Notes to the Instructor

These notes, which are not included in the student edition, offer additional information that is not included elsewhere in the introductory material. We highlight additional resources and provide a more in-depth overview of the topics covered in the text to assist you in choosing and organizing your curriculum to suit your audience.

Instructional Aids

The following resources are provided to make your linear algebra course more enjoyable and effective for you and your students.

Interactive Textbook and Figures

The *Sixth Edition* of the text is available in an interactive electronic format within MyLab Math. By migrating the e-Text from Wolfram CDF Player to the Wolfram Cloud, we provide you with additional tools for lecturing and group work. We've included over 200 interactive figures that give you and your students the ability to play with mathematics. You can now look at multiple examples with the click of a button, leading to the formulation of conjectures that you will later prove. The geometry of linear algebra comes also comes alive as you show moving pictures illustrating concepts and definitions as they are discussed.

Instructor's Solutions Manual

The *Instructor's Solutions Manual* contains detailed solutions for all exercises, along with teaching notes for many sections. The manual is available electronically for download at *www.pearson.com* and in MyLab Math.

MyLab Math—Online Homework and Resources

Support for the *Sixth Edition* is offered through MyLab Math (*pearson.com/mylab/math*). Students submit homework online for instantaneous feedback, support, and assessment. Take the time at the beginning of the semester to help your students see the online homework as a learning tool, rather than just a means of assessment.

You may want to spend some time in the first week discussing with students how to use this program effectively. Go over the help features, how to enter their work into the computer, and how many times a student can attempt an exercise before losing points. If you set the number of attempts to unlimited, students can continue to work on an exercise until they have mastered the technique or idea.

Additionally, encourage your students to have a workbook where they can show their work and make notes about what they are learning. You may even want to ask students to hand in their notebooks occasionally so that you can give them tips on what sorts of things to include in these notebooks.

Past editions of MyLab Math worked particularly well for supporting computation-based skills. In the *Sixth Edition*, more sophisticated exercises can be assigned through MyLab Math and graded online by yourself or a teaching assistant. This new exercise style allows you to go completely paperless while still requiring students to set up their work, create short proofs, and articulate relationships between concepts.

Study Guide

The challenge of teaching linear algebra is heightened by the need to teach students how to learn mathematics. The *Study Guide* is designed to help you do this, with new advice to students about how to use MyLab Math effectively. The *Guide* will reinforce many comments you are likely to make in class about the importance of reading the text carefully, working the exercises promptly, and the like. Appendixes in some sections explain mathematical concepts such as induction and logical implication.

The *Study Guide* can save you time in class. You can refer students to some solutions in the *Guide* instead of working all the exercises in class. The *Guide* also answers many questions that students might otherwise ask during office hours. For courses using technology, the *Guide* is a "lab manual" that gives keystroke instructions and command references for the most popular matrix programs.

The *Guide* has two unique pedagogical features:

1. The *Study Guide* supplies a detailed solution to about a third of the odd-numbered problems (which includes most key exercises) and a solution to every odd-numbered writing problem for which the text's answer is only a *Hint*. A *Hint* shows students that they must work harder on a problem before reading the solution. (The *Study Guide* is separate from the text to prevent quick access to the solution.) Nearly 100 conceptual exercises have this design.
2. A series of subsections in the *Study Guide*, entitled "Mastering Linear Algebra Concepts," shows students how to make review sheets that identify connections between concepts in a way that builds a robust mental "image" of each concept. Such mental images are a key to true comprehension and retention of ideas.

The *Sixth Edition* of the textbook has marginal notes at appropriate points that direct the student to go to a "Mastering" subsection of the *Study Guide*, which is included in MyLab Math. You can train your students to make effective review sheets by requiring that each student submit review sheets, at least for the first two or three key concepts. Some instructors glance at each sheet and add notes when key features are missing, but do not grade the sheets other than to record that an effort was made. You may choose to make a review section part of the notebook you ask students to create as part of doing MyLab Math exercises. If review sheets are not required, the students who need them most may never create them.

Online Chapter

Chapter 10, listed in the Table of Contents and posted online, is designed to be part of a second course in linear algebra that focuses on mathematics and applications of interest to students in mathematics, computer science, and engineering. The chapter offers opportunities for greater depth and breadth, in both the mathematics and the applications, than can be achieved in a single course.

Chapter 10 was written and class tested by Thomas Polaski (Winthrop University, Rock Hill, South Carolina). We are grateful to Tom for providing such interesting material. The online chapter is available for download at `bit.ly/2nj1HhO`.

Designing the Course

This text is written to have the flexibility to meet the opportunity to offer everything from a two-credit intensive course to two courses worth four credits each. The Topical Organization Table on the following page shows how the text can accommodate a variety of courses. The core topics are listed in the first column, and in a short course, these may be all you have time to cover. Some of the supplementary topics in the second column will be considered as core topics by some instructors, but they are optional in the technical sense that subsequent sections do not depend on them. All applications listed in the third column are optional, but interesting. Finally, the fourth column lists a condensed alternative to covering the core topics listed in the same boxed row.

A section-by-section commentary starting on page xv identifies connections between various sections. The balance between supplementary topics and applications will vary according to your own course objectives. Extensive class testing for more than a decade has verified that each core section and supplementary section can usually be covered in one 50-minute period, with time for some discussion of homework. To stay within this timeline, instructors can focus on the main points of each section and allow students to supplement lectures through careful reading of the text. Longer application sections have subsections that can be covered in one period or less. Some instructors plan an additional day per chapter to allow more in-depth coverage of selected topics.

A Note on Vector Spaces

The text offers instructors two different approaches to vectors spaces and subspaces:

1. For an in-depth coverage of the topic, skip Sections 2.8 and 2.9 and cover Chapter 4. We are hearing from industry partners as well as engineering and computer science colleagues that understanding abstract vectors spaces has become more important because of its applications in signal processing, data analytics, and artificial intelligence. Chapter 4 has been revised to touch on infinite dimensional vector spaces, with digital signal processing added as a new application.
2. For a more condensed course, or one that focuses more on applications later in the text, the instructor may choose to cover Sections 2.8 and 2.9 but then skip Chapter 4. There is enough information in these two sections to comfortably cover the material in the remaining chapters.

Topical Organization of *Linear Algebra and Its Applications*

Core Topics	Supplementary Topics	Applications	Condensed Alternative
1.1 Systems of Linear Equations			
1.2 Row Reduction and Echelon Forms			
1.3 Vector Equations			
1.4 The Matrix Equation $A\mathbf{x} = \mathbf{b}$			
1.5 Solution Sets of Linear Systems		1.6 Applications of Linear Systems	
1.7 Linear Independence			
1.8 Introduction to Linear Transformations			
1.9 The Matrix of a Linear Transformation		1.10 Linear Models in Business, Science, and Engineering	
2.1 Matrix Operations			
2.2 The Inverse of a Matrix	2.4 Partitioned Matrices	2.6 The Leontief Input–Output Model	
2.3 Characterizations of Invertible Matrices	2.5 Matrix Factorizations	2.7 Applications to Computer Graphics	
3.1 Introduction to Determinants	3.2 Properties of Determinants	3.3 Cramer's Rule, Volume, and Linear Transformations	
4.1 Vector Spaces and Subspaces			2.8 Subspaces of $\mathbb{R}^n$
4.2 Null Spaces, Column Spaces, Row Spaces, and Linear Transformations			2.9 Dimension and Rank
4.3 Linearly Independent Sets; Bases			
4.4 Coordinate Systems	4.6 Change of Basis	4.7 Digital Signal Processing	
4.5 The Dimension of a Vector Space		4.8 Applications to Difference Equations	
5.1 Eigenvectors and Eigenvalues	5.4 Eigenvectors and Linear Transformations	5.6 Discrete Dynamical Systems	
5.2 The Characteristic Equation	5.5 Complex Eigenvalues	5.7 Applications to Differential Equations	
5.3 Diagonalization	5.8 Iterative Estimates for Eigenvalues	5.9 Applications to Markov Chains	
6.1 Inner Product, Length, and Orthogonality	6.4 The Gram–Schmidt Process		
6.2 Orthogonal Sets	6.5 Least-Squares Problems	6.6 Machine Learning and Linear Models	
6.3 Orthogonal Projections	6.7 Inner Product Spaces	6.8 Applications of Inner Product Spaces	
7.1 Diagonalization of Symmetric Matrices	7.3 Constrained Optimization		
7.2 Quadratic Forms	7.4 The Singular Value Decomposition	7.5 Applications to Image Processing and Statistics	
8.1 Affine Combinations			
8.2 Affine Independence	8.4 Hyperplanes		
8.3 Convex Combinations	8.5 Polytopes	8.6 Curves and Surfaces	
9.2 Linear Programming—Geometric Method		9.1 Matrix Games	
9.3 Linear Programming—Simplex Method	9.4 Duality		

Sample Syllabi

The first three syllabi are for semester-length courses, with at least 39 lecture periods of 50 minutes, including time for homework discussions and occasional quizzes, but not counting days for hour exams or the final examination. The text has been class tested many times in each of these courses. The 39-day schedule is based on a 15-week semester, with three or four exams and two or three days' slack time to allow for unplanned delays or digressions that often occur when a textbook is used for the first time. Extra review days, if any, should be planned within the 39-day schedule.

