REVIEW FOR EXAM II

The exam covers sections 3.4–3.6, the part of 3.7 on Markov chains, and 4.1–4.3.

1. The LU factorization:

   An $n \times n$ matrix A has an LU factorization if $A = LU$, where L is lower triangular ($\ell_{i,j} = 0$ when $i < j$, i.e. every entry above the diagonal is zero) such that each $\ell_{i,i} = 1$ (i.e. every entry on the diagonal is 1) and U is an upper triangular matrix ($u_{i,j} = 0$ when $i > j$, i.e. every entry below the diagonal is zero). Not every A has an LU factorization. When A has an LU factorization and is invertible, then the L and the U are unique.

   Every A can be put in REF by a sequence of row operations in which one never does a row multiplication $R_i \to kR_i$. If A can be put in REF by a sequence of row operations in which one never does a row multiplication $R_i \to kR_i$ and never does a row exchange $R_i \leftrightarrow R_j$, i.e. one uses ONLY row operations of the form $R_i \to R_i + kR_j$, then A has an LU factorization. The U is just the REF. If $\rho_1, \rho_2, \ldots, \rho_m$ is the sequence of row operations which takes A to U, and $E_1, E_2, \ldots, E_m$ is the corresponding sequence of elementary matrices, then we have $E_m \ldots E_2 E_1 A = U$, and so $A = (E_m \ldots E_2 E_1)^{-1} U = E_1^{-1} E_2^{-1} \ldots E_m^{-1} U$. Each E_i^{-1} is a lower triangular matrix with all diagonal entries equal to one. Since the product of any string of matrices of this type is also of this type we have that $L = E_1^{-1} E_2^{-1} \ldots E_m^{-1}$.

   The argument just given shows that under these circumstances the LU factorization exists and shows how to compute U. DO NOT USE IT TO COMPUTE L. There is a far easier and quicker method. First of all use STRICT Gauss reduction to put A in REF U. This may not work if you don't use STRICT Gauss reduction. Each row operation has the form $R_i \to R_i - kR_j$. The number k is called the multiplier. The result of the row operation is to change the entry in the i^{th} row and j^{th} column of your current matrix to a zero. To compute L start with the identity matrix I. For each of the row operations change the corresponding entry in the i^{th} row and j^{th} column of I to k. This will be your L.

   For your final answer YOU MUST WRITE DOWN A SPECIFIC L AND U (WITH NUMBERS IN THEM) SIDE BY SIDE WITH L ON THE LEFT AND U ON THE RIGHT.

   There is a generalization of the LU factorization called the $\mathrm{P^T}$LU factorization. Every matrix A has a $\mathrm{P^T}$LU factorization. You first do a sequence of row operations to put A in REF. List all the row exchanges $R_i \leftrightarrow R_j$ that you do, ignoring all the others. Then start over by applying this sequence of row exchanges, in order, to A. This gives $PA = P_t \ldots P_2 P_1 A$, where the P_i are the elementary matrices that perform these row exchanges. Now find the LU factorization of this matrix $PA = LU$. It turns out that for this sort of P we have $P^{-1} = P^T$, so we have $A = P^T LU$.

2. Subspaces:

A subset S of $\mathbb{R}^n$ is a subspace of $\mathbb{R}^n$ if for *every* $\mathbf{u}$ and $\mathbf{v}$ in S and *every* real number r we have (1) $\mathbf{u}+\mathbf{v}$ is in S, and (2) $r\mathbf{u}$ is in S. To determine whether or not S is a subspace you must first know *exactly* what it means for a vector $\mathbf{x}$ to be in S. This will be given by some criterion stated in the definition of S. If you suspect that S *is* a subspace, then you must write down $\mathbf{u}$ and $\mathbf{v}$ using algebraic vector notation as ordered n-tuples of variables, write down what it means for $\mathbf{u}$ and $\mathbf{v}$ to be in S, then verify (usually algebraically) that $\mathbf{u}+\mathbf{v}$ and $r\mathbf{u}$ also satisfy this criterion. If you suspect that S is not a subspace, then the task is usually much easier; you just have to find specific numerical examples which violate (1) or (2). (One of these two is enough.) Note that if you are instructed on the exam that you do not have to give explanations, then you merely have to write down "yes" or "no"; you would then not have to write down the details mentioned above, although you would probably want to briefly think them through to yourself.

