Finger Exercises 11

Lights, the Four Vectors, and Shadow Rays

$\approx$ 60 minutes, plus a 10-minute bonus
CS 453 — Computer Graphics

How to do it: On **paper** first, no computer, squared paper beside you. Sketch the ray *before* you compute, and keep the scene picture on page 2 open beside you — almost every part uses it. Conventions are the lecture's: x right, y up, bold letters are points and arrows, plain letters are single numbers. If an item takes longer than 5 minutes, skip it, finish the rest, come back. Nothing here is an assignment task — this sheet is Lecture 11 and nothing else.

Short on time? Do Parts E, F, G, H and I first (about 24 minutes). They are the spine of the lecture: find $\mathbf{l}$ and r, find the four vectors at a hit, write the shadow ray down, run the test, and know what ϵ is for.

Part A • Do You Remember? RECALL $\approx$ 6 min

A1. Our model has exactly three kinds of light: ____________, ____________ and ____________. The scene keeps them all in ____________.

A2. Shading at a hit needs four vectors: $\mathbf{x}$ = ____________, $\mathbf{n}$ = ____________, $\mathbf{l}$ = ____________, $\mathbf{v}$ = ____________.

A3. For a point light at $\mathbf{p}$, shading a point $\mathbf{x}$: $\mathbf{l}$ = ____________ and r = ____________. Both fall out of the same ____________.

A4. With viewing ray $\mathbf{e} + t\mathbf{d}$ hitting at t: $\mathbf{x}$ = ____________ and $\mathbf{v}$ = ____________. Why the minus sign?

A5. $\mathbf{n}$ is the ____________ at the hit. Floor: $\mathbf{n}$ = ________. Circle of centre $\mathbf{c}$, radius R, hit at $\mathbf{x}$: $\mathbf{n}$ = ____________.

A6. A shadow ray is ____________. Its origin is ____________ and its direction is ____________. It is the same formula as a viewing ray, with a ____________ job.

A7. For a point light the shadow ray searches $t \in$ ____________. The far end is r because ____________.

A8. ϵ is there because ____________. Rule of thumb: bigger than ____________, smaller than ____________.

A9. Directional light: search $t \in$ ____________, because ____________. Ambient light: ____________.

A10. A hit inside the interval means ____________, and that light contributes ________. The point is still not black, because ____________.

Lecture 11 on one card — check Part A against it, then keep it open

Three lights: a *point* light has a position $\mathbf{p}$, so $\mathbf{l}$ and r differ at every shaded point; a *directional* light has only a direction, the same $\mathbf{l}$ everywhere and no r; *ambient* light arrives equally from everywhere and has no $\mathbf{l}$ at all.
Adding up: each light computes its own share at $\mathbf{x}$; the colour is the sum over the scene's list of lights. How big a share is, is the shading model — later.
The four vectors at a hit: $\mathbf{x}$ the hit point, $\mathbf{n}$ the surface normal, $\mathbf{l}$ toward the light, $\mathbf{v}$ toward the eye. Ambient light needs none of them.
Where each comes from: $\mathbf{x} = \mathbf{e} + t\mathbf{d}$ from the viewing ray; $\mathbf{n}$ from the surface that was hit; $\mathbf{l} = (\mathbf{p} - \mathbf{x})/\|\mathbf{p} - \mathbf{x}\|$ and $r = \|\mathbf{p} - \mathbf{x}\|$ from the light; $\mathbf{v} = -\mathbf{d}/\|\mathbf{d}\|$ from the viewing ray, flipped.

Normals: floor $(0, 1)$; circle $(\mathbf{x} - \mathbf{c})/R$. Each shape returns its own $\mathbf{n}$ with the hit.
The shadow ray: $\mathbf{x} + t\mathbf{l}$ — start at the surface, aim at the light. Same formula as a viewing ray, different question: not *which* object, just *is there one.*
The interval: point light $[\epsilon, r]$; directional $[\epsilon, \infty)$; ambient, no ray. Stop at r: nothing past the light can block it. Start at ϵ: otherwise the ray hits the surface it started on at $t \approx 0$ and every point calls itself shadowed — *shadow acne.*
The verdict: a hit in the interval $\Rightarrow$ shadow, this light adds 0. No hit $\Rightarrow$ add this light's share. One shadow ray per light, per shaded point.

The scene on this sheet. Every part from B onwards lives in it. Three lights, one ball, one floor, one eye, and three floor points to shade. Copy it onto squared paper once; you will draw on it a lot.

