308 PROOF EXAMPLES

JOSH SWANSON

Abstract. This document contains some example proofs. It attempts to provide examples which (1) allow you to separate "the idea" behind a proof from the proof itself and (2) allow you to coherently explain that idea rigorously. Most propositions are significantly more advanced and lengthy than you would be asked to prove on an exam.

1. Chapter 1: Linear equations, echelon forms

Theorem 1. *A linear system with (at least) two solutions has infinitely many solutions.*

Idea. Several possibilities.

(1) Geometrically, a linear system is the intersection of a bunch of hyperplanes, which is itself either empty or a hyperplane, and hyperplanes have 1 or infinitely many points. This would be hard to formalize at this stage, so let's find another route.
(2) Taking the average of two solutions gives another solution. Weighted averages also work.
(3) Consider echelon form systems. If they're consistent, the number of solutions is governed by the number of free variables—infinitely many if there's more than one free variable, and exactly 1 if there's no free variables.

□

Proof. Approach (3) is fleshed out in Holt as Theorem 1.2. It's quite clean and brief, but it uses the machinery of Gaussian elimination, so it might be considered overcomplicated. We'll use approach (2), which has no prerequisites.

Suppose $(x_1, \dots, x_n)$ and $(y_1, \dots, y_n)$ are distinct solutions of a linear system. Using our standard notation, the ith row for these two solutions reads

$$a_{i1}x_1 + \cdots + a_{in}x_n = b_i$$

and

$$a_{i1}y_1 + \cdots + a_{in}y_n = b_i.$$

We claim that $(tx_1 + (1-t)y_1, \dots, tx_n + (1-t)y_n)$ is also a solution of the system, for all scalars t. Indeed, if we multiply the first equation by t, multiply the second equation by $1-t$, and add the two equations, we get

$$(ta_{i1}x_1 + \cdots + ta_{in}x_n) + ((1-t)a_{i1}y_1 + \cdots + (1-t)a_{in}x_n) = tb_i + (1-t)b_i.$$

Date: April 6, 2015.

The right-hand side is just b_i, and the left-hand side is

$$a_{i1}(tx_1 + (1-t)y_1) + \cdots + a_{in}(tx_n + (1-t)y_n) = b_i.$$

This says exactly that our proposed solution $(tx_1 + (1-t)y_1, \ldots, tx_n + (1-t)y_n)$ is in fact a solution.

As we let t vary, we get infinitely many solutions—why? Since $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_n)$ are distinct, they must differ in some coordinate, say $x_j \neq y_j$. Then $tx_j + (1-t)y_j$ gives infinitely many distinct values for the jth coordinate of our solutions, so there are infinitely many solutions. □

Comments. This argument is much briefer in matrix notation: the heart of it is

$$\begin{aligned} A(t\mathbf{x} + (1-t)\mathbf{y}) &= t(A\mathbf{x}) + (1-t)(A\mathbf{y}) \\ &= t\mathbf{b} + (1-t)\mathbf{b} \\ &= \mathbf{b}. \end{aligned}$$

□

Corollary 2. *A linear system has either* 0*,* 1*, or infinitely many solutions.*

Idea. Intuitively rather clear from the theorem. □

Proof. We do a proof by cases. Either there are 0 solutions or there are not 0 solutions. If there are 0 solutions, great, we're done. If there are not 0 solutions, there are more than 0 solutions. In this case, either there is 1 solution or there is more than 1 solution. In the former case, we're again done, and in the latter case, there are at least 2 solutions. In that case, by the theorem, there are infinitely many solutions, so again, we're done. □

Comments. With practice, this becomes quite short: "either there are 0, 1, or more than one solutions, and in the last case there are infinitely many solutions by the theorem." □

Proposition 3. *A triangular system has exactly one solution.*

Idea. Back substitution allows us to successively compute the coordinates of the solution from right to left. □

Proof. A triangular system has the same number of variables and equations, and so is of the form

$$\begin{aligned} a_{11}x_1 + a_{21}x_2 + \cdots + a_{n1}x_n &= b_1 \\ a_{22}x_2 + \cdots + a_{n2}x_n &= b_2 \\ \ddots \qquad &\vdots \\ a_{nn}x_n &= b_n \end{aligned}$$

