UCLA COURSE FIELD NOTES

MATH 33A

Linear Algebra and Applications

Systems, vector spaces, transformations, eigenstructure, orthogonality, and least squares

How these notes are written. I wrote each chapter the way I wish the material had been explained the first time: the real idea, the point where I usually get stuck, the simplest mental model that still stays honest, and a worked decision instead of a polished answer appearing from nowhere. Every chapter closes with practice and a fully written solution. If a sentence sounds impressive but does not help me solve or explain something, it does not belong here.

Daniel Mastick

Curriculum map, working notes, practice, and solutions

Course map

Week	Core thread	What should be solid by the end
1	**Linear systems**	Row reduction, echelon form, pivots, free variables, and consistency.
2	**Matrix algebra**	Inverse matrices, factorization, and structural identities.
3	**Vector spaces**	Subspaces, span, independence, basis, dimension, and coordinates.
4	**Linear transformations**	Kernel, image, matrix representations, and composition.
5	**Determinants**	Volume scaling, invertibility, orientation, and cofactor structure.
6	**Eigenvalues**	Invariant directions, characteristic polynomials, and eigenspaces.
7	**Diagonalization**	Powers, recurrences, dynamical systems, and matrix exponentials.
8	**Orthogonality**	Projections, Gram-Schmidt, orthonormal bases, and QR.
9	**Least squares**	Normal equations, data fitting, and residual geometry.
10	**Symmetric matrices**	Spectral theorem and quadratic forms.

Scope anchor: https://catalog.registrar.ucla.edu/course/2025/MATH33A. The sequence is aligned to the official catalog description and expanded into a practical ten-week study plan.

How I would use this packet

Treat each numbered section as one chapter. Read the formal explanation, pause at the student interpretation, and see whether the easy version still says the same thing. Rework the example without looking, then cover the solution in the chapter practice box. The cumulative set at the end is deliberately mixed: knowledge becomes durable when I have to choose the tool instead of being told which chapter I am in.

Contents

CHAPTER 1 Systems and row reduction

How I actually think about it

Row reduction preserves the solution set while making the constraints easier to read. Pivots tell me which variables are controlled; free variables measure the remaining degrees of freedom.

The easy version

Simplify the equations without changing which points satisfy all of them.

A linear system is one object with three views: equations, an augmented matrix, and an intersection of geometric constraints. Elementary row operations preserve the solution set.

Reduced row echelon form reveals:

- pivot columns and basic variables;
- nonpivot columns and free variables;
- inconsistency through a row $[0 \ \cdots \ 0 \mid b]$ with $b \neq 0$.

Worked example

If reduction gives

$$\left[\begin{array}{ccc|c} 1 & 0 & 2 & 3 \\ 0 & 1 & -1 & 4 \end{array}\right],$$

let $x_3 = t$. Then $(x_1, x_2, x_3) = (3 - 2t, 4 + t, t)$, or

$$\mathbf{x} = \begin{bmatrix} 3 \\ 4 \\ 0 \end{bmatrix} + t \begin{bmatrix} -2 \\ 1 \\ 1 \end{bmatrix}.$$

The affine part chooses one solution; the direction spans the null space.

The idea

For $A\mathbf{x} = \mathbf{b}$, every solution is one particular solution plus a null-space solution. This pattern returns in differential equations, optimization, and numerical methods.

Rebuild the chapter from first principles

The claim I need to own is this: For $A\mathbf{x} = \mathbf{b}$, every solution is one particular solution plus a null-space solution. This pattern returns in differential equations, optimization, and numerical methods. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of systems and row reduction. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach systems and row reduction on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from the course foundations to matrix algebra and invertibility. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

Solved chapter practice

Practice. Solve $x + y + z = 2$, $x + 2y + 3z = 5$, $2x + 3y + 4z = 7$.
Solution. Subtract the first from the second: $y + 2z = 3$. Subtract twice the first from the third: $y + 2z = 3$, the same equation. Let $z = t$; then $y = 3 - 2t$ and $x = 2 - y - z = t - 1$. Solutions

are $(t-1, 3-2t, t)$.

Mastery practice A. Reconstruct systems and row reduction from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a systems and row reduction problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 2 Matrix algebra and invertibility

Rebuild the chapter from first principles

The claim I need to own is this: The model in Matrix algebra and invertibility controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of matrix algebra and invertibility. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.

- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach matrix algebra and invertibility on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from systems and row reduction to vector spaces, basis, and dimension. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: a matrix represents a linear action, and multiplication means composition. Invertibility has many equivalent signals because they all say the action loses no direction and reaches every target.

The easy version

An invertible matrix is a reversible machine: no inputs collapse together and every output can be reached.

Matrix multiplication composes linear actions. Order matters because performing B then A is AB.

For an $n \times n$ matrix, the following are equivalent:

- A is invertible;
- $A\mathbf{x} = \mathbf{b}$ has exactly one solution for every $\mathbf{b}$;
- the columns form a basis of $\mathbb{R}^n$;
- $\operatorname{rref}(A) = I$;
- $\det A \neq 0$;
- 0 is not an eigenvalue.

The inverse is a structural certificate, not usually the best computational method. To solve one system, elimination is often cheaper and more stable than explicitly forming A^{-1}.

Solved chapter practice

Practice. Determine whether $A = \begin{pmatrix} 1 & 2 \\ 2 & 4 \end{pmatrix}$ is invertible and explain geometrically.

Solution. $\det A = 4 - 4 = 0$, so A is not invertible. Its columns are dependent; both lie on the same line, so the transformation collapses the plane onto a line and cannot be reversed.

Mastery practice A. Reconstruct matrix algebra and invertibility from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a matrix algebra and invertibility problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 3 Vector spaces, basis, and dimension

Rebuild the chapter from first principles

The claim I need to own is this: The model in Vector spaces, basis, and dimension controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of vector spaces, basis, and dimension. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach vector spaces, basis, and dimension on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from matrix algebra and invertibility to linear transformations. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

A vector space is a collection closed under the linear moves I am allowed to make. A basis is a minimal coordinate system: it spans everything without redundancy, and dimension counts its independent directions.

The easy version

A basis is the smallest set of building blocks that can make every vector in the space exactly once.

A subspace contains 0 and is closed under addition and scalar multiplication. The span of vectors is the smallest subspace containing them. Independence means the only solution of

$$c_1\mathbf{v}_1 + \cdots + c_k\mathbf{v}_k = 0$$

is all $c_i = 0$.

A basis is independent and spanning. It is a lossless coordinate system for the space. Dimension is the number of basis vectors.

The move

To find a basis for a column space, row-reduce to locate pivot columns, then take those columns from the original matrix. Row operations preserve dependencies among columns but do not preserve the actual column space.

Rank-nullity for $A \in \mathbb{R}^{m\times n}$ is

$$\boxed{\operatorname{rank}(A) + \operatorname{nullity}(A) = n.}$$

It accounts for every input direction: each either survives into the image or disappears into the kernel.

Solved chapter practice

Practice. Find a basis for the span of $(1,0,1),(0,1,1),(1,1,2)$.
Solution. The third vector is the sum of the first two, so it adds no direction. The first two are independent, hence $\{(1,0,1),(0,1,1)\}$ is a basis and the span has dimension 2.

Mastery practice A. Reconstruct vector spaces, basis, and dimension from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a vector spaces, basis, and dimension problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 4 Linear transformations

Rebuild the chapter from first principles

The claim I need to own is this: The model in Linear transformations controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of linear transformations. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substi-

tuting values. This is where I make the assumption visible.

3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach linear transformations on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from vector spaces, basis, and dimension to determinants. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

Linearity means the transformation respects addition and scaling. Once I know what happens to basis vectors, I know what happens to every vector because every vector is a linear combination of that basis.

The easy version

Tell me where the coordinate axes go, and linearity fills in the rest of the map.

$T : V \to W$ is linear when it preserves addition and scalar multiplication. The kernel is inputs mapped to zero; the image is reachable outputs.

Once bases are chosen, a linear map is a matrix. Its columns are the images of basis vectors. Composition becomes multiplication, and a change of basis becomes a similarity transform.

Worked example

Reflection across the x-axis maps $\mathbf{e}_1 \mapsto \mathbf{e}_1$ and $\mathbf{e}_2 \mapsto -\mathbf{e}_2$, so its standard matrix is

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

Rotation by θ has columns equal to the rotated standard basis:

$$\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}.$$

Solved chapter practice

Practice. If $T(1,0) = (2,1)$ and $T(0,1) = (-1,3)$, find the matrix of T and $T(4,-2)$.

Solution. The images of basis vectors are the columns, so $[T] = \begin{pmatrix} 2 & -1 \\ 1 & 3 \end{pmatrix}$. Then $T(4,-2) = (10,-2)$.