The fourth syllabus, for short courses, describes how to use Sections 2.8 and 2.9 to cover six core topics in a week or less and get quickly to eigenvalues. The final syllabus shows how to incorporate the topics recommended by the original Linear Algebra Curriculum Study Group into a semester course.

Course 1. A 39-lecture class for freshmen and sophomores, with a wide variety of majors.

Such a course could include all 25 core topics from Chapters 1–7 shown in the table and perhaps seven of the supplementary topics. In addition to these 32 class days, you can probably fit in four or more applications and still have three review days. We recommend one application from each of Chapters 1, 2, 4, and 5. They are spread out to provide welcome relief from the theory. Early introduction of applications seems especially beneficial in a course for younger students. Some courses, such as the original course at the University of Maryland, are four credits instead of three. In this case, another five or more supplementary and application sections can be covered.

Course 2. A 39-lecture course for juniors who have already seen determinants, had brief glimpses of matrix algebra in other courses, and have had multivariable calculus.

This course can go a little faster through Sections 1.1–1.3, covering all three sections in two days, with students reading a lot on their own. (Later, the class will also speed through Section 6.1.) Early applications are not as important for students who are already aware that linear algebra is useful. Also, Chapter 3 can be skipped entirely and the brief treatment in Section 5.2 used instead. As a result, the 25 core sections can be covered in 24 days, and perhaps nine of the supplementary sections can be covered, still leaving time for four or five days on the applications in Chapters 4–7. Section 5.7 (Applications to Differential Equations) was written with this upper-level course in mind. In a 14-week course, a class can usually get as far as the singular value decomposition (Section 7.4).

Course 3. A 39-lecture course for juniors and seniors in the social and management sciences who are aiming at graduate school and will likely take a course later in multivariate statistics.

This course covers all core sections. In Chapter 4, you should focus on $\mathbb{R}^n$, with some work on $\mathbb{P}_n$. For supplementary topics, you might select partitioned matrices (Section 2.4), inner product (Examples 1-6 of Section 6.7), constrained optimization

(Section 7.3), and perhaps one or two other sections. That should leave room for applications from Sections 1.10 (Examples 1 and 3), 2.6, 5.9, 6.6, 6.8 (Examples 1 and 2), and 7.5, plus two days for review.

Course 4. A 27-lecture short course.

Short courses typically focus on linear algebra in $\mathbb{R}^n$, with perhaps only a brief mention of more general vector spaces. Such courses can take advantage of Sections 2.8 and 2.9. All vector space concepts needed for Chapters 5–7 can be covered in two to three days with the sections. The basic schedule is to cover the first eleven core sections, insert Sections 2.8 and 2.9, and then cover core sections 5.1–5.3, 6.1, 6.2, 7.1, and 7.2. Total time: 19–22 days (not counting exams). There should be time to choose a few supplementary topics and one or two applications. (According to users, the application with the broadest appeal is Section 2.7, Computer Graphics.)

A single lecture on general vector spaces could be added at the end of the short course. If you wish to cover the five core sections of Chapter 4 instead of Sections 2.8 and 2.9, you will need about three extra days. Omit most of Section 4.4 (except perhaps Theorem 7).

Course 5. A course that covers the basic syllabus recommended by the original Linear Algebra Curriculum Study Group (LACSG).

The LACSG "core" topics are all included in any course that covers all 25 topics listed from Chapters 1 to 7 in this text as core topics, plus the following sections:

- 2.4 Partitioned Matrices
- 2.5 Matrix Factorizations
- 3.1 Introduction to Determinants
- 3.2 Properties of Determinants (plus Theorem 7 and Example 1 from Section 3.3)
- 6.4 The Gram-Schmidt Process
- 6.6 Applications to Linear Models (Example 1 only)

This 31-lecture schedule is about three days longer than the LACSG time estimates, but the 31 days include time for linear transformations, general vector spaces, and positive definite matrices. Another five or six sections can be selected to fill out some of the supplementary topics and applications suggested by the LACSG. If you are using technology, Section 5.8 would be a natural choice, or you could spend extra time on special projects.

Section Commentaries

The following commentaries point out connections between a given section and those that follow it. Obvious connections are usually not mentioned. Some subsections of core material are identified as optional when their omission will cause no serious problem later. The commentaries also list exercises that are particularly valuable pedagogically, anticipate future discussions, or treat an optional supplementary topic.

In each core section, we urge you to assign True/False questions. Students will probably spend from 5 to 15 minutes on the questions because they will have to read the text. That is the point of these exercises! Many instructors have used them enthusiastically for several years.

SECTION 1.1
Systems of Linear Equations

This section and the next form the foundation for the row reduction algorithm used throughout the text. The existence and uniqueness questions introduced here are of fundamental importance and appear at several points later in the text. The exercises focus only on the existence question. Systems with many solutions are treated in Section 1.2.

Key Exercises: 7, 23–26, 35. The heat flow problem in Exercise 43 is also discussed in exercises in Sections 1.10 and 2.5.

SECTION 1.2
Row Reduction and Echelon Forms

The pivot positions in a matrix will play an essential role in several theorems—on spanning, linear independence, invertibility, the determinant, and rank. Theorem 1 is fundamental. Example 3 is helpful but need not be presented in class. The parametric descriptions of solution sets are written in vector form in Section 1.5. The solution method is reviewed in Section 4.2 and is used throughout Chapters 5 and 7.

Key Exercises: 1–24, 35–40. Students should work at least four or five of Exercises 7–14 in preparation for Section 1.5. Interpolating polynomials in Exercises 33 and 34 will appear again in exercises at the end of Chapter 2.

SECTION 1.3
Vector Equations

The traditional arrow representation of a vector is a handicap to understanding and visualizing Span $\{\mathbf{u}, \mathbf{v}\}$ as well as the set $\{\mathbf{p} + t\mathbf{v} : t \text{ in } \mathbb{R}\}$, discussed in Section 1.5. The set of free arrows (actually equivalence classes) in $\mathbb{R}^3$ is discussed in Section 4.1. Until then, a vector is a list of numbers, represented geometrically by a point. Sometimes an arrow is added to draw attention to the point.

Part or all of Section 6.1 could be inserted here, of course, along with more vector geometry, but the danger of doing so is that students will start to perceive the course as largely computational. The students' opinions of the course are set somewhere in the first two weeks, and they need to feel the conceptual emphasis early. Forcing students to deal with Span $\{\mathbf{v}_1, \dots, \mathbf{v}_p\}$ has that effect. Example 7 is optional, but it is used in Example 6 of Section 1.8.

Key Exercises: 11–14, 17–22, 33, 34. A discussion of Exercise 25 will help students understand $[\,\mathbf{a}_1 \quad \mathbf{a}_2 \quad \mathbf{a}_3\,]$, $\{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\}$, and Span $\{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\}$.

SECTION 1.4
The Matrix Equation $A\mathbf{x} = \mathbf{b}$

First impressions are often lasting. That is why the connection between $A\mathbf{x}$ and a linear combination of vectors must be made immediately. The value of this definition is far greater than one might expect. We urge you to try it and to wait until later to give the modern definition of matrix multiplication, as in Section 2.1.

Nothing in the section is optional. Theorem 4 is extremely important. Students may have difficulty with its proof. (Example 3 will help.) The connection between Theorem 4(a) and an *existence* question could be made here, but it could also be made when reviewing for an exam.

Key Exercises: 1–20, 37, 38, 41, 42. Students should be writing explanations as part of their answers.

Note: The term *rank* can be introduced here, if you wish, but that adds unnecessary terminology at this point. This chapter and the next will focus on the *locations* of pivot positions, which give more information than simply the number of pivot positions.