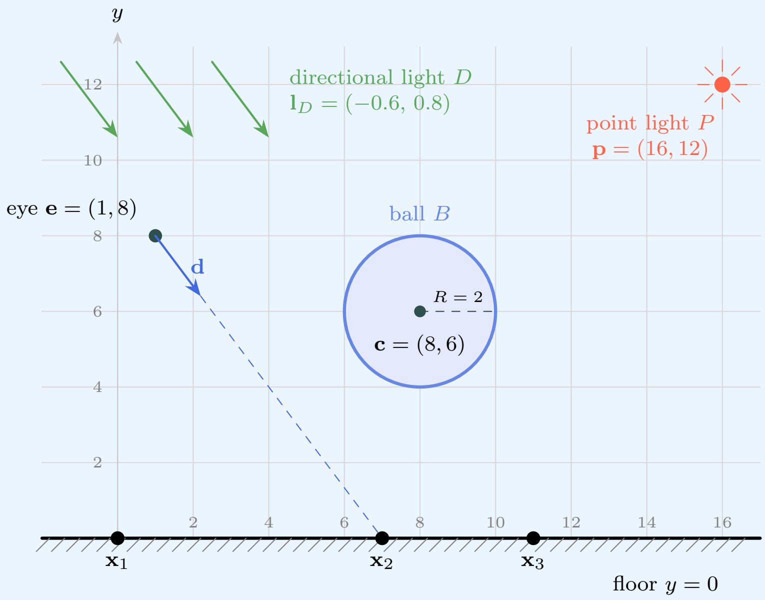

Floor: the line $y = 0$. **Ball** B: centre $\mathbf{c} = (8, 6)$, radius $R = 2$. **Point light** P at $\mathbf{p} = (16, 12)$. **Directional light** D, $\mathbf{l}_D = (-0.6, 0.8)$. **Ambient light** A.

Eye $\mathbf{e} = (1, 8)$, viewing direction $\mathbf{d} = (1.2, -1.6)$. **Floor points to shade**: $\mathbf{x}_1 = (0, 0)$, $\mathbf{x}_2 = (7, 0)$, $\mathbf{x}_3 = (11, 0)$.

Part B • Three Lights, Three Behaviours

CLASSIFY ≈ 4 min

Fill the table from memory, then check it against the card. Write **yes**, **no** or the interval.

light	has a position?	same $\mathbf{l}$ everywhere?	finite r?	needs a shadow ray?	interval
point	_______	_______	_______	_______	__________
directional	_______	_______	_______	_______	__________
ambient	_______	_______	_______	_______	__________

B1. Our scene has three lights. At $\mathbf{x}_2$, how many shadow rays are traced? How many would be traced if we added two more point lights?

B2. Which kind of light would you add to stop shadowed floor from coming out pure black, and why does adding it cost no extra rays?

Part C • Where Does Each One Come From?

CLASSIFY ≈ 4 min

C1. Join each quantity to the one thing that produces it: $\mathbf{x}$, $\mathbf{n}$, $\mathbf{l}$, r, $\mathbf{v}$ versus the viewing ray, the surface that was hit, the light, the viewing ray flipped.

C2. The light P moves to $(20, 9)$. The eye, the ball and the floor stay put. Which of $\mathbf{x}, \mathbf{n}, \mathbf{l}, \mathbf{v}, r$ change at $\mathbf{x}_2$?

C3. Now put the light back and move the *eye* instead. Which of the five change?

C4. Is $\mathbf{l}$ stored in the light or computed per shaded point? Answer for P and for D.

Part D • Normals by Hand

COMPUTE ≈ 4 min

The ball B has centre $\mathbf{c} = (8, 6)$ and radius $R = 2$.

D1. A viewing ray lands on the floor. What is $\mathbf{n}$ there, and does the answer depend on where along the floor it landed?

D2. The ball is hit at $(9.2,\ 7.6)$, then at $(6, 6)$, then at $(8, 4)$. Give $\mathbf{n}$ for each.

D3. Check one of your three: what should $\|\mathbf{n}\|$ be, and what does dividing by R have to do with it?

D4. The third one points straight *down*. Is that a mistake?

D5. Where on this ball is $\mathbf{n} = (0, 1)$, and where is $\mathbf{n} = (-0.6,\ 0.8)$?

Part E • $\mathbf{l}$ and r, Point by Point

COMPUTE ≈ 5 min

The point light P is at $\mathbf{p} = (16, 12)$. One subtraction gives you both answers. Lengths are one line: the root, then the number.

E1. At $\mathbf{x}_1 = (0, 0)$: compute $\mathbf{p} - \mathbf{x}_1$, then r, then $\mathbf{l}$.

E2. At $\mathbf{x}_2 = (7, 0)$: same three answers.

E3. At $\mathbf{x}_3 = (11, 0)$: same again. Leave $\mathbf{l}$ as a fraction if it is not tidy.

E4. Three points, three different $\mathbf{l}$ and three different r. Which property of a point light is this, in one sentence?

E5. Now the directional light D. Write $\mathbf{l}$ at $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$, and write r at each.

E6. Confirm the three arrows you wrote for P are unit arrows without doing three square roots.

