Optical fields, instruments and observations

Trace an optical field through a physical instrument. Calculate the signal a detector measures and test the approximations on which the interpretation rests.

Contents

1. **What an objective must collect to resolve a specimen**
   A microscope can transmit plenty of light and still erase fine detail. Treat the specimen as a weak grating, follow its spatial harmonics through the objective, and determine why changing illumination coherence changes the transfer of contrast.
2. **From a slowly varying wave to a curved ray in a graded index**
   A ray bends continuously in a glass whose index changes across its width. Derive where that ray equation comes from, predict the focusing length of a graded-index rod, and locate the points at which a smooth ray description ceases to determine an optical amplitude.
3. **Measure rotation with two waves following the same path**
   A stationary optical loop gives equal return times for opposite propagation directions in a reciprocal medium. When the loop rotates the beams meet their starting beamsplitter at different times. Derive the measurable differential phase and compare a passive interferometer with a ring resonator.
4. **Dispersive pulse compression and a detector's view of a tilted pulse**
   An ultrashort pulse can broaden because its spectral components acquire different group delays or because different parts of its beam arrive at different times. Derive the two mechanisms separately, predict when added glass can shorten an already chirped pulse, and decide which measurements distinguish them.
5. **Which polarisation measurements identify an optical element?**
   A retarder rotates the state of a coherent beam, while unresolved mixtures can reduce its measured degree of polarisation. Derive the passage from Jones fields to a Mueller matrix, reconstruct a simple instrument from measured intensities, and identify claims the available input states cannot support.
6. **Photon counting, bunching and the time resolution of a coincidence test**
   An intensity trace can be steady while individual detector clicks show correlations. Derive the observable of a two-detector Hanbury Brown–Twiss measurement, quantify how background and

bin width change it, and decide whether a dip near zero delay belongs to the source or to a detector.

7. Build a time lens from dispersion and a quadratic temporal phase

A temporal imaging system can stretch the time separation of two optical features without requiring them to arrive more slowly at the source. Derive the imaging condition from delay and frequency, calculate the sign of time magnification, and estimate the finite temporal aperture that limits which features can be recovered.

8. Recover an optical phase from one off-axis intensity frame

A camera measures intensity and discards field phase. A known tilted reference wave can shift a cross term away from the intensity background so the missing complex field becomes accessible. Derive how much camera bandwidth that costs, and convert the recovered phase into an optical-path measurement with a stated ambiguity.

9. Resolve two nearby wavelengths by watching fringes disappear

A Michelson interferometer shows very fine bright and dark fringes as one mirror moves. Two narrow spectral lines produce a much slower variation of the fringe contrast. Use the mirror travel between contrast minima to determine the separation of the lines.

10. Where a thick lens focuses and which aperture an instrument actually sees

The distance from a lens mount to a sharply focused sensor is often shorter than the focal length printed on the barrel. Two principal planes, together with the image of the physical aperture stop, reconcile these measurements and determine which rays an instrument accepts.

11. From measured wavefront error to coma, astigmatism and focus

A point source can become a comet-shaped patch or two perpendicular line foci. Describing these observations by a single blur diameter loses the distinction between aberrations and suggests the wrong correction. The gradient of a pupil wavefront predicts the ray displacement at the detector.

12. Why light can generate its second harmonic and then lose it again

A nonlinear polarisation can oscillate twice during one period of the applied field. Whether the generated wave grows through a crystal depends on how successive source layers arrive in phase. This spatial phase relation makes crystal length and dispersion as important as local nonlinearity.

13. How a material changes the speed, attenuation and direction of light

A detector behind a sample loses light both when matter absorbs energy and when light is scattered outside its acceptance. A microscopic oscillator connects refraction to absorption. A

radiating dipole then shows how a transparent particle can also reduce the measured direct beam.

Optical fields, instruments and observations · 1. What an objective must collect to resolve a specimen

A microscope can transmit plenty of light and still erase fine detail. Treat the specimen as a weak grating, follow its spatial harmonics through the objective, and determine why changing illumination coherence changes the transfer of contrast.

Diffract a sinusoidal object into orders

Use a thin specimen with amplitude transmission $t(x) = 1 + \epsilon\cos(2\pi f x)$, $|\epsilon| \ll 1$, and uniform normally incident coherent light. Write the cosine as half the sum of two exponentials. Immediately after the sample the field has spatial frequencies zero and $\pm f$. A lens maps transverse spatial frequency into angle, and its finite numerical aperture $\mathrm{NA} = n\sin\theta_{\max}$ accepts $|f| \le \mathrm{NA}/\lambda_0$ for coherent normal illumination, where λ_0 is the vacuum wavelength. Beyond that limit both first orders are blocked, leaving only a constant image field and no sinusoidal contrast.

If the zero order and both sidebands pass through a clear ideal pupil, inverse Fourier transformation reproduces $1 + \epsilon\cos(2\pi f x)$. Its intensity is approximately $1 + 2\epsilon\cos(2\pi f x)$. If only one sideband and the zero order pass, their interference still produces a feature at f, but the first-order intensity contrast is halved for this object and the image field acquires a spatially varying phase. If the zero order is suppressed as well, the intensity contrast follows a different mixing law. These possibilities show why the statement that one needs two interfering orders must specify which orders and which image quantity is measured.

Oblique coherent illumination shifts the object's spatial spectrum by the incident transverse frequency f_i. A sideband initially outside the pupil can enter while one from the opposite direction exits. With illumination aperture $\mathrm{NA}_{\mathrm{illum}}$ and objective aperture $\mathrm{NA}_{\mathrm{obj}}$, the extreme accessible object frequency from angular coverage can reach $(\mathrm{NA}_{\mathrm{illum}} + \mathrm{NA}_{\mathrm{obj}})/\lambda_0$. That support limit alone does not guarantee uniform contrast or isotropic information. The illumination distribution and reconstruction determine how well any frequency inside the support can be measured.

Derive the incoherent transfer from a pupil autocorrelation

For spatially incoherent emitters, intensities from different object points add. If h is the coherent amplitude point-spread function, the incoherent intensity point-spread function is $|h|^2$, and the intensity image is its convolution with object brightness. Fourier transforming the product $|h|^2$ produces an autocorrelation of the pupil P. A clear circular pupil gives a transfer value proportional to the overlap of two equal disks displaced by a spatial-frequency-dependent distance. At zero shift the disks overlap fully; at a diameter of shift their overlap vanishes. The incoherent cutoff is therefore $f_c = 2\mathrm{NA}/\lambda_0$, twice the coherent on-axis field cutoff.

For $u = |f|/f_c \le 1$, the normalised overlap area is

$$\mathrm{MTF}(u) = \frac{2}{\pi}\left[\cos^{-1} u - u\sqrt{1-u^2}\right].$$

To derive it divide the overlap lens into two symmetric disk segments. Each segment is a circular sector of angle $2 \arccos u$ with its isosceles triangle subtracted. Summing and dividing by the full area yields the expression above. It declines continuously and gives about 0.391 at half the incoherent cutoff. The coherent transfer of a clear pupil has unit field transmission inside its support; it operates on complex amplitude. These two curves describe different imaging experiments. Replacing a coherent field by its intensity before applying a coherent pupil would use the wrong linear system.