From the nth equation, we have $x_n = b_n/a_{nn}$. Note that we can divide by a_{nn} since leading terms are by assumption non-zero. The $n-1$st equation is $a_{n-1,n-1}x_{n-1} + a_{n,n-1}x_n = b_{n-1}$, which gives $x_{n-1} = (b_{n-1} - a_{n,n-1}x_n)/a_{n-1,n-1}$, and since $x_n = b_n/a_{nn}$, this determines the value of x_{n-1} in terms of the system's constants. Continuing in this manner, the variables

$x_1, \ldots, x_n$ are each forced to have a single value, so there is at most one solution. Moreover, the nth, $n-1$st, etc. equations are each satisfied in turn by these values, so they do actually form a solution. Hence there is exactly one solution. □

Comments. A completely formal version of this argument uses "mathematical induction," though we will avoid such rigor in this class and allow somewhat informal "continue in this way" arguments. In this particular case, I stopped after the $n-1$st equation since the later cases are essentially no more complex than that case. Stopping after the nth equation makes it less clear to me that you really understand the argument, and your rough goal when writing your own proofs is to convince me you know what you're talking about. □

Proposition 4. *If S can be obtained from T by a sequence of elementary row operations, then T can be obtained from S by a sequence of elementary row operations.*

Idea. Each ERO can be undone; to get from T to S by ERO's, do the ERO's you used from S to T but in the opposite order and reversing each one. □

Proof. First suppose S is obtained from T by a single ERO. We have three cases:

(1) The ERO interchanges row i and j, i.e. $R_i \Leftrightarrow R_j$. Apply this ERO again to undo it.
(2) The ERO scales row i by $c \neq 0$, i.e. $cR_i \Rightarrow R_i$. Divide by c to undo it, i.e. apply $(1/c)R_i \Rightarrow R_i$.
(3) The ERO adds c times row j to row i, i.e. $R_i + cR_j \Rightarrow R_i$. Subtract c times row j from row i to undo it, i.e. apply $R_i - cR_j \Rightarrow R_i$.

Now suppose that S is obtained from T by a sequence of ERO's, resulting in $T = T_0, T_1, \ldots, T_k = S$. Since T_i is obtained from T_{i-1} by an ERO, the above reasoning says that T_{i-1} can be obtained from T_i by an ERO. Hence we can successively go from $S = T_k$ to T_{k-1} to $\cdots$ to T_1 to T_0 by ERO's. □

Comments. Cases (1)-(3) above could have been written more formally. We could do (3) as follows. Let S_k denote the kth row of S and T_k the kth row of T. In case (3), by definition $T_i = S_i + cS_j$ and $T_k = S_k$ for $i \neq k$. Applying the suggested operation results in a matrix R with $R_i = T_i - cT_j$ and $R_k = T_k$ for $k \neq i$. But then $R_k = T_k = S_k$ for $k \neq i$ and $R_i = T_i - cT_j = (S_i + cS_j) - cS_j = S_i$, so $R = S$. This is a very clear argument, though it's also quite lengthy, and I at least am entirely convinced by the shorter, less detailed version above. □

Theorem 5. *Suppose a linear system is in echelon form. Having chosen particular values for each non-leading variable, there is a unique solution of the system with those values for those non-leading variables.*

Moreover, in every solution, a leading variable x_j is a sum of constants (depending only on the system) times either a non-leading variable x_i with $i > j$ or some constant b_k from the right-hand side of one of the linear system's equations.

Idea. Back substitution continues to allow us to successively compute the coordinates of the solution from right to left, just with a bit more freedom coming from non-leading variables because of the "stair step" pattern of echelon form. □

Proof. Say there are n variables. Suppose throughout that we've fixed values for the non-leading variables. We must show that there are unique values for the leading variables which yield a solution of the system. We'll compute them right to left.

If x_n is a leading variable, it must be so in the bottom-most equation, which is then of the form $a_{mn}x_n = b_m$ for $a_{mn} \neq 0$. Since $a_{mn} \neq 0$, this forces x_n to have a unique value, namely b_m/a_{mn}. On the other hand, if x_n is not a leading variable, we've already fixed its value. In either case, there is a unique value of x_n.