Part F • $\mathbf{x}$ and $\mathbf{v}$ from the Viewing Ray

COMPUTE ≈ 5 min

The eye is at $\mathbf{e} = (1, 8)$ and the viewing ray is $\mathbf{e} + t\mathbf{d}$ with $\mathbf{d} = (1.2,\ -1.6)$. It first hits something at $t = 5$.

F1. Find $\|\mathbf{d}\|$, then the hit point $\mathbf{x}$.

F2. How far is that hit from the eye? (Careful: it is not 5.)

F3. Give $\mathbf{v}$.

F4. A classmate writes $\mathbf{v} = -\mathbf{d} = (-1.2,\ 1.6)$. Are they pointing the right way? What exactly is wrong?

F5. Write down all four vectors at this hit, using the light P for $\mathbf{l}$.

F6. Here $\mathbf{l}$ and $\mathbf{v}$ are mirror images. What does that say about where the eye and the light are standing, and would it still hold if we shaded $\mathbf{x}_1$ instead?

Part G • Write the Shadow Ray Down

COMPUTE ≈ 4 min

No intersection tests yet — just the ray and the interval, in the form $\mathbf{x} + t\mathbf{l}$, $t \in [\ ,\ ]$.

G1. The shadow ray from $\mathbf{x}_2$ toward P: write the origin, the direction and the interval.

G2. Evaluate it at $t = 0$, at $t = 5$ and at $t = 15$. What sits at each?

G3. The shadow ray from $\mathbf{x}_1$ toward P: origin, direction, interval. Then evaluate it at $t = 10$ and say what you find there.

G4. The shadow ray from $\mathbf{x}_3$ toward the *directional* light D: origin, direction, interval.

G5. Why can we say "$t = 15$ is the light" at all — what had to be true of $\mathbf{l}$ first?

Part H • Blocked or Clear? TRACE ≈ 5 min

Now run the test. The only object that can block anything here is the ball. You have not been taught ray–circle intersection yet, so use this week-2 shortcut, which is all projection:

Ray $\mathbf{o} + t\mathbf{l}$ with $\mathbf{l}$ unit, circle of centre $\mathbf{c}$, radius R. Put $\mathbf{m} = \mathbf{c} - \mathbf{o}$.

1. $s = \mathbf{m} \cdot \mathbf{l}$ — how far along the ray the closest approach happens. If $s \le 0$ the circle is not ahead of you.
2. The leftover $\mathbf{m} - s\,\mathbf{l}$ is the part of $\mathbf{m}$ the ray never covers; call its length h, the *closest approach*.
3. $h \ge R$: the ray misses. Otherwise the first hit is at $t = s - \sqrt{R^2 - h^2}$.

H1. $\mathbf{x}_1$ toward P: find s, the leftover and h. What is special about this ray?

H2. Same ray: where is the first hit, and what is the verdict for light P at $\mathbf{x}_1$?

H3. $\mathbf{x}_2$ toward P: find s, the leftover and h, and give the verdict.

H4. $\mathbf{x}_3$ toward the directional light D: same three numbers, then the first hit, then the verdict.

H5. $\mathbf{x}_1$ toward D: compute s and stop there. What does the answer tell you, and what is the verdict?

Part I • Why ϵ? TRACE ≈ 5 min

The floor is the line $y = 0$, and all three shaded points sit on it. That is the whole problem.

I1. Take the shadow ray at $\mathbf{x}_2$: $(7, 0) + t(0.6,\ 0.8)$. Solve for the t at which it meets the floor $y = 0$.

I2. With the naive interval $[0,\ r]$, is that t inside? What verdict does every floor point in the scene then get?

I3. Round-off leaves $\mathbf{x}_2$ a hair below the floor, at $(7,\ -0.000001)$. Now solve for the t at which the ray meets $y = 0$ again. Is it inside $[0, r]$? Inside $[0.001,\ r]$?

I4. Complete: ϵ must be bigger than ____________________ and smaller than ________________________.

I5. Someone worries about it and sets $\epsilon = 9$. What verdict does $\mathbf{x}_1$ get for light P now, and is it right?

I6. Why does the *far* end of the interval need no such nudge?

Part J • One Shadow Ray per Light CLASSIFY ≈ 5 min

Fill the verdict table from Parts H and I, then answer the questions. Write **shadow** or **lit** — and for ambient, whatever is true.

	point light P	directional light D	ambient A
$\mathbf{x}_1 = (0, 0)$	____________	____________	____________
$\mathbf{x}_2 = (7, 0)$	____________	____________	____________
$\mathbf{x}_3 = (11, 0)$	____________	____________	____________

J1. Two of the three points are in shadow — but not from the same light. What does that tell you about where the test belongs in the code?

J2. $\mathbf{x}_1$ is in P's shadow. Is it black? Write its colour as a sum of three terms.

J3. Suppose some object crossed $\mathbf{x}_2$'s shadow ray to P at $t = 18$. Does $\mathbf{x}_2$ go into shadow?

J4. Same object, same place — but now the light along that direction is *directional* instead. Verdict?