For $\lambda_0 = 0.550\,\mu\mathrm{m}$ and $\mathrm{NA} = 0.80$, the coherent cutoff is 1.45 cycles per micrometre and the incoherent cutoff 2.91 cycles per micrometre. An incoherent sinusoidal fluorescent brightness pattern at $f = 2.0$ cycles per micrometre has $u = 0.6875$ and ideal transmitted contrast about 0.20 of its original. A perfectly transmitting objective cannot recover the missing higher frequencies merely by magnifying the image farther. Magnification can, however, change camera sampling and therefore recover detail that the objective passed but undersampled camera pixels discarded.

Separate diffraction, sampling and noise in a resolution estimate

A detector with physical pixel pitch p at magnification M samples the specimen at p/M. The Nyquist frequency is $M/(2p)$. For $p = 6.5\,\mu\mathrm{m}$, $M = 40$ this gives 3.08 cycles per micrometre, just above the illustrative incoherent cutoff. Finite pixel area also averages over a square patch and multiplies the axial-frequency contrast by a sinc factor $\sin(\pi f p/M)/(\pi f p/M)$. At 2.0 cycles per micrometre that factor is about 0.835. The optical and pixel factors multiply, leaving less measured contrast than an optical-only transfer curve predicts. A sinusoidal target in a diagonal orientation sees the product of the factors in the two camera axes.

A cutoff is a support statement. Discrimination in real noise requires a contrast or likelihood criterion, a measured background and a known photon count. Deconvolving a frequency where the transfer is 0.1 multiplies noise in that component by roughly ten when dividing by the transfer. At a true zero no such division can restore the missing datum without additional assumptions or measurements. A point-pair Rayleigh separation, the highest sinusoidal frequency transmitted, and precision in locating the centre of an isolated bright emitter answer three different quantitative questions.

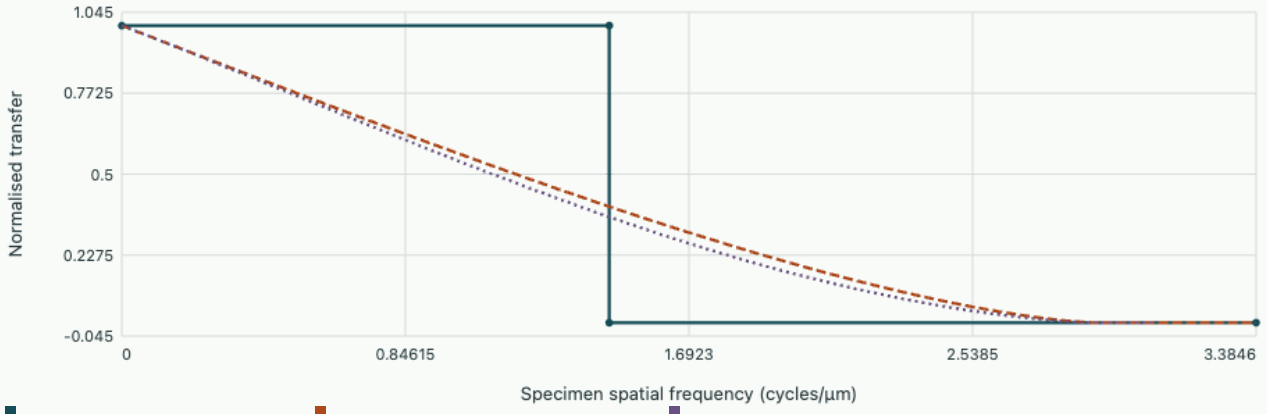
Compare the pupil models directly

The live graph plots the normalised coherent field pupil and incoherent intensity modulation transfer on the same specimen-frequency axis, with a separately identified pixel-averaged incoherent curve. Change numerical aperture while holding camera sampling fixed, then change wavelength. Predict whether optical cutoff or camera Nyquist is limiting before reading the returned values. The rectangular coherent curve should be interpreted as amplitude transmission for normally incident coherent light, whereas the smoother curves act on brightness contrast of incoherent emitters.

Open this experiment online

Try the calculation

Numerical aperture	0.8
Vacuum wavelength (nm)	550
Magnification	40



Coherent field cutoff 1.45, incoherent optical cutoff 2.91, camera Nyquist 3.08 cycles/µm. The sampled curve ends at Nyquist.

Measure transfer with a sinusoidal calibration target

For an incoherent test pattern of mean brightness I_0 and intrinsic modulation m_o, measure the image modulation as $m_i = (I_{\max} - I_{\min})/(I_{\max} + I_{\min})$ after correcting the additive detector background. In a linear shift-invariant imaging regime the ratio m_i/m_o estimates the magnitude of the total transfer at the target frequency. A constant stray background reduces the measured modulation even if the pupil is unchanged. It increases the denominator without increasing the numerator, so subtract its independently measured intensity before comparing with a diffraction formula.

For a target of modulation 0.60 producing corrected maximum and minimum signals 550 and 450 counts, the image modulation is 0.10, so the inferred transfer is 0.167. The earlier optical and pixel factors at two cycles per micrometre predict approximately this value. A target with sharp bars contains many harmonics rather than one pure frequency. Its peak-to-trough contrast cannot generally be inserted into the same single-frequency formula without accounting for those harmonics. A sinusoidal target or a measured Fourier coefficient gives a cleaner comparison with the pupil calculation.

Compare diffraction conventions across the pupil model

When checking a spatial transfer against a diffraction calculation, keep the vacuum wavelength, refractive index and numerical aperture together: index already enters through numerical aperture. Inserting the wavelength in the medium a second time in a formula written with vacuum wavelength

would count that index twice. The distinction matters most for immersion objectives and for comparisons of cutoffs reported from different beam geometries.

The open TU Delft diffraction chapter develops the wave-optical treatment of apertures and optical systems.

Tutorial on resolution claims

A normal-incidence coherent image of a grating with period $0.50\ \mu\mathrm{m}$ is uniform under the example objective. Can changing only illumination angle help?

▼ **Work through the solution**

The spatial frequency is 2.0 cycles per micrometre, beyond the on-axis coherent limit 1.45. An illumination shift of order 0.55 cycles per micrometre can bring one first order inside while retaining the reference order, restoring some contrast. The shift corresponds to illumination NA about 0.30 at the stated wavelength. Other orientations can respond differently.

At $20\times$ magnification with the same camera, which limit applies to an incoherent 2.0-cycle-per-micrometre pattern?

▼ **Work through the solution**

The specimen-plane sampling pitch is $0.325\ \mu\mathrm{m}$, giving Nyquist 1.54 cycles per micrometre. Optics can transmit the pattern, but the camera aliases it. Increasing magnification can help until sampling ceases to be limiting; it cannot raise the optical transfer at a fixed specimen frequency.

In an ideal coherent microscope a weak pure-phase grating has $t(x) \approx 1 + i\epsilon\cos(2\pi f x)$. Does passing all three orders guarantee first-order intensity contrast in focus?

