UCLA COURSE FIELD NOTES

MATH 32A

Calculus of Several Variables

Geometry, vector motion, partial derivatives, gradients, and constrained optimization

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	**Vectors and geometry**	Dot products, cross products, projections, lines, planes, and distance.
2	**Vector-valued functions**	Space curves, velocity, acceleration, arc length, and curvature.
3	**Multivariable functions**	Graphs, level sets, domains, and limits in more than one direction.
4	**Partial derivatives**	First and higher partials; differentiability and local change.
5	**Tangent planes**	Linearization, differentials, and multivariable error propagation.
6	**Chain rules**	Dependency trees, implicit differentiation, and coordinate changes.
7	**Gradient**	Directional derivatives, steepest ascent, and normal vectors.
8	**Unconstrained extrema**	Critical points and the second derivative test.
9	**Constrained extrema**	Lagrange multipliers and boundary-aware comparison.
10	**Synthesis**	Mixed geometry, approximation, and optimization.

Scope anchor: https://catalog.registrar.ucla.edu/course/2025/MATH32A. 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 Vectors turn geometry into algebra

Rebuild the chapter from first principles

The claim I need to own is this: The model in Vectors turn geometry into algebra 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 vectors turn geometry into algebra. 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 vectors turn geometry into algebra 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 vector-valued motion. 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 is a directed change, and dot and cross products answer different geometric questions. I use the dot product for alignment and projection; I use the cross product when I need a perpendicular direction or an oriented area.

The easy version

The dot product measures how much two arrows point together; the cross product builds an arrow perpendicular to both.

A vector carries magnitude and direction. In $\mathbb{R}^3$, $\mathbf{v} = \langle v_1, v_2, v_3\rangle$ and $\|\mathbf{v}\| = \sqrt{\mathbf{v}\cdot\mathbf{v}}$.

The dot product measures alignment:

$$\mathbf{u}\cdot\mathbf{v} = \|\mathbf{u}\|\|\mathbf{v}\|\cos\theta, \qquad \operatorname{proj}_{\mathbf{v}}\mathbf{u} = \frac{\mathbf{u}\cdot\mathbf{v}}{\mathbf{v}\cdot\mathbf{v}}\mathbf{v}.$$

The cross product produces a perpendicular vector with magnitude equal to parallelogram area:

$$\|\mathbf{u}\times\mathbf{v}\| = \|\mathbf{u}\|\|\mathbf{v}\|\sin\theta.$$

Worked example

The plane through $P = (1, 0, 2)$ with normal $\mathbf{n} = \langle 2, -1, 3\rangle$ is

$$2(x-1) - y + 3(z-2) = 0,$$

or $2x - y + 3z = 8$. A normal vector is the plane's compressed representation: one point plus one perpendicular direction determines the whole plane.

Distance from point Q to plane $\mathbf{n}\cdot(\mathbf{x}-\mathbf{p}) = 0$ is

$$d = \frac{|\mathbf{n}\cdot(\mathbf{q}-\mathbf{p})|}{\|\mathbf{n}\|}.$$

Solved chapter practice

Practice. Find the projection of $\mathbf{v} = (3, 4, 0)$ onto $\mathbf{u} = (1, 1, 0)$.
Solution. $\operatorname{proj}_{\mathbf{u}}\mathbf{v} = \frac{\mathbf{v}\cdot\mathbf{u}}{\mathbf{u}\cdot\mathbf{u}}\mathbf{u} = \frac{7}{2}(1, 1, 0) = (7/2, 7/2, 0)$.

Mastery practice A. Reconstruct vectors turn geometry into algebra 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 vectors turn geometry into algebra 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 Vector-valued motion

Rebuild the chapter from first principles

The claim I need to own is this: The model in Vector-valued motion 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-valued motion. 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-valued motion 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 vectors turn geometry into algebra to functions of several variables. 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 vector-valued function gives position one coordinate at a time, but velocity and acceleration are geometric vectors, not three unrelated derivatives. Speed is the magnitude of velocity; it is not the same object as velocity.