J5. A scene has 6 point lights and 1 ambient light. How many shadow rays are traced for one shaded point? For a 400×300 image in which every viewing ray hits something?

Part K • Put the Steps in Order

TRACE ≈ 3 min

Here is everything that happens when the ray tracer shades *one* viewing ray, in the wrong order. Number the steps 1 to 7.

_____ add this light's share to the running colour
_____ start the colour at the ambient share
_____ find the nearest hit along the viewing ray
_____ if the shadow ray hit something, add nothing for this light
_____ if nothing was hit at all, use the background colour and stop
_____ take the next light from the scene's list
_____ send a shadow ray from $\mathbf{x}$ along $\mathbf{l}$, over the light's interval

K1. Which steps repeat, and how many times?

K2. At which step is each of $\mathbf{x}$, $\mathbf{n}$, $\mathbf{v}$, $\mathbf{l}$, r first known?

K3. A viewing ray that hits nothing: how many shadow rays does it cause?

Part L • Find the Bug

FIX ≈ 5 min

Each snippet is one small change away from right. `first_hit(o, dir, t_min, t_max)` reports anything the ray meets with t in that range. Name the bug and give the fix.

L1. Symptom: every floor point comes out shadowed, and the ball's surface is speckled with dark dots.

```
blocked = first_hit(x, l, 0.0, r)
```

Bug & fix: ____

L2. Symptom: a point light. Some points come out shadowed although nothing at all stands between them and the lamp.

```
blocked = first_hit(x, l, EPS, INFINITY)
```

Bug & fix: ____

L3. Symptom: no shadows anywhere, although the test is certainly running.

```
l = normalize(x - p)
```

Bug & fix: ____

L4. Symptom: as soon as one light is blocked, the point turns pure black — the other lights and the ambient term vanish with it.

```
colour = ambient_share(x)
for light in scene.lights:
    if blocked(x, light):
        return BLACK
    colour = colour + light_share(x, light)
```

Bug & fix: ____

L5. Symptom: things far beyond the lamp cast shadows, and the fix from bug 2 did not help.

```
d = p - x
blocked = first_hit(x, d, EPS, r)
```

Bug & fix: ______________________________

Part M • One Sentence Each

EXPLAIN ≈ 4 min

M1. A shadow ray and a viewing ray use the same formula and the same intersection code. What is the one difference in the question they ask?

M2. Why is a shadowed point not black in our model?

M3. Why does the shadow test need only $\mathbf{x}$ and $\mathbf{l}$, and not $\mathbf{n}$ or $\mathbf{v}$?

M4. Why does a directional light's shadow ray have no far end?

M5. In our model the edge of a shadow is perfectly sharp. Why, and what kind of light would blur it?

M6. Why do shadows make a picture easier to read, even when they add no new objects?

Part N • Bonus: The Plank's Shadow ★

DESIGN ≈ 10 min

For the brave. The lecture ended by reading the shadow test backwards: instead of standing on the floor and asking whether the light can be seen, stand at the light and watch where the rays land. Same test, other direction — and it turns a shadow into a similar-triangles problem you can do in your head.

Add one object to the scene: a straight plank from (10, 3) to (11, 6). Everything else stays exactly as it was.

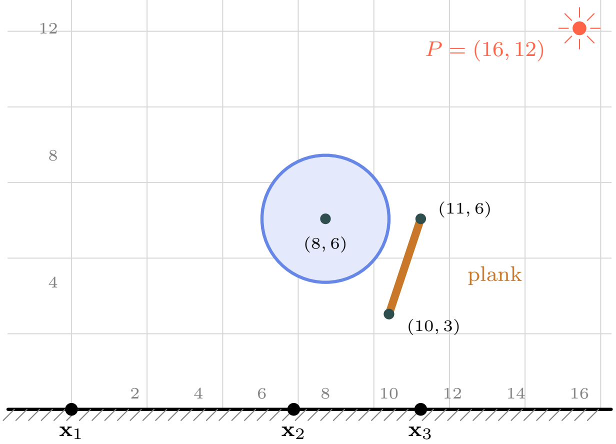

N1. Leave P along the ray that passes through the plank's upper end $(11, 6)$. One step of that ray is $(11, 6) - (16, 12) = (-5, -6)$. How many steps to fall from $y = 12$ to the floor, and where does it land?

N2. Same for the lower end $(10, 3)$.

N3. So the plank's shadow on the floor is the strip ________ $\leq x \leq$ ________. Which of $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$ falls inside it, and what does that do to the verdict you wrote in Part J?

N4. Check it the other way round, from the floor. The plank crosses $\mathbf{x}_2$'s shadow ray to P at $t = 6$. Does that agree? Which of the two methods would you use to shade one pixel, and which to understand a whole picture?

N5. Does the plank change $\mathbf{x}_2$'s ambient share? Its share from the directional light D? Give a reason for each.

