Math 642:550 — Summer 2009
MTTh 6:00–8:30 PM Hill 425
Prof. Bumby

Supplement 1, Prerequisites

1. Introduction

A motivation for the subject of **Linear Algebra** is to organize the techniques for solving systems of linear equations. Since you should have had a course in which this was done, it suffices to show how some simple examples fit into the theory.

1.1 One equation in three unknowns

What are **all** solutions of

$$x - 2y + 5z = 17? \tag{Q}$$

The equation can be written as

$$x = 17 + 2y - 5z, \tag{A}$$

which can be interpreted as: **allow** y and z to **take any values** and use (A) to find a value of x. **Every such assignment** of values to (x, y, z) satisfies (Q), and **every solution** of (Q) can be obtained in this way.

1.2 An extension

Suppose that we learn that equation (Q) was **expected to contain** a variable w. Then the solutions (w, x, y, z) are **still** given by (A), although **we must now allow** y, z, **and** w to take arbitrary values.

Although the **method** of solution is **the same**, the **answers** to these two questions are **different**. This serves as a reminder that a complete statement of the equation (Q) **must contain** a list of the variables that are allowed to appear in the equation. A solution will give values to **all** of these variables in terms of some **parameters**. The method of solution shows that these parameters can be chosen from the variables **declared** in the problem (even if they are not **seen** in any equation).

In order to have any theoretical results, **the variables must be declared in advance** and **only the declared variables will be allowed to appear in equations**. However, we do not demand that every declared variable be present in the equations. We will see that the method of solution has such an easy way of dealing with this possibility that we should be ashamed of feeling that it was a matter for concern.

We make the convention that the left side of each equation is a combination of the variables and the right side is a constant.

1.3 Two equations

Suppose that our original equation (Q) is joined by the equation

$$w + 3y - 2z = -5 \tag{Q'}$$

to form a system of two equations. Solving (Q) for x, as before, and solving (Q') for w, gives expressions for all solutions w, x, y, z in terms of y and z.

This is too easy. Each equation has one special variable that appears nowhere else, and we simply use that equation to determine the special variable. The other variables are not restricted in any way: they **parameterize** the solutions.

We are more likely to meet something like

$$\begin{aligned} w + x + y + 3z &= 12 \\ 7w + 2x + 17y - 4z &= -1 \end{aligned}$$

If we just had the first equation, we would solve it for w in terms of x, y, z. This solution is still valid when we meet the second equation, so we can **substitute** $12 - x - y - 3z$ for w in the second equation to find out what the second equation says about the parameters x, y, z in the solution of the first equation. It turns out to be (Q)! This allows us to use $17 + 2y - 5z$ in place of x in our expression for w to get the same expression for w in terms of the **free variables** y and z that was found by solving (Q'). This system has the same solution as (Q) and (Q').

While the above solution is correct, **it is never used in practice**. While it retains the variables that are valuable in **interpreting the answer**, it requires the use of those variables to identify like terms, and the process of shifting terms from one side of an equation to the other introduces changes of sign. This approach to algebra is unsuitable for humans because it is **tedious** and **difficult to do accurately**, and unsuitable for computers because the terms must be repeatedly examined to find how expressions can be simplified. Fortunately, **there is a better way**.

1.4 Elimination

The substitution step is **completely equivalent** to subtracting 7 times the first equation from the second in our system. That is, the left side of the first equation minus 7 times the left side of the second is set equal to the right side of the first equation minus 7 times the right side of the second Why 7? It leads to an expression in which w does not appear. As long as we do the same operation on both sides of our equations, we get relations between our variables w, x, y, z that are consequences of the relations in the given system. When we are done, we can **check** that all our results satisfy the original equations. However, in some cases, this is only a partial check of the accuracy of our computation since it only shows that we haven't introduced values that fail to satisfy the given system, but it cannot detect an incorrect **exclusion** of some solutions. One aim of a systematic method of solution is to assure that this can only happen because of a mistake in computation and not because of a flaw in the solution algorithm. While the computations in the **Gaussian elimination** method are usually easy to do correctly, it is so easy to check that the answers satisfy the given equations that there is no reason to accept answers without a check. The method of solution that we will obtain later actually includes a complete check that we have all solutions of a given system.