SECTION 1.5
Solution Sets of Linear Systems

The geometry helps students understand Span $\{\mathbf{u}, \mathbf{v}\}$, in preparation for later discussions of subspaces. The parametric vector form of a set will be used throughout the text to construct bases for null spaces and eigenspaces. Figure 6 will appear again in Section 4.8.

Key Exercises: 1–14, 41–46.

SECTION 1.6
Applications of Linear Systems

This optional section can provide a short homework assignment (on one or two of the applications) if your class needs a pause in the theoretical development of the chapter. You might omit the application to economics if you plan to cover Section 2.6.

SECTION 1.7
Linear Independence

The second half of Theorem 7 will be used several times later in the text. Exercise 30 could be stated as a theorem, in parallel with Theorem 4, but students *must* know the proof to avoid confusing the two results. The *Study Guide* will help your students learn about linear independence.

Key Exercises: 9–20, 29–36.

Note: Now that your students know about spanning and linear independence, there may be a temptation to introduce vector space concepts such as *basis* and *subspace*. However, piling on too many concepts before the fundamentals are mastered is a mistake most other texts make.

SECTION 1.8
Introduction to Linear Transformations

The connection with existence and uniqueness problems is made here and in Section 1.9. Introducing linear transformations here raises the level of the first part of the course and simplifies the work later in the course. Linear transformations appear in the following sections in the next three chapters:

2.1 The definition of AB is motivated by linear transformations.

2.3 Invertible transformations are discussed.

2.5 Matrix factorization is explained via linear transformations.

2.7 Linear transformations are at the heart of computer graphics.

3.3 The volume-change effect of a linear transformation is measured by a determinant.

4.2 General linear transformations are introduced without fanfare, because students are already familiar with concrete matrix transformations.

4.4 The coordinate mapping is a one-to-one linear mapping of V onto $\mathbb{R}^n$.

4.5 The coordinate mapping provides a transparent proof of the key theorem needed to discuss the dimension of a vector space.

4.7 Applying linear transformations to infinite vectors (signals) illustrates that digital signal processors and linear transformations are one and the same.

4.8 This section is a showcase of the power of linear algebra, mainly because linear transformations can be used in the explanations.

Linear transformations appear frequently, too, in Chapters 5 and 7. Exercises 19 and 20 segue into Section 1.9. Exercise 33 is used in Section 2.7, Computer Graphics.

Key Exercises: 17–20, 33, 39. The *Study Guide* offers help for Exercise 31.

SECTION 1.9
The Matrix of a Linear Transformation

In this section, you will want to cover at least Theorem 10 and a few geometric examples. The notions of *one-to-one* and *onto* appear in the Invertible Matrix Theorem (Section 2.3), in the discussion of the coordinate mapping (Sections 2.8 and 4.4) Section 5.4 generalizes the construction in Theorem 10. The proof of Theorem 11 applies to any linear transformation, in case you want to discuss this later in Chapter 4.

Exercises 33–36 and 39–42 offer fairly easy writing practice. Exercises 39, 40, and 43 provide important links to earlier material.

SECTION 1.10
Linear Models in Business, Science, and Engineering

This section gives the students a break from the concentrated material in the previous sections. If only one application is covered, you should have plenty of time to get a good start on Section 2.1.

The nutrition problem is not mentioned later. The discussions of electrical circuits here and in Sections 2.5 and 5.7 are basically independent of one another. However, if you plan to treat Markov chains later in Section 5.9, then you should present the population movement examples. By considering population sizes here and population distributions in Section 5.9, you can sidestep the discussion of probabilities. You might also use the difference equations here to prepare for Chapter 5, although the introduction to that chapter will suffice.

SECTION 2.1
Matrix Operations

The definition here of a matrix product AB is not only well motivated, but it also gives the proper view of AB for nearly all matrix calculations. The dual fact about the rows of A and the rows of AB is seldom used, mainly because vectors here are usually written as columns. Although the section is long, it is easy and quite a bit can be assigned as outside reading.

Key Exercises: 13, 25–30. Exercises 31 and 32 are cited in the proof of Theorem 8 in Section 2.3. Exercises 35 and 36 are optional, but they are mentioned in Example 4 of Section 2.4.

SECTION 2.2
The Inverse of a Matrix

The proof of Theorem 5 is important; students need to see that both uniqueness and existence must be proved. Elementary matrices are used in Section 2.5 (Matrix Factorizations) and in Section 3.2. The algorithm for finding A^{-1} is popular because it is so familiar and leads to easy exam questions. However, in most courses, time is better spent on the LU factorization in Section 2.5, for instance. (The LU factorization takes longer to present.)

Key Exercises: 11–34, 45. Exercise 8 is referenced in Section 2.3, after the proof of Theorem 8. Exercise 12 is useful and indicates how matrix products involving inverses are actually computed in practice. This exercise can be used in Section 4.6, to compute a change-of-coordinates matrix, and the exercise is mentioned in a numerical note in Section 5.4. Exercises 33 and 34 are cited in the proof of Theorem 8.

SECTION 2.3
Characterizations of Invertible Matrices

The Invertible Matrix Theorem ties together most of the concepts studied thus far. Additional statements are added to the theorem in Sections 2.9, 3.2, 4.6, 5.2, and 7.4. The subsection on invertible linear transformations is optional.

Key Exercises: 23–32.

SECTION 2.4
Partitioned Matrices

This section should be part of any modern course in linear algebra because partitioned matrix notation is widely used today in disciplines that employ linear algebra. It is listed as optional only because it is not yet a standard topic in many linear algebra courses. Block matrix multiplication is used to obtain the spectral decomposition of a symmetric matrix in Section 7.1 and the reduced singular value decomposition in Section 7.4.

Key Exercises: 15, 16, and 18. Exercises 1–10 provide excellent practice with matrix algebra and the Invertible Matrix Theorem. Exercises 21 and 22 are mentioned at the end of Section 2.5.

SECTION 2.5
Matrix Factorizations

Modern algorithms in numerical linear algebra are often described using matrix factorizations. For practical work, this section is more important than Sections 4.6 and 5.4, even though matrix factorizations can be explained nicely in terms of change of bases. See Exercises 24–28.

SECTION 2.6
The Leontief Input–Output Model

This section is independent of Section 1.10. The material here makes a good backdrop for the series expansion of $(I - C)^{-1}$, because this formula is actually used in some practical economic work. Exercise 8 gives an interpretation to entries of an inverse matrix that could be stated without the economic context.

SECTION 2.7
Applications to Computer Graphics

This section has had universal appeal. A five- or ten-minute video demonstration of computer graphics will heighten student interest. The section is fun to teach, and it provides practice with composition of linear transformations. A common student mistake is to reverse the order of the matrices.

SECTION 2.8
Subspaces of $\mathbb{R}^n$

This section and the next extract everything you need from Chapter 4 to discuss the topics in Chapters 5–7 (except for the general inner product spaces in Sections 6.7 and 6.8). Omit Sections 2.8 and 2.9 if you plan to cover Chapter 4.

Theorem 12 and Example 6 are important for eigenvector calculations in Chapter 5. Examples 7 and 8 and Theorem 13 are used mainly for the Rank Theorem in Section 2.9 and then for a discussion of the fundamental subspaces determined by a matrix, in Sections 6.1 and 7.4.

Key Exercises: 5–20, 31–34.

SECTION 2.9
Dimension and Rank

The concept of a coordinate vector is needed for Section 5.4. It also helps students to understand the change of variable $\mathbf{x} = P\mathbf{y}$ in Sections 5.6, 5.7, and 7.2. The notions of *dimension* and *rank* are important for later work in many sections. The Basis Theorem is crucial for the theory of homogeneous difference and differential equations (in Sections 5.6 and 5.7), and it is used in the Gram–Schmidt process (Section 6.4). Also, you will need the Basis Theorem if you plan to cover Section 4.8 later in the course.

Key Exercises: 9–16.

SECTION 3.1
Introduction to Determinants

Because Chapter 3 is mainly computational, it gives students time to absorb the theory from earlier chapters before they plunge into Chapter 4. The use of Theorem 1 in Section 3.1 enables you to get into the subject quickly. This is important, because a typical first course cannot afford to spend more than a week on determinants. Exercises 33–36 provide the first step in the inductive proof of Theorem 3, in the next section.

