VMS Electromagnetism — Student Workbook

Symbols (Aligned with Bridge V2)

$S_0 = \hbar$ — anchor constant

κ_t — torsion constant | κ_s — shear constant

A_d — display-area (projected orthogonal cross-section, $A\_d = \pi r^2$)

n — closure density factor (shadow: refractive index)

θ — angle from the surface normal

Q — conserved closure measure from loop orientation $\sigma \in \{+1, -1\}$ (shadow: charge $q = \sigma e$, with sign from relative handedness preserved by A3 expansion).

Axioms (Plain Language)

A1 — Voids move like shadows: when a source changes, the "missing" part of space moves outward, like a shadow sliding across a wall. A2 — Space has finite tension: like a trampoline, it resists bending and stretching (not zero, not infinite). A3 — Space can expand: like a balloon skin that can stretch or shrink; this expansion preserves handedness ($\sigma = +1$ or -1).

Constants Sidebar (Student Awareness)

Constants are current and stable (Bridge p.6–8), suitable for 2025 and beyond: • $\varepsilon_0 = 8.854 \times 10^{-12}$ F/m   • $k = 1/(4\pi \varepsilon_0) \approx 8.987 \times 10^9$ N·m^2/C^2 • Hydrogen Hα (Balmer) = 656.281 nm (NIST ASD v5.11) These won't change in a way that affects your homework or lab checks.

Section 1 — Inverse-Square from Closure Geometry (FULL Algebra, Projected A_d)

Figure: sphere around a compact oriented loop source; display-area taken as projected disk A_d.

Derivation — Every Step (No Flux Import)

Goal: Find $E(r)$ at distance r from a compact oriented loop source with conserved closure Q (σ gives the sign).

1) Conservation: total closure fixed by orientation: $Q = \sigma \cdot |Q|$, with $\sigma \in \{+1, -1\}$.

2) Display-area (Bridge A_d): $A(r) = \pi r^2$ (projected disk, orthogonal cross-section per bridge definition).

3) Density: $\rho(r) := Q / A(r) \to \rho(r) = Q / (\pi r^2)$.

4) Field definition with calibration: $E(r) := \alpha \cdot \rho(r) \Rightarrow E(r) = \alpha \cdot (Q / (\pi r^2))$.

5) Simplify: $E(r) = (\alpha Q) / (\pi r^2)$.

VMS Result with A_d: $E(r) = (\alpha Q) / (\pi r^2)$ (σ encodes the sign).

Cross-ref (Bridge p.2,4): Far-field $F \propto (\sigma_1 \sigma_2)/r^2$ from loop orientation and deficit propagation (A2–A3 finite tension/expansion).

Unit Bridge (α to SI)

Set $\alpha = 4/\varepsilon_0$ to account for projected $A\_d = \pi r^2$ vs. full flux conventions. Then

$E(r) = Q / (4\pi \varepsilon_0 r^2) = k Q / r^2$, with $k = 1/(4\pi \varepsilon_0)$.

Pause & Reflect

What did we just learn? The field gets weaker as $1/r^2$ because the same 'stuff' spreads over a bigger disk $A\_d = \pi r^2$. Why did we set $\alpha = 4/\varepsilon_0$? Because our VMS area uses $A\_d = \pi r^2$; this restores the standard classroom number $E = Q/(4\pi \varepsilon_0 r^2)$.

Worked Examples (Numbers, Step-by-Step)

Constants: $\varepsilon_0 = 8.854 \times 10^{-12}$ F/m $\to$ $k = 1/(4\pi \varepsilon_0) \approx 8.987 \times 10^9$ N·m^2/C^2.

Ex.1 $Q = 2.0$ µC, $r = 0.50$ m $\to$ $E = (8.987 \times 10^9)(2.0 \times 10^{-6})/0.25 = 7.19 \times 10^4$ V/m.

Ex.2 $Q = 1.0$ µC, $r = 1.00$ m $\to$ $E = (8.987 \times 10^9)(1.0 \times 10^{-6})/1.00 = 8.99 \times 10^3$ V/m.

Scaling r: $r_1 = 1.0$ m $\to$ $E_1 = 8.99 \times 10^3$; $r_2 = 2.0$ m $\to$ $E_2 = 2.25 \times 10^3$; $E_2/E_1 = 1/4$.

Bench replication: two e at 1 nm $\to$ $F = k e^2 / r^2 \approx 2.3 \times 10^{-10}$ N (repulsion; sign via σ).

Teaching Moments (Plain Language)

• Same stuff over more area $\to$ weaker. Double r $\Rightarrow$ area ×4 $\Rightarrow$ field ÷4.

• This is geometry, not 'flux lines.' We never counted lines crossing a surface.

Practical Use Case — Field Near a Probe (Engineering)

Why this matters: Engineers estimate E around electrodes, sensors, or charged probes for safety and performance.

VMS explains the inverse-square shape (geometry). Classical SI form (Coulomb/Gauss) gives the exact number once ε_0 is fixed.

Shadow / Classical Anchor + Use

Classical (textbook): $\oint E \cdot dA = Q/\varepsilon_0$ and $E = k\,Q\,/\,r^2$. Use: field mapping, capacitor edge estimates, sensor safety distances.

Bridge: VMS gives the $1/r^2$ from σ-orientation + A2–A3 expansion; classical provides calibrated SI numbers once ε_0 is fixed.

Guided Exercises (Self-Check)

1) With $Q = 3.0$ μC and $r = 2.0$ m, compute E using $E = Q/(4\pi\,\varepsilon_0\,r^2)$. Show every step and the units.

2) Compare E at $r = 1$ m and $r = 3$ m for $Q = 1.0$ μC. What is the ratio $E(3\text{ m})/E(1\text{ m})$?

3) (Challenge) If σ flips sign for one source, what happens to the force direction between two charges? Explain using $\sigma_1\sigma_2$.

Section 2 — Refraction as Ratio-First Closure (FULL Algebra, Caustics Tie-In)

Figure: boundary with normal; show θ_1 in medium 1, θ_2 in medium 2; draw tangential component matching across the boundary.

Derivation — Every Step (No Fermat Import)

Idea: At the boundary, nothing can pile up sideways; transverse oscillations must match.

Cross-ref (Bridge p.3): Emerges from Void waveform at caustics/folds; "sideways" continuity from transverse oscillations matching at boundary (no surplus A_d).

1) Geometry: θ_1 in medium 1 (n_1), θ_2 in medium 2 (n_2).

2) Sideways (tangential) continuity: $n_1 \sin\theta_1 = n_2 \sin\theta_2$.

3) Solve: $\sin\theta_2 = (n_1/n_2) \sin\theta_1$ → $\theta_2 = \arcsin((n_1/n_2) \sin\theta_1)$.

VMS Refraction: $n_1 \sin\theta_1 = n_2 \sin\theta_2$

Pause & Reflect

Why is sideways (tangential) part the same? Nothing can teleport sideways at the interface; the wavefront must meet perfectly along the edge. Bigger n means the ray bends toward the normal on entry; smaller n means it bends away on exit.

Numeric Angle Examples — Step by Step

Air→Glass n_1=1.000, n_2=1.500, θ_1=30° → $\sin\theta_2$=(1/1.5)·0.500=0.333 → θ_2≈19.5°.

Glass→Air n_1=1.500, n_2=1.000, θ_1=30° → $\sin\theta_2$=1.5·0.500=0.750 → θ_2≈48.6°.

Water→Air n_1=1.333, n_2=1.000, θ_1=45° → $\sin\theta_2$=1.333·0.707=0.943 → θ_2≈70.0°.

Practical Use Case — Fiber Optics (Critical Angle)

Critical angle θ_c when θ_2→90° ($\sin\theta_2$=1): $n_1 \sin\theta_c = n_2$ → $\sin\theta_c = n_2/n_1$.

Worked: n_core=1.48, n_clad=1.44 → $\sin\theta_c$=0.97297 → θ_c≈76.7°.

Validity: Plane wave, sharp interface (n jump); ties to bridge caustics where J=0 forces TIR-like evasion.

Bridge to Classical + Side-by-Side

Classical anchor (Snell): $n_1 \sin\theta_1 = n_2 \sin\theta_2$. Use: lenses, eyeglasses, microscopes, prisms, fiber-optic routing.

Side-by-Side: air→glass, n_1=1.000, n_2=1.500, θ_1=40° → $\sin\theta_2$=(1/1.5)·sin40°=0.4285 → θ_2≈25.4° (VMS and Classical match).

VMS: boundary continuity of transverse A_d (tangential continuity / no sideways pile-up at the boundary). Classical: Snell from experiment/Fermat. Both agree numerically.

Guided Exercises (Self-Check)

1) Water→air: n_1=1.33, n_2=1.00, θ_1=45°. Compute θ_2. Show sin and arcsin steps.

2) Air→glass: n_1=1.00, n_2=1.50, θ_1=20°. Compute θ_2. Then reverse (glass→air) and compare.

3) (Challenge) Prove that if $n_1 = n_2$ then $\theta_2 = \theta_1$ from the equation $n_1 \sin\theta_1 = n_2 \sin\theta_2$.

Section 3 — Interference (Revised, VMS → Classical)

Figures: (1) double-slit geometry (d, z, y_m); (2) thin-film stack; (3) Michelson paths.

Concept First (Plain Language)

In classic terms interference happens when two coherent contributions reach the same spot and add. Sometimes they add up (bright), sometimes they cancel (dark). In VMS language, phases are comparisons of closure along two routes; the key control knob is the path difference ΔL. They don't "interfere" with each other they just take different paths because the space they cross has been impacted by nearby mass of the slit (closed void loops)

- **Teaching Moment:** "Coherent" = same color (wavelength) and a steady phase relationship.

- **Quick rules:** Constructive when $\Delta L = m\lambda$; destructive when $\Delta L = (m+1/2)\lambda$

Derivation — Every Step

1) Two contributions arrive at point P from paths of lengths L_1 and L_2.

2) Phase difference: $\Delta\phi = k(L_2 - L_1) = (2\pi/\lambda)\cdot\Delta L$.

3) Add complex amplitudes: $U \propto 1 + e^{i\Delta\phi}$.

4) Intensity: $I \propto |U|^2 = |1 + e^{i\Delta\phi}|^2 = 2(1 + \cos\Delta\phi)$.

Conditions: Constructive when ΔL = mλ; Destructive when ΔL = (m+1/2)λ (m ∈ ℤ)

Double-Slit Geometry (Small Angles)

Path difference: $\Delta L \approx d\sin\theta$. Fringe angles: $d\sin\theta_m = m\lambda \Rightarrow \sin\theta_m = m\lambda/d$.

Screen distance z (small θ): $y_m \approx z\tan\theta_m \approx z\sin\theta_m \approx z\,m\,\lambda/d$.

Fringe spacing: Δy ≈ zλ/d (small-angle approximation)

Finite Slit Width — Envelope Matters

A single slit of width a produces an envelope $I_1(\theta) \propto \text{sinc}^2(\beta)$ with $\beta = (\pi a\sin\theta)/\lambda$.