The easy version

Track a moving point with an arrow; differentiate the arrow to get velocity and again to get acceleration.

A space curve is $\mathbf{r}(t) = \langle x(t), y(t), z(t)\rangle$. Velocity is $\mathbf{r}'(t)$, speed is $\|\mathbf{r}'(t)\|$, and acceleration is $\mathbf{r}''(t)$.

Arc length is accumulated speed:

$$L = \int_a^b \|\mathbf{r}'(t)\| dt.$$

The unit tangent is $\mathbf{T} = \mathbf{r}'/\|\mathbf{r}'\|$. Curvature measures turning per unit distance:

$$\kappa = \frac{\|\mathbf{T}'(t)\|}{\|\mathbf{r}'(t)\|} = \frac{\|\mathbf{r}' \times \mathbf{r}''\|}{\|\mathbf{r}'\|^3}.$$

Sanity check

Velocity depends on parameter timing; the geometric path does not. Reparameterizing a curve can change speed without changing the traced curve.

Solved chapter practice

Practice. For $\mathbf{r}(t) = (\cos t, \sin t, t)$, find velocity, speed, and acceleration.
Solution. $\mathbf{v} = (-\sin t, \cos t, 1)$, so speed is $\sqrt{\sin^2 t + \cos^2 t + 1} = \sqrt{2}$. Acceleration is $\mathbf{a} = (-\cos t, -\sin t, 0)$.

Mastery practice A. Reconstruct vector-valued motion 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-valued motion 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 Functions of several variables

Rebuild the chapter from first principles

The claim I need to own is this: The model in Functions of several variables 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 functions of several variables. 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 functions of several variables 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-valued motion to partial derivatives and differentiability. 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 multivariable function is a landscape: the input is a point in a plane or space, and the output is height or another scalar. Domain and level sets are my first sketches because they expose constraints before differentiation.

The easy version

Feed a point into the function and imagine the output as the height of terrain above that point.

A function $f(x, y)$ is visualized as a surface $z = f(x, y)$ or as level curves $f(x, y) = c$. Level curves are often more informative: they show how outputs are organized across the plane.

For $f(x, y) = x^2 + 4y^2$, levels are ellipses. Higher c means larger ellipses. The gradient will later point across them, toward fastest increase.

Limits need every path

$$\lim_{(x,y)\to(a,b)} f(x, y) = L$$

requires the same behavior along every path. Two paths producing different values prove nonexistence. Agreement along a few paths proves very little.

Worked example

For $f(x, y) = \frac{x^2-y^2}{x^2+y^2}$ near $(0, 0)$, the path $y = 0$ gives 1, while $x = 0$ gives -1. Therefore the limit does not exist.

To prove a limit, inequalities and polar form are useful. If $|f(x, y)| \le C(x^2 + y^2)^{p/2}$ with $p > 0$, then $f \to 0$.

Solved chapter practice

Practice. Describe the domain and level curves of $f(x, y) = \sqrt{9 - x^2 - y^2}$.
Solution. The radicand must be nonnegative, so the domain is the disk $x^2 + y^2 \leq 9$. A level $f = c$ satisfies $x^2 + y^2 = 9 - c^2$ for $0 \leq c \leq 3$, giving concentric circles.

Mastery practice A. Reconstruct functions of several variables 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 functions of several variables 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 Partial derivatives and differentiability

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: a partial derivative freezes all but one input, so it gives slope along a coordinate direction. Differentiability is stronger: one tangent plane must approximate the surface for every small direction at once.

The easy version

Walk east while holding north fixed to get one partial; walk north while holding east fixed to get the other.

The partial derivative f_x changes x while freezing y; f_y does the opposite. Higher partials describe changing rates. When the second partials are continuous near a point, $f_{xy} = f_{yx}$.