SECTION 3.2
Properties of Determinants

The development here is more efficient than some presentations, because the hard work is all packed into the proof of Theorem 3. The characterization of det A via pivots gives a concrete interpretation to the somewhat bizarre calculations with cofactors. Theorems 4 and 6 are important for later work.

SECTION 3.3
Cramer's Rule, Volume, and Linear Transformations

There are several independent topics here from which to choose. Some faculty tend to prefer the geometric interpretation of the determinant, because of its connection with calculus. But who can resist presenting the wonderful proof of Cramer's Rule as developed by Roger Horn and

Charles Johnson? See page 21 of their book, *Matrix Analysis* (Cambridge University Press, Cambridge, 1986).

SECTION 4.1
Vector Spaces and Subspaces

The space $\mathbb{S}$ of signals is used in Sections 4.7 and 4.8. The spaces of polynomials, $\mathbb{P}_n$, are used in many sections of Chapters 4 and 6. This section was designed to avoid the standard exercises in which a student is asked to check ten axioms on an array of sets (some of which may be artificially weird). In fact, Theorem 1 provides the main homework tool in this section for showing that a set is a subspace. Students could be taught how to check the closure axioms, of course, but we think the time for that is better spent elsewhere in a first course.

Key Exercises: 1–18. The exercises here and in the next few sections emphasize $\mathbb{R}^n$, to give students time to absorb the abstract concepts.

SECTION 4.2
Null Spaces, Column Spaces, Row Spaces, and Linear Transformations

This section provides a review of Chapter 1 using the new terminology. Linear transformations are introduced without fuss, because the students are already comfortable with the idea for $\mathbb{R}^n$. The comments for Section 1.8 list the sections in Chapter 4 where linear transformations are used.

Key Exercises: 3–6, 17–24. They are simple but helpful. The idea in Exercises 7–14 is to use Theorem 1 or Theorem 2 to show that a given set is a subspace.

SECTION 4.3
Linearly Independent Sets; Bases

The definition of *basis* is given initially for subspaces because this emphasizes that the basis elements must be in the subspace. Students often overlook this point when the definition is given only for a vector space. The subsection on bases for Nul A and Col A is essential for Section 4.5.

Key Exercises: 1–20.

SECTION 4.4
Coordinate Systems

Section 4.6 depends heavily on this section, as does Section 5.4. It is possible to cover the $\mathbb{R}^n$ parts of the two later sections, however, if the first half of Section 4.4 (and perhaps Example 7) is covered. The linearity of the coordinate mapping is used in Section 5.4 to find the matrix of a transformation relative to two bases. The change-of-coordinates matrix appears in Theorem 8.

Exercise 29 is used in Section 4.6 to show that the change-of-coordinates matrix is invertible.

SECTION 4.5
The Dimension of a Vector Space

Theorem 10 is true because a vector space isomorphic to $\mathbb{R}^n$ has the same algebraic properties as $\mathbb{R}^n$. A proof may not be needed to convince the class. If you skipped Theorem 9 in Section 4.4, you could give a proof such as the one on page 111 of *Introduction to Linear Algebra*, 2nd ed., by Serge Lang (Springer-Verlag, New York, 1986). The Basis Theorem is used in Sections 4.8 and 6.4.

SECTION 4.6
Change of Basis

The row reduction algorithm that produces $\underset{\mathcal{C}\leftarrow\mathcal{B}}{P}$ can also be deduced from Exercise 22 in Section 2.2, by row reducing $[P_{\mathcal{C}} | P_{\mathcal{B}}]$ to $[I \,|\, P_{\mathcal{C}}^{-1} P_{\mathcal{B}}]$. The change-of-coordinates matrix here will be interpreted in Section 5.4 as the matrix of the identity transformation relative to two bases.

SECTION 4.7
Digital Signal Processing

In this section, we work with infinite dimensional vector spaces, subspaces, and linear transformations in the context of digital signals.

SECTION 4.8
Applications to Difference Equations

This is an important section for engineering students and worth two days of class time, if you can arrange that. To spend only one day, you could cover up through Example 5, but assign Example 3 for reading (because it takes some time to present in class). Example 3 is the background for Exercise 34 in Section 6.5.

SECTION 5.1
Eigenvectors and Eigenvalues

The subsection of eigenvectors and difference equations anticipates the discussion that follows in Section 5.2. Also, see Exercises 41 and 42. The T exercises do not involve the full power of an eigenvalue command because students need to make the connection between an eigenspace of A and a null space of $A - \lambda I$. After Section 5.3, students are encouraged to use their matrix programs to generate both eigenvalues and eigenvectors.

SECTION 5.2
The Characteristic Equation

Theorem 3 is review for students who have studied Chapter 3. For others, the theorem lists the facts they need to know about determinants. The idea of similarity, introduced here, is discussed again in Section 5.4, along with eight more exercises. The subsection on dynamical systems provides a short introduction to the subject. The calculations here eliminate the need to supply details for the eigenvector decomposition (2) in Section 5.6.

Exercises 9–14 can be omitted, unless you want your students to have some facility with such problems. If you covered partitioned matrices in Chapter 2, you might connect Exercise 16 in Section 2.4 with Supplementary Exercises 34–36 in Chapter 5.

SECTION 5.3
Diagonalization

This section is essential for the diagonalization of symmetric matrices, in Section 7.1. Theorem 7 is needed for Practice Problem 3 and Exercises 29–32, but is not used elsewhere. A number of exercises in later sections are stated for diagonalizable matrices, because the proofs are simpler for this class of matrices. (See Exercise 26 in Section 5.4, for example.)

SECTION 5.4
Eigenvectors and Linear Transformations

Students should be encouraged to review Section 4.4 before the lecture here. To simplify this section, the focus is on the matrix for a linear transformation relative to a single basis, and no special notation is given to the matrix for a transformation that acts between two different spaces.

SECTION 5.5
Complex Eigenvalues

Although some applications, such as in electrical engineering, naturally involve complex vector spaces, the majority of applications accessible to undergraduates require only real vector spaces. For that reason, we have chosen to discuss complex eigenvalues of only real matrices. A discussion of eigenvalues for complex matrices would take more time than most first courses have available. Moreover, in such discussion, the important case of a real matrix acting on $\mathbb{R}^n$ seldom receives the attention it deserves.

Even if there is no time for this section, students can be encouraged to look at the pictures! Figure 5 in Section 5.6 could be examined, too.

Exercises 27 and 28 provide the proof that a real symmetric matrix has only real eigenvalues. These exercises, together with the (real) Schur factorization described in Supplementary Exercise 34 in Chapter 6 and Exercise 38 in Section 7.1, give you pieces from which you can construct a proof of the spectral theorem for symmetric matrices, if you wish.

SECTION 5.6
Discrete Dynamical Systems

The results here are easier to describe than those for differential equations, and the discrete case deserves more attention in the undergraduate curriculum than it commonly now receives.

A brief treatment of this material could cover Example 1, the discussion following it, and some of Exercises 1–6. The last subsection on the spotted owls could be added to this, requiring only a little hand-waving if you have not discussed complex eigenvalues.

For classes that follow Course Syllabus 2, leave the ecology for outside reading and spend one day on the graphical description of trajectories, change of variable, and complex eigenvalues. If you have some engineers in your class who have already seen state-space methods, their comments may help to spark considerable enthusiasm for this material, particularly the change of variable that decouples a system.

SECTION 5.7
Applications to Differential Equations

This section presumes that students have seen the differential equation $y' = ky$. Sections 5.6 and 5.7 offer an opportunity to contrast discrete and continuous models, which is something engineering texts on signal processing, for example, often do. Students enjoy both sections, although they tend to confuse conditions such as $|\lambda| < 1$ and $\operatorname{Re}\lambda < 0$ when both sections are covered.

SECTION 5.8
Iterative Estimates for Eigenvalues

Example 5 of Section 5.2 and the eigenvector decomposition in Section 5.6 both flow naturally into the power method. When discussing the power method, be careful to *avoid* the following incorrect statements: (1) for sufficiently large k, the vector $A^k\mathbf{x}$ is a good approximation to an eigenvector of A corresponding to the strictly dominant eigenvalue, λ_1; and (2) for large k, the vector $A(A^k\mathbf{x})$ is a good approximation to $\lambda_1(A^k\mathbf{x})$.