Mastery practice A. Reconstruct linear transformations from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a linear transformations problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 5 Determinants

Rebuild the chapter from first principles

The claim I need to own is this: The model in Determinants controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of determinants. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach determinants on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from linear transformations to eigenvalues and eigenspaces. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

The determinant is an oriented scale factor for area or volume. Its algebraic rules make more sense when I remember that swapping rows flips orientation and dependent columns flatten volume to zero.

The easy version

The determinant tells how much a linear map stretches space and whether it flips orientation.

$|\det A|$ is volume scale; its sign records orientation. Key laws:

$$\det(AB) = \det A \det B, \quad \det(A^T) = \det A, \quad \det(A^{-1}) = (\det A)^{-1}.$$

Swapping rows negates the determinant, scaling a row scales it, and adding a multiple of one row to another leaves it unchanged.

Sanity check

The determinant is one number. It answers invertibility and volume questions, but it does not expose the null space or solve the system by itself.

Solved chapter practice

Practice. Find the area scale factor of $A = \begin{pmatrix} 3 & 1 \\ 2 & 2 \end{pmatrix}$.

Solution. $\det A = 3 \cdot 2 - 1 \cdot 2 = 4$. Areas are multiplied by $|4| = 4$, and the positive sign means orientation is preserved.

Mastery practice A. Reconstruct determinants from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a determinants problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 6 Eigenvalues and eigenspaces

Rebuild the chapter from first principles

The claim I need to own is this: The model in Eigenvalues and eigenspaces controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of eigenvalues and eigenspaces. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach eigenvalues and eigenspaces on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from determinants to diagonalization and dynamics. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

An eigenvector is a direction the transformation does not turn away from itself; it only stretches, shrinks, or reverses it. I solve $(A - \lambda I)v = 0$ because an eigenvalue makes that system acquire a nonzero nullspace.

The easy version

Most arrows turn under a matrix. Eigenvectors are the special arrows that stay on their own line.

An eigenvector keeps its direction under A:

$$A\mathbf{v} = \lambda\mathbf{v}.$$

Eigenvalues satisfy $\det(A - \lambda I) = 0$. For each λ, the eigenspace is $\ker(A - \lambda I)$.

Trace is the sum of eigenvalues and determinant is their product, counted with algebraic multiplicity.

Worked example

For $A = \begin{bmatrix} 2 & 1 \\ 0 & 3 \end{bmatrix}$, eigenvalues are $2, 3$. For $\lambda = 2$, an eigenvector is $\langle 1, 0 \rangle$. For $\lambda = 3$, one is $\langle 1, 1 \rangle$.

Solved chapter practice

Practice. Find eigenvalues and eigenvectors of $A = \begin{pmatrix} 2 & 1 \\ 0 & 3 \end{pmatrix}$.

Solution. The triangular entries give eigenvalues $2, 3$. For $\lambda = 2$, $(A - 2I)v = 0$ gives $y = 0$, so $\mathrm{span}(1, 0)$. For $\lambda = 3$, $-x + y = 0$, so $\mathrm{span}(1, 1)$.

Mastery practice A. Reconstruct eigenvalues and eigenspaces from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a eigenvalues and eigenspaces problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 7 Diagonalization and dynamics

Rebuild the chapter from first principles

The claim I need to own is this: The model in Diagonalization and dynamics controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of diagonalization and dynamics. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach diagonalization and dynamics on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from eigenvalues and eigenspaces to orthogonality and projection. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: diagonalization changes to an eigenvector coordinate system where repeated multiplication becomes simple scaling. The matrix P is not a magic accessory; its columns are the directions that make the transformation simple.

The easy version

Rotate the coordinate system to the matrix's favorite directions, where the action becomes independent stretching.

A is diagonalizable if it has a basis of eigenvectors. Put them into P and eigenvalues into D:

$$A = PDP^{-1}, \qquad A^k = PD^kP^{-1}.$$

Distinct eigenvalues guarantee independent eigenvectors, but repeated eigenvalues require checking geometric multiplicity.

This turns repeated application into scalar powers. In $\mathbf{x}_{k+1} = A\mathbf{x}_k$, long-term behavior is dominated by eigenvalues of largest magnitude, unless the initial state has no component in those eigendirections.

Common miss

Algebraic multiplicity belongs to the characteristic polynomial. Geometric multiplicity is the dimension of the eigenspace. Diagonalization requires enough eigenvectors, not merely enough eigenvalues written with repetition.