N6. Sketch four or five more rays leaving P and landing on the floor. State the rule that connects "a ray from the light lands here" with the verdict a shadow ray would give there.

Stuck on more than three items? Re-read the Lecture 11 slides — everything above Part N comes straight from them. Start with "$\mathbf{l}$ points from $\mathbf{x}$ toward the light", "$\mathbf{v}$ points back toward the eye", "$\mathbf{n}$: each shape knows its own normal", "The shadow ray must stop at the light", "Start the shadow ray at ϵ" and "One shadow ray per light". Most items here are one of those six slides with new numbers.

Answer Key

Check your work *after* attempting everything. If an answer surprises you, that item is worth re-reading in the slides.

A1. **Point**, **directional** and **ambient**. The scene keeps them in **one list**; adding a light is one more entry in it, and the colour at a point is the sum over that list.

A2. $\mathbf{x}$ the **hit point**, $\mathbf{n}$ the **surface normal** there, $\mathbf{l}$ the direction **toward the light**, $\mathbf{v}$ the direction **back toward the eye**. Three of them are unit arrows; $\mathbf{x}$ is a point.

A3. $\mathbf{l} = (\mathbf{p}-\mathbf{x})/\|\mathbf{p}-\mathbf{x}\|$ and $r = \|\mathbf{p}-\mathbf{x}\|$ — both from the same **subtraction** $\mathbf{p}-\mathbf{x}$. $\mathbf{l}$ says which way to look, r says how far.

A4. $\mathbf{x} = \mathbf{e} + t\mathbf{d}$ and $\mathbf{v} = -\mathbf{d}/\|\mathbf{d}\|$. Minus because $\mathbf{d}$ runs *from* the eye *to* the surface, and $\mathbf{v}$ has to point the other way, back at the eye. The division makes it unit length.

A5. The **unit direction the surface faces**. Floor: $(0, 1)$. Circle: $\mathbf{n} = (\mathbf{x}-\mathbf{c})/R$ — from the centre out through the hit, divided by the radius to make it unit. Each shape computes its own and hands it back with the hit.

A6. $\mathbf{x} + t\mathbf{l}$. Origin: the **surface point being shaded**. Direction: $\mathbf{l}$, **toward the light**. Same formula, **new** job — a viewing ray asks *what do I see*, a shadow ray asks *can this light see me*.

A7. $t \in [\epsilon,\ r]$. The far end is r because the light sits at $t = r$ (as $\mathbf{l}$ is unit, t is distance), and an object *past* the light cannot block light that has already left it.

A8. Because the ray starts *on* a surface, so with $t = 0$ allowed it immediately hits the surface it left — **shadow acne**. ϵ must be bigger than the **round-off wobble** in $\mathbf{x}$, and smaller than the **distance to the nearest genuine blocker**.

A9. $[\epsilon,\ \infty)$, because the light has no position, so there is no finite r to stop at — anything along $\mathbf{l}$ blocks it. Ambient: **no shadow ray at all**, so no test.

A10. **Shadow**; that light contributes $\mathbf{0}$. Not black, because the ambient share is added with no test, and any *other* light with a clear shadow ray still adds its own share.

B. Point: yes / no / yes / yes / $[\epsilon, r]$. Directional: no / yes / no / yes / $[\epsilon, \infty)$. Ambient: no / — (there is no $\mathbf{l}$) / no / no / none. Read the table down the last two columns: that is the whole shadow section in six boxes.

B1. **Two** — one for P, one for D, none for the ambient light A. With two more point lights, **four**. One shadow ray per light that has a direction, per shaded point.

B2. **Ambient**. It arrives equally at every point from every direction, so there is no direction to trace along and no test to run — it is simply added to every shaded point, shadowed or not.

C1. $\mathbf{x}$ — the viewing ray, evaluated at the hit t. $\mathbf{n}$ — the surface that was hit. $\mathbf{l}$ and r — the light (both from $\mathbf{p}-\mathbf{x}$). $\mathbf{v}$ — the viewing ray flipped. Note $\mathbf{l}$ and r share a source, which is why one subtraction gives you both.

C2. Only $\mathbf{l}$ and r. $\mathbf{x}$ came from the viewing ray, $\mathbf{n}$ from the floor, $\mathbf{v}$ from the eye — none of them has heard of the light.

C3. Potentially **all** of them. A new eye means a new viewing ray, so it lands somewhere else: $\mathbf{x}$ moves, and once $\mathbf{x}$ moves, $\mathbf{l}$ and r are recomputed from the new point; $\mathbf{v}$ changes with the new direction; and $\mathbf{n}$ changes too if the ray now lands on the ball instead of the floor.

C4. For the point light P it is **computed per shaded point**, from $\mathbf{p}-\mathbf{x}$ — that is what having a position means. For the directional light D it is **stored**: one arrow, the same at every point, with nothing to compute.

D1. $\mathbf{n} = (0, 1)$, straight up, the same everywhere on the floor — a flat surface faces one way at every point.