The elimination method aims to find a description of **all** solutions of a system of m equations in n unknowns. Since we have agreed to declare all variables, we have identified n unknowns whether or not they appear in an equation. Since we no longer need an equation to introduce a variable, we can allow an extreme case of a system with **no equations** in the given list of variables. The solution of such a system allows arbitrary values for **all** variables.

Here is one version of the method, organized as an induction on m. The basis is $m = 0$: a system of no equations whose solution is given by allowing each variable to be assigned an arbitrary value. The induction step uses elimination to show how to solve a system of m equations with $m > 0$ assuming that we are able to solve a system of $m - 1$ equations.

Since we have at least one equation, we **select** one, which we call the **pivot** equation. The aim is to use that equation to complete the solution after solving a modified form of the system of the remaining equations. However, there are several cases.

Consider the variables that appear with **nonzero** coefficients in that equation. On possibility is that there are no such variables. Then, the left side reduces to zero, and we look at the right side. If it is zero, the equation is the **tautology** $0 = 0$ that places no restriction on the solution. The equation may be removed from the system, leaving a system of $m - 1$ equations in n variables, completing the induction step. The second possibility is that the right side is not zero. In this case, we have proved that, **if there are any solutions to the system**, then $0 = 1$. We take this as a proof that the equation has no solutions and stop.

Otherwise, there are variables with nonzero coefficients and we **select** one, which we call the **pivot** variable. To complete the induction step, we form a system of $m - 1$ equations in $n - 1$ variables consisting of modified forms of the nonpivot equations in which the pivot variable does not appear. To do this, use

the elimination step to replace each nonpivot equation by an equation not containing the pivot variable that is equivalent to it provided that the pivot equation is satisfied. This gives the new system. Since the pivot variable has a nonzero coefficient in the pivot equation, the pivot equation may be solved for the pivot variable in terms of the other variables to complete the solution.

In practice, this method of **Gaussian elimination** proceeds in two steps. The first part is the succession of pivot steps to reduce to a system of no equations. Any remaining variables are **free variables** that may take any value. The method leaves a succession of pivot equations to be solved for the pivot variables **in the reverse order of their identification**. We work back through the list to replace references to the other pivot variables by their expressions in terms of the free variables. This is called **back substitution**, but it can be done by additional **elimination** steps. (It turns out that it is more efficient for all but very small systems to wait until the free variables have been identified before eliminating other variables from our answers.)

Although the inductive description of the method uses systems of different numbers of equations, the usual implementations use the same values of m and n throughout. The pivot equations are set aside for use in the back substitution phase, but each array that is written represents a system of equations with the same solutions. In machine computation, this allows the work to be done **in place**. For hand computation, one has a sequence of m by n arrays that give a record of the computation that can be use to localize any mistake that is detected. These arrays are only written when necessary to summarize a number of computations that are done in parallel (called a **pivot step**), although theoretical discussions often emphasize the individual **row operations**.. Thus, theoretical discussion and different forms of computation may differ in details although they share a common strategy.

In this method, the variables only serve to arrange their coefficients in columns and the order of columns is fixed throughout the solution. Similarly, for hand computation, the rows are kept in a fixed order that trace their history back to a particular equation in the original list although the other equations are used to modify the information on that row (unless it is **absolutely necessary** to change the order of the rows). The calculation can now be done with a rectangular array (or **matrix**) of coefficients.

2. Matrices

Strictly speaking, the array of coefficients that we have described is the **augmented matrix** of the system. The last column represents the right side of the equation, while the remaining columns are coefficients of particular variables. This is usually indicated by separating the last column from the others by a dotted line. This dotted line corresponds to the equals sign in the equations: columns to the left of the line are coefficients of the variables; those to the right are the constants on the right sides of the equations. Typically, one is solving a single system of equations, so there is one column to the right of the dotted line. However, there could be **any number** of columns on the right representing different right sides that will be considered in the same computation. One important special case is the **homogeneous system** in which all right sides are zero. The operations in the solution preserve a zero column, so it is customary to omit a zero column when solving a homogeneous system. This leads to systems with **no** columns in the right side of the augmented matrix, giving another example of a vacuous case. Another case that most books treat separately is the computation of the **inverse** that will be described below.