▼ **Work through the solution**

No first-order intensity term survives because the perturbation is in quadrature with the reference. The field modulation is present but the intensity changes only at higher order in this approximation. Introduce a controlled reference phase shift, defocus, or another phase-sensitive measurement to convert the phase modulation into brightness contrast. Distinguish absence of an order from absence of interference in the detected quadrature.

TU Delft's Advanced Optical Imaging course derives coherent image formation from the pupil.

Optical fields, instruments and observations · 2. From a slowly varying wave to a curved ray in a graded index

A ray bends continuously in a glass whose index changes across its width. Derive where that ray equation comes from, predict the focusing length of a graded-index rod, and locate the points at which a smooth ray description ceases to determine an optical amplitude.

Expand the wave equation in inverse optical frequency

In a transparent isotropic nonmagnetic medium, a scalar monochromatic field can satisfy $\nabla^2 U + k_0^2 n(\mathbf{r})^2 U = 0$, with $k_0 = 2\pi/\lambda_0$. Assume n and a slowly varying envelope change over distances long compared with a local wavelength. Seek $U = A(\mathbf{r}) e^{ik_0 S(\mathbf{r})}$. Substitution gives a leading term $k_0^2 A[n^2 - |\nabla S|^2]$, a next term $ik_0[2\nabla A \cdot \nabla S + A\nabla^2 S]$, and $\nabla^2 A$. Vanishing of the leading term yields the eikonal equation $|\nabla S| = n$. The gradient of optical path is the local normal to a phase front.

For a real leading amplitude the next term gives $\nabla \cdot (A^2 \nabla S) = 0$: flux through a thin ray tube is conserved to this order. In a one-dimensional monotone geometry with rays along z, $S' = \pm n$ and the travelling-wave amplitude scales as $n^{-1/2}$. Its field intensity parameter $|A|^2$ falls as index rises while the transported flux proportional to $n|A|^2$ stays constant. At an abrupt interface a reflected wave is generally needed to satisfy matching conditions, so the single-branch envelope law cannot determine the reflected power from a sudden index step.

The locally one-dimensional adiabatic criterion can be written $|k'|/k^2 \ll 1$ with $k = k_0 n$, accompanied by adequate smoothness of higher derivatives. Near a turning point the relevant longitudinal or transverse component of k can vanish even when n itself is finite. Then the WKB amplitude proportional to that component's inverse square root diverges. A finite-wave treatment gives a reflected and penetrating field. Thus merely comparing vacuum wavelength with the overall lens diameter cannot certify the ray approximation at every point.

Derive the graded-index ray equation and paraxial solution

A ray makes the optical path $\int n\, ds$ stationary. With longitudinal axis z and transverse coordinate $x(z)$, the integrand is $n(x, z)\sqrt{1 + x'^2}$. Its Euler–Lagrange equation is

$$d[n x'/\sqrt{1 + x'^2}]/dz = n_x \sqrt{1 + x'^2}.$$

For small slopes and a smooth index $n(x) \simeq n_0[1 - g^2 x^2/2]$, retain the leading transverse restoring term. The equation becomes $x'' + g^2 x = 0$, with g in inverse metres. The axis has the highest index, and rays curve back toward it. Solving gives

$$x(z) = x_0 \cos(gz) + \alpha_0 \sin(gz)/g,$$

where $\alpha_0 = x'(0)$ is the internal paraxial slope.

A full spatial pitch is $2\pi/g$, and a quarter-pitch length $\pi/(2g)$ transforms entrance position into angle and entrance angle into exit position. Collimated input rays with $\alpha_0 = 0$ cross the axis together at a quarter pitch. Rays launched from an axial point $x_0 = 0$ instead emerge parallel there since $x' = \alpha_0 \cos(gz)$ vanishes. These are reciprocal ray maps. At a flat end face outside the rod, refraction changes the external slope according to the index ratio, so an internal angle-to-position scaling should not be directly applied to the angle of an external beam without that conversion.

For $g = 200\,\mathrm{m}^{-1}$ and $n_0 = 1.50$, the quarter-pitch length is $7.85\,\mathrm{mm}$. A ray launched at $x_0 = 0.40\,\mathrm{mm}$ parallel to the axis reaches it with slope $-gx_0 = -0.080$, compatible with a paraxial starting estimate. The assumed index depression at that entrance radius is

$$n_0 g^2 x_0^2/2 = 0.0048.$$

If a measured rod reaches only 0.002 depression at the same radius, either g was overstated or the parabolic model is inaccurate. An independent transverse index map and a measured pitch provide complementary checks.

Connect ray turning points to confined wave modes

For a guide uniform in z, seek $U(x,z) = u(x)e^{i\beta z}$. The scalar Helmholtz equation gives

$$u'' + [k_0^2 n^2(x) - \beta^2]u = 0.$$

The transverse local wave number $k_x = \sqrt{k_0^2 n^2 - \beta^2}$ becomes imaginary outside the classical turning points. Away from them an oscillatory WKB mode has amplitude $1/\sqrt{k_x}$ and phase $\int k_x dx$; the evanescent tail decays as $\exp[-\int |k_x| dx]$. Matching through two simple smooth turns gives the approximate condition $\int_{x_-}^{x_+} k_x dx = (m + 1/2)\pi$. The half-integer arises from the turning-point phase shifts and differs from the phase shifts at two rigid reflecting walls.

Under $n^2(x) \simeq n_0^2(1 - g^2x^2)$, the transverse equation takes harmonic-oscillator form with quadratic confinement. Direct polynomial-Gaussian normalisability yields

$$k_0^2 n_0^2 - \beta_m^2 = (2m+1)k_0 n_0 g.$$

For small mode index relative to $k_0 n_0/g$,

$$\beta_m \simeq k_0 n_0 - (m + 1/2)g.$$

Adjacent propagation constants then differ by nearly g, so their relative phases refocus at $2\pi/g$, exactly the ray pitch. A real finite core and a departure from parabolic index spoil exact equal spacing and can broaden a refocused image or pulse.

Estimate the wave scale of a smooth turning region

Near a transverse turning point x_t suppose $k_x^2(x) \simeq -\gamma(x - x_t)$, where $\gamma > 0$ has units of inverse length cubed and the oscillatory side is $x < x_t$. The transverse equation becomes $u'' - \gamma(x - x_t)u = 0$. Rescale $X = (x - x_t)\gamma^{1/3}$; the resulting dimensionless equation is $u_{XX} - Xu = 0$, whose decaying Airy solution matches a sinusoid with a definite quarter-phase shift on the oscillatory side. Thus the width of the turning region is $l_t = \gamma^{-1/3}$. The wave does not reflect at an infinitely sharp geometric point, even for a perfectly smooth deterministic index profile.

If k_x^2 changes by $10^{10}\,\mathrm{m}^{-2}$ over $10\,\mu\mathrm{m}$, the slope scale is $10^{15}\,\mathrm{m}^{-3}$ and l_t is roughly $10\,\mu\mathrm{m}$. A detector averaging over a larger transverse spot cannot establish the detailed penetration profile from a direct image. Increasing the index gradient makes the turning region narrower only as its inverse cube root. A ray sketch may still predict the trajectory well over millimetres while missing a micrometre-scale reflected-wave phase or modal shift. Compare the length scale of the particular observable with the turning region before replacing the local wave equation by a specular bounce.