Two slits → interference × envelope: $I(\theta) \propto \cos^2(\pi d\sin\theta/\lambda) \cdot \text{sinc}^2(\beta)$. Zeros of the envelope at $a\sin\theta = m\lambda$ (m≠0) suppress outer fringes.

Worked Examples (Line by Line)

Example 1 — Fringe Spacing:

Given λ = 532 nm = 5.32×10^{-7} m, d = 0.50 mm = 5.0×10^{-4} m, z = 1.5 m.

Use $\Delta y = z\lambda/d$.

Compute numerator: $z\lambda = 1.5 \times 5.32 \times 10^{-7} = 7.98 \times 10^{-7}$.

Divide by d: $7.98 \times 10^{-7} / 5.0 \times 10^{-4} = 1.596 \times 10^{-3}$ m = 1.596 mm.

Answer: $\Delta y \approx$ 1.596 mm per bright fringe.

Example 2 — First Six Bright Fringes (m = 0...5):

$y_0 = m \times \Delta y = 0 \times 1.596$ mm = 0.000 mm

$y_1 = m \times \Delta y = 1 \times 1.596$ mm = 1.596 mm

$y_2 = m \times \Delta y = 2 \times 1.596$ mm = 3.192 mm

$y_3 = m \times \Delta y = 3 \times 1.596$ mm = 4.788 mm

$y_4 = m \times \Delta y = 4 \times 1.596$ mm = 6.384 mm

$y_5 = m \times \Delta y = 5 \times 1.596$ mm = 7.980 mm

Example 3 — Envelope Value at m=4 with a = d/5:

a = d/5 = 0.50 mm / 5 = 0.10 mm = 1.0×10^{-4} m.

Small angle: $\sin\theta \approx m\lambda/d$.

Compute $\sin\theta \approx 4 \times 5.32 \times 10^{-7} / 5.0 \times 10^{-4} = 4.256 \times 10^{-3}$.

$\beta = \pi a \sin\theta / \lambda = \pi \times 1.0 \times 10^{-4} \times 4.256 \times 10^{-3} / 5.32 \times 10^{-7} \approx 2.51$.

$\text{sinc}(\beta) = \sin(\beta)/\beta \approx 0.234$. Envelope factor = $\text{sinc}^2(\beta) \approx 0.055$.

Observation: Outer fringes are dimmed by the envelope.

Thin-Film Anti-Reflection (Normal Incidence)

At normal incidence, two reflections (air→film, film→glass) can cancel if they are π out of phase. Choose quarter-wave thickness so the second reflection picks up an extra half-cycle inside the film.

Condition (with one phase flip): $2 n_f t = (m + 1/2) \lambda_0 / n_f$; most common m=0 $\Rightarrow$ t $\approx \lambda_0/(4 n_f)$.

Example: λ_0 = 550 nm, n_f = 1.38 (MgF_2) $\Rightarrow$ t $\approx \lambda_0/(4n_f)$ = 550/(4×1.38) $\approx$ 99.6 nm.

Practical Use: Camera lenses and displays use quarter-wave MgF_2 to reduce glare near green light.

Michelson Interferometer — Distance from Fringe Count

Path difference is Δ = 2ΔL. Each full bright→bright cycle corresponds to Δ increasing by λ.

So m = Δ/λ = 2ΔL/λ $\Rightarrow$ ΔL = (m λ)/2.

Worked: λ = 632.8 nm; if you count m = 10 fringes, then ΔL = (10 × 632.8 nm)/2 = 3,164 nm.

Rule of Thumb: One fringe shift corresponds to ΔL = λ/2.

ⓘ Teaching Moments

• Path difference → phase difference. Integer multiples of λ add; half-multiples cancel.

• The envelope comes from the finite width a: the slit itself diffracts, shaping the interference pattern.

• Thin films engineer destructive interference to kill reflections at a target color.

• Michelson turns color into distance by counting fringes.

Classical Anchor & Bridge

Double-slit: $y_m \approx m\lambda z/d$. Single-slit minima: a sinθ = mλ. Thin-film AR: t $\approx \lambda_0/(4n_f)$. Michelson: m = 2ΔL/λ.

VMS expressed the same conditions via closure phases without importing wave ideas first; the classroom equations above are the calibrated shadows used in labs.

Practice Problems

1) λ=500 nm, d=0.20 mm, z=2.0 m. Compute Δy and y_3.

2) a=50 μm, d=0.25 mm, λ=633 nm. Is the 6th bright fringe strongly suppressed by the envelope? Show β and $sinc^2(\beta)$.

3) Thin-film AR: n_f=1.30, target λ_0=600 nm. Find t.

4) Michelson: λ=633 nm, observed 12 fringes. Find ΔL.

Worked Solutions

1) $\Delta y = z\lambda/d = 2.0\times500\times10^{-9} / 2.0\times10^{-4} = 5.00$ mm; $y_3 = 3\Delta y \approx 15.00$ mm.

2) $sin\theta \approx m\lambda/d \Rightarrow \beta=\pi\, a\, sin\theta/\lambda \approx \pi\, a\, (m/d)$. Here $\beta \approx \pi\times50\times10^{-6}\times(6/2.5\times10^{-4})\approx3.77$. $sinc^2\approx(-0.156)^2\approx0.024$ (suppressed).

3) $t\approx\lambda_0/(4n\_f)=600/(4\times1.30)\approx115.4$ nm.

4) $\Delta L=(m\lambda)/2=(12\times633$ nm$)/2\approx3798.0$ nm.

Section 4 — Propagation in Varying Media (Revised, Eikonal + Rays)

Figures: (1) S-contours with normals (rays); (2) GRIN slab bending; (3) mirage profile.

Concept First (Plain Language)

Rays don't "decide" anything. In VMS terms they follow the steepest climb of S — a map of how space has been shaped by nearby stuff. Bigger n makes S steeper, so paths bend toward higher n. You control bending by changing n(x): gentle gradients give gentle curves; a sharp step is refraction (the Snell result you already derived).

• Teaching Moment: Think of S as a contour map; rays always cross the contours at right angles (they're normals to S).

Derivation — Every Step

Start: $U(x) = A(x) \cdot exp(i\, k_0\, S(x))$.

Compute gradients explicitly:

$\nabla U = (\nabla A + i\, k_0\, A\, \nabla S)\, e^{i k_0 S}$

$\nabla^2 U = [\nabla^2 A + 2\, i\, k_0\, (\nabla S\cdot\nabla A) + i\, k_0\, A\, \nabla^2 S - k_0^{\,2}\, A\, |\nabla S|^2]\, e^{i k_0 S}$

Insert into wave equation (Helmholtz): $\nabla^2 U + k_0^{\,2}\, n^2\, U = 0$.

Balance powers of k_0:

- $O(k_0^2)$: $-k_0^2 A |\nabla S|^2 + k_0^2 n^2 A = 0 \Rightarrow |\nabla S|^2 = n^2$ (eikonal)

- $O(k_0^1)$: $2 \nabla S \cdot \nabla A + A \nabla^2 S = 0$ (transport)

$|\nabla S|^2 = n^2(x)$ and $2 \nabla S \cdot \nabla A + A \nabla^2 S = 0$

Rays — Equations of Motion

Rays follow normals to S: $r'(s) \parallel \nabla S$. In isotropic media: $d/ds\,(n\,t) = \nabla n$, where $t = r'(s)$.

1D gradient $n = n(x)$: $d\theta/ds = (1/n)\,\partial n/\partial x$ (bend toward higher n).

Ray bend rule: $d\theta/ds = (1/n)\,\partial n/\partial x$

Worked Example 1 — Parabolic GRIN (Step-by-Step)

Take $n(y) = n_0 (1 - (g^2/2)\,y^2)$. For small y, and small angles, the ray equation reduces to $y'' + g^2 y = 0$.

Solve: $y(s) = y_0 \cos(g\,s) + (\theta_0/g) \sin(g\,s)$.

Pitch: $P = 2\pi/g \Rightarrow$ with $g = 20.0\ \mathrm{m}^{-1}$, $P \approx 2\pi/20 \approx 0.314$ m. A 0.25-pitch element length $\approx$ 0.0785 m focuses collimated input at the exit.

Result: $y'' + g^2 y = 0 \Rightarrow y(s) = y_0 \cos(gs) + (\theta_0/g) \sin(gs)$

Worked Example 2 — Linear Gradient (Mirage, Numbers)

Let $n(z) = n_0 + \beta z$ with $n_0 \approx 1.000$. For a shallow ray, $d\theta/ds \approx (1/n)\,\partial n/\partial z = \beta/n_0$.

Choose $\beta = 0.010\ \mathrm{m}^{-1}$, $n_0 \approx 1.000$. Over $s = 100$ m, $\Delta\theta \approx (\beta/n_0)\,s \approx 1.00$ rad $\approx 57.3°$ upward.

Interpretation: Hot air near the road lowers n near the surface; rays bend upward, showing a fake 'pool.'

ⓘ Teaching Moments

- Eikonal builds the 'direction map' S; rays follow the steepest climb of S.

• Where n is larger, S changes faster, so rays turn toward that region.

• Transport keeps brightness reasonable: if rays bunch up, intensity grows; if they spread, it drops.

Classical Anchors + Practical Uses

• Eikonal/transport underpin geometric optics and commercial ray-tracing (lens design, microscopes, AR/VR).

• GRIN lenses enable compact imaging; the quarter-pitch rule is used in fiber collimators.

• Atmospheric gradients explain mirages and ducting; same math helps laser-beam propagation near hot surfaces.

Practice Problems

1) For $n(x)=n_0+\alpha x$ with $\alpha=0.02$ m^{-1}, estimate the bend over 50 m starting at $\theta_0=0$. Use $d\theta/ds\approx\alpha/n_0$.

2) GRIN: $g=15$ m^{-1}. Compute the pitch P and the length for a 0.25-pitch element.

3) Sketch S-contours for a lens with higher n at the center. Indicate ray directions (normals to S).

Worked Solutions

1) $\Delta\theta\approx(\alpha/n_0)s=0.02\times50\approx1.00$ rad $\approx 57.3°$ toward +x.

2) $P=2\pi/g\approx2\pi/15\approx0.419$ m; 0.25-pitch≈0.105 m.

3) S-contours are denser at the center (higher n). Rays (normals) bend inward toward the center.

Section 5 — Diffraction (Kirchhoff → Fresnel → Fraunhofer)

Angles are measured from the surface normal.

Symbols: λ (wavelength), $k=2\pi/\lambda$, a (slit width), d (slit spacing), D (aperture diameter), z (screen distance), y (screen height), m (order).

Concept — What is Diffraction?

An opening (aperture) decides which routes are allowed from the source to the screen. Different places across the opening are slightly different routes, so they pick up slightly different phase. When those contributions add at one spot on the screen, you get bright or dim depending on how the phases line up.