The idea

Partials are axis-aligned probes. Differentiability is stronger: one plane must approximate the function well in every nearby direction.

Rebuild the chapter from first principles

The claim I need to own is this: Partials are axis-aligned probes. Differentiability is stronger: one plane must approximate the function well in every nearby direction. 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 partial derivatives and differentiability. 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 partial derivatives and differentiability 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 functions of several variables to chain rules and dependency trees. 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.

For $z = f(x, y)$, the tangent plane at (a, b) is

$$\boxed{z = f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b).}$$

This is also the linearization $L(x, y)$.

Worked example

For $f(x,y) = \sqrt{4+x-y}$ near $(0,0)$,

$$f(0,0) = 2, \quad f_x(0,0) = \frac{1}{4}, \quad f_y(0,0) = -\frac{1}{4}.$$

Thus $f(0.04,-0.02) \approx 2 + 0.01 + 0.005 = 2.015$.

The total differential is $df = f_x dx + f_y dy$. It is a local accounting system for propagated input changes.

Solved chapter practice

Practice. Find the tangent plane to $z = x^2y + e^y$ at $(1,0,1)$.
Solution. $f_x = 2xy$ and $f_y = x^2 + e^y$. At $(1,0)$ these are 0 and 2. Thus $z-1 = 0(x-1)+2(y-0)$, so $z = 1+2y$.

Mastery practice A. Reconstruct partial derivatives and differentiability 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 partial derivatives and differentiability 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 Chain rules and dependency trees

Rebuild the chapter from first principles

The claim I need to own is this: The model in Chain rules and dependency trees 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 chain rules and dependency trees. 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 chain rules and dependency trees 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 partial derivatives and differentiability to gradient and directional change. 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: with several variables, the chain rule is best drawn as a dependency graph. Every path from an independent variable to the output contributes a product of rates, and those path contributions add.

The easy version

Follow every route by which the input can affect the output; multiply along a route and add across routes.

If $z = f(x, y)$ and $x = x(t)$, $y = y(t)$, then

$$\frac{dz}{dt} = f_x \frac{dx}{dt} + f_y \frac{dy}{dt}.$$

If x and y depend on (s, t), there is one equation per input. A dependency tree keeps every path visible. For an implicit surface $F(x, y, z) = 0$ with $z = z(x, y)$,

$$z_x = -\frac{F_x}{F_z}, \qquad z_y = -\frac{F_y}{F_z},$$

when $F_z \neq 0$.

Common miss

Every path from an independent variable to the output contributes. Missing one branch is the multivariable version of forgetting an inner derivative.

Solved chapter practice

Practice. If $z = x^2 y$, $x = s + t$, and $y = st$, find $\partial z/\partial s$ at $(s, t) = (1, 2)$.
Solution. $z_x = 2xy$, $z_y = x^2$, $x_s = 1$, and $y_s = t$. At $(1, 2)$, $x = 3, y = 2$, so $z_s = z_x x_s + z_y y_s = 12(1) + 9(2) = 30$.

Mastery practice A. Reconstruct chain rules and dependency trees 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 chain rules and dependency trees 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 Gradient and directional change

Rebuild the chapter from first principles

The claim I need to own is this: The model in Gradient and directional change 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 gradient and directional change. 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 gradient and directional change 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 chain rules and dependency trees to unconstrained optimization. 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 gradient packages all first partials into the direction of steepest increase. A directional derivative is its projection onto a unit direction, so the dot product is measuring how much of that steepest climb I actually follow.

The easy version

The gradient is the uphill arrow; dot it with the direction you walk to get your slope.

$$\nabla f = \langle f_x, f_y, f_z \rangle.$$

For a unit direction $\mathbf{u}$, the directional derivative is

$$D_{\mathbf{u}} f = \nabla f \cdot \mathbf{u}.$$

The maximum rate is $\|\nabla f\|$, achieved in the gradient direction. The gradient is perpendicular to the level curve or level surface through the point.