The steps in the solution of a system is determined entirely by the left side of its augmented matrix. If you want to solve several systems with the same left side, all right sides could be written as columns to the right of the dotted line in a large augmented matrix. A solution of one of these systems is found by looking in its column on the right side of the final augmented matrix. When you have solved a system with one right side, the system whose right side is a fixed multiple of that will have a solution that is the same multiple of the solution of the original system. When you have solved systems with two right sides, the system whose right side is the sum of these will have a solution that is the sum of the solutions of the separate equations. This allows the solution of a small number of systems with a common left side to determine the solution of a large set of related systems. An important special case is that the **general** solution of a system is the sum

of a single **particular** solution with the general solution of the **homogeneous** equation.

Although the computations are done entirely with arrays of numbers, the variables and equations are never far away. Multiplying the entry in a column by the variable belonging to that column in the original system and adding these terms gives an expression whose equality to the value in the last column is asserted by the system being solved. This way of combining elements of matrices is the definition of **matrix multiplication**. Given a matrix M and a column v with the **same number of entries** as M has columns (equivalently, the number of entries in a row of M) Mv is found by, for each i, adding the product of the jth entry of the ith row of M with the jth entry of v to form the ith entry of the product. This allows systems of equations to be abbreviated as a matrix equation $Mx = b$.

This can be extended to give a product of matrices A and B — **in that order**, matrix multiplication is not generally commutative — if the number of **columns** of A is **the same** as the number of **rows** of B with each column of the product being the product of A with the corresponding column of B.

In matrix multiplication, the order of terms in a product is important. If AB is defined, BA may not be. If both are defined, they may not even be the same size. Even if they are the same size, they need not be equal. However, the **associative law** holds. That is, if either of the products $A(BC)$ or $(AB)C$ make sense, then both do, and the results are equal. One proof is to simply write expressions for the elements of these product and show that they are the same, but other approaches to matrices lead to more conceptual proofs. This explains why, when X is the solution of $MX = B$, the solution of $My = Bc$ is $y = Xc$.

An n by n matrix I like

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

with entries that are zero except 1 in every place on the **main diagonal**, has the property that multiplying by I (where defined) gives a product equal to the other factor. Such a matrix is called an **identity matrix**. If X is a matrix with $AX = I$, X is said to be a (right) **inverse** of A. Thus, one method to find an inverse (or to show that one doesn't exist) is to write an augmented matrix with A on the left and I on the right and apply Gaussian reduction, which agrees with our general method for solving systems of equations applied to the equation $AX = I$. This is the typical computation of A^{-1} found in textbooks. When A has the same number of rows and columns (called a **square** matrix), an analysis of this computation shows that, whenever it exists, a right inverse is also a left inverse. In this case, an equation $Ax = b$ has a solution that can be written $x = A^{-1}b$. Although we may **write** this, it should not be taken as an instruction to find A^{-1} and multiply this matrix with b. The solution should be found using an augmented matrix with **only** b on the right side.

2.1 Row operations

Another role of matrices is to express the steps of Gaussian elimination as row operations that can be expressed by multiplying the augmented matrix of the system **on the left** by a suitable matrix. In all cases, since the operation applies to all matrices, this matrix is obtained by applying the intended operation to an identity matrix.

The elementary row operations of Gaussian reduction (including the rarely used row interchange) are represented by square matrices that have unique inverses giving row operations that **undo** that operation. Insisting that the number of rows and columns remain fixed throughout the computation allows the work to be described in terms of matrix multiplication.

In the first part of Gaussian elimination, M is reduced to a triangular (or **echelon**) form U. The role of b is to **accept these row operations** so that each step in the process gives the augmented matrix of a system $Ux = c$ with the same solutions as $Mx = b$.