■ **BOX: Big idea: A pattern on the screen is just 'many routes with phases' adding together.**

① Teaching Moment — 'Fourier transform' in plain words

Think of the aperture shape as a 'recipe' that tells you how much of each tiny sideways tilt you have. Far away (Fraunhofer), each tilt sends light to a particular angle. So the far-field pattern is a map of 'how much of each tilt' the aperture recipe contains.

→ If the opening is wider → you have more 'narrow tilts' → the central bright spot gets tighter.

→ If the opening is narrower → you have more 'wide tilts' → the light spreads out more.

Bridge to Classical (what equations we use)

Start with the Helmholtz equation (steady waves): $\nabla^2 U + k^2 U = 0$. Use the Kirchhoff integral to add contributions from each point in the opening. Keep the quadratic terms in the phase → Fresnel (near/intermediate). Drop them in the far field → Fraunhofer, where the pattern is essentially the Fourier transform of the aperture.

Derivation — Step by Step (what each step does & why)

→ Model: $\nabla^2 U + k^2 U = 0$ (Why: one color, uniform medium).

→ Green function $G(P,Q)=e^{ikR}/(4\pi R)$ (Why: how a point radiates to a point).

→ Kirchhoff integral over the opening Σ (Why: sum all allowed routes through the aperture).

→ Paraxial expansion of distance R (Why: small angles let us keep just the important terms).

→ Fresnel: keep quadratic phase → curved 'Fresnel' fringes; Fraunhofer: drop it → straight, evenly spaced fringes.

Canonical Results (memorize these shapes)

■ **BOX: Single slit (width a): $I(\theta) = I_0 (\sin\beta/\beta)^2$, $\beta = (\pi a / \lambda) \sin\theta$; minima when $a \sin\theta = m\lambda$ ($m \neq 0$).**

■ **BOX: Double slit (width a, spacing d): $I(\theta) = I_0 \cos^2(\pi d \sin\theta/\lambda) \cdot (\sin\beta/\beta)^2$; small angles: $\Delta y \approx \lambda z / d$.**

■ **BOX: Transmission grating (period d): $m\lambda = d \sin\theta\_m$ (m integer).**

■ **BOX: Circular aperture (diameter D, Airy): $I(\theta)=I_0[2J_1(\pi D \sin\theta/\lambda)/(\pi D \sin\theta/\lambda)]^2$; first dark at $\sin\theta \approx 1.22 \lambda/D$.**

■ **BOX: Fresnel number: $N\_F = a^2/(\lambda z)$. If $N\_F \gtrsim 1$ → Fresnel; if $N\_F \ll 1$ → Fraunhofer.**

Worked Examples (numbers shown, no skips)

Example 1 — Single slit, first minimum and central-lobe width

Given: $\lambda = 632.8$ nm $= 6.328\times10^{-7}$ m, $a = 100$ μm $= 1.00\times10^{-4}$ m, $z = 1.5$ m.

First minimum at $a \sin\theta \approx \lambda \Rightarrow \sin\theta \approx \lambda/a$ (small angle $\theta \approx \sin\theta$).

$\theta_1 \approx \lambda/a = 6.328\times10^{-7} / 1.00\times10^{-4} = 6.328\times10^{-3}$ rad.

Distance on screen: $y_1 \approx z \theta_1 = 1.5 \times 6.328\times10^{-3} \approx 9.49$ mm.

■ **BOX: Central bright width (first-to-first minima): $\approx 2 y_1 \approx 18.98$ mm.**

→ Check: making the slit wider (bigger a) makes y_1 smaller → sharper central spot.

Example 2 — Double slit, fringe spacing and envelope

Use the small-angle spacing $\Delta y \approx (\lambda z)/d$.

$\Delta y = (6.328\times10^{-7} \times 1.5) / (3.00\times10^{-4}) = 3.16$ mm.

→ If each slit has width a, the single-slit envelope (sinc^2) dims the outer fringes.

Example 3 — Circular aperture (Airy disk), first dark ring

Small angles: $\sin\theta \approx y/z \Rightarrow y \approx z \cdot 1.22\,\lambda/D$.

Given $D = 5$ mm $= 5.0\times10^{-3}$ m, $z = 2.0$ m: $y \approx 2.0 \times 1.22 \times 6.328\times10^{-7} / 5.0\times10^{-3} \approx 0.31$ mm.

■ **BOX: Rayleigh tip: two point sources are 'just resolved' when one's peak sits on the other's first dark ring.**

Example 4 — Transmission grating, first order angle

Period $d = 2.0$ μm, $\lambda = 500$ nm $\Rightarrow \sin\theta_1 = \lambda/d = 0.25 \Rightarrow \theta_1 \approx 14.48°$.

→ More lines per mm (smaller d) spread colors more — used in spectrometers.

Example 5 — Which regime am I in? (Fresnel number)

$N\_F = a^2/(\lambda z) = (5.0\times10^{-5})^2 / (5.32\times10^{-7} \times 0.5) \approx 0.01 \rightarrow$ Fraunhofer (far-field).

→ If you move the screen closer (smaller z) or use a larger aperture a, N_F grows and you slide toward Fresnel patterns.

Teaching Moments (recap)

→ Wider slit (bigger a) → tighter central spot; narrower slit → broader pattern.

→ Double-slit spacing set by d; outer fringes fade because of the single-slit envelope (a).

→ Circular aperture: smaller D → larger Airy disk (worse angular resolution).

→ Gratings: smaller period d → larger angular spread between orders.

Practice (try these, then check solutions)

1) Single slit: $\lambda=500$ nm, $a=80$ μm, $z=1.0$ m. Find y_1 (first minimum) and central-lobe width ($\approx 2y_1$).

2) Double slit: $\lambda=633$ nm, $d=250$ μm, $z=1.2$ m. Find Δy. If $a=50$ μm, explain what happens to the outer fringes.

3) Grating: $d=2.0$ μm, $\lambda=450$ nm. Find θ_1 and θ_2 (if they exist).

4) Circular aperture: $D=3$ mm, $\lambda=550$ nm, $z=2.5$ m. Find Airy first-dark radius y.

5) Fresnel/Fraunhofer: a=60 μm, λ=532 nm, z=0.40 m. Compute N_F and name the regime.

6) Design: You want $\Delta y \approx 2.0$ mm at z=1.0 m using λ=532 nm. What should d be?

Worked Solutions

1) $y_1 \approx z(\lambda/a) = 1.0 \times (5.00 \times 10^{-7}/8.0 \times 10^{-5}) = 6.25 \times 10^{-3}$ m=6.25 mm; width$\approx 2y_1 \approx 12.50$ mm.

2) $\Delta y = (\lambda z)/d = (6.33 \times 10^{-7} \times 1.2)/(2.50 \times 10^{-4}) = 3.04$ mm. With a=50 μm the sinc^2 envelope dims outer fringes.

3) $\sin\theta_1 = \lambda/d = 0.225 \Rightarrow \theta_1 \approx 13.00°$. For m=2, $\sin\theta_2 = 2\lambda/d = 0.45 \Rightarrow \theta_2 \approx 26.74°$ (exists since ≤1).

4) $y \approx z \cdot 1.22\lambda/D = 2.5 \times 1.22 \times 5.50 \times 10^{-7}/3.0 \times 10^{-3} \approx 0.56$ mm.

5) $N\_F = a^2/(\lambda z) = (6.0 \times 10^{-5})^2/(5.32 \times 10^{-7} \times 0.40) \approx 0.02 \rightarrow$ Fraunhofer (far-field).

6) $\Delta y \approx \lambda z/d \Rightarrow d \approx (\lambda z)/\Delta y = (5.32 \times 10^{-7} \times 1.0)/0.002 \approx 266.00$ μm.

Section 6 — Polarization (Jones → Stokes → Mueller)

Angles are measured from each element's reference axis (e.g., a polarizer's transmission axis). Symbols are defined up front. Teen-friendly, step-by-step, VMS-first.

Symbols

• λ — wavelength; $k = 2\pi/\lambda$; $i^2 = -1$

• $E = [E\_x ; E\_y]$ — Jones vector (complex amplitudes)

• θ — angle of a polarizer's transmission axis (from +x)

• δ — retardance (phase delay between x and y)

• Δn — birefringence; L — thickness; $\delta = 2\pi \Delta n L / \lambda$

• $S = (S_0, S_1, S_2, S_3)$ — Stokes vector (intensity domain)

Concept — What is Polarization?

VMS Math — how the x/y numbers are built (routes → sums)

Idea: many allowed routes r reach the same point on the screen. Each route picks up a phase S_r (a path-like number) and carries some 'closure weight' into x and y.

For each transverse axis u ∈ {x,y}:

$E\_u \propto \Sigma\_r \ \sqrt{\rho\_u(r)} \cdot e^{\{ i\ k_0\ S\_r \}}$ with $k_0 = 2\pi/\lambda$

Uniform medium & narrow beam: the sum behaves like one effective term per axis →

$E\_x = A\_x\ e^{\{i\ \phi\_x\}}$, $E\_y = A\_y\ e^{\{i\ \phi\_y\}}$, $\Delta\phi = \phi\_y - \phi\_x$

Bridge to classroom (Jones vector): identify $E\_x \leftrightarrow U\_x$, $E\_y \leftrightarrow U\_y$, so $J = [E\_x\ ;\ E\_y]$. Stokes (intensities): $S_0 = |E\_x|^2 + |E\_y|^2$, $S_1 = |E\_x|^2 - |E\_y|^2$, $S_2 = 2\text{Re}(E\_x E\_y^*)$, $S_3 = 2\text{Im}(E\_x E\_y^*)$.

▣ Teaching Moment: Polarization = the balance (A_x vs A_y) and the phase gap ($\Delta\phi$) between the two axes.

VMS → Polarizer (why it's a projection)

At a linear polarizer with axis n̂(θ), the 'no sideways pile-up' boundary rule drains closure perpendicular to n̂. Only the component aligned with n̂ continues downstream; closure perpendicular to n̂ is culled (no sideways pile-up).

Math gate from that rule: $P = \hat{n}\ \hat{n}^T$. With $\hat{n} = [\cos\theta, \sin\theta]$, $P = [[\cos^2\theta, \cos\theta\cdot\sin\theta],[\cos\theta\cdot\sin\theta, \sin^2\theta]]$.

→ This is exactly the $J\_pol(\theta)$ matrix you use in the classroom derivation (now you know why it's a projector).

VMS → Retarder (why a phase delay appears)

If the medium prefers one transverse direction (shear anisotropy), the 'slowness' differs along x and y, so phases accumulate at different rates.

Route integral for the phase gap: $\Delta\phi = k_0 \int (n\_y - n\_x)\ ds$

For a uniform plate of thickness L: $\Delta\phi = k_0 (n\_y - n\_x) L = 2\pi\ \Delta n\ L / \lambda$ (this is the classroom formula for a waveplate).

▣ Teaching Moment: Quarter-wave ($\Delta\phi = \pi/2$) turns equal x/y into circular; Half-wave ($\Delta\phi = \pi$) mirrors the angle around the fast axis.