In Example 1, the points $A^k\mathbf{x}$ lie on a line parallel to the eigenspace for $\lambda = 2$. So statement (1) is false. In fact, $\mathbf{x} = (.1)\mathbf{v}_1 + \mathbf{v}_2$, where $\mathbf{v}_2$ is an eigenvector for $\lambda = 1$. Thus $A^k\mathbf{x} = 2^k(.1)\mathbf{v}_1 + \mathbf{v}_2$ and $A^{k+1}\mathbf{x} - 2A^k\mathbf{x} = \mathbf{v}_2$ for all k. Hence statement (2) is false.

One or both statements essentially appear in a majority of the leading linear algebra texts that cover the power method. (However, their final algorithms for the power method itself are correct.) The statements are true when the second largest eigenvalue is less than 1 in magnitude. But, in general, the sequence $\{A^k\mathbf{x}\}$ *must* be scaled to a sequence of unit vectors (in some convenient norm). The scaling is not simply a matter of keeping numbers from getting too large or too small.

SECTION 5.9
Applications to Markov Chains

The migration matrix is examined in Section 5.2, where an eigenvector decomposition shows explicitly why the sequence of state vectors $\mathbf{x}_k$ tends to a steady-state vector. The discussion in Section 5.2 leads into the discussion in Section 5.9.

SECTION 6.1
Inner Product, Length, and Orthogonality

The general material on orthogonal complements is essential for later work. Theorem 3 is an important general fact, but it is needed only for Supplementary Exercise 31 at the end of the chapter and in Section 7.4. The optional material on angles is not used later. Exercises 35–39 concern facts that are used later.

SECTION 6.2
Orthogonal Sets

The nonsquare matrices in Theorems 6 and 7 are needed for the QR factorization, in Section 6.4. It is important to emphasize that the term *orthogonal matrix* applies only to certain *square* matrices. The subsection on orthogonal projections not only sets the stage for the general case in Section 6.3, it also provides what is needed for the orthogonal diagonalization exercises in Section 7.1, because none of the eigenspaces there have dimension greater than 2. For this reason, the Gram–Schmidt process (Section 6.4) is not really needed in Chapter 7. Exercises 13 and 14 prepare for Section 6.3.

SECTION 6.3
Orthogonal Projections

Example 1 seems to help students understand Theorem 8. Theorem 8 is needed for the Gram–Schmidt process (but only for a subspace that itself has an orthogonal basis). Theorems 8 and 9 are needed for the discussions of least squares in Sections 6.5 and 6.6. Theorem 10 is used with the QR factorization to provide a good numerical method for solving least-squares problems, in Section 6.5.

Key Exercises: 19 and 20. They lead naturally into the Gram–Schmidt process.

SECTION 6.4
The Gram–Schmidt Process

The QR factorization encapsulates the essential outcome of the Gram–Schmidt process, just as the LU factorization describes the result of a row reduction process. For practical use of linear algebra, the factorizations are more important than the algorithms that produce them. In fact, the Gram–Schmidt process is *not* the appropriate way to compute the QR factorization. (See the Numerical Notes.) For that reason, one should consider deemphasizing the hand calculation of the Gram–Schmidt process, even though it provides easy exam questions.

The Gram–Schmidt process is used in Sections 6.7 and 6.8, in connection with various sets of orthogonal polynomials. The process is mentioned in Sections 7.1 and 7.4, but the one-dimensional projection constructed in Section 6.2 will suffice. The QR factorization is used in an optional subsection of Section 6.5, and it is needed in Supplementary Exercise 23 of Chapter 7 to produce the Cholesky factorization of a positive definite matrix.

SECTION 6.5
Least-Squares Problems

This is a core section, but it need not take a full day. Each example provides a stopping place. Theorem 13 and Example 1 are all that is needed for Section 6.6. Theorem 15, however, gives an illustration of why the QR factorization is important. Example 4 is related to Exercise 17 in Section 6.6.

SECTION 6.6
Machine Learning and Linear Models

All science students will benefit from Example 1. The general linear model and the subsequent examples are aimed at students who may take a multivariate statistics course. In a general sophomore class, that may include more students than one might expect.

SECTION 6.7
Inner Product Spaces

The three types of inner products described here are matched by applications in Section 6.8. It is possible to spend just one day on selected portions of both sections. Example 1 matches the weighted least squares in Section 6.8. Examples 2–6 are applied to trend analysis in Section 6.8. This material is aimed at students who have not had much calculus or who intend to take more than one course in statistics.

For students who have seen some calculus, Example 7 is needed for the Fourier series in Section 6.8. Example 2 is used to motivate the inner product on $C[a, b]$. The Cauchy–Schwarz and triangle inequalities are not used later, but they should be part of the training of every mathematics student.

SECTION 6.8
Applications of Inner Product Spaces

The connections with Section 6.7 have already been identified. For a junior-level course (see Syllabus 2), one might spend three days on the following topics: Theorems 13 and 15 in Section 6.5, plus Examples 1, 3, and 5; Example 1 in Section 6.6; Examples 2 and 7 in Section 6.7, with the motivation for the definite integral; and Fourier series in Section 6.8.

SECTION 7.1
Diagonalization of Symmetric Matrices

Theorems 1 and 2 and the calculations in Examples 2 and 3 are important for the sections that follow. Note that *symmetric matrix* means *real symmetric matrix*, because all matrices in the text have real entries, as mentioned at the beginning of Chapter 7.

Theorem 2 is easily proved for the 2×2 case:

$$\text{If } A = \begin{bmatrix} a & b \\ b & d \end{bmatrix},$$
$$\text{then } \lambda = \frac{1}{2}\left(a + d \pm \sqrt{(a-d)^2 + 4b^2}\right)$$

If $b = 0$, there is nothing to prove. Otherwise, there are two distinct eigenvalues, so A must be diagonalizable. In this case, an eigenvector for λ is $(d - \lambda, -b)$.

SECTION 7.2
Quadratic Forms

This section can provide a good conclusion to the course, because the mathematics here is widely used in applications. For instance, Exercises 31 and 32 can be used to develop the second derivative test for functions of two variables. However, if time permits, some interesting applications still lie ahead. Theorem 4 is used to prove Theorem 6 in Section 7.3, which in turn is used to develop the singular value decomposition.

SECTION 7.3
Constrained Optimization

Theorem 6 is the main result needed in the next two sections. Theorem 7 is mentioned in Example 2 of Section 7.4. Theorem 8 is needed at the very end of Section 7.5. The economic principles in Example 6 may be familiar to students who have had a course in macroeconomics.

SECTION 7.4
The Singular Value Decomposition

This section presents a modern topic of great importance in applications, particularly in computer calculations. Moreover, the singular value decomposition explains much about the structure of matrix transformations. The SVD does for an arbitrary matrix almost what an orthogonal diagonalization does for a symmetric matrix.

SECTION 7.5
Applications to Image Processing and Statistics

The application here has turned out to be of interest to a wide variety of students, including engineers. It is included in Course Syllabus 3, described above, but other classes may only have time to mention the idea briefly.

SECTION 8.1
Affine Combinations

This section introduces special kinds of linear combinations used to describe geometric objects involving two or more variables, including affine sets and translations of subspaces. They set the stage for later applications to computer graphics and can be used in courses on linear programming.

SECTION 8.2
Affine Independence

Affine independence is a restricted type of linear dependence; it provides useful background for a course in computer graphics.

SECTION 8.3
Convex Combinations

Convex combinations are the basic objects studied in Section 8.5 (Polytopes), and they also play a role in computer graphics.

SECTION 8.4
Hyperplanes

This section generalizes the earlier material on lines and planes in $\mathbb{R}^3$ to higher dimensions, in which a hyperplane divides $\mathbb{R}^n$ into two regions.

SECTION 8.5
Polytopes

Polytopes are a class of compact convex sets used to study optimization problems in a variety of fields, including engineering design, linear programming, and business management.

SECTION 8.6
Curves and Surfaces

This section moves beyond lines and planes to curves and surfaces used in engineering and CAD (computer-aided design).

SECTION 9.1
Matrix Games

The two-person "zero-sum" games in this section introduce mathematical strategies for making rational decisions in a simplified but competitive environment. Optimal strategies here depend on geometric properties of convex sets. These results form the foundation for more complex decision problems in business and industrial environments, discussed in the remainder of the chapter.

SECTION 9.2
Linear Programming—Geometric Method

The simplified problems here at first involve only two variables. A solution is found by constructing a *feasible set* in the plane, with a boundary formed by three or more line segments. Each intersection of two lines identifies a point that may be on the boundary of the feasible set. The solution of the problem is given by the coordinates of one of these boundary points.