The names "square", "diagonal", "triangular", and "echelon" are intended to suggest the visual appearance of a matrix. If is customary to refer to the (i, j) entry of a matrix as the number in row i (from the top)

and column j (from the left). With this notation, the **main diagonal** consists of the (i, i) entries, and only this diagonal is usually considered significant. A matrix is called **upper triangular** if all of its **nonzero** entries are on or above the main diagonal, so that everything below the diagonal is zero. That is, if $i > j$, the (i, j) element is zero. Similarly, a **lower triangular** matrix is zero above the main diagonal. In an **echelon form**, the first nonzero element of row $i + 1$ is to the right of the first nonzero element of row i. In particular, an echelon form is upper triangular, but

$$\begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$$

is upper triangular but not in echelon form.

The left multiplications expressing the steps in the first part of Gaussian elimination involve matrices in which nonzero elements can only appear on or below the main diagonal. The product of such matrices also has the same property.

2.2 The LU factorization The reduction of an m by n matrix M to echelon form, when it is possible without row interchanges, can then be summarized by a factorization $M = LU$ where L is a m by m **lower** triangular matrix whose **inverse** gives the product of the operations and U is the m by n **upper** triangular matrix that is the result of the operations. (We insist that U be in echelon form, but only the weaker property of being upper triangular is reflected in its name.) The matrix L can also be found in the record of the computation (this is described in detail in the textbook — see index entries for "pivot", "LU factorization", and "LDU factorization"). If the original matrix is augmented by a column b, the echelon form U will be augmented by a column c with $Lc = b$. Note that we have required L to be a square matrix with 1 on its main diagonal. Since L is invertible, the original system $LUx = b$ has the same solutions as $Ux = c$. The matrix U has the same shape as M. These are the conventions of this text — other conventions may be used elsewhere. Note that the entries of the matrix L describe the row operations used to transform M to U and b to c.

It is easy to see from the echelon form whether the system has solutions, and to determine **qualitative** information about the nature of the set of all solutions. Back substitution is required to **find** the solution. As we have already noted, back substitution is done by additional row operations.

If there are rows of zeros at the bottom of U, then a test for $Mx = b$ to have a solution, is that c also have zero in these positions, since such a row represents an equation whose left side is zero for all values of the variables, and $Mx = b$ is **equivalent** to $Ux = c$.

2.3 Uniqueness If no row of U consists entirely of zero, then this is the **only** LU factorization of M. We prove this by induction. The first **nonzero column** of an **echelon form** has a nonzero entry in its first position and the rest of the column zero. Similarly, the first **row** of a **lower** triangular matrix will be zero except perhaps for its first position. Since we have insisted that L have all diagonal entries equal to 1, its first row is completely determined. The first row of a product is found by multiplying this first row of L with the whole second factor U, so the first rows of U and M are identical. Likewise, a column of a product is the product of the whole first factor with the corresponding column of the second factor. If a column of U is zero, so is the corresponding column of M, and conversely. The first nonzero column of U is zero except for its first entry, so the product with L is a multiple of the first column of L. The first row and column of L, the first row of U, and the first nonzero column of U are uniquely determined by M.

After the pivot step determined by the first column of L, both M and L have been reduced to matrices in which all rows except the first have 0 in the first position. If we now remove the first row and column from L, U, and M, we have an LU factorization of the replacement for M. If we can continue until the rows of U are exhausted, every element of the factorization will have been characterized by the calculation that we used.

On the other hand, if the last row of U consists entirely of zeros, the elements in the last column of L are multiplied by zero when L and U are multiplied to give M. Since we normalized L to have 1 on its diagonal, the result is still unique. However, if **two or more** rows of U are zero, there will be entries of L that do not contribute to the product LU. The factorization will not be unique in this case.

Uniqueness can be restored by removing the rows of zeros at the bottom of U and the columns of L that multiply those quantities. A brief discussion of this construction — leading to a **reduced** LU factorization, denoted $\underline{L}\,\underline{U}$ — was in earlier editions of the textbook, but has been removed from the fourth edition. This seems to have been a step backward since we will see that the reduced factorization is the key to finding the **four fundamental subspaces** that we discuss in Section 4.1 below.

In exercises aiming at this factorization, we will emphasize that we want the reduced factorization since the significance of the underlining in this notation often escapes notice.