VMS parameters in this story (light touch):

• κ_s (shear budget) sets the anisotropy → different n_x, n_y (birefringence).

• κ_t (torsion budget) twists the preferred axes along the path (optical activity). For small angles, the net twist is the rotation you measure.

Light can wiggle more in one direction than another. That "wiggle direction" is polarization.

In VMS language, a source and the medium set preferred transverse directions; different routes can keep phase along x and y differently, so one direction dominates at the screen.

Classroom shadow: we represent this with components along x and y and track how elements (polarizers, plates) change those components.

◙ Teaching Moment — Two ways to describe polarization

→ Jones: track complex amplitudes (good when the light is fully coherent/polarized).

→ Stokes/Mueller: track intensities (works for partial/real-world beams and lossy elements).

Part A — Jones Calculus (amplitudes & phases)

WHAT: Write the electric field as a 2-component column E = [E_x ; E_y]. Linear optical elements are 2×2 matrices that act on E.

WHY: In a uniform medium with one color (single λ), each element mixes or delays components linearly, so matrix math fits.

A1) Linear Polarizer at angle θ

Algebra (no skips):

• Rotation matrix $R(\theta) = [[\cos\theta, -\sin\theta], [\sin\theta, \cos\theta]]$; $R(-\theta) = [[\cos\theta, \sin\theta], [-\sin\theta, \cos\theta]]$.

• Ideal pass in its own frame $P_0 = [[1, 0], [0, 0]]$.

• Compose in the lab frame $J\_pol(\theta) = R(-\theta) \cdot P_0 \cdot R(\theta)$.

• Multiply: $J\_pol(\theta) = [[\cos^2\theta, \cos\theta\cdot\sin\theta], [\cos\theta\cdot\sin\theta, \sin^2\theta]]$.

→ This projects any input onto the axis at angle θ.

Malus from the matrix:

Let input be linear at angle φ, $E\_in = [\cos\phi, \sin\phi]^T$ (unit amplitude). Then

$E\_out = J\_pol(\theta)\cdot E\_in \Rightarrow E\_out,x = \cos\theta\cdot\cos(\phi-\theta)$, $E\_out,y = \sin\theta\cdot\cos(\phi-\theta)$.

Intensity $I\_out = |E\_out|^2 = E\_out,x^2 + E\_out,y^2 = \cos^2(\phi-\theta)$.

□ Box: Malus' Law $I\_out = I\_in \cdot \cos^2(\phi - \theta)$ (unpolarized input → first polarizer gives ½ I_in).

Idea: keep the component along its transmission axis, kill the perpendicular component.

Matrix recipe (rotation into the polarizer frame, ideal pass along x, rotate back):

$J\_pol(\theta) = R(-\theta) \cdot [[1, 0], [0, 0]] \cdot R(\theta)$

with $R(\theta) = [[\cos\theta, -\sin\theta], [\sin\theta, \cos\theta]]$.

Worked Example A — Two Polarizers (Malus' Law)

Given unpolarized light of intensity I_in. First ideal polarizer passes half: I_1 = ½ I_in.

A second polarizer at relative angle Δ transmits I_out = $I_1 \cos^2\Delta$ = (½ I_in) $\cos^2\Delta$.

Numbers: I_in = 1.00 → at Δ = 30°, I_out = 0.5 × $\cos^2 30°$ ≈ 0.5 × 0.75 = 0.375.

A2) Linear Retarder (Waveplate) — phase delay δ

Algebra (retarder at angle φ):

• In plate axes (fast=x): J_ret(δ) = [[1, 0], [0, e^{iδ}]].

• In lab axes: J_ret(δ,φ) = R(−φ) · J_ret(δ) · R(φ).

• Physical link: δ = 2π·Δn·L / λ.

Quarter-wave to circular (worked):

Input linear at 45°: E_in = [1/√2, 1/√2]. QWP δ=π/2 with fast axis along x (φ=0):

J_ret = [[1, 0], [0, i]] ⇒ E_out = [1/√2, i/√2].

Equal magnitudes; y lags x by 90° → circular polarization (handedness set by sign of δ).

Half-wave rotates linear by 2× angle difference (derivation sketch):

HWP: δ=π so J = [[1,0],[0,−1]] in plate axes. For input linear at ψ and plate at φ,

E_out = R(φ) · [[1,0],[0,−1]] · R(−φ) · [cosψ, sinψ]^T.

Carrying out the multiplication gives a linear output at angle ψ' = 2φ − ψ.

☐ Box: HWP rule — output angle = 2φ − input angle (mirror around the fast axis).

Matrix (fast axis along x): J_ret(δ) = [[1, 0], [0, e^{iδ}]]   (global phase ignored).

Physical link: δ = 2π Δn L / λ (birefringence Δn over thickness L).

Orientation φ: rotate the fast axis with R(φ) on both sides.

Worked Example B — Quarter-Wave Plate makes circular

Input: linear at 45°: E_in ∝ [1 ; 1].

Through a QWP (δ = π/2) with fast axis along x: E_out ∝ [1 ; e^{iπ/2}] = [1 ; i] → equal magnitudes with 90° phase → circular polarization.

Check: intensities equal; the y-component lags by 90°, so the tip of E rotates.

▣ Teaching Moment — Half-Wave and Quarter-Wave plates

→ Half-wave ($\delta=\pi$) flips the component along y: great for rotating linear polarization.

→ Quarter-wave ($\delta=\pi/2$) shifts by 90°: turn 45° linear into circular, or circular back to linear.

Part B — Stokes & Mueller (intensity domain; great for real beams)

"Polarized = has a favorite direction; Unpolarized = no favorite direction."

Angles are measured from each element's reference axis (e.g., a polarizer's transmission axis).

VMS Bridge — why we need Stokes/Mueller

Many allowed routes reach the same spot on the screen. Each route adds with a phase. When the environment jiggles those phases differently (lamp light, rough paths), the x and y parts don't stay perfectly in step. We still have light (big S_0), but the 'in-step' part shrinks (smaller S_1, S_2, S_3).

→ VMS one-liner (idea only): for axis $u\in\{x,y\}$, a route-sum builds a complex number; random phase scatter reduces the cross-term average.

→ So we switch calculators: Stokes (S) and Mueller (M) track intensity balances even when phases jitter.

WHAT & WHY (in one screen)

WHAT:

Use $S = (S_0,S_1,S_2,S_3)$ to say how bright the beam is (S_0) and how aligned its polarization is (S_1,S_2,S_3).

WHY:

Works for real beams (lamps/sunlight) and lossy elements; systems act via 4×4 Mueller matrices S_out = M S_in.

▣ BOX — Mapping note (VMS ↔ classroom)

We build two transverse components from routes (VMS). In classroom math they're written as E_x and E_y. Stokes just converts those components into four easy intensity numbers (S_0..S_3).

B1) Stokes parameters from components (step by step)

Start with two components (classroom symbols): E_x and E_y (can be complex).

Algebra (no skips):

• $S_0 = |E\_x|^2 + |E\_y|^2$ (total brightness)

• $S_1 = |E\_x|^2 - |E\_y|^2$ (more x than y?)

• $S_2 = 2\,\mathrm{Re}(E\_x\,E\_y^*)$ (x vs y at 45°)

• $S_3 = 2\,\mathrm{Im}(E\_x\,E\_y^*)$ (right vs left circular)

→ Use: S says 'how polarized' without needing to track the exact phase of every route.

Examples (plug numbers):

• Linear at angle ψ: E = [cosψ, sinψ] (real) $\Rightarrow$ S_0=1, S_1=cos2ψ, S_2=sin2ψ, S_3=0.

• Right-circular: E = [1/$\sqrt{2}$, i/$\sqrt{2}$] $\Rightarrow$ S_0=1, S_1=0, S_2=0, S_3=+1.

▣ **BOX — Teen check: if S_1=S_2=S_3=0 (but S_0>0), the light is unpolarized — bright, but no preferred direction.**

B2) Mueller essentials (use without getting lost)

▣ **BOX — How to use a Mueller matrix (4 steps)**

1) Name your state: S_in = (S_0,S_1,S_2,S_3).

2) Pick the element matrix M (polarizer, rotator, etc.).

3) Multiply once: S_out = M · S_in.

4) Say what it means in plain words (e.g., 'the analyzer kept the part aligned with Δ').

Rotation (twist axes by θ):

S_1' = S_1 cos2θ + S_2 sin2θ; S_2' = −S_1 sin2θ + S_2 cos2θ; S_0, S_3 unchanged.

→ Use: models optical activity/Faraday — 'turn the polarization map by 2θ' in (S_1,S_2).

Ideal linear polarizer at angle θ (Mueller form):

M_pol(θ) = ½ · [[1, cos2θ, sin2θ, 0],

     [cos2θ, cos^{2}2θ, cos2θ·sin2θ, 0],

     [sin2θ, cos2θ·sin2θ, sin^{2}2θ, 0],

     [0, 0, 0, 0]]

→ Use: halves unpolarized light; for polarized light, output depends on alignment with θ.

Worked Example — Malus with Stokes (numbers, no skips)

"Alignment rule: more mis-alignment → dimmer."

Start: unpolarized S_in = (I,0,0,0). Through first polarizer (any θ): S_1 = (I/2, I/2, 0, 0).

Analyzer at relative angle Δ (second polarizer): S_out = (I/2)(1, cos2Δ, 0, 0).

Intensity after analyzer: I_out = (I/2) $\cos^2\Delta$ (same Malus result you got with Jones).

◙ **BOX — Plain words: more mis-alignment → less makes it through.**

B3) Birefringence & Faraday — what the numbers mean (VMS tie-back)

Birefringence (waveplate): one transverse direction is 'faster' than the other (shear anisotropy). Phase gap δ grows with thickness L: $\delta = 2\pi \Delta n L / \lambda$.

→ Quarter-wave (δ=π/2) turns equal x/y into circular; half-wave (δ=π) mirrors the linear angle around the plate axis.

Faraday/optical activity (torsion): tiny twist per meter adds up: $\theta \approx \int \kappa_t \, ds$ (classroom anchor: θ_F = V B L).

→ Use: analyzer converts tiny θ into a measurable brightness change.

Practical anchors (why this matters)

• Glare-cutting sunglasses (block horizontal polarization).

• Screens/cameras (LCDs use polarizers + retarders).

• Stress in plastics/fibers (measure Δn via δ).

• Magnetometry (read tiny θ with polarimeters).

Practice (then check solutions)

1) Two polarizers: Δ = 25°. Unpolarized I_in = 1.0. Find I_out using Stokes.

2) Quarter-wave plate: input linear at 45°, QWP (δ=π/2) fast axis along x. Describe the output polarization in words and with S.

3) Half-wave plate: input linear at 30°, HWP (δ=π) with fast axis at 15°. What is the output linear angle?

4) Birefringent plate: $\Delta n = 1.2\times10^{-5}$ at λ=633 nm. What thickness L gives a half-wave plate? (δ=π).