A second reflection from a nearby turning point can overlap this Airy region. The usual two independent turning phases are then no longer justified. Solve the transverse wave equation with both sides at once or compare a matched model with the actual index profile and verify normalisable boundary conditions.

Check the paraxial assumptions against a measured ray map

Measure entrance and exit positions for several collimated beams launched at progressively larger distances from the axis. The ideal paraxial solution gives a straight line of exit versus entrance position with slope $\cos(gL)$. Curvature in that plot can reveal aberration or departures from parabolic index. Comparing opposite launch offsets separates symmetric aberration from an off-centre guide or tilted coordinate origin.

Tutorial on a graded-index instrument

A guide with $g = 200\,\mathrm{m}^{-1}$ has length $5.0\,\mathrm{mm}$. A ray enters at $x_0 = 0.30\,\mathrm{mm}$ with zero internal slope. Predict its exit position and slope.

▼ Work through the solution

The phase advance is $gz = 1.00$. Thus

$$x = 0.30\cos 1 = 0.162\,\mathrm{mm},$$

$$x' = -200(0.00030)\sin 1 = -0.0505.$$

The ray is still on its original side of the axis and moving inward. Treating this rod as an exact quarter-pitch element would predict a premature focus.

Using $\lambda_0 = 0.60\,\mu\mathrm{m}$ and the example n_0, g, estimate the approximate Gaussian ground-mode width scale $w_0 = (k_0 n_0 g)^{-1/2}$.

▼ **Work through the solution**

Here $k_0 = 1.047 \times 10^7\,\mathrm{m}^{-1}$, so $w_0 = 17.8\,\mu\mathrm{m}$. The transverse field is proportional to $\exp[-x^2/(2w_0^2)]$ in this convention. A quoted intensity $1/e^2$ radius uses a different numerical factor. The computed width should be much smaller than the parabolic region for this ideal mode to be valid.

The ray-tube transport equation predicts unbounded amplitude at the common ray focus. Which measurement shows what the approximation omitted?

▼ **Work through the solution**

Measure the finite transverse focal spot while changing wavelength and aperture. Wave diffraction limits resolution even when all geometric rays meet at one point. The leading transport approximation conserves flux as ray-tube area shrinks but cannot describe interference and phase near a caustic. A wave propagation or mode superposition through the same index map should produce a finite profile and permit a direct comparison with those measurements.

TU Delft derives optical rays and geometrical imaging.

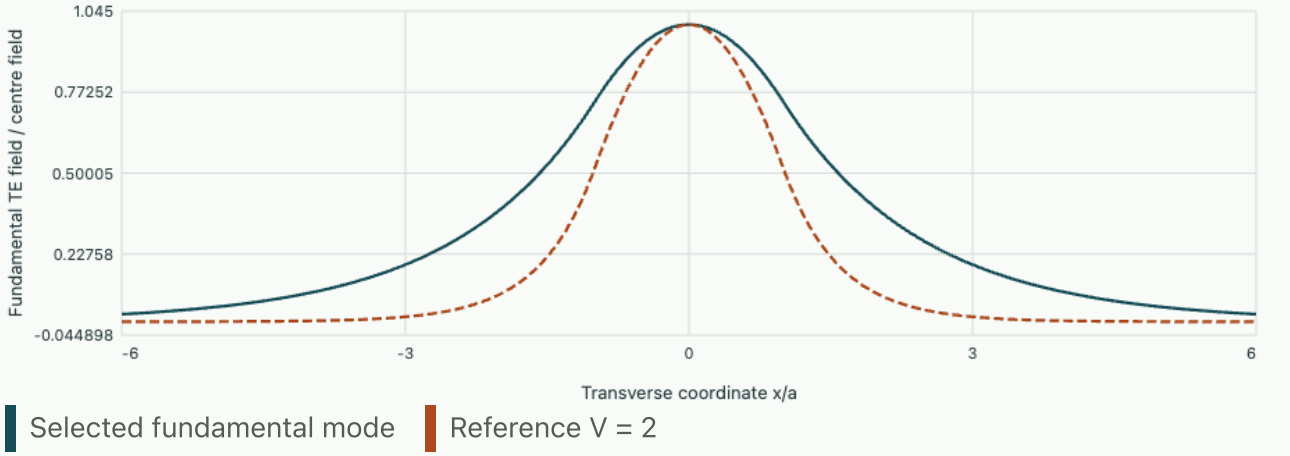
Compare smooth index turning points with modes confined at an index jump.

Open this experiment online

Try the calculation

Normalised half-width V

1



Root u = 0.7391, cladding exponent w = 0.6736. Field amplitude decays by 1/e over 1.485 half-widths outside the core.

Optical fields, instruments and observations · 3. Measure rotation with two waves following the same path

A stationary optical loop gives equal return times for opposite propagation directions in a reciprocal medium. When the loop rotates the beams meet their starting beamsplitter at different times. Derive the measurable differential phase and compare a passive interferometer with a ring resonator.

Calculate the return times in a rotating circular loop

Take a vacuum circular loop of radius R rotating at constant angular speed Ω in its plane, with a co-moving source, beamsplitter and receiver. In an inertial frame the receiver travels at $v = \Omega R$. A beam propagating in the direction of rotation must catch up with the moving receiver, while the oppositely directed beam meets it early. For an ideal constrained circular optical path, $ct_+ = 2\pi R + vt_+$ and $ct_- = 2\pi R - vt_-$. Thus $t_\pm = 2\pi R/(c \mp v)$, and their difference is $4\pi Rv/(c^2 - v^2)$. To first order in v/c this is $4\Omega A/c^2$, with $A = \pi R^2$.

The times just calculated are laboratory coordinate times. A clock on the rotating detector accumulates proper time with factor $\sqrt{1 - v^2/c^2}$; this changes the first-order result only at higher order. The local vacuum speed remains c for either beam. The effect comes from comparing their arrival at a moving part of a closed apparatus. There is no requirement to assign a different local fundamental light speed to the clockwise and anticlockwise beams.

A general co-moving loop at low rotational speed gives an extra time of flight $dt_{\mathrm{corr}} = \mathbf{v}(\mathbf{r}) \cdot d\mathbf{l}/c^2$ along one orientation. Reversing orientation doubles its contribution to the difference. With $\mathbf{v} = \boldsymbol{\Omega} \times \mathbf{r}$,

$$\Delta t = \frac{2}{c^2} \oint (\boldsymbol{\Omega} \times \mathbf{r}) \cdot d\mathbf{l} = \frac{4}{c^2} \boldsymbol{\Omega} \cdot \mathbf{A},$$

$$\mathbf{A} = \frac{1}{2} \oint \mathbf{r} \times d\mathbf{l}.$$

The oriented area handles tilted loops and folded paths. A figure-eight winding with two equal oppositely oriented lobes can enclose a large total unsigned area yet have zero net response to uniform rotation. Winding a fibre around the same area N times in the same orientation multiplies the directed area by N.