The number of nonzero rows in U is the **rank** of M, usually indicated by r. Thus, $\underline{L}$ is an m by r matrix and $\underline{U}$ is an r by m matrix.

2.4 Example Consider the 5 by 6 matrix

$$M = \begin{bmatrix} 2 & 0 & 3 & 5 & -1 & 2 \\ 4 & 0 & 9 & 12 & -2 & 9 \\ 6 & 0 & 6 & 12 & 1 & 4 \\ -2 & 0 & -3 & -7 & 9 & 4 \\ 10 & 0 & 6 & 20 & -9 & -8 \end{bmatrix}$$

The LU factorization of M is

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 & 0 \\ 3 & -1 & 1 & 0 & 0 \\ -1 & 0 & 2 & 1 & 0 \\ 5 & -3 & -1 & x & 1 \end{bmatrix} \begin{bmatrix} 2 & 0 & 3 & 5 & -1 & 2 \\ 0 & 0 & 3 & 2 & 0 & 5 \\ 0 & 0 & 0 & -1 & 4 & 3 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

You are invited to work through the eliminations steps to derive this result and note that the $(5, 4)$ element of L, denoted by x here, is not determined by the process. You may also check the result by multiplying the factors, which will show that the product is M for any value of x.

Since the last two rows of U are zero, the rank is 3 and the reduced $\underline{L}\,\underline{U}$ factorization is

$$\begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & -1 & 1 \\ -1 & 0 & 2 \\ 5 & -3 & -1 \end{bmatrix} \begin{bmatrix} 2 & 0 & 3 & 5 & -1 & 2 \\ 0 & 0 & 3 & 2 & 0 & 5 \\ 0 & 0 & 0 & -1 & 4 & 3 \end{bmatrix}$$

3. Row interchanges

The process for finding an LU factorization by Gaussian elimination doesn't always work. In the inductive description of the process in the previous section, we may be led to a system in which the first variable doesn't appear in the first equation. If it appears in no **nonpivot** equation, we may skip this column when finding the echelon form. However, if it appears in **some** equation, it must be made a pivot variable by selecting a different equation to solve for the first variable. The simplest way to do this is to exchange two equations in our list. This has no effect on the solution of the system, but at the matrix level, this means interchanging two rows in the matrix.

This difficulty only arises rarely, so its role in solving systems by hand can be treated in an *ad hoc* manner, although a systematic treatment, using theoretical considerations to be given later in this course, can be given. However, in numerical work, there are often reasons to perform row interchanges to assure that the calculation can be performed accurately, so it must be made part of the method. This appears in Chapter 7 of the textbook and will also be treated later in the course.

For now, we assume that we can find an LU factorization.

4. Vector spaces

Having introduced matrices to facilitate the solution of linear systems, we are led to look at intrinsic properties of matrices that go beyond that application. This leads to the abstract idea of a **vector space**. A key example is the collection of all **columns** of n real numbers, which is denoted $\mathbb{R}^n$. It has an operation of **addition** defined by putting the sum of corresponding entries in that place of the sum. It is also possible to multiply by a number (called a **scalar** in this context) by multiplying each entry by that scalar. These operations have the expected algebraic properties.

Any collection of objects with operations of addition and multiplication by scalars satisfying the usual algebraic properties is called a **vector space**. An introductory course on Linear Algebra explores the consequences of this definition in order to apply it to a large number of examples. One property that is easy to overlook is that a vector space **must contain a zero vector**, so the simplest vector space is one that consists only of a zero vector.

We have seen that $v \in \mathbb{R}^n$ can be multiplied by an m by n (the first number always counts the rows and the second counts the columns) matrix M to give $Mv \in \mathbb{R}^m$. Another interpretation of this product is that each column of M is multiplied by a scalar taken from the corresponding position in v and the results are added together. This gives a **linear combination** of the columns of M. The collection of all such linear combinations is a vector space called the **column space** of M.

In this language, $Mx = b$ has a solution if and only if b lies in column space of M.