Given a finite set $T = \{\mathbf{w}_1, \ldots, \mathbf{w}_k\}$ of vectors in $\mathbb{R}^n$ the span of T is the set span(T) of all linear combinations $c_1\mathbf{w}_1 + \cdots + c_k\mathbf{w}_k$ of the elements of T as $c_1, \ldots, c_k$ range over all the real numbers. It is easy to see that this is a subspace of $\mathbb{R}^n$.

3. Bases and dimension:

A subset B of a subspace S is a basis for S if it spans S and is linearly independent. This is equivalent to the statement that every vector in S can be expressed as a *unique* linear combination of these vectors.

Every basis for S has the same number of elements. This number is called the dimension of S, written $\dim(S)$.

Suppose $\dim(S) = d$. Let $T = \{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_p\}$ be a subset of S.

If $p < d$, then T cannot span S.

If $p > d$, then T cannot be linearly independent, i.e. it must be linearly dependent.

If $p = d$, the T spans S if and only if it is linearly independent, i.e. either both properties hold (and so T is a basis for S) or neither of them holds.

If T spans S, then T can be reduced to a basis B for S by discarding any vector in T that is a linear combination of the others to obtain a smaller set T', then repeating the process until one gets a set for which no vector is a linear combination of the others.

If T is a linearly independent subset of S, then T can be extended to a basis B for S by taking a known basis B' for S, taking the union $T \cup B'$, and then discarding vectors from $T \cup B'$ as above, being careful to never discard an element of T.

4. Coordinate vectors:

   Suppose $B = \{\mathbf{b}_1, \ldots, \mathbf{b}_k\}$ is a basis for S. Let $\mathbf{v}$ be an element of S. Because B spans S we have that $\mathbf{v} = c_1\mathbf{b}_1 + \cdots + c_k\mathbf{b}_k$ for some numbers c_i. Because B is linearly independent, these c_i are the *only* coefficients for which this holds. They are the coordinates of $\mathbf{v}$ with respect to the basis B. The coordinate vector of $\mathbf{v}$ with respect to B is

$$[\mathbf{v}]_B = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_k \end{bmatrix}$$

   If you are given $\mathbf{v}$ and need to find $[\mathbf{v}]_B$ you must solve the equation $\mathbf{v} = c_1\mathbf{b}_1 + \cdots + c_k\mathbf{b}_k$ for the c_i and then assemble them into the vector $[\mathbf{v}]_B$ in $\mathbb{R}^k$.

   If you are given $[\mathbf{v}]_B$ and need to find $\mathbf{v}$ all you need to do is compute $\mathbf{v} = c_1\mathbf{b}_1 + \cdots + c_k\mathbf{b}_k$.

5. Subspaces associated to matrices:

   There are three main subspaces associated to an $m \times n$ matrix A. Bases for them can be found by examining a REF U as detailed below.

   The row space $\text{row}(A)$ is the subspace of $\mathbb{R}^n$ spanned by the rows of A. The non-zero rows of U form a basis for $\text{row}(A)$. Since the number of non-zero rows of U is the rank of A we have that the dimension of $\text{row}(A)$ is equal to the rank of A.

   The column space $\text{col}(A)$ is the subspace of $\mathbb{R}^m$ spanned by the columns of A. Locate the columns of U which contain leading entries of the rows of U. Say, for example, that these are the first, third, and fourth columns of U. In general THESE DO NOT FORM A BASIS FOR THE COLUMN SPACE OF A. Instead, take the corresponding columns of A. In this example you would take the first, third, and fourth columns of A. The number of these columns is equal to the number of leading entries of U, which is equal to the number of non-zero rows of U which is the rank of A, so the dimension of $\text{col}(A)$ is equal to the rank of A.

   The nullspace $\text{null}(A)$ is the set of all solutions $\mathbf{x}$ of $A\mathbf{x} = \mathbf{0}$. You find a basis for this in the usual way. (See the review for Exam I.) The dimension of $\text{null}(A)$ is called the nullity of A. It is equal to the number of free variables, which is equal to n minus the number of bound varibles, which is equal to n minus the rank of A, so we have

$$\text{rank}(A) + \text{nullity}(A) = n$$

6. Linear transformations:

   A function $T : \mathbb{R}^n \to \mathbb{R}^m$ is a linear transformation if for *every* $\mathbf{u}$ and $\mathbf{v}$ in $\mathbb{R}^n$ and *every* real number r we have (1) $T(\mathbf{u}+\mathbf{v}) = \mathbf{T}(\mathbf{u}) + \mathbf{T}(\mathbf{v})$, and (2) $T(r\mathbf{u}) = rT(\mathbf{u})$. Note that this looks somewhat like a subspace problem. If you suspect that T *is* a linear transformation, then you would need to write down $\mathbf{u}$ and $\mathbf{v}$ in algebraic notation and prove that (1) and (2) hold. If you suspect that T is *not* a linear transformation, then it suffices to find specific vectors for which at least one of these equations is not true. In many cases you should be able to guess quickly from the form of T what to suspect. Again, if you are told that you do not have to give explanations, then you need merely write down "yes" or "no".