Worked example

For $f(x, y) = x^2 + xy + y^2$, at $(1, 2)$,

$$\nabla f = \langle 2x + y, x + 2y \rangle = \langle 4, 5 \rangle.$$

Toward $\mathbf{v} = \langle 3, 4 \rangle$, use the unit vector $\mathbf{u} = \langle 3/5, 4/5 \rangle$:

$$D_{\mathbf{u}} f = \frac{12 + 20}{5} = \frac{32}{5}.$$

Solved chapter practice

Practice. For $f = x^2 + 2y^2$, find the maximum rate of change at $(1, -1)$ and its direction.
Solution. $\nabla f = (2x, 4y)$, so at $(1, -1)$ it is $(2, -4)$. The maximum rate is its magnitude $\sqrt{20} = 2\sqrt{5}$, in unit direction $(2, -4)/(2\sqrt{5}) = (1/\sqrt{5}, -2/\sqrt{5})$.

Mastery practice A. Reconstruct gradient and directional change 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 gradient and directional change 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 Unconstrained optimization

Rebuild the chapter from first principles

The claim I need to own is this: The model in Unconstrained optimization 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 unconstrained optimization. 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 unconstrained optimization 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 gradient and directional change to lagrange multipliers. 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

Setting the gradient to zero only creates candidates. The Hessian test reads local curvature, and just like one-variable calculus it can be inconclusive. I still inspect the function and domain when the determinant test returns zero.

The easy version

Find flat points, then use curvature to decide whether the surface is bowl-shaped, hill-shaped, or saddle-shaped there.

Interior critical points satisfy $\nabla f = \mathbf{0}$ or have undefined partials. For two variables, define

$$D = f_{xx}f_{yy} - (f_{xy})^2.$$

At a critical point:

- $D > 0$ and $f_{xx} > 0$: local minimum.
- $D > 0$ and $f_{xx} < 0$: local maximum.
- $D < 0$: saddle point.
- $D = 0$: inconclusive.

On a closed bounded region, absolute extrema require interior critical points plus the entire boundary. Parameterize boundary pieces and reduce them to one-variable problems.

Solved chapter practice

Practice. Classify the critical point of $f(x, y) = x^2 + 4xy + y^2$.
Solution. $f_x = 2x + 4y$ and $f_y = 4x + 2y$, giving only $(0, 0)$. Here $D = f_{xx}f_{yy} - f_{xy}^2 = 2 \cdot 2 - 4^2 = -12 < 0$, so the point is a saddle.

Mastery practice A. Reconstruct unconstrained optimization 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 unconstrained optimization 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 Lagrange multipliers

Rebuild the chapter from first principles

The claim I need to own is this: The model in Lagrange multipliers 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 lagrange multipliers. 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 lagrange multipliers 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 unconstrained optimization 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

At a constrained optimum, I cannot move along the constraint in a direction that increases the objective. That means the objective gradient is perpendicular to the constraint curve, parallel to the constraint gradient.

The easy version

On the boundary path, the best point occurs where the uphill arrow points directly across the path rather than along it.

To optimize f subject to $g = c$, solve

$$\boxed{\nabla f = \lambda \nabla g, \qquad g = c.}$$

At a constrained extremum, the level curve of f is tangent to the constraint, so the normal vectors align.

Worked example

Maximize xy on $x^2 + y^2 = 1$.
Solution.

$$\langle y, x\rangle = \lambda\langle 2x, 2y\rangle.$$

The nonzero solutions satisfy $x^2 = y^2$. With the constraint, $|x| = |y| = 1/\sqrt{2}$. The maximum $1/2$ occurs when signs match; the minimum $-1/2$ when signs differ.

Common miss

Lagrange multipliers produce candidates. Compare objective values, and separately inspect points where $\nabla g = 0$ or the constraint is not smooth.