Every **subspace** of $\mathbb{R}^m$, i.e., a sub**set** of $\mathbb{R}^m$ that is a vector space with the operations of addition and multiplication by scalars inherited from $\mathbb{R}^m$, can be realized as the column space of some matrix with m rows.

The set of all linear combinations of a set $\mathcal{S}$ of vectors is a vector space called the **span** of $\mathcal{S}$. If two different linear combinations have the same sum, their difference is a nontrivial representation of the zero vector. A set allowing a nontrivial linear combination whose sum is the zero vector is called **linearly dependent**. In particular, a set containing the zero vector is always linearly dependent. Otherwise the set is said to be **linearly independent**. If a set is linearly dependent, a dependence relation can always be used to write one vector in the set as a linear combination of the others (including as a special case the empty expression for the zero vector). This expression can be used in place the vector, so the vector can be removed without changing the span. A linearly independent spanning set of vector space is called a **basis** of the space. It can be shown that all bases of a space have the same number of elements. The number of elements in a basis is called the **dimension** of the space.

4.1 The four fundamental subspaces

In Section 2.4 of the textbook, Strang summarizes this part of Linear Algebra by describing four subspaces associated with a matrix, and applying the LU factorization to the computation of these spaces. His task would have been **simpler**, **more natural**, and **symmetric** in its treatment of rows and columns if he had used the reduced $\underline{L}\,\underline{U}$ factorization defined in Section 2.3 of this supplement and illustrated in Section 2.4 of this supplement, so we will include that characterization as well.

A basis for the **column space** is taken to be those columns of **the original matrix** M, that contain pivot elements of the first part of Gaussian elimination process. Since the vectors are columns of M, there is no doubt that they lie in the column space. The process used to solve equations with this matrix of coefficients shows that they form a basis. Note that Gaussian elimination is used to identify the columns that are needed,

but they are found in the original matrix. Many students want to see a more **visible** role of the elimination process and seek to use columns of U. This is unlikely to work because **row** operations don't preserve **columns**. However, the columns of $\underline{L}$ do form a basis for the column space of M. Here, it is important that one use $\underline{L}$, since L was constructed to be a basis for all of $\mathbb{R}^m$. The linear independence of these columns is **visible**, since the matrix is triangular.

The **nullspace** is the set of all vectors v with $Mv = 0$. In particular, each vector gives the coefficient in a linear combination of the columns of M that evaluates to the zero vector. The vectors that result from back substitution form a special basis for the nullspace that shows how to express the non-pivot columns of M in terms of earlier columns. These columns are produced from the rows of U (or $\underline{U}$) by back substitution.

We need to be a little careful about the phrase "row space". It it tempting to think of it as a space of rows, but it has become standard to think of "vector" as only denoting **column** vectors. Thus, the row space of M is defined to be the **column space** of M^T. We will still relate it to rows of M, and use only the reduced $\underline{L}\,\underline{U}$ factorization of M to find it. It is not necessary to do separate elimination steps on M^T.

The easiest basis for the **row space** of M will be built from the rows of $\underline{U}$, which are the nonzero rows of U. Although these are **not necessarily** rows of M, they are certainly in the row space, and it is easily seen that they are linearly independent and span the row space of M. It is tempting to look for actual rows of M to form the basis. They certainly belong to the row space, but one needs to show that they are **linearly independent**. If row interchanges were used in the Gaussian elimination process, one needs extra care to relate rows of M with rows of U, so it is better to use the rows of U that are **visibly** linearly independent and belong to the rowspace by the **result** of the elimination process.

That leaves the **left nullspace**. Since the nonzero rows of U are linearly independent, if $0 = v^\mathsf{T}M = v^\mathsf{T}LU$, all entries of $v^\mathsf{T}L$ that multiply the nonzero part of U must be zero. These are the first $n-r$ columns of L, which is $\underline{L}$. Thus, a basis for the left nullspace can be found using **back substitution** on the transpose of $\underline{L}$. The restores the symmetry between the nullspace and the left nullspace.

The back substitution calculations can be checked by showing that the vectors v in the nullspace satisfy $Uv = 0$, and similarly for the left nullspace: a vector u in the left nullspace satisfies $u^\mathsf{T}L = 0$.