When three variables are involved, the feasible set is in $\mathbb{R}^3$, bounded by planes that create "faces" on the sides of the feasible set. Each face is a polygonal region that has three or more boundary points, and the geometric situation becomes complicated. So this section only provides practice setting up such a linear programming problem in preparation for Section 9.3.

SECTION 9.3
Linear Programming—Simplex Method

Once the feasible set is in $\mathbb{R}^n$ for $n > 3$, an algebraic method of solution becomes essential. The feasible set is described by two inequalities, $A\mathbf{x} \leq \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$. Suppose $\mathbf{x}_0$ is an extreme point on the feasible set, F. The next step is to examine all other extreme points of F that are connected by one edge to $\mathbf{x}_0$. Of these points, choose one that maximizes the value of the objective function and call it $\mathbf{x}_1$. Then repeat the process and obtain, if possible, another extreme point $\mathbf{x}_2$ at which the value of F increases. Continue the search until no further points are available. This is the Simplex Method, and Section 9.3 has the details that organize the search for the best possible extreme point.

SECTION 9.4
Duality

To each maximization problem in linear programming there corresponds a minimization problem, called the dual problem. Section 9.4 describes this dual problem and how to solve it. The section also shows how *any matrix game* can be solved using the primal and dual versions of a suitable linear programming problem.

SECTION 10.1 (`bit.ly/2nj1HhO`)
Introduction and Examples

This section reviews the terminology from Section 5.9 and introduces additional concepts needed to study Markov chains. New examples of Markov chains are developed for later use, including errors in a multi-stage process (such as signal transmission) and diffusion of gases between two compartments. Random walks on graphs are used as examples of unexpected behavior of Markov chains. Exercises include mathematical models for tennis and volleyball matches.

SECTION 10.2 (`bit.ly/2nj1HhO`)
The Steady-State Vector and Google's PageRank

The steady-state vector for a Markov chain, introduced in Section 5.9, is studied here in more detail. Theorem 1 describes a Markov chain that has a unique steady-state vector. Random internet surfing is then modeled as a random walk on an immense directed graph. The surfing behavior is altered somewhat to allow Theorem 1 to apply to the Markov chain. The steady-state vector for the resulting Markov chain is essentially the vector that Google's PageRank algorithm uses to rank the popularity of web pages.

SECTION 10.3 (`bit.ly/2nj1HhO`)
Communication Classes

This section studies the ways in which Markov chains can fail to have a unique steady-state vector. This study involves partitioning the set of states of a Markov chain into subsets called communication classes. An important special case is an irreducible Markov chain, which has only one communication class. This section also introduces the mean return time for a state—a key idea that is the expected number of steps needed for the chain to return to its starting state.

SECTION 10.4 (`bit.ly/2nj1HhO`)
Classification of States and Periodicity

The states and communication classes of a Markov chain can be classified by whether the chain has a positive probability of never visiting that state (or class) after starting there. States and classes that have this property are called transient; those that do not are called recurrent. Each state or class of a Markov chain also has a period, which is a number that specifies how often the chain may return to that state or class. The recurrence or transience of the states in a Markov chain and their periods determine when the chain has a unique steady-state vector and how to interpret it when it does.

SECTION 10.5 (`bit.ly/2nj1HhO`)
The Fundamental Matrix

The fundamental matrix is derived from the transition matrix of a Markov chain that has both recurrent and transient classes. The entries in the fundamental matrix provide the expected number of visits to a given transient state before the chain is absorbed into a recurrent class. When the chain starts at a transient state, the sum of the entries in a column of the fundamental matrix gives the expected number of time steps until absorption. The fundamental matrix can be used to compute other important quantities, such as the expected transit time between any two states in an irreducible Markov chain.

SECTION 10.6 (`bit.ly/2nj1HhO`)
Markov Chains and Baseball Statistics

This case study shows how to model the offensive action in baseball as a Markov chain. The model uses 28 states that depend on which bases are occupied and how many outs have occurred; appropriate historical statistics are used to compute the transition probabilities. The results in Section 10.5 are then used to predict the number of earned runs a team will score in an inning and to compare the offensive abilities of different players. Exercises suggest how this model can be used to investigate baseball strategies, such as deciding whether to attempt a sacrifice or to attempt to steal a base.

Chapter 10 APPENDIX 1 (`bit.ly/2nj1HhO`)
Proof of Theorem 1

This appendix provides a proof of Theorem 1 in Section 10.2.

Chapter 10 APPENDIX 2 (`bit.ly/2nj1HhO`)
Probability

This appendix provides the background in probability theory that is needed to develop a formal definition of a Markov chain. Proofs are provided for some results that are only stated or justified informally in Chapter 10.

Preface

The response of students and teachers to the first five editions of *Linear Algebra and Its Applications* has been most gratifying. This *Sixth Edition* provides substantial support both for teaching and for using technology in the course. As before, the text provides a modern elementary introduction to linear algebra and a broad selection of interesting classical and leading-edge applications. The material is accessible to students with the maturity that should come from successful completion of two semesters of college-level mathematics, usually calculus.

The main goal of the text is to help students master the basic concepts and skills they will use later in their careers. The topics here follow the recommendations of the original Linear Algebra Curriculum Study Group (LACSG), which were based on a careful investigation of the real needs of the students and a consensus among professionals in many disciplines that use linear algebra. Ideas being discussed by the second Linear Algebra Curriculum Study Group (LACSG 2.0) have also been included. We hope this course will be one of the most useful and interesting mathematics classes taken by undergraduates.

What's New in This Edition

The *Sixth Edition* has exciting new material, examples, and online resources. After talking with high-tech industry researchers and colleagues in applied areas, we added new topics, vignettes, and applications with the intention of highlighting for students and faculty the linear algebraic foundational material for machine learning, artificial intelligence, data science, and digital signal processing.

Content Changes

- Since matrix multiplication is a highly useful skill, we added new examples in Chapter 2 to show how matrix multiplication is used to identify patterns and scrub data. Corresponding exercises have been created to allow students to explore using matrix multiplication in various ways.
- In our conversations with colleagues in industry and electrical engineering, we heard repeatedly how important understanding abstract vector spaces is to their work. After reading the reviewers' comments for Chapter 4, we reorganized the chapter, condensing some of the material on column, row, and null spaces; moving Markov chains to the end of Chapter 5; and creating a new section on signal processing. We view signals

as an infinite dimensional vector space and illustrate the usefulness of linear transformations to filter out unwanted "vectors" (a.k.a. noise), analyze data, and enhance signals.

- By moving Markov chains to the end of Chapter 5, we can now discuss the steady state vector as an eigenvector. We also reorganized some of the summary material on determinants and change of basis to be more specific to the way they are used in this chapter.
- In Chapter 6, we present pattern recognition as an application of orthogonality, and the section on linear models now illustrates how machine learning relates to curve fitting.
- Chapter 9 on optimization was previously available only as an online file. It has now been moved into the regular textbook where it is more readily available to faculty and students. After an opening section on finding optimal strategies to two-person zero-sum games, the rest of the chapter presents an introduction to linear programming—from two-dimensional problems that can be solved geometrically to higher dimensional problems that are solved using the Simplex Method.

Other Changes

- In the high-tech industry, where most computations are done on computers, judging the validity of information and computations is an important step in preparing and analyzing data. In this edition, students are encouraged to learn to analyze their own computations to see if they are consistent with the data at hand and the questions being asked. For this reason, we have added "Reasonable Answers" advice and exercises to guide students.
- We have added a list of projects to the end of each chapter (available online at `bit.ly/30IM8gT` and in MyLab Math). Some of these projects were previously available online and have a wide range of themes from using linear transformations to create art to exploring additional ideas in mathematics. They can be used for group work or to enhance the learning of individual students.
- Free-response writing exercises have been added to MyLab Math, allowing faculty to ask more sophisticated questions online and create a paperless class without losing the richness of discussing how concepts relate to each other and introductory proof writing.
- The electronic interactive textbook has been changed from Wolfram CDF to Wolfram Cloud format. This allows faculty and students to interact with figures and examples on a wider variety of electronic, devices, without the need to install the CDF Player.
- PowerPoint lecture slides have been updated to cover all sections of the text and cover them more thoroughly.