Convert the time shift into optical phase

A continuous source of vacuum wavelength λ_0 accumulates an optical phase difference $\Delta\phi = \omega_0 \Delta t = 8\pi \boldsymbol{\Omega} \cdot \mathbf{A}/(\lambda_0 c)$ at leading order. For a co-moving nondispersive dielectric loop the reciprocal static propagation time increases with index, but the standard first-order geometric Sagnac phase at a fixed external optical frequency retains this scale; applying relativistic velocity addition to propagation in the moving medium accounts for the cancellation. Flow of the medium relative to the loop, time-dependent index disturbances and nonreciprocal phase shifts can introduce additional effects.

For a hundred turns of radius $R = 0.10\,\mathrm{m}$, total directed area is $3.142\,\mathrm{m}^2$. At $\lambda_0 = 1.55\,\mu\mathrm{m}$ and $\Omega = 0.10\,\mathrm{rad/s}$ perpendicular to the loop, the time difference is only $1.40 \times 10^{-17}\,\mathrm{s}$, yet the phase is $0.0170\,\mathrm{rad}$. Direct electronic time-of-flight timing would be exceptionally demanding, while interference converts this small delay into an intensity change. Tilting the normal by 60° relative to rotation halves the signal and reversing the winding or rotation reverses its sign.

With equal returning intensities, a detector output can be written $I = I_{\mathrm{mean}}[1 + \mathcal{V}\cos(\Delta\phi + \phi_b)]$, where $\mathcal{V}$ is visibility and ϕ_b a controlled bias. Near $\phi_b = 0$ a small rotation changes intensity only quadratically and does not identify its sign. At quadrature, $\phi_b = \pi/2$, the change is approximately $-I_{\mathrm{mean}}\mathcal{V}\Delta\phi$. Alternating the applied bias or using both complementary outputs can reject a common power drift. A measured percentage intensity change only directly represents phase when the visibility and operating point are known.

Derive a ring resonator's beat instead of a passive phase

In an active vacuum ring laser or a tracked passive ring cavity, each propagation direction can have a separate resonant frequency satisfying an integer round-trip phase. In the simple nondispersive vacuum case let perimeter be P, unperturbed time $T = P/c$, and resonance frequency $\nu = m/T$. A small directional time difference Δt gives a resonance frequency splitting of magnitude $\Delta\nu \simeq \nu\Delta t/T = 4A\Omega_\perp/(\lambda_0 P)$. This readout measures a beat frequency. The passive interferometer measured a single-pass phase difference; the dimensions and perimeter dependence of the two formulas distinguish them.

A square vacuum ring of side $1.0\,\mathrm{m}$ has $A = 1.0\,\mathrm{m}^2$, $P = 4.0\,\mathrm{m}$. At $633\,\mathrm{nm}$ with $\Omega_\perp = 10^{-5}\,\mathrm{rad/s}$, the predicted directional splitting is $15.8\,\mathrm{Hz}$. Double every side at fixed wavelength and angular speed: area quadruples but perimeter doubles, so the beat doubles. In a physical active ring, backscattering can couple the counterpropagating modes and produce a dead band near low rotation. A zero observed beat does not uniquely show a stationary loop unless this coupling and other bias effects have been characterised.

Estimate a rotation sensitivity from counted photons

At a balanced quadrature readout with visibility $\mathcal{V}$, a small phase changes the mean differential count divided by total count by roughly $\mathcal{V}\Delta\phi$. For independently detected photons and a shot-noise-limited balanced measurement the standard deviation of that normalised difference is approximately $1/\sqrt{N}$, giving a phase uncertainty near $1/(\mathcal{V}\sqrt{N})$. This estimate assumes the photons used in the signal actually reached the detectors and that gain mismatch, intensity noise and nonreciprocal drift do not dominate. It is a count budget, separate from the geometric scale relating phase to rotation.

For $\mathcal{V} = 0.80$ and $N = 10^8$ detected photons the ideal phase uncertainty is $1.25 \times 10^{-4}\,\mathrm{rad}$. The hundred-turn example then has projected-rate uncertainty about $7.4 \times 10^{-4}\,\mathrm{rad/s}$ over that acquisition interval. Repeating statistically independent acquisitions improves the random part by the square root of their number, but cannot reduce a persistent reference bias in the same way. Rotating the coil normal through a known angle changes the physical signal as a cosine and can help identify

additive biases. The calibration must use the actual normal orientation and total directed area, rather than assuming the apparatus remains level.

A longer fibre winding increases area at fixed turn radius while also increasing attenuation and susceptibility to distributed disturbances. The ideal phase scale grows with turn count; the number of detected photons and long-time noise can change at the same time. Compare those quantities independently when evaluating a different winding.

Tutorial on rejecting false rotations

What happens if a rigid loop translates at constant velocity without rotating in the first-order geometric calculation?

▼ **Work through the solution**

Its velocity is constant around the path, and $\oint \mathbf{v} \cdot d\mathbf{l} = \mathbf{v} \cdot \oint d\mathbf{l} = 0$. There is no corresponding Sagnac area signal. A moving segment of a nonrigid loop can have a different line integral; the stated cancellation is for common translation of the entire closed path.

A passive interferometer with visibility 0.80 reports a quadrature change $\Delta I / I_{\mathrm{mean}} = -0.0080$ for the hundred-turn geometry above. Estimate the projected rotation.

▼ **Work through the solution**

At quadrature the phase change is $0.0100\,\mathrm{rad}$. The example scale is $0.170\,\mathrm{rad}$ of phase per radian per second, giving $\Omega_{\perp} \simeq 0.0589\,\mathrm{rad/s}$. Use the actual directed area and vacuum wavelength. A visibility correction or a bias angle error propagates directly into the rotation estimate.

Can a drifting temperature distribution mimic rotation in a geometrically reciprocal fibre loop?

▼ **Work through the solution**

Two counterpropagating waves encounter a given fibre segment at different times. A static reciprocal phase shift cancels to leading order, but a changing local index can give different integrated phases on the two traversals. Compare runs with reversed physical rotation, symmetric thermal conditions and a calibrated artificial phase bias. Winding arrangement and timing of disturbances help distinguish a true oriented-area response from a transient nonreciprocal measurement error.

The Sagnac interferometer chapter hosted by Rutgers derives the phase shift of the two counterpropagating beams.

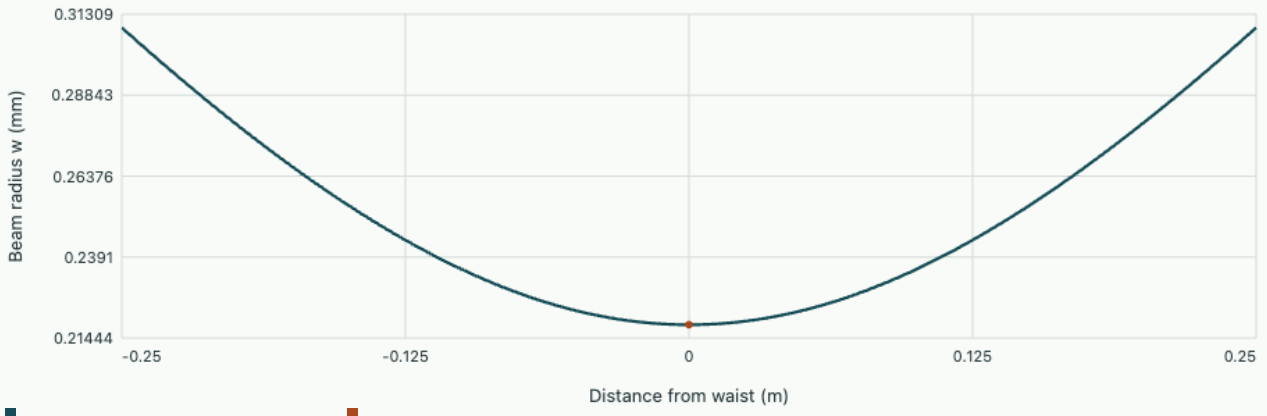
Compare a resonator's mode conditions with the directional frequency readout.