5) Faraday: $V = 3.8\times10^{-5}$ rad·T^{-1}·m^{-1}, B = 0.25 T, L = 0.50 m. Compute θ_F in degrees.

Worked solutions (what we computed & why)

1) First polarizer halves to 0.5. Analyzer Δ=25°: I_out = 0.5 × $\cos^2 25°$ ≈ 0.5 × 0.821 = 0.410 (alignment).

2) Equal magnitudes + 90° phase gap → circular; Stokes ≈ (S_0,0,0,±S_0).

3) HWP mirrors angle: ψ' = 2φ − ψ = 2×15° − 30° = 0° (aligned with x).

4) L = λ/(2Δn) = 633e−9 /(2×1.2e−5) ≈ 0.0264 m = 26.4 mm (thick if Δn is tiny).

5) θ_F = VBL = 3.8e−5 × 0.25 × 0.50 = 4.75e−6 rad ≈ 2.72e−4°.

Pause & Reflect (say it out loud)

→ Project / delay / rotate — then do the numbers.

→ Stokes/Mueller are just the intensity shadow of the VMS route story.

→ Bigger mis-alignment → dimmer after an analyzer.

"Companion: Math Appendix § Polarization for edge cases"

Section 7 — Imaging (Paraxial / ABCD) and Invariants

Words to Know (1-liners)

Polarized = has a favorite transverse direction. Unpolarized = no favorite direction.

Image plane = a plane where rays from one object point meet again (sharp).

Magnification m = image size / object size (negative m means flipped/inverted).

Symbols

x — ray height above the optical axis (meters)

θ — small angle the ray makes with the axis (radians)

n — medium factor (classroom: refractive index)

q := n·θ — n-weighted slope (lets invariants survive across interfaces)

M = [[A, B], [C, D]] — ABCD matrix acting on a 2-vector

Units sanity: B in meters; C in 1/meters; A and D are unitless.

VMS Bridge — why imaging works the way it does

In VMS, 'no sideways pile-up' plus smooth propagation enforce a conserved cross-flow for small angles. With q := n·θ, the product H := x·q stays constant through lossless, symmetric elements.

▣ Teaching Moment — The invariant H = x·q

→ If you squeeze width x smaller, q grows (angles spread) so the product stays the same.

→ This is why pupils and f-numbers exist: you can't shrink everything at once.

Core Relation (how we move a ray)

For one color in a uniform n, each element mixes (x, θ) linearly:

```
[ x2 ; θ2 ] = [ [A, B] ; [C, D] ] · [ x1 ; θ1 ]      with   det = AD −
BC = 1
```

If any element changes n (air→glass), switch to q := n·θ and use the 2-vector [x ; q]. This keeps det = 1 and preserves H = x·q across the system (nice and clean).

```
[ x2 ; q2 ] = [ [A, B] ; [C, D] ] · [ x1 ; q1 ]      (determinant still
1)
```

▣ What this lets me compute

→ Given any start (x1,θ1) or (x1,q1), I can find where the ray goes next.

→ If B = 0 at a plane, that plane is an image plane for an on-axis object point (rays reconverge there).

▣ Signs (keep it simple)

→ Positive f: converging (brings parallel rays to a focus).

→ Negative f: diverging (spreads rays).

→ Negative magnification m: image is inverted; |m| is the size factor.

Element Matrices (memorize these)

Free space (length L, uniform n): moves x forward by L·θ, leaves θ alone.

```
F(L) = [ [1,  L] ; [0,  1] ]     ⇒   [x2 ; θ2] = [x1 + L·θ1 ; θ1]
```

Thin lens (f): bends θ by −x/f, leaves x alone.

```
L(f) = [ [1,  0] ; [−1/f,  1] ]   ⇒   [x2 ; θ2] = [x1 ; θ1 − x1/f ]
```

Interfaces that change n: use q = n·θ so matrices still multiply with det = 1 and H = x·q stays constant.

▣ Check Yourself

→ B has units of meters; C has units of 1/meters.

→ Order matters: F(L)·L(f) ≠ L(f)·F(L).

Worked Example 1 — Derive the thin-lens formula 1/f = 1/d_o + 1/d_i

System: free space from object to lens (distance d_o), then a thin lens f, then free space to the image plane (distance d_i).

Matrix (light goes left→right): M = F(d_i) · L(f) · F(d_o)

```
F(d) = [ [1, d] ; [0, 1] ],      L(f) = [ [1, 0] ; [−1/f, 1] ]
```

Multiply (track B because 'B = 0' means 'this plane is in focus'):

```
K = L(f) ·F(d_o) = [ [1, d_o] ; [−1/f, 1 − d_o/f] ]
```

```
M = F(d_i) ·K = [ [1, d_i] ; [0, 1] ] ·K  ⇒  B = d_o + d_i − (d_o d_i)/f
```

Focus condition: set B = 0 (rays from one point meet again at that plane). Solve:

```
B = 0  ⇒  d_o + d_i = (d_o d_i)/f  ⇒  1/f = 1/d_o + 1/d_i   ← BOX
```

▣ Plain words

→ 'B = 0' is the 'no blur' condition for that plane.

→ If d_o is huge (very far object), $1/d_o \approx 0$ ⇒ the sharp plane is at $d_i \approx f$ (the focal plane).

Numbers: take d_o = 0.30 m, f = 0.050 m. Then $1/d_i = 1/0.050 − 1/0.30 = 20 − 3.33... = 16.67...$ ⇒ d_i ≈ 0.060 m (≈ 6.0 cm).

Worked Example 2 — Multiply ABCD and read magnification

Choose d_o = 0.10 m, f = 0.050 m, d_i = 0.10 m (symmetric around the lens). Compute M and read off A,B:

```
K = L(f) ·F(0.10) = [[1, 0],[−20, 1]] ·[[1, 0.10],[0, 1]] = [[1,
0.10],[−20, −1]]
```

```
M = F(0.10) ·K = [[1, 0.10],[0, 1]] ·[[1, 0.10],[−20, −1]] = [[−1,
0.0],[−20, −1]]
```

At an image plane (B = 0), the transverse magnification is m = A. Here A = −1 ⇒ same size, inverted (classic 1:1 imaging).

▣ Teen tip

→ If B ≠ 0 at your sensor plane, it will look blurry. Slide the sensor until B hits 0 for the object distance you care about.

Invariant (why q = n·θ matters when n changes)

When the beam passes through different media (air → glass → air), switching to q = n·θ keeps the determinant 1 and preserves H = x·q.

```
H_out = x2 ·q2 = (A x1 + B q1)(C x1 + D q1)  with  AD − BC = 1  ⇒
H_out = H_in  (no loss, symmetric optics)
```

▣ Pause & Reflect — What this buys you

→ You can trade beam width for angle spread, but the product stays fixed.

→ Smaller aperture (stop down) → narrower bundle x → angles spread more → more depth of field but less light.

Mini-Example — Mixed media (air → glass → air)

If there's a 2 cm glass gap (n≈1.50) between a lens and a sensor, track a ray with [x ; q] so the interface steps are clean. You'll keep det = 1, and your magnification and focus checks still work the same way (just use q).

Practice (then solutions)

1) A camera has f = 35 mm. An object is at d_o = 2.0 m. Where is the sharp image plane d_i? Is the image upright or inverted? (state m's sign)

2) Build M for: free space L1 = 0.20 m → thin lens f = 0.10 m → free space L2 = 0.30 m. Compute B. Is the final plane an image plane?

3) You reduce the aperture diameter by 2× (stop down). What happens qualitatively to x, q, and depth of field (use H = x·q)?

4) A slab of glass (n=1.50) of thickness 0.020 m sits between lens and sensor. Explain in one line why using q keeps det = 1.

Worked Solutions (what & why)

1) $1/d_i$ = 1/0.035 − 1/2.0 = 28.571 − 0.500 = 28.071 ⇒ d_i ≈ 0.036 m; m ≈ -0.018 (negative ⇒ inverted).

2) K = L(f)·F(L1) = [[1, 0.20], [−10, -1.00]]; M = F(L2)·K ⇒ B = -0.10. Since B ≠ 0, the final plane is NOT an image plane.

3) Halving aperture shrinks x (beam width). To keep H = x·q constant, q (angle spread) grows → larger depth of field but less light reaches the sensor.

4) At each interface, θ rescales with n, but q = n·θ stays continuous. Using [x ; q] keeps each element's matrix in the same det=1 family, so overall det=1.

Common Gotchas (read before lab)

• Unit slip: B must be in meters; C in 1/meters. If not, re-check your F(L) and L(f) order.

• Steep rays: if θ is not small, paraxial formulas break. Draw it; if it looks steep, don't use ABCD for that part.

• B=0 only at the in-focus plane. Don't force magnification rules at planes where B≠0 — that's blur, not image.

Section 8 — Stops, Pupils, f-Number, NA, and Etendue (with Depth of Field & Resolution) — Expanded Student Insert

WHAT & WHY (one screen)

WHAT: Define stop, pupils, f-number (f/#), numerical aperture (NA), and etendue; connect them to depth of field (DOF) and diffraction-limited resolution.

WHY: Real systems are limited by how much 'size × angle' can pass (etendue) and by diffraction of a finite pupil.

VMS Core Idea (kid-friendly)

No sideways pile-up across apertures: transverse flow continues smoothly. A smaller opening cuts off large angles, so the transported closure cannot over-concentrate. In paraxial optics this appears as a conserved product of beam width and weighted angle.

▣ The invariant

→ Let $q := n \cdot \theta$. Then $H := x \cdot q$ stays constant in lossless, symmetric systems.

→ You can't squeeze both width x and angle q at the same time.

Bridge to classical

ABCD matrices linearly mix height and angle. With $q = n \cdot \theta$, the 2-vector [x ; q] is mapped by a 2×2 matrix with determinant 1, so $H = x \cdot q$ is conserved.

Radiometry says the same thing in 3D: etendue $G = n^2 \cdot A \cdot \Omega$ (area × solid angle) is conserved in passive optics.

Vocabulary (plain English)

Stop — the tightest bottleneck in the train of optics (sets the bundle).

Entrance pupil — the stop as seen from object space (after front optics).

Exit pupil — the stop as seen from image space (before back optics).

f-number (f/#) — focal length divided by entrance-pupil diameter: $f/\# = f / D_{ent}$.

Numerical aperture (NA) — $n \cdot \sin u \approx n \cdot u$ (small angles), half-cone angle at the focus.

Etendue (G) — how much 'light phase-space' you own: $G = n^2 \cdot A \cdot \Omega$.

Algebra A — f-Number, Cone Angle, and NA (no skips)

Goal: turn geometry (f, D_ent) into angle u and NA.

1) Cone angle at focus (small angles):

    Draw a right triangle from the focal point to the pupil edge.
Half-angle u satisfies tan u = (D_ent/2)/f.

    Paraxial → small angle → tan u ≈ u, so u ≈ (D_ent/2)/f =
D_ent/(2f).