7. Given an $m \times n$ matrix A, the function $T(\mathbf{x}) = A\mathbf{x}$ is easily seen to be a linear transformation. As we will show below every linear transformation has this form.

   The range $T(\mathbb{R}^n)$ of T is the set of all vectors $\mathbf{b}$ in $\mathbb{R}^m$ for which there is some $\mathbf{x}$ in $\mathbb{R}^n$ with $T(\mathbf{x}) = \mathbf{b}$. This is a subspace of $\mathbb{R}^m$. When T is represented by a matrix A this is just the column space of A.

   The kernel $\ker(T)$ is the set of all vectors $\mathbf{x}$ in $\mathbb{R}^n$ such that $T(\mathbf{x}) = \mathbf{0}$. This is a subspace of $\mathbb{R}^n$. When T is represented by a matrix A this is just the nullspace of A.

   If $B = \{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n\}$ is a basis for $\mathbb{R}^n$ and $\mathbf{x} = c_1\mathbf{b}_1 + c_2\mathbf{b}_2 + \cdots + c_n\mathbf{b}_n$, then $T(\mathbf{x}) = c_1T(\mathbf{b}_1) + c_2T(\mathbf{b}_2) + \cdots + c_nT(\mathbf{b}_n)$. So the function T is completely determined by the values of $T(\mathbf{b}_1), T(\mathbf{b}_2), \ldots, T(\mathbf{b}_n)$.

   In particular, if $E = \{\mathbf{e}_1, \ldots, \mathbf{e}_n\}$ is the standard basis for $\mathbb{R}^n$ consisting of the vectors

$$\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \ldots, \mathbf{e}_n = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}$$

   then T is completely determined by the $m \times n$ matrix

$$[T] = \left[\ T(\mathbf{e}_1) \quad T(\mathbf{e}_2) \quad \ldots \quad T(\mathbf{e}_n)\ \right]$$

   This is called the standard matrix of T. It has the property that for every $\mathbf{x}$ in $\mathbb{R}^n$ we have $T(\mathbf{x}) = [T]\mathbf{x}$. Thus every linear transformation has the form $T(\mathbf{x}) = A\mathbf{x}$.

   If $T : \mathbb{R}^n \to \mathbb{R}^m$ and $S : \mathbb{R}^m \to \mathbb{R}^p$ are linear transformations, then the composition $S \circ T : \mathbb{R}^n \to \mathbb{R}^p$ given by $(S \circ T)(\mathbf{x}) = S(T(\mathbf{x}))$ is also a linear transformation. The standard matrices satsify $[S \circ T] = [S][T]$, i.e. the matrix of the composition is the product of the matrices.

When $n = m = p$ and $(S \circ T)(\mathbf{x}) = \mathbf{x}$ and $(T \circ S)(\mathbf{x}) = \mathbf{x}$ for all $\mathbf{x}$ in $\mathbb{R}^n$ we say that T is invertible with inverse $T^{-1} = S$. In this case $[T^{-1}] = [T]^{-1}$.

8. Markov chains:

   A stochastic matrix is a square matrix with the properties that all the entries are non-negative and for each column the sum of the entries in that column is 1. These matrices are used to model processes in which a system repeatedly changes from one state to another and the new state depends only on the state immediately preceding it (a Markov chain). The main skill to have here is the ability to translate a problem into a statement about stochastic matrices. See Example 3.64 in the book (page 236 of the 3rd edition, page 230 of the 4th edition) or the example about the population of California presented in the notes on 3.7.

9. Eigenvalues and eigenvectors:

   Let A be an $n \times n$ matrix. A number λ is an eigenvalue of A if there is a vector $\mathbf{x} \neq \mathbf{0}$ such that $A\mathbf{x} = \lambda\mathbf{x}$. In this case $\mathbf{x}$ is called an eigenvector of A belonging to (or associated to) λ. You should be able to tell from this definition whether or not a given vector is an eigenvector of a given matrix and, if it is, what the eigenvalue is, without doing any involved computations.