D2. $(\mathbf{x}-\mathbf{c})/R$ each time: $(1.2, 1.6)/2 = (0.6,\ 0.8)$; $(-2, 0)/2 = (-1,\ 0)$; $(0, -2)/2 = (0,\ -1)$.

D3. 1. Every point of the circle is exactly R from the centre, so $\mathbf{x}-\mathbf{c}$ always has length R; dividing by R is exactly what makes it unit. No square root needed.

D4. No. $(8, 4)$ is the bottom of the ball, and the surface there faces down — away from the centre, as always. The normal has no idea where the light or the floor is.

D5. $\mathbf{n} = (0, 1)$ at the top, $\mathbf{c} + 2(0, 1) = (8, 8)$. $\mathbf{n} = (-0.6, 0.8)$ at $\mathbf{c} + 2(-0.6, 0.8) = (6.8,\ 7.6)$, up and to the left. Going the other way, $\mathbf{x} = \mathbf{c} + R\,\mathbf{n}$.

E1. $\mathbf{p}-\mathbf{x}_1 = (16, 12)$; $r = \sqrt{16^2 + 12^2} = 20$; $\mathbf{l} = (16, 12)/20 = (0.8,\ 0.6)$.

E2. $\mathbf{p}-\mathbf{x}_2 = (9, 12)$; $r = \sqrt{9^2 + 12^2} = 15$; $\mathbf{l} = (9, 12)/15 = (0.6,\ 0.8)$.

E3. $\mathbf{p}-\mathbf{x}_3 = (5, 12)$; $r = \sqrt{5^2 + 12^2} = 13$; $\mathbf{l} = (5, 12)/13 \approx (0.38,\ 0.92)$.

E4. All its light leaves **one position**, so the direction toward it — and the distance to it — depend on where you are standing. $\mathbf{l}$ and r are per shaded point, not per light.

E5. $\mathbf{l} = (-0.6,\ 0.8)$ at all three — it is stored in the light, not computed. There is **no** r: the light has no position, so there is no distance to it. That is the one fact the shadow interval will need.

E6. Each was built by dividing a vector by its own length, so each is unit by construction. If you want one check: $0.8^2 + 0.6^2 = 1$. The point of dividing is that t along the shadow ray then measures distance.

F1. $\|\mathbf{d}\| = \sqrt{1.2^2 + 1.6^2} = 2$. $\mathbf{x} = (1, 8) + 5(1.2, -1.6) = (7,\ 0)$ — which is $\mathbf{x}_2$, on the floor.

F2. $t\,\|\mathbf{d}\| = 5 \times 2 = 10$. $\mathbf{d}$ was not a unit arrow, so t counts copies of $\mathbf{d}$, not units of distance — Lecture 10's warning, still true here.

F3. $\mathbf{v} = -\mathbf{d}/\|\mathbf{d}\| = -(1.2, -1.6)/2 = (-0.6,\ 0.8)$. Up and to the left — back toward the eye at $(1, 8)$, which is where it should point.

F4. The *direction* is right; the *length* is not. $\mathbf{v}$ must be unit, so it has to be divided by $\|\mathbf{d}\| = 2$. $-\mathbf{d}$ is only correct when $\mathbf{d}$ was already normalized.

F5. $\mathbf{x} = (7, 0)$; $\mathbf{n} = (0, 1)$ (floor); $\mathbf{l} = (0.6,\ 0.8)$ (from Part E); $\mathbf{v} = (-0.6,\ 0.8)$. Three unit arrows and a point — everything shading asks for, all of it available the moment the hit is known.

F6. They lean by the same amount on opposite sides of the upward normal, so the eye and the light sit symmetrically about the vertical through $\mathbf{x}_2$. It is a coincidence of this point: at $\mathbf{x}_1$, $\mathbf{l} = (0.8, 0.6)$ while $\mathbf{v}$ would be computed from a different viewing ray entirely. $\mathbf{l}$ and $\mathbf{v}$ answer different questions and only agree by accident.

G1. $(7, 0) + t\,(0.6,\ 0.8)$, with $t \in [\epsilon,\ 15]$. Origin: the point being shaded. Direction: $\mathbf{l}$. Far end: r, from Part E.

G2. $t = 0$: $(7, 0)$, the shaded point itself — which is exactly why we do not start there. $t = 5$: $(10,\ 4)$, a point in mid-air on the way. $t = 15$: $(16,\ 12)$, the light. The far end of the interval is the light, every time.

G3. $(0, 0) + t\,(0.8,\ 0.6)$, $t \in [\epsilon,\ 20]$. At $t = 10$: $(8,\ 6)$, which is $\mathbf{c}$ — this shadow ray runs straight through the centre of the ball. Useful for Part H.

G4. $(11, 0) + t\,(-0.6,\ 0.8)$, $t \in [\epsilon,\ \infty)$. No r exists, so there is no far end to stop at.