Now suppose we've found unique values of the rightmost p variables $x_n, x_{n-1}, \ldots, x_{n-p+1}$ which satisfy the equations whose leading variable is one of $x_n, \ldots, x_{n-p+1}$. If x_p is not a leading variable, we've again already fixed its value, and we can safely say the rightmost $p+1$ variables $x_n, \ldots, x_{n-p}$ have unique values which satisfy the equations whose leading variable is one of $x_n, \ldots, x_{n-p}$. On the other hand, if x_p is a leading variable, then it's in a row of the form

$$a_{i,p}x_p + a_{i,p+1}x_{p+1} + \cdots + a_{i,n}x_n = b_i. \tag{1}$$

Since $x_{p+1}, \ldots, x_n$ have already been fixed and $a_{i,p} \neq 0$, we can solve equation (1) for x_p uniquely, and again we can extend to the $p+1$ case. Contining in this manner, eventually all variables will have been processed, at which point there is a unique value of the n variables $x_n, \ldots, x_1$ which satisfy the equations.

In the above procedure, leading variables depended only on the value of later variables—see equation (1). The rightmost leading variable then depends only on later variables, which must be non-leading variables. The next leading variable to its left similarly depends only on later variables, including free variables and the rightmost leading variable. Replacing the rightmost leading variable with free variables, the next leading variable also depends only on free variables to its right. Continuing in this way gives the second claim. □

Comments. In the triangular system variant of this proposition, we were able to write out the procedure for just the last two rows and appeal to an ability to "continue in this manner." Here, though, the procedure is more complex and depends on the precise location of the pivots. To write a coherent explanation, we used a standard organizational method. We made a series of claims involving the rightmost p variables, going through the $p = 1$ case explicitly and then describing the procedure for going from the p case to the $p+1$ case in general. Applying it repeatedly gives us the $1, 2, 3, \ldots$ cases of our claim, which is really just an algorithmic form of induction. □

Theorem 6. *Given just the solution set of a homogeneous linear system with m equations in n variables whose augmented matrix is in reduced row echelon form, one can compute the underlying linear system.*

Idea. The non-pivot variables can be used as free variables and the pivot variables can be computed in terms of the free variables, which abstractly gives the entire solution set. To

identify the pivots from the solutions, roughly note that the free variables can be completely arbitrary as we range over all solutions, but the pivot variables cannot be so arbitrary. Compute the remaining entries by plugging in 0 for every free variable but one and examining the resulting solutions. The details of this proof are somewhat technical. □

Proof. The linear system is homogeneous, so the rightmost column of its augmented matrix is zero. We may delete this column, so every column of the matrix corresponds to a variable in the system, including the new rightmost column. If the matrix has zero rows, remove the corresponding trivial equations (all of the form $0 = 0$, so the solutions are unaffected) from the linear system in what follows. In this way we may assume the linear system is in echelon form and the preceding theorem may be applied to it.

Suppose $i_1, \ldots, i_k$ are the columns in which pivots do not occur. The variables $x_{i_1}, \cdots, x_{i_k}$ are free variables which can be chosen completely arbitrarily. That is, for any choice of values for these k variables, there is a unique solution to the system with those values in those variables, by the preceding theorem. If we choose values for only some of these free variables, there are still solutions, though they are no longer unique.

We now identify pivots and free variables by working our way right to left. If the rightmost column contains a pivot, that pivot's row is of the form $(0, \ldots, 0, 1)$, which corresponds to the equation $x_n = 0$, and in particular any solution whatsoever has nth coordinate 0. On the other hand, if the rightmost column does not contain a pivot, then x_n is a free variable, so in particular there is some solution with $x_n = 1$. Hence we can detect a pivot in the rightmost column by looking for solutions with $x_n = 1$.

Supposing we can identify the pivots among the last p variables $x_{n-p+1}, \ldots, x_n$ (i.e. in the rightmost p columns), we can determine if the next variable x_{n-p} is a pivot as follows. Set all the free variables amongst the last p variables to zero and look for solutions with $x_{n-p} = 1$. There are two cases:

- If x_{n-p} is a free variable, then there will be a solution of this form.
- Suppose x_{n-p} is a pivot. We're considering solutions where the free variables among the last p variables are 0. By the preceding theorem, a pivot variable among the last p variables is a sum of the free variables to its right, which are all 0, so all of the last p variables must be 0 in such a solution. By the same reasoning, x_{n-p} must also be zero, so there is no solution of the suggested form.