Distinctive Features

Early Introduction of Key Concepts

Many fundamental ideas of linear algebra are introduced within the first seven lectures, in the concrete setting of $\mathbb{R}^n$, and then gradually examined from different points of view. Later generalizations of these concepts appear as natural extensions of familiar ideas, visualized through the geometric intuition developed in Chapter 1. A major achievement of this text is that the level of difficulty is fairly even throughout the course.

A Modern View of Matrix Multiplication

Good notation is crucial, and the text reflects the way scientists and engineers actually use linear algebra in practice. The definitions and proofs focus on the columns of a matrix rather than on the matrix entries. A central theme is to view a matrix–vector product $A\mathbf{x}$ as a linear combination of the columns of A. This modern approach simplifies many arguments, and it ties vector space ideas into the study of linear systems.

Linear Transformations

Linear transformations form a "thread" that is woven into the fabric of the text. Their use enhances the geometric flavor of the text. In Chapter 1, for instance, linear transformations provide a dynamic and graphical view of matrix–vector multiplication.

Eigenvalues and Dynamical Systems

Eigenvalues appear fairly early in the text, in Chapters 5 and 7. Because this material is spread over several weeks, students have more time than usual to absorb and review these critical concepts. Eigenvalues are motivated by and applied to discrete and continuous dynamical systems, which appear in Sections 1.10, 4.8, and 5.9, and in five sections of Chapter 5. Some courses reach Chapter 5 after about five weeks by covering Sections 2.8 and 2.9 instead of Chapter 4. These two optional sections present all the vector space concepts from Chapter 4 needed for Chapter 5.

Orthogonality and Least-Squares Problems

These topics receive a more comprehensive treatment than is commonly found in beginning texts. The original Linear Algebra Curriculum Study Group has emphasized the need for a substantial unit on orthogonality and least-squares problems, because orthogonality plays such an important role in computer calculations and numerical linear algebra and because inconsistent linear systems arise so often in practical work.

Pedagogical Features

Applications

A broad selection of applications illustrates the power of linear algebra to explain fundamental principles and simplify calculations in engineering, computer science, mathematics, physics, biology, economics, and statistics. Some applications appear in separate sections; others are treated in examples and exercises. In addition, each chapter opens with an introductory vignette that sets the stage for some application of linear algebra and provides a motivation for developing the mathematics that follows.

A Strong Geometric Emphasis

Every major concept in the course is given a geometric interpretation, because many students learn better when they can visualize an idea. There are substantially more drawings here than usual, and some of the figures have never before appeared in a linear algebra text. Interactive versions of these figures, and more, appear in the electronic version of the textbook.

Examples

This text devotes a larger proportion of its expository material to examples than do most linear algebra texts. There are more examples than an instructor would ordinarily present in class. But because the examples are written carefully, with lots of detail, students can read them on their own.

Theorems and Proofs

Important results are stated as theorems. Other useful facts are displayed in tinted boxes, for easy reference. Most of the theorems have formal proofs, written with the beginner student in mind. In a few cases, the essential calculations of a proof are exhibited in a carefully chosen example. Some routine verifications are saved for exercises, when they will benefit students.

Practice Problems

A few carefully selected Practice Problems appear just before each exercise set. Complete solutions follow the exercise set. These problems either focus on potential trouble spots in the exercise set or provide a "warm-up" for the exercises, and the solutions often contain helpful hints or warnings about the homework.

Exercises

The abundant supply of exercises ranges from routine computations to conceptual questions that require more thought. A good number of innovative questions pinpoint conceptual difficulties that we have found on student papers over the years. Each exercise set is carefully arranged in the same general order as the text; homework assignments are readily available when only part of a section is discussed. A notable feature of the exercises is their numerical simplicity. Problems "unfold" quickly, so students spend little time on numerical calculations. The exercises concentrate on teaching understanding rather than mechanical calculations. The exercises in the *Sixth Edition* maintain the integrity of the exercises from previous editions, while providing fresh problems for students and instructors.

Exercises marked with the symbol **T** are designed to be worked with the aid of a "matrix program" (a computer program, such as MATLAB, Maple, Mathematica, MathCad, or Derive, or a programmable calculator with matrix capabilities, such as those manufactured by Texas Instruments).

True/False Questions

To encourage students to read all of the text and to think critically, we have developed over 300 simple true/false questions that appear throughout the text, just after the computational problems. They can be answered directly from the text, and they prepare students for the conceptual problems that follow. Students appreciate these questions-after they get used to the importance of reading the text carefully. Based on class testing and discussions with students, we decided not to put the answers in the text. (The *Study Guide* tells the students where to find the answers to the odd-numbered questions.) An additional 150 true/false questions (mostly at the ends of chapters) test understanding of the material. The text does provide simple T/F answers to most of these supplementary exercises, but it omits the justifications for the answers (which usually require some thought).

Writing Exercises

An ability to write coherent mathematical statements in English is essential for all students of linear algebra, not just those who may go to graduate school in mathematics. The text includes many exercises for which a written justification is part of the answer. Conceptual exercises that require a short proof usually contain hints that help a student get started. For all odd-numbered writing exercises, either a solution is included at the back of the text or a hint is provided and the solution is given in the *Study Guide,* described below.

Projects

A list of projects (available online at `bit.ly/30IM8gT`) have been identified at the end of each chapter. They can be used by individual students or in groups. These projects provide the opportunity for students to explore fundamental concepts and applications in more detail. Two of the projects even encourage students to engage their creative side and use linear transformations to build artwork.

Reasonable Answers

Many of our students will enter a workforce where important decisions are being made based on answers provided by computers and other machines. The Reasonable Answers boxes and exercises help students develop an awareness of the need to analyze their answers for correctness and accuracy.

Computational Topics

The text stresses the impact of the computer on both the development and practice of linear algebra in science and engineering. Frequent Numerical Notes draw attention to issues in computing and distinguish between theoretical concepts, such as matrix inversion, and computer implementations, such as LU factorizations.

Acknowledgments

David Lay was grateful to many people who helped him over the years with various aspects of this book. He was particularly grateful to Israel Gohberg and Robert Ellis for more than fifteen years of research collaboration, which greatly shaped his view of linear algebra. And he was privileged to be a member of the Linear Algebra Curriculum Study Group along with David Carlson, Charles Johnson, and Duane Porter. Their creative ideas about teaching linear algebra have influenced this text in significant ways. He often spoke fondly of three good friends who guided the development of the book nearly from the beginning—giving wise counsel and encouragement—Greg Tobin, publisher; Laurie Rosatone, former editor; and William Hoffman, former editor.

Judi and Steven have been privileged to work on recent editions of Professor David Lay's linear algebra book. In making this revision, we have attempted to maintain the basic approach and the clarity of style that has made earlier editions popular with students and faculty. We thank Eric Schulz for sharing his considerable technological and pedagogical expertise in the creation of the electronic textbook. His help and encouragement were essential in bringing the figures and examples to life in the Wolfram Cloud version of this textbook.

Mathew Hudelson has been a valuable colleague in preparing the *Sixth Edition;* he is always willing to brainstorm about concepts or ideas and test out new writing and exercises. He contributed the idea for new vignette for Chapter 3 and the accompanying project. He has helped with new exercises throughout the text. Harley Weston has provided Judi with many years of good conversations about how, why, and who we appeal to when we present mathematical material in different ways. Katerina Tsatsomeros' artistic side has been a definite asset when we needed artwork to transform (the fish and the sheep), improved writing in the new introductory vignettes, or information from the perspective of college-age students.

We appreciate the encouragement and shared expertise from Nella Ludlow, Thomas Fischer, Amy Johnston, Cassandra Seubert, and Mike Manzano. They provided information about important applications of linear algebra and ideas for new examples and exercises. In particular, the new vignettes and material in Chapters 4 and 6 were inspired by conversations with these individuals.