A final check is to verify that $M = \underline{L}\,\underline{U}$.

The original problem of solving a system of equations is essentially the problem of finding a nullspace, but the use of matrices reveals that this is only one of four spaces that should be considered. This also obscures the symmetry between rows and columns of a matrix.

4.2 Intersection of subspaces

Having a record of the reduction of M to the echelon form U, one can solve a system $Mx = b$ by augmenting the matrices appearing in the reduction with those determined by b. The process is the same even if the record contains matrices that were previously augmented since only the part to the left of the dotted line is used to find the LU factorization. Having provided the details of the computation for two columns b_0 and b_1, to obtain U augmented by columns c_0 and c_1, there is a shortcut to the c corresponding to a linear combination $b = \alpha_0 b_0 + \alpha_1 b_1$. At each stage, one would get the same linear combination of columns, so it is not necessary to write them all — one can skip to the last step and write $c = \alpha_0 c_0 + \alpha_1 c_1$. This is useful if neither $Mx = b_0$ nor $Mx = b_1$ has a solution because c_0 and c_1 have nonzero entries where U has a zero row, since we can immediately describe those c which have zero in this position and use the coefficients α_i to identify the corresponding b. This is much simpler than multiplying by L.

This leads to a method for finding the intersection of subspaces V and W. Begin by using a **basis** for V as the columns in the left side of an augmented matrix M and a **basis** for W as the columns in the right side. Then begin Gaussian elimination to try to express each vector in the given basis of W in terms of the given basis of V. This leads to $M = LU$. Here L is a square matrix and U has the same shape as M and may be considered as an augmented matrix in the same way. If the left side of U has a zero row, any column c in the right side not having zero in this position corresponds to a vector in the basis for W that

does **not** belong to V. Typically, none of the vectors in the **basis** for W will belong to V, but there may be other **combinations** of these vectors that give elements of the intersection. Specifically, a combination belongs to the intersection if it has zero in all positions corresponding to the zero rows on the left side. The coefficients in this combination can be found by solving the homogeneous system whose left side consists of the right sides of all rows of M whose left sides are zero. This is equivalent to doing back substitution steps for the part of the **right side** of U corresponding to the zero rows on the left side.

Note that the back substitution has been done only for the rows of U with zero left side, the expressions for vectors of the intersection in terms of the basis of W has been revealed. This allows a basis of the intersection to be written. Further row operations will give expressions in terms of the basis of V, but this only serves as a check of the answer. Example 6 of Appendix A (of the fourth edition of the text) gives a brief description of the use of an LU factorization of this augmented matrix to find the intersection of spaces given as column spaces, but it fails to notice that a basis for the intersection can be found using only **one** of the spaces being intersected. This oversight removes this problem from the realm of useful exercises in the text. We correct that defect here with an illustrative example and an exercise.

The above description may not be clear, so here is a simple example.

4.3 Example

Let

$$A = \begin{bmatrix} 3 & -1 & 2 \\ 9 & -1 & 11 \\ 6 & -4 & 2 \\ -3 & 5 & 20 \end{bmatrix} \qquad B = \begin{bmatrix} 4 & -3 \\ 19 & -10 \\ 0 & -3 \\ 8 & 15 \end{bmatrix}$$

Find the intersection of the space V spanned by the columns of A and the space W spanned by the columns of B.

Solution. Begin by forming the augmented matrix AB, then reduce A to echelon form (pivots will be shown in bold type).

$$\begin{bmatrix} \mathbf{3} & -1 & 2 & \vdots & 4 & -3 \\ 9 & -1 & 11 & \vdots & 19 & -10 \\ 6 & -4 & 2 & \vdots & 0 & -3 \\ -3 & 5 & 20 & \vdots & 8 & 15 \end{bmatrix}$$

$$\begin{bmatrix} 3 & -1 & 2 & \vdots & 4 & -3 \\ 0 & \mathbf{2} & 5 & \vdots & 7 & -1 \\ 0 & -2 & -2 & \vdots & -8 & 3 \\ 0 & 4 & 22 & \vdots & 12 & 12 \end{bmatrix}$$