Solved chapter practice

Practice. Using the previous matrix, compute $A^n(1,1)$.
Solution. $(1,1)$ is already an eigenvector for $\lambda = 3$. Therefore each application multiplies it by 3, so $A^n(1,1) = 3^n(1,1)$.

Mastery practice A. Reconstruct diagonalization and dynamics from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a diagonalization and dynamics problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 8 Orthogonality and projection

Rebuild the chapter from first principles

The claim I need to own is this: The model in Orthogonality and projection controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of orthogonality and projection. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach orthogonality and projection on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from diagonalization and dynamics to least squares. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: orthogonality makes coordinates independent in a geometric sense. Projection finds the closest point in a subspace because the leftover error is perpendicular to every allowable direction.

The easy version

Drop a perpendicular shadow onto the subspace; that shadow is the closest vector in it.

Orthogonal vectors satisfy $\mathbf{u} \cdot \mathbf{v} = 0$. An orthonormal basis makes coordinates immediate:

$$\mathbf{x} = \sum_i (\mathbf{x} \cdot \mathbf{q}_i)\mathbf{q}_i.$$

Projection onto a subspace with orthonormal basis columns in Q is

$$\boxed{\operatorname{proj}_W \mathbf{b} = QQ^T\mathbf{b}.}$$

Gram-Schmidt builds an orthogonal basis by subtracting previous projections. Normalizing yields Q and supports a QR factorization $A = QR$.

Solved chapter practice

Practice. Project $b = (2, 1, 3)$ onto the line spanned by $u = (1, 0, 1)$.
Solution. $\operatorname{proj}_u b = \frac{b \cdot u}{u \cdot u} u = \frac{5}{2}(1, 0, 1) = (5/2, 0, 5/2)$.

Mastery practice A. Reconstruct orthogonality and projection from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a orthogonality and projection problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 9 Least squares

Rebuild the chapter from first principles

The claim I need to own is this: The model in Least squares controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of least squares. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach least squares on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from orthogonality and projection to symmetric matrices and quadratic forms. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: when $Ax = b$ is inconsistent, least squares does not pretend the equations all hold. It chooses Ax closest to b, forcing the residual to be perpendicular to the column space and producing the normal equations.

The easy version

If no exact fit exists, choose the prediction whose error vector points straight away from the model space.

When $A\mathbf{x} = \mathbf{b}$ is inconsistent, choose $\hat{\mathbf{x}}$ minimizing $\|A\mathbf{x} - \mathbf{b}\|$. The residual must be orthogonal to the column space:

$$A^T(A\hat{\mathbf{x}} - \mathbf{b}) = 0.$$

Thus the normal equations are

$$\boxed{A^T A\hat{\mathbf{x}} = A^T\mathbf{b}.}$$

If columns are independent, $A^T A$ is invertible.

Least-squares fitting is projection: $A\hat{\mathbf{x}}$ is the closest point in $\mathrm{Col}(A)$ to the data vector $\mathbf{b}$.

Solved chapter practice

Practice. Fit a line $y = a + bx$ through $(0, 1), (1, 2), (2, 2)$ by least squares.

Solution. $A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{pmatrix}$, $y = (1, 2, 2)^T$. Normal equations are $\begin{pmatrix} 3 & 3 \\ 3 & 5 \end{pmatrix}(a, b)^T = (5, 6)^T$. Solving gives $b = 1/2$, $a = 7/6$.

Mastery practice A. Reconstruct least squares from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a least squares problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 10 Symmetric matrices and quadratic forms

Rebuild the chapter from first principles

The claim I need to own is this: The model in Symmetric matrices and quadratic forms controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of symmetric matrices and quadratic forms. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach symmetric matrices and quadratic forms on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from least squares to cumulative practice. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: symmetric matrices are unusually well behaved: real eigenvalues and orthogonal eigenvectors let them be diagonalized without distorting angles. A quadratic form becomes a sum of signed squares in that eigenbasis.

The easy version

A symmetric matrix has perpendicular favorite directions, so its curved surface can be rotated to line up with clean axes.

The spectral theorem says a real symmetric matrix has an orthonormal eigenbasis:

$$A = QDQ^T.$$

For a quadratic form $q(\mathbf{x}) = \mathbf{x}^T A\mathbf{x}$, orthogonal coordinates remove cross terms. The signs of eigenvalues classify the form: all positive gives positive definite; mixed signs gives indefinite.

