm{sol}(T)$.

Proof. For clarity, let's say S is the following system of linear equations;

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2 \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m \end{aligned}$$

To prove the theorem it suffices to show;

(1) If $r \in \mathbb{R}^n$ is a solution to S then r is also a solution to T.
(2) There is a row-operation which, when applied to T, produces S.

(1) proves that $\mathrm{sol}(S) \subseteq \mathrm{sol}(T)$; then, (2) and (1) prove together that $\mathrm{sol}(T) \subseteq \mathrm{sol}(S)$.

Let's prove (1) first. Let $(s_1, \ldots, s_n) \in \mathbb{R}^n$ be a solution to S. This means we have

$$\begin{aligned} a_{11}s_1 + a_{12}s_2 + \cdots + a_{1n}s_n &= b_1 \\ &\vdots \\ a_{m1}s_1 + a_{m2}s_2 + \cdots + a_{mn}s_n &= b_m \end{aligned}$$

To fully prove (1) we'd have to handle three cases, according to which type of row operation we applied to S to get T; I'll just do the row-combination case, since that's the hardest. So for some $i \neq j$ and some $\lambda \in \mathbb{R}$, T is the system

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ &\vdots \\ (a_{i1}x_1 + \cdots a_{in}x_n) + \lambda(a_{j1}x_1 + \cdots a_{jn}x_n) &= b_i + \lambda b_j \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m \end{aligned}$$

Since $(s_1, \ldots, s_n)$ is a solution to S, we have $a_{i1}s_1 + \cdots a_{in}s_n = b_i$ and $a_{j1}s_1 + \cdots a_{jn}s_n = b_j$, so

$$(a_{i1}s_1 + \cdots a_{in}s_n) + \lambda(a_{j1}s_1 + \cdots a_{jn}s_n) = b_i + \lambda b_j$$

The other equations in T are the same as those in S. Hence $(s_1, \ldots, s_n)$ is a solution to T.

Now we prove (2). This comes down to another proof by cases, based on which type of row operation we applied to get T from S. I'll write the row operation down along with its reverse.

- To reverse a swap $\rho_i \leftrightarrow \rho_j$, we just apply the same swap again; $\rho_i \leftrightarrow \rho_j$.
- To reverse a row combination $\rho_i \to \rho_i + \lambda\rho_j$ we apply the row combination $\rho_i \to \rho_i - \lambda\rho_j$. (It's implicit that $i \neq j$, here. Why is this important?)
- To reverse a scaling operation $\rho_i \to \lambda\rho_i$, we apply $\rho_i \to \frac{1}{\lambda}\rho_i$. (Here it's implicit that $\lambda \neq 0$.)

This completes the proof of the theorem.

$\square$

Now I'll describe the *Gaussian elimination* algorithm. Let S be the following system.

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2 \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m \end{aligned}$$

Perform the following operations for each i from 1 to m in that order, or until I tell you to stop.

- Find the row, below row i, with the leftmost leading term among those rows. If there is no such thing (ie, if $a_{jk} = 0$ for all $j \geq i$ and $1 \leq k \leq m$), stop. Otherwise, swap this row with row i.
- Say the leading coefficient in row i is in column j. For each $k > i$, perform the row combination $\rho_k \to \rho_k - (a_{kj}/a_{ij})\rho_i$.
- Repeat.

After the first i steps of this algorithm, we've ensured that the leading terms of the first i rows go from left to right (this is easily proven by induction on i). If we go through every row, then we've ensured that the leading terms are ordered this way throughout the whole matrix. If we stop at row i, then none of the rows below have leading terms. Either way, after finishing, the system is in echelon form.

So what do we do with this? First, let's consider the equations without leading terms. They look something like this;

$$0x_1 + 0x_2 + \cdots + 0x_n = b$$

This equation is either true or false outright, depending on whether $b = 0$ or $b \neq 0$. If $b \neq 0$ then we call the equation *inconsistent*. If, after performing Gaussian elimination to bring a system into echelon form, we find an inconsistent equation at the bottom, then we've determined that it has no solutions, since no system with an inconsistent equation can have a solution. But what if there are no inconsistent equations at the bottom? Do we have solutions in this case?

Back-substitution is the algorithm which shows that indeed, the answer is yes. To perform back-substitution we perform the following operations for each i, starting with the largest i such that row i has a leading variable, and going down to 1.

- Solve equation i for the leading variable, in terms of all the other variables in equation i.
- In your formula for the leading variable, for each $j > i$, replace the leading variable of row j with formula you found for it previously.

After this process ends, you've found a formula for each leading variable in the system, in terms of only the free variables. Any choice of values for the free variables then results in a solution to the given system of linear equations.