Open this experiment online

Try the calculation

Mirror separation / curvature radius **1**

Wavelength (nm) **600**



Gaussian beam radius Waist

Rayleigh length 0.25 m, waist radius 0.2185 mm, mirror spot radius 0.309 mm.

Optical fields, instruments and observations · 4. Dispersive pulse compression and a detector's view of a tilted pulse

An ultrashort pulse can broaden because its spectral components acquire different group delays or because different parts of its beam arrive at different times. Derive the two mechanisms separately, predict when added glass can shorten an already chirped pulse, and decide which measurements distinguish them.

Derive the propagation of a spectral envelope

Write the positive-frequency electric field as the real carrier with complex envelope $A(z,t)$, using the Fourier convention $A(t) \propto \int \tilde{A}(\Omega) e^{-i\Omega t} d\Omega$,

$$\Omega = \omega - \omega_0.$$

Propagation through a transparent homogeneous length L multiplies the spectrum by $e^{ik(\omega)L}$. Expand $k = k_0 + k_1\Omega + k_2\Omega^2/2 + \cdots$. The constant is an overall phase, the linear term shifts the pulse by group delay $k_1 L$, and the quadratic term gives group-delay dispersion $D = k_2 L$ in seconds squared. The group-delay difference between spectral components separated by $\Delta\Omega$ is $D\Delta\Omega$.

Take a Gaussian spectral amplitude $\tilde{A}(\Omega) = A_0 \exp[-\tau_0^2\Omega^2/2]$, and give it an initial quadratic spectral phase $\phi_2\Omega^2/2$. The effective quadratic coefficient after propagation is $\Phi = \phi_2 + D$. Fourier transforming the Gaussian at the retarded time $T = t - k_1 L$ gives an envelope proportional to $(\tau_0^2 - i\Phi)^{-1/2} \exp[-T^2/(2(\tau_0^2 - i\Phi))]$. Its intensity is Gaussian with temporal scale $\tau = \tau_0\sqrt{1 + (\Phi/\tau_0^2)^2}$, where intensity is proportional to $\exp[-T^2/\tau^2]$. The intensity full width at half maximum is $2\sqrt{\ln 2}\,\tau$; keeping this convention explicit prevents confusing amplitude widths, intensity widths and rms durations.

The intensity spectrum is unchanged by quadratic phase in a lossless element. Temporal broadening reduces the peak at fixed integrated pulse energy by τ_0/τ. Its duration depends on Φ^2, so one measured width at a single setting cannot determine the sign of chirp. Measuring the width after a known added dispersion does break that degeneracy. The full temporal phase contains a quadratic term giving instantaneous frequency that changes approximately linearly through the pulse. A transform-limited Gaussian has $\Delta t_{\mathrm{FWHM}} \Delta\nu_{\mathrm{FWHM}} = 2\ln 2/\pi \simeq 0.441$ when frequency is in cycles per second.

Calculate an imperfect compressor setting

Start with a transform-limited duration $30.0\,\mathrm{fs}$, so $\tau_0 = 18.0\,\mathrm{fs}$ and $\tau_0^2 = 325\,\mathrm{fs}^2$. A spectral phase of $+400\,\mathrm{fs}^2$ broadens this pulse to about $47.6\,\mathrm{fs}$, with peak intensity about 0.630 of its transform-limited value. Passing it through a component that contributes $-300\,\mathrm{fs}^2$ leaves $\Phi = 100\,\mathrm{fs}^2$ and duration $31.4\,\mathrm{fs}$. The same component applied to a pulse initially at $-400\,\mathrm{fs}^2$ instead gives $-700\,\mathrm{fs}^2$ and

substantially more broadening. A quoted dispersion accuracy without the initial sign does not establish compression.

Experimental scans are simplest in squared width. Vary known D and fit $\Delta t^2 = \Delta t_0^2[1 + ((D + \phi_2)/\tau_0^2)^2]$. The parabola has a minimum at $D = -\phi_2$. A non-Gaussian spectrum or a cubic spectral phase can produce pedestals, asymmetric profiles or a minimum width larger than the transform limit inferred from spectrum. In an intensity autocorrelation a Gaussian pulse has a correlation width $\sqrt{2}$ times its intensity width. Apply the correct deconvolution for the assumed shape, or measure a spectrally resolved temporal diagnostic when the shape itself is unknown.

Derive pulse-front tilt and the spatially averaged width

Suppose the centre arrival time varies across one transverse coordinate as $T_c(x) = sx$. The slope s has units of time per length. A grating or misaligned dispersive optic can correlate frequency, propagation angle and position. For a local pulse of intensity $\exp[-(t - sx)^2/\tau^2]$ and beam weighting $\exp[-x^2/w^2]$, a detector integrating the whole beam sees the temporal convolution of the local profile with the distribution of arrival centres. Independent Gaussian variances add, giving measured temporal scale $\tau_\mathrm{d}^2 = \tau^2 + s^2 w^2$. A single-pixel trace at fixed x remains narrow and shifts in time. A whole-beam trace broadens even when no additional spectral phase changes along any one local ray.

With $w = 0.50\,\mathrm{mm}$ in this intensity convention and $s = 40\,\mathrm{fs/mm}$, the extra scale is $20\,\mathrm{fs}$. A locally transform-limited $30.0\,\mathrm{fs}$ pulse has $\tau_0 = 18.0\,\mathrm{fs}$, so the detector width becomes

$$2\sqrt{\ln 2}\sqrt{18.0^2 + 20^2} = 44.8\,\mathrm{fs}.$$

Changing the aperture alters the broadened trace by changing the sampled transverse weighting. A hard slit does not produce a Gaussian weighting, so the variance addition remains applicable for independent position and local pulse shape but the full width at half maximum need not follow the same Gaussian conversion.

Angular dispersion can create such a delay gradient. Near the reference propagation axis a spectral plane wave carries transverse phase $k_x(\omega)x$. In the inverse transform the term $(dk_x/d\omega)\Omega x$ shifts the local envelope time by $(dk_x/d\omega)x$, giving $s = dk_x/d\omega$ under the adopted conventions. Keeping only $k_x(\omega_0)$ accounts for a tilted carrier phase but misses the tilted pulse front. Mapping cross-correlation delay versus transverse position separates this linear slope from uniform material dispersion.