Solved chapter practice

Practice. Maximize xy subject to $x^2 + y^2 = 8$ in the first quadrant.
Solution. $\nabla(xy) = (y, x) = \lambda(2x, 2y)$. Symmetry or division gives $x = y$. The constraint gives $2x^2 = 8$, so $x = y = 2$, and the maximum product is 4.

Mastery practice A. Reconstruct lagrange multipliers 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 lagrange multipliers 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 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 lagrange multipliers 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. Find the plane through $(1, 2, 0)$, $(0, 1, 1)$, and $(2, 0, 1)$.
2. For $\mathbf{r}(t) = \langle \cos t, \sin t, t\rangle$, find speed and arc length on $[0, 2\pi]$.
3. Show that $\lim_{(x,y)\to(0,0)} \frac{x^2y}{x^2+y^2} = 0$.
4. Find the tangent plane to $z = e^{xy}$ at $(1, 0, 1)$.
5. If $z = x^2y$, $x = s + t$, and $y = st$, find $\partial z/\partial s$.
6. Find the unit direction of fastest increase of $f = x^2 - yz$ at $(1, 2, 3)$ and the maximum rate.
7. Classify all critical points of $f = x^3 - 3x + y^2$.
8. Find the extrema of $f = x + y$ on $x^2 + 4y^2 = 4$.

CHAPTER 10 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. Plane

Two directions are $\mathbf{u} = \langle -1, -1, 1\rangle$ and $\mathbf{v} = \langle 1, -2, 1\rangle$. Their cross product is $\langle 1, 2, 3\rangle$. Thus

$$(x-1) + 2(y-2) + 3z = 0,$$

or $x + 2y + 3z = 5$.

2. Helix

$\mathbf{r}' = \langle -\sin t, \cos t, 1\rangle$, so speed is $\sqrt{2}$. Therefore $L = \int_0^{2\pi} \sqrt{2}dt = 2\pi\sqrt{2}$.

3. Limit bound

Since $|y| \le \sqrt{x^2+y^2}$ and $x^2 \le x^2+y^2$,

$$\left|\frac{x^2y}{x^2+y^2}\right| \le \sqrt{x^2+y^2} \to 0.$$

The squeeze theorem gives the limit 0.

4. Tangent plane

$f_x = ye^{xy}$ and $f_y = xe^{xy}$. At $(1,0)$ these are 0 and 1. Hence $z = 1 + (y-0) = 1 + y$.

5. Chain rule

Either expand or use paths:

$$z_x = 2xy, \quad z_y = x^2, \quad x_s = 1, \quad y_s = t.$$

Thus $z_s = 2xy + x^2t$. Substituting $x = s+t$, $y = st$ gives $2(s+t)st + (s+t)^2t$.

6. Gradient

$\nabla f = \langle 2x, -z, -y\rangle = \langle 2, -3, -2\rangle$. Maximum rate is $\sqrt{17}$, in unit direction $\langle 2, -3, -2\rangle/\sqrt{17}$.

7. Critical points

$f_x = 3x^2 - 3$ and $f_y = 2y$, so points are $(\pm 1, 0)$. Here $D = (6x)(2) - 0 = 12x$. At $(1,0)$, $D > 0$ and $f_{xx} > 0$: local minimum. At $(-1,0)$, $D < 0$: saddle.

8. Constrained extrema

Solve $\langle 1,1\rangle = \lambda\langle 2x, 8y\rangle$, so $x = 1/(2\lambda)$ and $y = 1/(8\lambda)$, hence $x = 4y$. The constraint gives $20y^2 = 4$, so $y = \pm 1/\sqrt{5}$ and $x = \pm 4/\sqrt{5}$ with matching signs. Values are $\pm\sqrt{5}$.

Closing frame

32A is where calculus stops being tied to one input but keeps the same architecture. The derivative becomes a linear map, the tangent line becomes a tangent plane, and the chain rule becomes a sum over dependency paths. The notation grows; the core idea stays restrained.