2) f-number definition:

    f/# ≡ f / D_ent ⇒ D_ent = f / (f/#).

3) Combine 1) and 2):

    u ≈ D_ent/(2f) = [f/(f/#)] /(2f) = 1 / (2·f/#).

4) Numerical aperture:

    NA ≡ n·sin u ≈ n·u (small u) ⇒ NA ≈ n /(2·f/#).

▣ Teaching Moment — sanity checks

→ Air (n≈1): f/4 → NA≈1/8=0.125; f/16 → NA≈1/32≈0.031.

→ Microscope oil (n≈1.5): the same geometry gives 1.5× larger NA.

Algebra B — From H = x·q to Etendue G (what the conservation really means)

1) Paraxial 'phase space' for one transverse axis is the rectangle Δx
by Δq. The matrix map [x2;q2] = M [x1;q1] has det M = 1 → area Δx·Δq
is preserved.

2) Because q = n·θ, we can write Δx·Δq = Δx·(n·Δθ).

3) In 2D (x and y), multiply the two axes: ΔA·(n²·ΔΩ) is preserved →
that's etendue G = n²·A·Ω.

4) Practical readout: shrinking A forces Ω to grow. A stop that cuts diameter in half must let angles double (if lossless).

▣ Teaching Moment — camera vs. telescope

→ Cameras trade angle for depth of field by stopping down (smaller A ⇒ larger angle spread after focus).

→ Telescopes chase big A (large D) to keep angles tiny and resolve fine detail.

Algebra C — Depth of Field (DOF) from similar triangles

Setup: a point that should focus at the image plane is actually formed Δz away from the sensor. The lens still forms a cone; the sensor slices it off as a disk (circle of confusion).

1) Blur diameter from similar triangles:

```
    At the image plane: cone half-angle u ≈ D_ent/(2 d_i). If the sensor
is |Δz| from focus, the disk radius r ≈ |Δz|·u.

    So diameter  c = 2r ≈ 2|Δz|·u ≈ 2|Δz|·(D_ent/(2 d_i)) =
(D_ent/d_i)|Δz|.
```

2) Allowable defocus for a chosen tolerance c_max:

```
    |Δz| ≤ (c_max · d_i) / D_ent.
```

3) Tie to f/# using D_ent = f/(f/#) and d_i ≈ f for distant objects:

```
    |Δz| ≲ c_max · (f/#).    ← super-useful rule of thumb.
```

▣ Teaching Moment — hyperfocal idea (kid version)

→ There's a far distance where everything to infinity looks acceptably sharp at a given f/#.

→ Stop down (bigger f/#) shrinks the cone → the same Δz makes a smaller blur.

Algebra D — Diffraction-Limited Resolution (why too much f/# gets soft)

```
1) Circular pupil far-field (Airy): first dark ring at sin θ_R ≈ 1.22·λ
/ D_ent.
```

```
2) Small angles: θ_R ≈ 1.22·λ / D_ent. Image-plane radius r ≈ f·θ_R.
```

```
3) Substitute D_ent = f/(f/#):  r ≈ 1.22·λ·(f/#).
```

▣ Teaching Moment — DOF vs Diffraction

→ DOF blur ∝ 1/(f/#) (bigger f/# helps). Diffraction blur ∝ (f/#) (bigger f/# hurts).

→ Pick f/# to balance your blur tolerance (c_max) and color (λ).

Finding the Pupils (step-by-step recipe)

1) Identify the physical stop (tightest bottleneck).

2) Look toward it from the front through any glass → the apparent diameter is the entrance pupil D_ent.

3) Look toward it from the back through the rest → the apparent diameter is the exit pupil D_exit.

4) Use D_ent to compute f/# (brightness seen from object space). Use D_exit to check sensor coverage and vignetting.

Worked Examples (numbers with steps)

Example 1 — f/number ↔ NA (air):

```
Given f/# = 4.  u ≈ 1/(2·f/#) = 1/8 = 0.125 rad ≈ 7.16°.   NA ≈ n·u ≈
1·0.125 = 0.125.
```

Example 2 — Resolution & DOF together:

```
λ = 550 nm, f/# = 8 → diffraction spot radius r ≈ 1.22·λ·(f/#) ≈
1.22·5.5e-7·8 ≈ 5.37 µm.
```

```
If c_max = 10 µm and d_i ≈ f = 50 mm with D_ent = f/(f/#) = 6.25 mm,
```

```
then from DOF rule:  |Δz| ≲ c_max·(f/#) ≈ 80 µm (rough, but useful).
```

Example 3 — Telescope resolution (angular):

```
D = 200 mm, λ = 500 nm → θ_R ≈ 1.22·λ/D ≈ 3.050e-06 rad ≈ 0.63 arcsec.
```

Practice Problems (with solutions)

1) A 50-mm lens at f/2.8 in air: compute u and NA.

2) Using the DOF rule, what f/# keeps |Δz| ≤ 120 µm when c_max = 6 µm?

3) A microscope objective in air has NA = 0.25. What f/# does that roughly match? (use NA ≈ 1/(2·f/#)).

4) D_ent seen from object space is 12.5 mm for a 100-mm focal length lens. What is f/# and the cone half-angle u?

5) If you halve D_ent while keeping f fixed, explain what must happen to Ω (solid angle) so G stays constant.

Worked Solutions

```
1) u ≈ 1/(2·2.8) ≈ 0.1786 rad ≈ 10.2°. NA ≈ 0.179.
```

```
2) f/# ≳ |Δz|/c_max = 120/6 = 20 → f/20.
```

```
3) f/# ≈ 1/(2·NA) = 1/(0.5) = f/2.
```

```
4) f/# = f/D_ent = 100/12.5 = 8 → f/8;  u ≈ 1/(2·8) = 1/16 = 0.0625
rad ≈ 3.58°.
```

```
5) Halving D_ent halves A; to keep G = n²·A·Ω, Ω must roughly double
(paraxial, small angles).
```

Common Gotchas (read before lab)

• DOF rule uses d_i ≈ f for distant objects; for closer objects use the exact d_i for tighter bounds.

• Paraxial means small angles; if u is big (say > ~10°), sines/tangents deviate and these approximations drift.

• Etendue conservation assumes lossless, passive optics; real stops and baffles reduce G.

Wrap-Up (what to remember)

Smaller f/# → brighter but shallower DOF and more chance of aberrations; larger f/# → deeper DOF but softer from diffraction.

Bigger D_ent → smaller Airy spot → finer detail (if the optics are good).

You can't beat etendue: squeeze size and the allowed angles swell to compensate.

Section 9 — Chromatic Dispersion & Achromats (Why Lenses Have Color and How We Fix It)

WHAT & WHY (one screen)

WHAT: Colors bend differently because the medium's slowness n depends on wavelength λ. We quantify that bending, predict color blur, and design a two-glass achromat that cancels it (for two colors).

WHY: Cameras, microscopes, telescopes, and VR lenses all fight color fringing. These steps let you estimate and fix it with numbers you can check.

VMS Core Idea (plain words first)

In VMS, routes carry closure through space. The space response (encoded in n) changes slightly with color (λ). At boundaries and inside a lens, 'no sideways pile-up' still holds, but each color uses its own n(λ). So blue and red follow slightly different bending → different best-focus distances unless we design around it.

▣ Teaching Moments — what actually changes with color?

→ Same geometry, but use n(λ) for each color. Larger n usually bends more (blue > red in normal glasses).

→ We can pair two glasses so the extra blue bend cancels the red shortfall → achromat.

Bridge to Classical (what math we'll use)

Thin-lens power at color λ: $\varphi(\lambda) = (n(\lambda) - 1) \cdot S$, where $S := (1/R_1 - 1/R_2)$ is the shape factor (thin lens).

Focal length: $f(\lambda) = 1 / \varphi(\lambda)$. Small color changes → small focus shifts.

Abbe number (dispersion fingerprint): $V = (n\_d - 1) / (n\_F - n\_C)$, using d=587.6 nm, F=486.1 nm, C=656.3 nm.

▣ What this lets me compute

→ Given glass data (n_d, V), I can estimate how much the focus changes between blue (F) and red (C).

→ Given a target f, I can choose two powers ϕ_1, ϕ_2 that cancel the blue–red focus spread.

Vocabulary

Axial chromatic aberration (ACA): different colors focus at different distances (front–back color).

Lateral chromatic aberration (LCA): magnification changes slightly with color (edge fringing).

Abbe number V: bigger V = lower dispersion (colors stay together better).

Achromat: two thin lenses in contact (or close) that cancel ACA at two colors (typically F and C).

Algebra A — Single-Lens Chromatic Focus Shift (no skips)

1) Power and focal length:

```
φ(λ)  =  (n(λ)  −  1) ·S  ;     f(λ)  =  1/φ(λ).
```

2) Small change from d-line (use a derivative):

```
Δφ  ≈  S  ·  Δn     because  φ  depends  linearly  on  n.

Since  f  =  1/φ,    df/dφ  =  −1/φ²  =  −f²   ⇒   Δf  ≈  −  f_d²  ·  Δφ   (for
small  changes).

Combine:   Δf  ≈  −  f_d²  ·  S  ·  Δn   .
```

▣ Teaching Moment — signs & sanity

→ If blue has larger n than red, then ϕ_blue > ϕ_red → f_blue < f_red (blue focuses closer).

→ Using the Abbe lines (F & C), the blue–red difference sets the worst-case focus spread.

Algebra B — Use V to estimate the blue–red focus spread

```
Abbe:   V  =  (n_d  −  1)/(n_F  −  n_C)   ⇒   (n_F  −  n_C)  =  (n_d  −  1)/V.
```

```
Blue-red power difference for one lens:  Δφ_FC = φ_F − φ_C = S·(n_F −
n_C) = S·(n_d − 1)/V.
```

```
Approximate blue-red focal shift (centered near f_d):   Δf_FC ≈ − f_d²
· Δφ_FC = − f_d² · S·(n_d − 1)/V.
```

```
Replace S using φ_d = (n_d − 1)S = 1/f_d  ⇒  S = 1/[(n_d − 1) f_d].
```

```
Then  Δf_FC ≈ − f_d²  · [ (1/((n_d − 1) f_d))  · (n_d − 1)/V ] = − f_d /
V.
```

◨ Result to remember

→ For a simple thin lens, the blue–red focus gap scale is roughly $|\Delta f\_FC| \approx f\_d / V$.

→ Lower V (more dispersive) → larger color focus shift.

Algebra C — Achromat Condition (derive it, don't just quote it)

Let the d-line powers be φ_1, φ_2 with Abbe numbers V_1, V_2. Two thin lenses in contact → total power $\varphi = \varphi_1 + \varphi_2$.

Goal: keep total blue-red power change ≈ 0 so focus stays put: $\Delta\varphi\_total \approx 0$.

Blue-red change per lens: $\Delta\varphi\_i \approx \varphi\_i / V\_i$ (from the V-based estimate above).