Predict a compression scan and perturb its detection area

The live temporal traces compare a local pulse and the second trace integrated across a known Gaussian beam, with the same normalised pulse energy. Predict the added dispersion that gives the narrowest local pulse, then turn up the transverse delay slope without moving the compressor. Which curve keeps its width? A larger minimum in an area-integrated scan is compatible with spatial averaging even when a local sample reaches its transform limit. Use a second spatial sampling condition before fitting that extra duration entirely to a higher-order material phase.

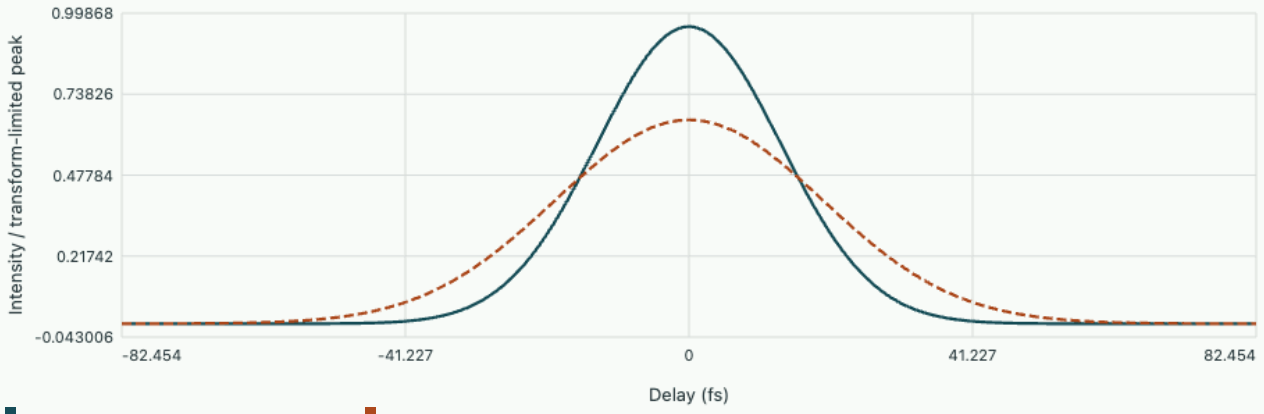
Open this experiment online

Try the calculation

Input spectral phase (fs²) **400**

Added dispersion (fs²) **-300**

Arrival slope (fs/mm) **40**



Local intensity FWHM 31.39 fs. Full-beam intensity FWHM 45.77 fs. Residual quadratic phase 100 fs².

Check the group delay with two narrow spectral slices

A pulse-width scan determines the magnitude of residual quadratic phase but can hide its sign at a single setting. Split off two narrow spectral slices centred at offsets $+\Delta\Omega/2$ and $-\Delta\Omega/2$, and compare their envelope delays against a common reference. To first order the arrival-time difference is $\Phi\Delta\Omega$. At $\Phi = +400\,\mathrm{fs}^2$ and an ordinary-frequency separation of $5.0\,\mathrm{THz}$, $\Delta\Omega = 0.0314\,\mathrm{rad/fs}$, giving a $12.6\,\mathrm{fs}$ delay difference. Under the adopted propagation convention the higher-frequency slice arrives later for positive dispersion.

The spectral filters themselves can contribute unequal group delays and broaden their selected pulses. Calibrate their differential delay using a reference path or a reference spectrum with independently characterised phase. Narrowing each slice improves its frequency selectivity while increasing its temporal duration, so a cross-correlation delay fit must retain adequate signal and temporal stability. Measuring several frequency centres permits a slope fit; visible curvature of delay against angular frequency indicates the need for a cubic or higher spectral phase. A camera position offset correlated with selected wavelength can add an apparent slope through pulse-front tilt, so compare at the same physical beam position as well as after spatial integration.

Tutorial on interpreting pulse widths

A $50\,\mathrm{fs}$ transform-limited Gaussian has a measured optical bandwidth of $5.0\,\mathrm{THz}$. Are the two measurements consistent if both are intensity full widths at half maximum?

▼ **Work through the solution**

Their product is 0.250, below the Gaussian transform-limited value 0.441. A $50\,\mathrm{fs}$ Gaussian requires about $8.82\,\mathrm{THz}$ of ordinary-frequency width. Check whether the reported bandwidth was angular frequency, an amplitude rather than intensity width, or clipped by the spectrometer. Pure spectral dispersion cannot reduce the time-bandwidth product below the transform limit.

The minimum duration occurs at $D = -250\,\mathrm{fs}^2$. What is the input quadratic phase, and what additional observation would show significant residual cubic phase?

▼ **Work through the solution**

The simple model gives

$$\phi_2 = +250\,\mathrm{fs}^2.$$

A minimum much longer than the transform limit for the measured spectrum, or asymmetric temporal satellites that persist near the minimum, motivates a higher-order spectral-phase measurement. Spatial averaging must be controlled separately before attributing every residual feature to cubic dispersion.

Two local cross-correlation peaks differ by $24\,\mathrm{fs}$ between transverse positions separated by $0.60\,\mathrm{mm}$. What extra rms time spread is expected across the $w = 0.50\,\mathrm{mm}$ Gaussian beam?

▼ **Work through the solution**

The slope is $40\,\mathrm{fs/mm}$. Since the transverse rms width is $w/\sqrt{2}$ for the stated intensity convention, the rms arrival spread is

$$40(0.50)/\sqrt{2} = 14.1\,\mathrm{fs}.$$

The corresponding Gaussian scale parameter $sw = 20\,\mathrm{fs}$ is larger by $\sqrt{2}$. Match the width convention to the formula used by the detector fit.

Kolner and Nazarathy connect dispersive evolution to temporal imaging.

Compare quadratic phase in temporal propagation with its spatial diffraction counterpart.

Optical fields, instruments and observations · 5. Which polarisation measurements identify an optical element?

A retarder rotates the state of a coherent beam, while unresolved mixtures can reduce its measured degree of polarisation. Derive the passage from Jones fields to a Mueller matrix, reconstruct a simple instrument from measured intensities, and identify claims the available input states cannot support.

Move from a field vector to a coherence matrix

For a monochromatic transverse field write the complex amplitudes as $\mathbf{e} = (E_x, E_y)^T$ with time convention $e^{-i\omega t}$. A deterministic optical element gives

$$\mathbf{e}_{\mathrm{out}} = J\mathbf{e}_{\mathrm{in}}.$$

A detector averaging over unresolved time, frequency or spatial channels instead sees a coherence matrix $C = \langle \mathbf{e}\mathbf{e}^\dagger \rangle$, which remains a 2×2 Hermitian positive semidefinite matrix. A fixed Jones element acts as $C_{\mathrm{out}} = JC_{\mathrm{in}}J^\dagger$. An incoherent mixture of distinct elements acts as $\sum_a p_a J_a C_{\mathrm{in}} J_a^\dagger$, generally not representable by one common Jones matrix.

Use the convention $S_0 = C_{xx} + C_{yy}$, $S_1 = C_{xx} - C_{yy}$, $S_2 = 2\Re C_{xy}$, $S_3 = 2\Im C_{xy}$. This explicitly sets the sign of the circular component independent of competing right-handedness naming conventions. The determinant is $[S_0^2 - S_1^2 - S_2^2 - S_3^2]/4$. Positive eigenvalues require $P = \sqrt{S_1^2 + S_2^2 + S_3^2}/S_0 \leq 1$ for nonzero signal. A pure Jones field gives rank one and $P = 1$; a mixture can give smaller P. A fixed lossy polariser can convert an unpolarised input into a fully polarised transmitted field, so attenuation and loss of polarisation are separate observations.