Solved chapter practice

Practice. Classify $q(x, y) = 2x^2 + 2xy + 2y^2$ as positive definite, negative definite, or indefinite.

Solution. The symmetric matrix is $\begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$ with eigenvalues 3 and 1, both positive. Thus $q > 0$ for every nonzero vector and is positive definite.

Mastery practice A. Reconstruct symmetric matrices and quadratic forms from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a symmetric matrices and quadratic forms problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 11 Cumulative practice

Rebuild the chapter from first principles

The claim I need to own is this: The model in Cumulative practice controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of cumulative practice. List the known quantities, the unknown, the units or types, and any hidden constraint.

2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.

3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.

4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach cumulative practice on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from symmetric matrices and quadratic forms to written solutions. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

1. Solve $x + y + z = 2$, $2x + y - z = 1$, $x + 2y + 3z = 4$.
2. Find bases for the null space and column space of $A = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \end{bmatrix}$.
3. Find the matrix of $T(x, y) = (x + 2y, 3x - y)$ and determine whether T is invertible.
4. Find $\det \begin{bmatrix} 1 & 2 & 0 \\ 0 & 3 & 4 \\ 0 & 0 & -2 \end{bmatrix}$ and interpret its magnitude.
5. Diagonalize $A = \begin{bmatrix} 4 & 1 \\ 0 & 2 \end{bmatrix}$.
6. Orthogonally project $\mathbf{b} = \langle 2, 1, 3 \rangle$ onto $\operatorname{span}\{\langle 1, 1, 0 \rangle\}$.
7. Find the least-squares line $y = a + bx$ through $(0, 1), (1, 2), (2, 2)$.
8. Classify $q(x, y) = 3x^2 + 4xy + 3y^2$.

CHAPTER 12 Written solutions

Rebuild the chapter from first principles

The claim I need to own is this: The model in Written solutions controls the work. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of written solutions. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach written solutions on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from cumulative practice to the cumulative course model. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

1. System

Elimination gives $z = 1$, $y = 0$, $x = 1$. Direct substitution checks all three equations.

2. Fundamental subspaces

The second row is twice the first. Equation $x_1 + 2x_2 + 3x_3 = 0$ gives

$$\mathbf{x} = s\langle -2, 1, 0\rangle + t\langle -3, 0, 1\rangle.$$

These form a null-space basis. The only pivot column is the first original column, so $\{\langle 1, 2\rangle\}$ is a column-space basis.

3. Transformation

$[T] = \begin{bmatrix} 1 & 2 \\ 3 & -1 \end{bmatrix}$. Its determinant is $-7 \neq 0$, so T is invertible.

4. Determinant

For an upper triangular matrix, determinant is the product of diagonal entries: $1 \cdot 3 \cdot (-2) = -6$. Volumes scale by 6 and orientation reverses.

5. Diagonalization

Eigenvalues are 4 and 2. Corresponding eigenvectors may be $\langle 1, 0\rangle$ and $\langle 1, -2\rangle$. Thus

$$P = \begin{bmatrix} 1 & 1 \\ 0 & -2 \end{bmatrix}, \qquad D = \begin{bmatrix} 4 & 0 \\ 0 & 2 \end{bmatrix}, \qquad A = PDP^{-1}.$$

6. Projection

With $\mathbf{v} = \langle 1, 1, 0\rangle$,

$$\operatorname{proj}_{\mathbf{v}} \mathbf{b} = \frac{\mathbf{b} \cdot \mathbf{v}}{\mathbf{v} \cdot \mathbf{v}} \mathbf{v} = \frac{3}{2}\langle 1, 1, 0\rangle.$$

7. Least-squares line

Use $A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix}$ and $\mathbf{y} = \langle 1, 2, 2\rangle$. Normal equations give

$$\begin{bmatrix} 3 & 3 \\ 3 & 5 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} 5 \\ 6 \end{bmatrix},$$

so $a = 7/6$, $b = 1/2$.

8. Quadratic form

The symmetric matrix is $\begin{bmatrix} 3 & 2 \\ 2 & 3 \end{bmatrix}$, with eigenvalues 5 and 1, both positive. The form is positive definite.

Closing frame

Linear algebra is the study of what survives when a system is constrained to be linear. Kernels record what disappears, images record what can be reached, eigenvectors record what keeps its direction, and orthogonality gives the cleanest approximation when an exact solution is unavailable.