   Note that $A\mathbf{x} = \lambda\mathbf{x}$ is equivalent to $(A - \lambda I)\mathbf{x} = \mathbf{0}$, where I is the identity matrix.

10. If λ_i is an eigenvalue of A, then the set of *all* solutions $\mathbf{x}$ of $A\mathbf{x} = \lambda\mathbf{x}$ is called the eigenspace E_{λ_i} of λ. The eigenvectors belonging to λ_i are thus the non-zero elements of E_{λ_i}. Since E_{λ_i} is the same thing as the nullspace of $A - \lambda_i I$, it is a subspace of $\mathbb{R}^n$. To determine E_{λ_i} you just solve the equation $(A - \lambda_i I)\mathbf{x} = \mathbf{0}$ in the usual way. A basis B_{λ_i} for E_{λ_i} is called an eigenbasis corresponding to λ_i. Every element of E_{λ_i} can be expressed as a linear combination of the elements of B_{λ_i} with coefficients the parameters you use in solving the equation. The eigenvectors are precisely those linear combinations for which at least one parameter is non-zero. Be sure to read a problem carefully to see whether the required answer is the eigenspace, the eigenvectors, or the eigenbasis; these are not the same things. Note in particular that an eigenbasis is a finite list of specific numerical vectors; it never has parameters in it like s or t.

11. For the case $n = 2$ it is easy to find all the eigenvalues of A. If λ is an eigenvalue, then $(A - \lambda I)\mathbf{x} = \mathbf{0}$ will have a non-trivial solution.

    For $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ we have $A - \lambda I = \begin{bmatrix} (a-\lambda) & b \\ c & (d-\lambda) \end{bmatrix}$.

    We will have a non-trivial solution when the determinant $(a - \lambda)(d - \lambda) - bc = 0$. Cleaning this up we have the characteristic polynomial $\lambda^2 - (a+d)\lambda + (ad - bc) = 0$. So the eigenvalues are just the solutions of this quadratic equation.

In this course we will require all our eigenvalues to be real numbers, so if you get only non-real solutions to the characteristic equation we will say that A has no eigenvalues.

12. Determinants:

(a) First order determinants: Given a 1×1 matrix $A = [a]$ we define the determinant to be $\det(A) = a$.

(b) Second order determinants: Given a 2×2 matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, we define the determinant to be $\det(A) = \begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc$. Note: The vertical lines in the second expression do NOT mean absolute value.

(c) Third order determinants: Given a 3×3 matrix $A = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}$, we define the determinant to be $\det(A) = \begin{vmatrix} a & b & c \\ d & e & f \\ g & h & i \end{vmatrix} = a \begin{vmatrix} e & f \\ h & i \end{vmatrix} - b \begin{vmatrix} d & f \\ g & i \end{vmatrix} + c \begin{vmatrix} d & e \\ g & h \end{vmatrix}$.

Alternatively, this can be computed by the <u>criss-cross</u> method (also known as the <u>basket weave</u> method):

$$\left|\begin{array}{ccc|cc} a & b & c & a & b \\ d & e & f & d & e \\ g & h & i & g & h \end{array}\right. = aei + bfg + cdh - ceg - afh - bdi$$

(Multiply going northwest to southeast and add; multiply going northeast to southwest and subtract.)

(d) Cofactor expansion: Suppose A is an $n \times n$ matrix.

The <u>minor submatrix</u> $A_{i,j}$ is the $(n-1) \times (n-1)$ matrix obtained by deleting the i^{th} row and j^{th} column of A.

The <u>minor determinant</u> is $\det(A_{i,j})$.

The <u>cofactor</u> is $C_{i,j} = (-1)^{(i+j)} \det(A_{i,j})$.

Cofactor expansion along the i^{th} row gives

$$\det(A) = a_{i,1}C_{i,1} + a_{i,2}C_{i,2} + \cdots + a_{i,n}C_{i,n}$$

Cofactor expansion along the j^{th} column gives

$$\det(A) = a_{1,j}C_{1,j} + a_{2,j}C_{2,j} + \cdots + a_{n,j}C_{n,j}$$

The way to remember the signs in these formulas is to use the checkerboard matrices

$$\begin{bmatrix} + & - \\ - & + \end{bmatrix}, \begin{bmatrix} + & - & + \\ - & + & - \\ + & - & + \end{bmatrix}, \begin{bmatrix} + & - & + & - \\ - & + & - & + \\ + & - & + & - \\ - & + & - & + \end{bmatrix}, \text{etc.}$$

Since cofactor expansion along a row of A gives the same result as cofactor expansion along a column of A^T we have that $|A^T| = |A|$. So any of the statements made below about the rows of A apply to the columns of A as well.