Set the sum to zero: $\Delta\varphi\_1 + \Delta\varphi\_2 = 0$ ⇒ $\varphi_1/V_1 + \varphi_2/V_2 = 0$.

Also hit the target focal length f_target ($\varphi\_target = 1/f\_target$): $\varphi_1 + \varphi_2 = \varphi\_target$.

Solve the two equations (showing the elimination):

   From $\varphi_1 + \varphi_2 = \varphi\_target$ ⇒ $\varphi_2 = \varphi\_target - \varphi_1$.

   Plug into $\varphi_1/V_1 + \varphi_2/V_2 = 0$:

      $\varphi_1/V_1 + (\varphi\_target - \varphi_1)/V_2 = 0$ ⇒ $\varphi_1(1/V_1 - 1/V_2) = -\varphi\_target / V_2$

      $\varphi_1 = \varphi\_target \cdot [(-1/V_2) / (1/V_1 - 1/V_2)] = \varphi\_target \cdot V_1 / (V_1 - V_2)$

   Then $\varphi_2 = \varphi\_target - \varphi_1 = - \varphi\_target \cdot V_2 / (V_1 - V_2)$.

◨ Teaching Moments — signs and glass choices

→ Crown: high V (low dispersion), usually positive power.

→ Flint: lower V (higher dispersion), usually negative power.

→ Achromat: pick $\varphi_1 > 0$ in crown, $\varphi_2 < 0$ in flint, sized by the formulas so total φ hits the target and color shift cancels.

Worked Example — Design a 100 mm Achromat (BK7 + F2)

d-line data: BK7 ~ n_d=1.5168, $V_1 \approx 64.17$; F2 ~ n_d=1.6200, $V_2 \approx 36.37$.

Target: f = 0.100 m ⇒ φ_target = 10.00 D.

Formulas give: $\varphi_1 \approx 23.08$ D (crown, +), $\varphi_2 \approx -13.08$ D (flint, −).

Implied element focal lengths: $f_1 \approx 43.3$ mm, $f_2 \approx -76.4$ mm.

Checks: (i) $\varphi_1 + \varphi_2 \approx \varphi\_target$, (ii) $\varphi_1/V_1 + \varphi_2/V_2 \approx 0$ (blue-red shift cancels).

◉ Plain words

→ The positive crown does most of the focusing.

→ The weaker negative flint undoes the color spread while keeping the net focal length at 100 mm.

→ Spacing and shapes fine-tune the design in real lenses, but this gets you to 'no obvious rainbow' fast.

Algebra D — Lateral Chromatic (edge color fringing)

Magnification m depends on $\varphi(\lambda)$ and the conjugates. For small changes across color, $\Delta m/m \approx - \Delta f/f \approx - \Delta\varphi/\varphi$.

Achromatizing $\varphi'(color) \to 0$ shrinks LCA too, but best correction uses element spacing and shapes.

◉ Teaching Moment — how to spot LCA

→ Edge of a high-contrast object shows blue on one side and red on the other when m changes with color.

→ Stopping down reduces both ACA and LCA by narrowing the bundles, but doesn't change dispersion itself.

Extra Clarity (kid-level anchors)

◉ Signs & units quick list

→ φ in diopters (D = 1/m). Positive φ focuses; negative φ spreads.

→ V has no units. Bigger V = less dispersion.

→ If numbers explode, check: did you mix mm and m? Did you use the right V definition (F,d,C lines)?

◉ Myths to avoid

→ 'Big f/# fixes color' — it reduces the visible blur (smaller cone), but dispersion still exists.

→ 'All crowns/flints are the same' — V varies widely by glass type; look up n_d and V.

◉ Lab demo ideas

→ Shine white light through a single lens onto a screen: measure f for blue vs red filters.

→ Stack a weak negative lens with a positive lens: adjust spacing until the blue/red spots overlap.

Practice Problems (then solutions)

1) Single lens: S = 20 m^{-1}, n_d=1.50 → φ_d=?, f_d=?; n_F=1.51, n_C=1.49 → estimate Δf_FC using Δf ≈ − f_d^2·S·Δn.

2) Achromat target f=150 mm. Take V_1=60 (crown), V_2=30 (flint). Solve $φ_1$, $φ_2$ exactly.

3) If an old singlet shows |Δf|/f ≈ 0.8%, estimate V from |Δf|/f ≈ 1/V.

4) You need f=85 mm at d-line with V_1=62, V_2=35 and choose $φ_1$=+18 D. Solve $φ_2$ and check the achromat condition.

Worked Solutions (what & why)

```
1) φ_d=(1.50-1)·20=10 D ⇒ f_d=0.100 m. Δn=0.02 →
Δf≈-f_d²·S·Δn=-(0.1²)·20·0.02=-0.004 m (-4.0 mm).

2) φ_target=6.667 D. φ₁= 13.333 D,  φ₂= -6.667 D (negative). Check sums
and φ₁/V₁+φ₂/V₂≈0.

3) |Δf|/f ≈ 0.008 ⇒ V ≈ 1/0.008 ≈ 125 (very low-dispersion behavior).

4) φ_target≈11.765 D; with φ₁=+18 D ⇒ φ₂≈-6.235 D. Achromat check
φ₁/V₁+φ₂/V₂≈0.1122 (aim for ≈0; tweak with shapes/spacing).
```

Wrap-Up (remember this)

Colors bend differently because n(λ) changes → single lenses have ACA/LCA. Two-glass achromats cancel blue/red focus shift while keeping the target f.

Key calculator: solve $φ_1$+$φ_2$=φ_target and $φ_1$/V_1+$φ_2$/V_2=0. Use Abbe number V to estimate color shifts quickly.

Section 10 — Aberrations & Wavefront Error (why real images aren't perfect)

WHAT & WHY

WHAT: If the eikonal S deviates from the ideal quadratic that focuses rays, the slopes ∇S are off a little, and rays miss a perfect crossing → blur patterns (coma, spherical, astigmatism…).

WHY: Writing the wavefront error W(x,y) and taking slopes gives you spot size and shape. That points to fixes: stop down, tweak shapes, move elements, swap glass, or use aspheres.

VMS Core Idea (plain words first)

In VMS, rays follow the normals to equal-S surfaces. Ideal imaging means equal-S surfaces are perfect spheres aimed at the image point. Any extra wrinkle W in S tilts the normals → tilt the rays → miss the point.

▣ Teaching Moment — 'Wavefront' in one line

→ Equal-S surfaces are like contour lines on a hill. Rays are the arrows pointing straight downhill (normals). Extra bumps (W) tip the arrows.

Bridge to Classical

We split S into an ideal part S_id (does perfect focusing) and an error W (small). Using small-angle paraxial optics:

```
S(x,y) = S₀ + S_id(x,y) + W(x,y).  Rays tilt by ∇S; the extra tilt is
from ∇W.
```

Algebra A — From W(x,y) to ray tilt and image shift (no skips)

```
1) Extra angle (small):  δθ_x ≈ (1/n) ∂W/∂x,   δθ_y ≈ (1/n) ∂W/∂y.
```

What this computes: how much a ray's direction is nudged by the local wavefront slope.

```
2) Spot shift at the focal plane (distance f):  Δx ≈ f·δθ_x,  Δy ≈
f·δθ_y.

   Combine:  Δx ≈ (f/n) W_x,  Δy ≈ (f/n) W_y   with W_x≡∂W/∂x.
```

▣ Result — slope→spot map

→ Take derivatives of W, multiply by f/n. Bigger pupil (larger |x|,|y|) → bigger slopes → bigger spots.

Algebra B — Defocus as a quadratic wavefront (connect Δz to W)

Defocus is the simplest 'aberration': the image plane is shifted by Δz from best focus. That adds (or removes) curvature from S_id.

1) Ideal quadratic near the pupil (object at infinity): S_id ≈ $(x^2+y^2)/(2f)$.

2) Move the screen by Δz along z. The required added quadratic to refocus is proportional to r^2 with $r^2=x^2+y^2$.

   Write W_def(r) = K · r^2 and solve for K by matching the change in ray slope to the angular change from moving the plane.

3) A small axial move changes the angular convergence by roughly $\Delta(\theta)$ ≈ r·$\Delta(1/f)$ with $\Delta(1/f)\approx-\Delta z/f^2$ (paraxial), giving K on the order of $\Delta z/(2f)$.

   Takeaway: W_def ∝ $(\Delta z/f)\, r^2$. (Constants depend on convention; the key is: quadratic in r, proportional to Δz.)

▣ Why this matters

→ Defocus W ∝ r^2 explains why stopping down (shrinking r) reduces blur disks fast.

→ It also lets you convert 'turn the focus ring by Δz' into an equivalent W for calculations.

Algebra C — What common W terms do (Seidel/Zernike intuition)

• Spherical (on-axis): W ∝ r^4 → symmetric blur; best focus shifts with pupil size.

• Coma (off-axis, field height h): W ∝ h·r^3·cosϕ → comet tails pointing away from the field center.

• Astigmatism/Field curvature: W ∝ h^2·r^2 cos2ϕ (astig), and W ∝ h^2·r^2 (field) → different tangential/sagittal foci, curved best-focus surface.

• Distortion: terms that change image magnification with field (straight lines bend).

▣ Zernike snapshot (unit disk)

→ Z2,0 ∝ r^2 (defocus); Z3,±1 ∝ r^3 cosϕ,sinϕ (coma); Z4,0 ∝ r^4 (spherical); Z2,±2 ∝ r^2 cos2ϕ,sin2ϕ (astig).

Worked Example A — Coma spot scaling (explicit derivatives)

Model: W(x,y)=C·h·r^3 cosφ. In Cartesian, r^3 cosφ = (x^2+y^2)·x.

Derivatives: W_x=C·h·$(3x^2+y^2)$,   W_y=C·h·(2xy).

Edge (|r|=r_max) gives scale of slopes ~ C·h·r_max². Spot half-size ~ (f/n)·C·h·r_max².

▣ Plain words

→ Coma falls very fast when you stop down: r_max→r_max/2 → blur ~ ×1/4.

Worked Example B — Spherical aberration & best focus shift

```
Take W(r)=S·r⁴ (axisymmetric). Slope: W_r = 4S r³. Average slope across
the pupil drives a mean angular error.
```

```
Shift the sensor by Δz to cancel the average slope at focus; the
required Δz scales like S·r_max².
```

▣ Takeaway

→ Smaller pupils reduce both blur and the focus shift from spherical.

Worked Example C — Astigmatism: two focal planes

```
Let W=B·h²·r² cos2φ. Slopes differ along tangential vs sagittal
directions (φ=0 vs φ=π/2).
```

```
Compute along x (tangential line y=0): cos2φ=cos0=1 → W_x=2B·h²·x;
along y (sagittal x=0): cos2φ=cosπ=-1 → W_y=-2B·h²·y.
```

```
Changed slopes of opposite sign shift best focus to two z-locations
separated by an amount ∝ B·h²·r_max².
```

▣ Plain words

→ Horizontal lines and vertical lines focus at slightly different distances → one looks sharp while the other looks soft.