G5. $\mathbf{l}$ had to be **unit**. Then t along the ray is distance in world units, and the light, being $r = 15$ away, sits at $t = 15$. With an unnormalized direction the light would sit at $t = 1$ and the interval $[\epsilon, r]$ would be nonsense.

H1. $\mathbf{m} = (8, 6)$, $\mathbf{l} = (0.8, 0.6)$, so $s = 6.4 + 3.6 = 10$ and the leftover is $(8, 6) - 10(0.8, 0.6) = (0, 0)$, so $h = 0$. The ray goes **straight through the centre** — the closest approach is the centre itself.

H2. $t = s - \sqrt{R^2 - h^2} = 10 - 2 = 8$: it meets the ball one radius before the centre. Is 8 in $[\epsilon,\ 20]$? Yes — comfortably. **Shadow**: light P adds nothing at $\mathbf{x}_1$.

H3. $\mathbf{m} = (1, 6)$, $\mathbf{l} = (0.6, 0.8)$: $s = 0.6 + 4.8 = 5.4$; leftover $(1, 6) - 5.4(0.6, 0.8) = (-2.24,\ 1.68)$; $h = \sqrt{2.24^2 + 1.68^2} = 2.8$. Since $2.8 \geq R = 2$ the ray **misses** the ball — by 0.8. Nothing at all lies in $[\epsilon, 15]$, so $\mathbf{x}_2$ is **lit** by P.

H4. $\mathbf{m} = (-3, 6)$, $\mathbf{l} = (-0.6, 0.8)$: $s = 1.8 + 4.8 = 6.6$; leftover $(-3, 6) - 6.6(-0.6, 0.8) = (0.96,\ 0.72)$; $h = 1.2 < 2$, so it hits, at $t = 6.6 - \sqrt{4 - 1.44} = 6.6 - 1.6 = 5$. The interval is $[\epsilon, \infty)$, and 5 is in it: **shadow**.

H5. $\mathbf{m} = (8, 6)$, $\mathbf{l} = (-0.6, 0.8)$: $s = -4.8 + 4.8 = 0$. The closest approach is at $t = 0$, i.e. the ball is exactly *sideways* from $\mathbf{x}_1$ and never ahead of the ray. Nothing in $[\epsilon, \infty)$: **lit**. (If you carried on: $h = 10$, far outside the ball anyway.)

I1. $0 + 0.8t = 0$, so $t = 0$. The ray starts on the floor, so of course it meets the floor — right where it started.

I2. Yes, $t = 0$ is inside $[0, r]$. So every floor point reports a hit and calls itself **shadowed** — the floor goes uniformly dark, and no light ever reaches anything.

I3. $-0.000001 + 0.8t = 0$ gives $t = 0.00000125$. Tiny, but positive, so it *is* inside $[0, r]$ — this is **shadow acne**, and it speckles because the round-off differs from pixel to pixel. It is not inside $[0.001, r]$, so starting at ϵ skips it.

I4. Bigger than the round-off wobble in $\mathbf{x}$ (around 10^{-6} here), and smaller than the distance from the surface to the nearest thing that genuinely blocks it. Any value in that gap works, which is why a single small constant does the job for a whole scene.

I5. Its real hit was at $t = 8$, and $8 < 9$, so the search $[9, 20]$ skips it: $\mathbf{x}_1$ is reported **lit**. Wrong — the ball is plainly in the way. Too large an ϵ deletes real shadows, starting with the ones that matter most: the contact shadows right under an object.

I6. Because there is no surface at the far end to hit by accident — $t = r$ is the light, not geometry. The near end is special only because the ray starts on a surface.

J. $\mathbf{x}_1$: shadow (hit at $t = 8$) / lit / always arrives. $\mathbf{x}_2$: lit (missed by 0.8) / lit / always arrives. $\mathbf{x}_3$: lit (the ray to P passes well to the right of the ball — read it off the picture) / shadow (hit at $t = 5$) / always arrives. Ambient needs no row of reasoning at all: no direction, no ray, no test.

J1. Inside the per-light step, not once per point. Each light sits somewhere different, so each one needs its own shadow ray from the same $\mathbf{x}$. A point can be shadowed for one light and lit for another at the same instant.

J2. Not black: 0 (from P) + D's share + the ambient share. Only the blocked light drops out; the sum over the list carries on.

J3. No. The interval is $[\epsilon,\ 15]$ and $18 > 15$, so the object sits *behind the light* from $\mathbf{x}_2$'s point of view. It cannot block light that has already left the lamp.

J4. **Shadow**. The interval is $[\epsilon, \infty)$, and 18 is inside it. With no position there is no "behind the light" — anything along $\mathbf{l}$, at any distance, blocks it.

J5. 6 per shaded point. For the image: $400 \times 300 \times 6 = 720{,}000$ shadow rays, on top of the 120,000 viewing rays. Shadows are cheap per ray and expensive in total — one more reason the intersection test has to be fast.