$$\begin{bmatrix} 3 & -1 & 2 & \vdots & 4 & -3 \\ 0 & 2 & 5 & \vdots & 7 & -1 \\ 0 & 0 & \mathbf{3} & \vdots & -1 & 2 \\ 0 & 0 & 12 & \vdots & -2 & 14 \end{bmatrix}$$

$$\begin{bmatrix} 3 & -1 & 2 & \vdots & 4 & -3 \\ 0 & 2 & 5 & \vdots & 7 & -1 \\ 0 & 0 & 3 & \vdots & -1 & 2 \\ 0 & 0 & 0 & \vdots & 2 & 6 \end{bmatrix}$$

The elements of right side of the bottom row are not zero. This means that no column of B belongs to V. However, if 3 times the first column is subtracted from the second column, the result has has a zero in the last position. Performing this operation on the columns of the original matrix B gives

$$\begin{bmatrix} -15 \\ -67 \\ -3 \\ -9 \end{bmatrix}.$$

This is a basis for the intersection, i.e., the multiples of this vector give all elements of $V \cap W$.

Since the vector is a linear combination of columns of B, it belongs to W. Properties of the echelon form show that it also belongs to V.

When you have a basis for $V \cap W$, finding the descriptions of those vectors in terms of **both** the given basis of V and the given basis of W is equivalent to finding the nullspace of a matrix whose set of columns is the union of those bases (this observation was given in a previous edition of the textbook, but is one of the few things removed in the transition to the current edition). The method of these notes, while based on row operations on the same matrix, involves less computation since it only finds expressions with respect to the basis of one of the subspaces after formulating the property that a combination of the vectors in a basis for W belongs to V.

A variation on the method for finding both representations is to do the back substitution only with columns on the right side that come from the intersection. These are recognized in the echelon form by having zero in the positions where there is a row of zeroes on the left side of the augmented matrix. Thus, if you continue the solution with U augmented by

$$\begin{bmatrix} -3 \\ -1 \\ 2 \\ 6 \end{bmatrix} - 3 \begin{bmatrix} 4 \\ 7 \\ -1 \\ 2 \end{bmatrix} = \begin{bmatrix} -15 \\ -22 \\ 5 \\ 0 \end{bmatrix}$$

you will get the solution $\left({}^{-67}\!/_{6}, {}^{-91}\!/_{6}, {}^{5}\!/_{3} \right)$ giving the coefficients in the expression of this vector in terms of the columns of A.

Our example, and Exercise A below are just large enough to give an example of a nontrivial intersection. In general, the intersection can have any dimension, so many combination of columns may need to be found in order to arrive at a basis.

5. Exercises

A. Let

$$A = \begin{bmatrix} 1 & 3 & -2 \\ -2 & -5 & 3 \\ 3 & 9 & -5 \\ -2 & -4 & 3 \end{bmatrix} \qquad B = \begin{bmatrix} 5 & -3 \\ -12 & 7 \\ 15 & -6 \\ -11 & 13 \end{bmatrix}$$

(a) Find the intersection of the space V spanned by the columns of A and the space W spanned by the columns of B by forming an augmented matrix with A on the left and B on the right.

(b) Repeat with a matrix that has B on the left and A on the right. (This has **not** been constructed to avoid fractions in the computation. You may use of a calculator or computer giving decimal values for the matrix entries.)

(c) Interpret these results. Although the details are different, both claim to find the intersection of V and W. You should see that the spaces are the same.

Note. In simple cases, you can **see** that the spaces are the same. However, if the dimension of the intersection is greater than 1, different computations may give **very** different bases for the space, so another test must be used to recognize that the same space has been found.

B. Given

$$M = \begin{bmatrix} 1 & 3 & -2 & 1 & 5 \\ 5 & 15 & -9 & 6 & 23 \\ -2 & -6 & 3 & -3 & -8 \\ 3 & 9 & -4 & 5 & 11 \end{bmatrix},$$

find the **reduced** $\underline{L}\,\underline{U}$ factorization described in Section 2.4 of this supplement. Use this to find bases for the four fundamental subspaces of M as described in Section 4.1 of this supplement.

End of Supplement