Having now identified the pivots among the last $p+1$ variables, we repeat this process until we have identified all the pivots and free variables. The number of non-zero rows is equal to the number of pivots, so we know the number of zero rows, which gives us the location (both row and column) of each pivot. All that remains is to compute the entries to the right of each pivot. For that, focus on a particular pivot variable x_i which is in a row of the form $(0, \ldots, 0, 1, a_{j,i+1}, \ldots, a_{j,n})$. The corresponding equation for this row is

$$x_i = -a_{j,i+1}x_{i+1} - \cdots - a_{j,n}x_n. \tag{2}$$

We must compute $a_{j,k}$ for $i + 1 \leq k \leq n$ in terms of the solution set. Note that if x_k is a pivot, then $a_{j,k} = 0$ since position (j, k) is above the location of the pivot x_k. Hence we can restrict our attention to k for which x_k is a free variable. Moreover, all the non-zero terms

on the right-hand side of equation (2) involve free variables. Now consider solutions where all the (free) variables on the right-hand side of (2) are zero except for x_k, which is 1. The equation is then of the form $x_i = -a_{j,k}$, so we can read off $a_{j,k}$ from the ith coordinate of such a solution. □

Comments. I've repeatedly made claims without giving all the details. These claims are meant to be very specific and checkable. Some examples:

- "Note that if x_k is a pivot, then $a_{j,k} = 0$ since position (j, k) is above the location of the pivot x_k."
- "the number of zero rows ... gives us the location (both row and column) of each pivot"

In the first case, justifying this "geometrically obvious" statement would take a line or two and would mostly be a distraction. In the second case, I don't actually say how to compute the pivot locations from the given data. In both cases the key point is that I could tell you the missing details if you asked me, but I've decided they are not worth writing down. Your proofs will probably be heavier on details than mine, but sometimes a lack of detail actually makes for a clearer proof. For instance, sometimes several algebraic steps can be combined into one without sacrificing clarity.

(I've included this theorem despite the proof's relative complexity since Holt does not include a proof and I was not happy with any of the proofs I found in other sources. The argument is my own, though it can't possibly be new. None of the proofs I found worked with the underlying solution set, which is what I was after.) □

Corollary 7. *Every matrix is equivalent (i.e. obtainable by a sequence of elementary row operations) to a unique matrix in reduced row echelon form.*

Idea. The solutions of the underlying linear system determine the reduced row echelon form, these solutions are preserved by ERO's, and they determine the linear systems by the theorem. □

Proof. First apply Gauss–Jordan elimination to obtain a reduced row echelon form matrix M of the suggested form. Suppose some sequence of elementary row operations results in another reduced row echelon form matrix M'. We must show $M = M'$.

Thinking of matrices M and M' as homogeneous linear systems, their solutions are preserved by elementary row operations, so M and M' have the same solution set. By the theorem, the linear systems are equal, so $M = M'$. □

Remark 8. Theorem 6 is actually rather powerful. It can be used to prove the following statements, which we'll get to in more detail later:

- Every matrix has a pair of numbers associated to it called its rank and its nullity. In the present context, the nullity is the number of free variables coming from the reduced row echelon form of the matrix and the rank is the number of pivots. Hence the number of variables is equal to the rank plus the nullity.

- Given a homogeneous system, you can scale its equations and add them to get new equations which still must be zero for any solution. Consider the set of all equations obtainable in this way for a particular homogeneous system coming from a matrix M. This set is (essentially) the row space of M. Now also consider the set of solutions of the homogeneous system, which is called the null space of M. In fact, two matrices have the same null space if and only if they have the same row space.
- There is a "duality" between solutions of homogeneous systems and the equations of homogeneous systems themselves. The act of computing the solutions of a homogeneous system actually computes the "orthogonal complement" to the row space. This can be reversed: one can obtain a system of equations from the orthogonal complement whose solutions are precisely the row space of the original system.
- Reduced row echelon form matrices "parameterize" both the space of solution sets of systems of equations and the space of row spaces of matrices. This insight is used in the "Schubert calculus" to give geometric structure to sets of hyperplanes called "Schubert varieties." This is the starting point of an old but very active area of mathematical research.
- (Everything said so far also works over arbitrary fields. In particular, the duality above gives the usual symmetry of q-binomial coefficients. The one strange thing is that the row space and its orthogonal complement no longer have trivial intersection.)