Practice (with solutions later)

1) Show that defocus $W=A\,r^2$ produces a uniform extra slope $\propto r$; match its average to a small Δz to link A and Δz.

2) If r_max halves, by what factors do (a) coma blur and (b) spherical blur change? Justify from the r-powers in W.

3) For $W=k \cdot h \cdot x(x^2+y^2)$, compute W_x and W_y. Which way does the tail point for +h vs −h?

Section 11 — Fourier Optics & MTF/OTF (how aperture → detail)

WHAT & WHY

WHAT: The point-spread function (PSF) is the image of a point. In the focal plane under paraxial conditions, the field is (scaled) a Fourier transform of the pupil. For incoherent scenes, intensities add and the OTF/MTF describe contrast vs spatial frequency.

WHY: MTF predicts which details survive. It also shows when stopping down helps (cuts aberrations) and when it hurts (diffraction).

VMS Core Idea

Each route across the pupil contributes with a phase set by path length. At focus those phased contributions sum like a Fourier-type integral — the aperture 'selects' which tilts (spatial frequencies) pass.

▣ Teaching Moment — 'Aperture as recipe'

→ Think: each point in the pupil adds a tiny wave aimed a little differently. The sum of all those tiny waves draws the PSF.

Bridge to Classical (definitions we'll use)

Pupil function P(x,y): 1 inside the clear aperture (include a phase term e^{iφ(x,y)} if aberrated).

Coherent field at focus: U_img(ξ,η) ∝ FT{P}(ξ/(λ f), η/(λ f)). Intensity = |U_img|².

Incoherent imaging: OTF(f_x,f_y) = normalized autocorrelation of P; MTF = |OTF|.

▣ Cutoffs

→ Circular clear pupil (no aberration): angular cutoff θ_c≈D/(2f) (in radians).

→ Linear frequency cutoff at the sensor: f_c = D/(λ f).

Algebra A — From Fresnel to Fraunhofer (scaled FT)

Start from scalar diffraction (Fresnel integral). In the focal plane and for small angles, the phase becomes linear in x,y → a Fourier transform with scaling 1/(λ f).

Define spatial frequencies u=ξ/(λ f), v=η/(λ f). Then U_img(u,v) ∝ ∫∫ P(x,y) e^{-i2π(ux+vy)} dx dy.

What this computes: how pupil shape/phase maps to the amplitude pattern at focus.

Algebra B — Circular pupil: PSF & MTF formulae (write them down)

Amplitude PSF ∝ (2J₁(ka sinθ))/(ka sinθ), intensity PSF = Airy pattern. First dark ring at θ_R≈1.22 λ/D.

Incoherent MTF for a perfect circular pupil (radial frequency ν=f/f_c):

MTF_circ(ν) = (2/π)[cos^{-1}(ν) − ν √(1−ν²)],   for 0≤ν≤1; 0 beyond.

▣ What this lets me compute

→ Given D, f, and λ, compute f_c and then MTF at any fraction of f_c to estimate contrast loss.

Algebra C — 1-D rectangular pupil (kid-level triangular MTF)

```
Let the pupil be a 1-D rectangle of width D. Coherent field is sinc;
the incoherent OTF is the autocorrelation of a rectangle → triangle.
```

```
MTF_rect(f) = 1 - |f|/f_c  for |f|≤f_c;  0 beyond, with f_c=D/(λ f).
```

▣ Teaching Moment — 'slide and overlap'

→ The OTF is literally 'how much two copies overlap when one slides.' A rectangle sliding over a rectangle overlaps linearly → triangle.

Algebra D — Sampling & pixel MTF (don't ignore the sensor)

```
Pixel response (square pixel, pitch p) adds its own MTF:  MTF_px(f) =
|sinc(π f p)| along each axis.
```

```
System MTF ≈ MTF_optics · MTF_px · MTF_motion · MTF_demosaic …
(product of the major stages).
```

```
Nyquist frequency: f_N = 1/(2p). If f_c ≫ f_N, sampling limits detail
even if optics are great.
```

▣ Plain words

→ If optics beat the pixels, shrinking pixels (or binning differently) matters more than opening the lens further.

Worked Example — f/4 camera vs pixels (numbers)

```
f=100 mm, D≈25 mm, λ=550 nm → f_c ≈ 0.455 cycles/µm. Pixel 4 µm → f_N
≈ 0.125 cycles/µm.
```

Since f_c ≈ 0.455 cycles/µm and f_N ≈ 0.125 cycles/µm, pixels limit sharpness. Stopping down a little may still help if aberrations are present.

Practice (with solutions later)

1) For f=50 mm, f/2, λ=550 nm, compute f_c (cycles/mm). Compare to 3.2-µm pixel Nyquist.

2) Stop down from f/2.8 to f/8 (same f, λ). By what factor does f_c change? Why might images look sharper anyway?

3) A square pixel p=5 µm: what is MTF_px at 50 lp/mm (i.e., f=0.05 cycles/µm)?

Section 12 — Detectors, Noise & SNR (how dim can we see?)

WHAT & WHY

WHAT: Photon arrivals jitter (shot noise). Sensors add read noise (electronics) and sometimes dark current. The combined noise sets how faint you can measure. SNR gives you the 'is this measurement real?' number.

WHY: Lets you choose exposure time, f/#, and ISO intelligently — and know when stacking or a brighter lens actually helps.

VMS Core Idea

Counts over time are outcomes of many routes; randomness is fundamental (Poisson). Even with perfect optics, you still have shot noise.

Bridge to Classical (formulas we'll use)

```
Shot noise: variance = mean. If you collect N_e electrons, shot noise
≈ √N_e.

Read noise σ_r adds in quadrature; dark current adds N_dark
(electrons).

Total noise:  σ_tot = √(N_e + σ_r² + N_dark).   SNR = N_e / σ_tot.

Electrons gathered:  N_e = η · E_img · t · A_px, with η quantum
efficiency, E_img irradiance at the sensor, t time, A_px pixel area.
```

▣ Teaching Moment — f/# and light on sensor

→ For a uniform scene, $E_{img} \propto 1/(f/\#)^2$. Two stops (×4 in f/#) change electrons ×4 if time and ISO fixed.

Algebra A — Solve for exposure to hit a target SNR (full steps)

```
Given SNR₀ and read noise σ_r, solve SNR = N_e/√(N_e + σ_r²) for N_e:

SNR² (N_e + σ_r²) = N_e²  →  N_e² − SNR² N_e − SNR² σ_r² = 0.

N_e = ½ [ SNR² + √(SNR⁴ + 4 SNR² σ_r²) ].  (If σ_r→0, this → SNR².)

Convert N_e to time:  t = N_e/(η E_img A_px). With E_img ∝ 1/(f/#)²
you can trade aperture vs time.
```

▣ Plain words

→ If read noise is tiny, shoot until ~SNR^2 electrons. If read noise is big, front-load more light (open aperture or extend time) until $N_e \gg \sigma_r^2$.

Algebra B — Stacking exposures (averaging N frames)

Stack N equal frames (each with N_e electrons on average, read noise σ_r).

Signal sums: $N \cdot N_e$. Noise adds in quadrature: $\sqrt{N \cdot (N_e + \sigma_r^2)}$.

$SNR\_stack = (N \cdot N_e)/\sqrt{N \cdot (N_e + \sigma_r^2)} = \sqrt{N} \cdot N_e/\sqrt{N_e + \sigma_r^2}$.

Shot-noise regime ($N_e \gg \sigma_r^2$): $SNR \propto \sqrt{N \cdot N_e}$. Read-noise regime: need N large enough that $N_e \gtrsim \sigma_r^2$ per frame.

▣ Teaching Moment — When stacking helps

→ If you're read-noise-limited, fewer longer frames are better than many tiny ones.

→ If you're shot-noise-limited, many shorter frames stack almost as well as one long one.

Worked Example — Is this exposure enough? (expanded)

Pixel 4×4 μm² (A_px=16 μm²), η=0.6, σ_r=2 e⁻ rms, target SNR=20. Scene yields $E\_img=2.0\times10^{13}$ photons/(m²·s).

Electrons needed: $N_e \approx \frac{1}{2}[SNR^2 + \sqrt{SNR^4 + 4\ SNR^2\ \sigma_r^2}] \approx 404.0\ e^-$.

Electron rate $\dot{N}_e = \eta\ E\_img\ A\_px \approx 192.0\ e^-/s \rightarrow$ time t ≈ 2.104 s to reach SNR≈20.

If f/#→f/#·√2 (1 stop down), E_img halves → time doubles. If pixel is 8×8 μm (area ×4), time quarters for same SNR.

Practice (with solutions later)

1) σ_r=5 e⁻, N_e=400 e⁻ → SNR? (Use $\sigma\_tot = \sqrt{N\_e + \sigma_r^2}$.)

2) Target SNR=30, σ_r negligible. How many electrons do you need?

3) If you double pixel pitch (area ×4) and keep all else the same, how does SNR change in the shot-noise regime?

4) You stack N=9 frames, each SNR=10 (shot-noise-limited). What is the stacked SNR?

Short Answers (for practice)

10-1) W=A r² is quadratic; match its average slope to the slope from moving the screen by Δz to link A∝Δz/f.

10-2) r_max→r_max/2 ⇒ coma blur ×1/4 (∝r_max²), spherical slope terms ×1/8 (∝r_max³); spot area shrinks even faster.

10-3) $W\_x=\partial/\partial x[k\ h\ x(x^2+y^2)]=k\ h(3x^2+y^2)$; $W\_y=2k\ h\ x\ y$. Tail points away from field center; flipping h flips the tail.

11-1) $f\_c = D/(\lambda\ f)$ then compare to pixel Nyquist $1/(2p)$. Lower one limits detail; convert to cycles/mm as needed.

11-2) $f\_c \propto D \propto 1/(f/\#)$. f/2.8→f/8 cuts D by ~2.86× so f_c drops by 2.86×, but reduced aberrations can raise usable MTF.

11-3) MTF_px = $|\text{sinc}(\pi\ f\ p)|$. For f=0.05 cycles/μm and p=5 μm, argument=$\pi\cdot0.25$→MTF_px≈$|\sin(0.785)/0.785|$≈0.90.

12-1) SNR=$400/\sqrt{(400+25)}$≈19.4. 12-2) Need N_e=$(30)^2$=900 e⁻. 12-3) Area ×4 → electrons ×4 → SNR ×2. 12-4) √N scaling → $\sqrt{9}$×10=30.

Wrap-Up (what to remember)

Aberrations: compute slopes of W to get ray tilts → spot maps. Bigger pupil = bigger slopes; stopping down shrinks high-order terms fast.

Fourier optics: the pupil and its phase decide the PSF/MTF. Use f_c and closed-form MTFs to estimate contrast at frequencies you care about.

Detectors: SNR tells you if detail is usable. Solve for electrons→time, then decide: open aperture, expose longer, or stack frames.