(e) Row operations: Suppose we perform a row operation on a matrix A to obtain a matrix B.

Doing $R_i \leftrightarrow R_j$ gives $|B| = -|A|$.

Doing $R_i \to R_i + kR_j$ gives $|B| = |A|$.

Doing $R_i \to kR_i$ gives $|B| = k|A|$.

The most troublesome of these is the third one. You simply have to *think* about which matrix is your original matrix A and which is the matrix B obtained by doing $R_i \to kR_i$ to A. Sometimes you will use this procedure in reverse. You will have a matrix B one of whose rows contains a common factor k. Factoring out that common factor gives you a matrix A. The situation is exactly the same as above. You could get B back again by doing $R_i \to kR_i$ on A, so we still have $|B| = k|A|$, and hence $|A| = (1/k)|B|$. Don't just try to memorize this; *understand* what's going on.

Another troublesome situation is when you have $|kA|$. Well, each of the n rows of A has been multiplied by k. You have to factor out k a total of n times, so that you build up $k^n|A|$.

(f) Easy determinants:

If any of the following occur, then $|A| = 0 :$ a row of zeroes, two equal rows, two proportional rows. (Remember, the same statements hold for columns.)

If A is upper triangular (all entries below the diagonal being zero) or lower triangular (all entries above the diagonal being zero), or diagonal (both upper and lower triangular), then $|A|$ is just the product of the entries on the diagonal.

If you multiply an $n \times n$ matrix A by a number c to get a matrix cA, then $|cA| = c^n|A|$. (Every row of cA has a factor of c. Doing $R_1 \to (1/c)R_1$, $R_2 \to (1/c)R_2$, $\dots$, $R_n \to (1/c)R_n$ gives $|A| = |cA|/c^n$, so $|cA| = c^n|A|$.)

(g) Triangulation: The generally fastest way to compute $|A|$ is to put it in row echelon form U. Since U is upper triangular, we compute $|U|$ as above. Then we keep track of how the row operations change the determinant in order to get $|A|$. In particular, REF can be achieved without ever using $R_i \to kR_i$, and if you compute REF this way the determinant will at most change sign.

In practice one sometimes combines row and column operations with cofactor expansion to obtain shortcuts. This is not a good idea in general, but in some cases it saves a lot of time.

(h) Products and powers: $|AB| = |A||B|$, and so $|A^p| = |A|^p$.

(i) Inverses: A is invertible if and only if $|A| \neq 0$.

It follows from the formula for the determinant of a product and the fact that $|I| = 1$ that $|A^{-1}| = \dfrac{1}{|A|}$.

Suppose that $|A| \neq 0$.

If $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, then $A^{-1} = \dfrac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$. This is the fastest way to find the inverse of a 2×2 matrix.

For all n there is a similar formula. $A^{-1} = \operatorname{adj}(A)/\det(A)$. The adjoint $\operatorname{adj}(A) = C^T$, where C is the matrix of cofactors $C_{i,j}$. The $n = 2$ case has $\operatorname{adj}(A) = \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$ and $\det(A) = ad - bc$. For $n > 2$ this formula takes much longer to compute than using row operations to transform $[A|I]$ to $[I|A^{-1}]$.

(j) Cramer's rule: Given a system $A\mathbf{x} = \mathbf{b}$, where A is an $n \times n$ matrix with $|A| \neq 0$, let B_k be the matrix obtained by replacing the k^{th} column of A with $\mathbf{b}$. Then the k^{th} component of $\mathbf{x}$ is given by $x_k = \det(B_k)/\det(A)$.

(k) Transposes: $|A^T| = |A|$. This is because expansion along any row or column always gives the same result (Laplace's expansion theorem) and expanding along a row of A is the same computation as expanding along the corresponding column of A^T.

13. Finding eigenvalues and eigenvectors:

(a) A number λ is an eigenvalue of a matrix A if there is a vector $\mathbf{x} \neq \mathbf{0}$ such that $A\mathbf{x} = \lambda\mathbf{x}$. In this case $\mathbf{x}$ is called an eigenvector of A belonging to λ. (We may also say "corresponding to"or "associated to".)