K. Reading down the list as printed: $7, 3, 1, 6, 2, 4, 5$. In order: (1) nearest hit along the viewing ray; (2) nothing hit, background, stop; (3) start at the ambient share; (4) take the next light; (5) shadow ray over that light's interval; (6) hit, so add nothing for it; (7) no hit, so add its share — then back to (4) until the list is done.

K1. Steps 4 to 7 — once per light in the list. Everything before them happens once per viewing ray.

K2. $\mathbf{x}$ and $\mathbf{n}$ come out of step 1 with the hit; $\mathbf{v}$ can be had at step 1 too, since it is only the viewing direction flipped. $\mathbf{l}$ and r have to wait for step 4 — there is no such thing as "the" light direction until you have picked a light.

K3. None. Step 2 stops it. No hit, no point to shade, no light to test — the pixel just gets the background.

L1. The search starts at 0, so the ray finds the surface it started on, at $t \approx 0$ — shadow acne (Part I). Fix: `first_hit(x, l, EPS, r)` with a small positive `EPS`.

L2. No far end, so objects *past* the lamp are counted as blockers (J3). Fix: `first_hit(x, l, EPS, r)` with $r = \|\mathbf{p} - \mathbf{x}\|$. Only a directional light may use `INFINITY`.

L3. The subtraction is backwards, so $\mathbf{l}$ points *away* from the light and the shadow ray is fired into the empty half of the scene (usually straight into the ground, which `EPS` then skips). Fix: `l = normalize(p - x)`. $\mathbf{l}$ points *from* the surface *toward* the light.

L4. Being in shadow means *this light* adds zero, not that the point has no colour. The `return` throws away the ambient share already banked and never reaches the rest of the list. Fix: replace it with `continue` — add nothing for this light and carry on (J2).

L5. `d` was never normalized, so t counts copies of `d`, not distance: the lamp sits at $t = 1$, and the search runs to $t = r$, that is r times past it. Fix: normalize first — `l = d / length(d)`, then search $[\epsilon, r]$ — or keep `d` and search $[\epsilon, 1]$. Pick one and stay with it.

M1. A viewing ray asks *which* object is nearest and *where*, because the answer must be shaded. A shadow ray asks only *is there anything at all* in the interval — any hit will do, so it may stop at the first one it finds.

M2. Ambient light is added with no test at all, and any other light whose shadow ray is clear still contributes. Shadow removes one term from a sum; it does not zero the sum.

M3. The test is pure geometry — is the straight path from the point to the light clear? Which way the surface faces and where the eye stands change how bright the point looks, not whether the light can reach it. $\mathbf{n}$ and $\mathbf{v}$ are the shading model's business.

M4. It has no position, so there is no finite distance r at which the light sits and past which a blocker would be irrelevant. Everything along $\mathbf{l}$, however far, is between the point and the light.

M5. A point light is a single point, so the test is yes-or-no: either the one straight path is blocked or it is not, and neighbouring floor points flip from one to the other in an instant. A light with *size* — an area light — can be partly visible from a point, which is what gives a soft edge.

M6. They show which things touch which. A ball with no shadow could be resting on the floor or floating above it; the shadow, and where it meets the ball, settles the question. Shadows carry the spatial relationships between nearby objects.

N1. Each step falls 6, and we must fall 12, so **two** steps: $(16, 12) + 2(-5, -6) = (6,\ 0)$. No quadratic, no solving — just how many copies of the arrow reach the floor.

N2. One step is $(-6, -9)$, falling 9 each time, so $12/9 = 4/3$ steps: $(16, 12) + \frac{4}{3}(-6, -9) = (8,\ 0)$.

N3. $6 \le x \le 8$. That catches $\mathbf{x}_2 = (7, 0)$, which Part J had as *lit* by P. With the plank in the scene it is now in **shadow** for P. One new object changed a verdict without moving the light, the eye or the point.

N4. 6 is inside $[\epsilon,\ 15]$, so: shadow. The two methods agree, as they must — they walk the same line in opposite directions. For *one* point, the shadow ray: it answers exactly the question asked and stops. For the whole floor, the fan from the light: one sweep gives every edge at once. A renderer uses the first; the second is how you reason about what you are seeing.

N5. Ambient: no — it is added with no test, so no object can ever remove it. Directional: no — $\mathbf{l}_D = (-0.6, 0.8)$ points up and to the *left* from $(7, 0)$, while the plank stands to the right at x between 10 and 11, so that ray never meets it. $\mathbf{x}_2$ loses exactly one term, P's.

N6. A floor point receives a ray from P exactly when its own shadow ray to P is clear — it is the same straight line, walked the other way. The floor is in shadow precisely where no ray lands: the strip $6 \le x \le 8$ behind the plank, and the much wider band the ball hides (it contains $\mathbf{x}_1$, and reaches past the left edge of the picture). Two objects, two shadows, one test.