We are energized by Sepideh Stewart and the other new Linear Algebra Curriculum Study Group (LACSG 2.0) members: Sheldon Axler, Rob Beezer, Eugene Boman, Minerva Catral, Guershon Harel, David Strong, and Megan Wawro. Initial meetings of this group have provided valuable guidance in revising the *Sixth Edition.*

We sincerely thank the following reviewers for their careful analyses and constructive suggestions:

Maila C. Brucal-Hallare, *Norfolk State University*
Kristen Campbell, *Elgin Community College*
Charles Conrad, *Volunteer State Community College*
R. Darrell Finney, *Wilkes Community College*
Xiaofeng Gu, *University of West Georgia*
Jeong Mi-Yoon, *University of Houston–Downtown*
Michael T. Muzheve, *Texas A&M U.–Kingsville*
Iason Rusodimos, *Perimeter C. at Georgia State U.*
Rebecca Swanson, *Colorado School of Mines*
Casey Wynn, *Kenyon College*
Taoye Zhang, *Penn State U.–Worthington Scranton*
Steven Burrow, *Central Texas College*
J. S. Chahal, *Brigham Young University*
Kevin Farrell, *Lyndon State College*
Chris Fuller, *Cumberland University*
Jeffrey Jauregui, *Union College*
Christopher Murphy, *Guilford Tech. C.C.*
Charles I. Odion, *Houston Community College*
Desmond Stephens, *Florida Ag. and Mech. U.*
Jiyuan Tao, *Loyola University–Maryland*
Amy Yielding, *Eastern Oregon University*
Houlong Zhuang, *Arizona State University*

We appreciate the proofreading and suggestions provided by John Samons and Jennifer Blue. Their careful eye has helped to minimize errors in this edition.

We thank Kristina Evans, Phil Oslin, and Jean Choe for their work in setting up and maintaining the online homework to accompany the text in MyLab Math, and for continuing to work with us to improve it. The reviews of the online homework done by Joan Saniuk, Robert Pierce, Doron Lubinsky and Adriana Corinaldesi were greatly appreciated. We also thank the faculty at University of California Santa Barbara, University of Alberta, Washington State University and the Georgia Institute of Technology for their feedback on the MyLab Math course. Joe Vetere has provided much appreciated technical help with the *Study Guide* and *Instructor's Solutions Manual*.

We thank Jeff Weidenaar, our content manager, for his continued careful, well-thought-out advice. Project Manager Ron Hampton has been a tremendous help guiding us through the production process. We are also grateful to Stacey Sveum and Rosemary Morton, our marketers, and Jon Krebs, our editorial associate, who have also contributed to the success of this edition.

Steven R. Lay and Judi J. McDonald

MyLab Math for *Linear Algebra and Its Applications*

Lay, Lay, McDonald

(access code required)

MyLab Math features hundreds of assignable algorithmic exercises that mirror those in the text. It is tightly integrated with each author team's style, offering a range of author-created resources, so your students have a consistent experience.

eText with Interactive Figures

The eText includes **Interactive Figures**, created by author Judi McDonald, that bring the geometry of linear algebra to life. Students can manipulate figures and experiment with matrices to provide a deeper geometric understanding of key concepts and examples.

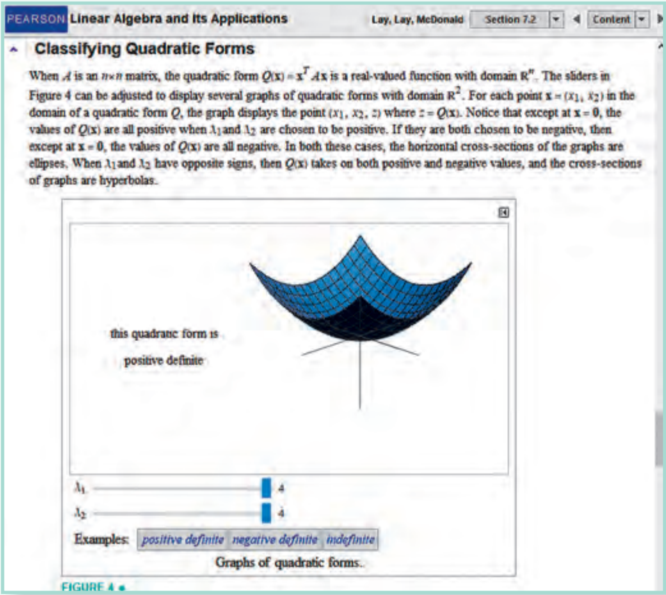

Teaching with Interactive Figures

Interactive Figure files are available within MyLab Math to use as a teaching tool for classroom demonstrations. Instructors can illustrate concepts that are difficult for students to visualize, leading to greater conceptual understanding.

pearson.com/mylab/math

Supporting Instruction

MyLab Math provides resources to help you assess and improve student results and unparalleled flexibility to create a course tailored to you and your students.

PowerPoint® Lecture Slides

Fully editable PowerPoint slides are available for all sections of the text. The slides include definitions, theorems, examples and solutions. When used in the classroom, these slides allow the instructor to focus on teaching, rather than writing on the board. PowerPoint slides are available to students (within the Video and Resource Library in MyLab Math) so that they can follow along.

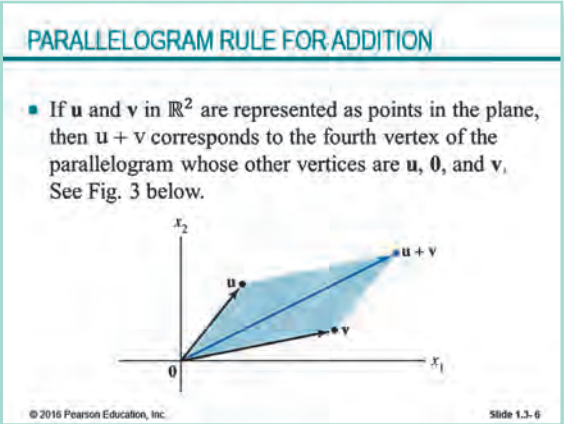

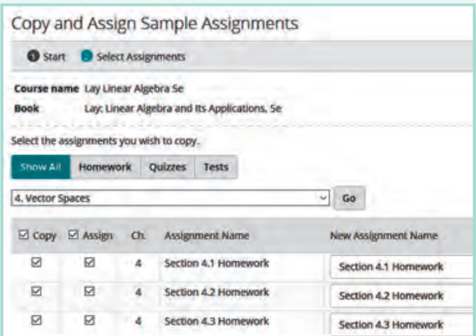

Sample Assignments

Sample Assignments are crafted to maximize student performance in the course. They make course set-up easier by giving instructors a starting point for each section.

Comprehensive Gradebook

The gradebook includes enhanced reporting functionality, such as item analysis and a reporting dashboard to enable you to efficiently manage your course. Student performance data are presented at the class, section, and program levels in an accessible, visual manner so you'll have the information you need to keep your students on track.

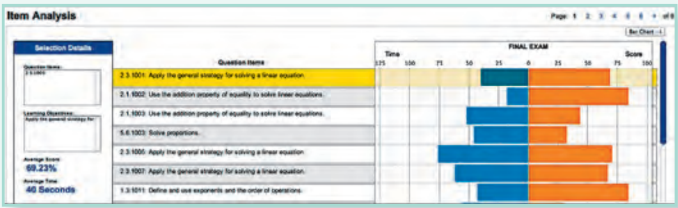

pearson.com/mylab/math

Resources for Success

Pearson MyLab

Instructor Resources

Online resources can be downloaded from MyLab Math or from www.pearson.com.

Instructor's Edition

ISBN 013588280X / 9780135882801
The instructor's edition includes brief answers to all exercises and provides instructor teaching and course structure suggestions, including sample syllabi.

Instructor's Solution Manual

Includes fully worked solutions to all exercises in the text and teaching notes for many sections.

PowerPoint® Lecture Slides

These fully editable lecture slides are available for all sections of the text.

Instructor's Technology Manuals

Each manual provides detailed guidance for integrating technology throughout the course, written by faculty who teach with the software and this text. Available For MATLAB, Maple, Mathematica, and Texas Instruments graphing calculators.

TestGen®

TestGen (www.pearsoned.com/testgen) enables instructors to build, edit, print, and administer tests using a computerized bank of questions developed to cover all the objectives of the text.

Student Resources

Additional resources to enhance student success. All resources can be downloaded from MyLab Math.

Study Guide

Provides detailed worked-out solutions to every third odd-numbered exercise. Also, a complete explanation is provided whenever an odd-numbered writing exercise has a Hint in the answers. Special subsections of the *Study Guide* suggest how to master key concepts of the course. Frequent "Warnings" identify common errors and show how to prevent them. MATLAB boxes introduce commands as they are needed. Appendixes in the *Study Guide* provide comparable information about Maple, Mathematica, and TI graphing calculators. Available within MyLab math and also available for purchase separately using ISBN 9780135851234.

Getting Started with Technology

A quick-start guide for students to the technology they may use in this course. Available for MATLAB, Maple, Mathematica, or Texas Instrument graphing calculators. Downloadable from MyLab Math.

Projects

Exploratory projects, written by experienced faculty members, invite students to discover applications of linear algebra.