You should be able to tell from this definition whether or not a given vector is an eigenvector of a given matrix and, if it is, what the eigenvalue is, without doing any involved computations. If $\mathbf{x} = \mathbf{0}$, then $\mathbf{x}$ is not an eigenvector. If $\mathbf{x} \neq \mathbf{0}$ then multiply A times $\mathbf{x}$ and see whether or not you get something of the form $\lambda\mathbf{x}$; if you do, then $\mathbf{x}$ is an eigenvector with eigenvalue λ.

(b) The characteristic polynomial $p(\lambda)$ of a matrix A is $\det(A - \lambda I)$. For a particular value λ_i of λ the following statements are equivalent:

i. λ_i is an eigenvalue of A.
ii. $A - \lambda_i I$ has a non-trivial nullspace.
iii. $A - \lambda_i I$ is singular.
iv. $p(\lambda_i) = 0$.
v. $(\lambda - \lambda_i)$ is a factor of $p(\lambda)$.

The highest power n_i such that $(\lambda - \lambda_i)^{n_i}$ is a factor of $p(\lambda)$ is called the algebraic multiplicity of the eigenvalue λ_i. If $\lambda_1, \ldots, \lambda_k$ are all of the *different* numbers which are eigenvalues of A, then $k \leq n$ and $n_1 + \cdots + n_k = n$. Note that $k = n$ if and only if each $n_i = 1$; in this case we say that A has n distinct eigenvalues.

(c) Finding eigenvalues: In this course we will be interested in only real numbers as eigenvalues. For 2×2 matrices the eigenvalues are easy to find since we have just a quadratic equation to solve. For larger matrices it may be more difficult to find the zeroes of $p(\lambda)$. While you are computing $p(\lambda)$ remember that your goal is to find its zeroes, and this amounts to finding its linear factors $(\lambda - \lambda_i)$, so you want to keep $p(\lambda)$ as factored out as you possibly can. Always be on the lookout for common factors as you compute $p(\lambda)$. Resist the temptation to expand $p(\lambda)$ out into some cubic or higher degree polynomial that will be difficult to handle.

If no common factors arise as you compute $p(\lambda)$, then you will have to expand it out. In this case look first for some easy factorization given by repeated patterns in the coefficients. If you cannot find such a pattern, then try the rational roots test. (See Appendix D of the book for the rational roots test. If $p(\lambda)$ has integer coefficients, then list all the factors, both positive and negative, of the constant term. Evaluate $p(\lambda)$ at each of these numbers. If for one of these numbers a, you get $p(a) = 0$, then you have found an eigenvalue. Then $p(\lambda)$ factors as $(\lambda - a)q(\lambda)$. The fastest way to both evaluate a polynomial and to divide it by $(\lambda - a)$ is by using synthetic division. The last number in the third row gives you $p(a)$. If $p(a) = 0$, then the numbers in the third row to the left of $p(a)$ give you the coefficients of $q(\lambda)$. (If you don't remember how to do synthetic division, look it up in a college algebra book. It may also be added to the notes for lecture 24 [October 22].) Next discard $p(\lambda)$ and examine $q(\lambda)$. It may be easy to factor or may be a quadratic, in which case you can quickly finish the problem. If not, then try the rational roots test again on $q(\lambda)$. Don't drag along all the candidates you had for $p(\lambda)$; start fresh with $q(\lambda)$.

(d) Finding eigenvectors, eigenspaces, and eigenbases: If λ_i is an eigenvalue of A, then the set of *all* solutions $\mathbf{x}$ of $A\mathbf{x} = \lambda\mathbf{x}$ is called the eigenspace E_{λ_i} of λ. The eigenvectors belonging to λ_i are thus the non-zero elements of E_{λ_i}. Since E_{λ_i} is the same thing as the nullspace of $A - \lambda_i I$, it is a subspace of $\mathbb{R}^n$. The dimension m_i of E_{λ_i} is called the geometric multiplicity of λ_i. One always has $1 \leq m_i \leq n_i$. To determine E_{λ_i} you just solve the equation $(A - \lambda_i I)\mathbf{x} = \mathbf{0}$ in the usual way. A basis B_{λ_i} for E_{λ_i} is called an eigenbasis corresponding to λ_i. Every element of E_{λ_i} can be expressed as a linear combination of the elements of B_{λ_i} with coefficients the parameters you use in solving the equation. The eigenvectors are precisely those linear combinations for which at least one parameter is non-zero. Be sure to read a problem carefully to see whether the required answer is the eigenspace, the eigenvectors, or the eigenbasis; these are not the same things.