Derive the Mueller matrix of a retarder and an unresolved mixture

Write $C = \frac{1}{2}\sum_j S_j \tau_j$, with τ_0 identity, $\tau_1 = \mathrm{diag}(1, -1)$, τ_2 the real off-diagonal Pauli matrix, and τ_3 chosen with $+i$ in its upper off-diagonal entry for the stipulated convention. Since $\mathrm{tr}(\tau_i \tau_j) = 2\delta_{ij}$,

$$S_i = \mathrm{tr}(\tau_i C).$$

Inserting $JCJ^\dagger$ gives the real Mueller matrix $M_{ij} = \frac{1}{2}\mathrm{tr}(\tau_i J \tau_j J^\dagger)$. A common phase on J cancels. These intensity measurements cannot determine a spatially uniform optical carrier phase, even if every Mueller entry is known.

For $J = \mathrm{diag}(1, e^{i\delta})$, C_{xy} is multiplied by $e^{-i\delta}$. Thus S_0, S_1 remain fixed while $S_2' = S_2\cos\delta + S_3\sin\delta$ and $S_3' = -S_2\sin\delta + S_3\cos\delta$. A 45° linear input has $(S_1, S_2, S_3)/S_0 = (0, 1, 0)$. A quarter-wave retardance sends it to $(0, 0, -1)$ in the stated sign convention; a half-wave retardance sends it to $(0, -1, 0)$. Measuring only horizontal and vertical powers would return equal readings for all three states. A diagonal analyser and a quarter-wave analyser are needed to distinguish them.

Now suppose half the unresolved beam spots undergo retardance $+\delta$ and half $-\delta$. Average their Mueller actions rather than averaging Jones fields. The S_2, S_3 cross terms cancel and both components acquire a factor $\cos\delta$. For $\delta = \pi/3$, the diagonal linear input becomes $P = 0.50$, while a horizontal input remains fully polarised. This is anisotropic depolarisation. Resolving the two spots would reveal two fully polarised local states. A coherent recombination of amplitudes with a controlled phase is a different physical experiment and requires adding fields before averaging.

Reconstruct an element from independent incident states

Put four independent known incident Stokes vectors into columns of a matrix W, and the corresponding measured output vectors into Y. A linear optical response obeys $Y = MW$, so $M = YW^{-1}$ if W is nonsingular. Horizontal, vertical, diagonal-linear and a circular state with calibrated sign are one possible independent set. Four purely linear input states have $S_3 = 0$, so their columns span at most a three-dimensional subspace. No amount of repetition of only those input states reveals the fourth response column. Additional independent preparations are more informative than merely increasing nominal data volume.

For normalised inputs $H = (1, 1, 0, 0)$, $V = (1, -1, 0, 0)$, $D = (1, 0, 1, 0)$, $C_+ = (1, 0, 0, 1)$, call measured outputs y_H, y_V, y_D, y_C. The first column is $(y_H + y_V)/2$, the second $(y_H - y_V)/2$, the third $y_D - M_{:0}$, and the fourth $y_C - M_{:0}$. These differences assume known equal incident intensities or a correction to a common monitored intensity. A source that dims during one preparation can look like diattenuation. Measure an independent input monitor and repeat the preparations in a different order to distinguish drift from element properties.

Suppose the outputs are $(1, \pm 1, 0, 0)$ for H, V, $(1, 0, 0.50, 0)$ for D, and $(1, 0, 0, 0.50)$ for C_+. The reconstructed matrix is diagonal with entries $1, 1, 0.50, 0.50$, as in the unresolved equal-retardance example. The matrix establishes the action on Stokes states but does not uniquely identify that microscopic mechanism. A temporal fluctuator and an unresolved spatial mixture can give the same matrix. Change acquisition time, spectral bandwidth or spatial aperture to test which ensemble the detector is averaging.

Recover the retardance of a real plate with an imperfect analyser

Send a unit 45° linear input through the diagonal retarder above. Measuring the two independent output components gives

$$S_2/S_0 = \cos\delta, \;\; S_3/S_0 = -\sin\delta.$$

Thus $\delta = \operatorname{atan2}(-S_3, S_2)$ modulo 2π. A fit using only the diagonal analyser, and hence only $\cos\delta$, cannot distinguish opposite retardances and becomes insensitive near a maximum or minimum. The circular analyser supplies the missing quadrature. If the beam is partially depolarised equally in these two components, the angle from their ratio can remain well determined while their magnitude shrinks.

Suppose calibrated output readings give

$$S_2/S_0 = 0.48,\ S_3/S_0 = -0.64,$$

with negligible S_1. The inferred retardance is $\operatorname{atan2}(0.64, 0.48) = 53.1^\circ$, while the degree of polarisation is 0.80. A single ideal Jones retarder acting on a pure input at one frequency could produce the angle but would preserve unit degree; something in the measured ensemble or instrument reduced it. A pure retarder model fitted only to S_2 would instead infer $\arccos 0.48 = 61.3^\circ$, wrongly absorbing the contrast loss into its phase.

An analyser with leakage can also reduce the measured Stokes components. A differential pair of orthogonal polarisers with identical leakage fraction η scales its normalised difference by $(1-\eta)/(1+\eta)$. Calibrate each channel using known pure preparations and monitor intensity drifts before assigning all contrast loss to the specimen. Repeating with narrower spectral bandwidth or a smaller beam spot then tests the ensemble dependence of the inferred retardance and depolarisation.

Separate common optical phase from relative retardance

Inserting an isotropic plate can add the same phase to both orthogonal components of every coherent input. Its ideal Mueller matrix is unchanged because J and $e^{i\alpha}J$ act identically on the coherence matrix. An interferometer with an external reference arm can nevertheless detect the added phase. This experiment supplies precisely the reference that a standalone polarimeter lacks. The two instruments can thus disagree about whether anything changed while measuring distinct physical observables accurately.

Tutorial on polarimetric inference

An analyser reports intensities $I_H = 0.60,\ I_V = 0.40,\ I_D = 0.75,\ I_A = 0.25$ in common units. Which Stokes entries follow, and what remains unknown?

▼ **Work through the solution**

The first pair gives

$$S_0 = 1.00, S_1 = 0.20;$$

the diagonal pair gives $S_2 = 0.50$. No circular analysis was performed, so S_3 remains unknown within $|S_3| \le \sqrt{1 - 0.20^2 - 0.50^2} = 0.843$. Assuming a pure state would set its magnitude to that bound but would still not determine its sign. A partially polarised beam can have any smaller magnitude.

An ideal deterministic $J = \operatorname{diag}(\sqrt{0.8}, \sqrt{0.2})$ transmits an unpolarised input with unit incident intensity. What are the output intensity and degree of polarisation?

▼ **Work through the solution**

The input coherence matrix is half the identity. The transmitted diagonal entries are 0.40 and 0.10, so $S_0 = 0.50$, $S_1 = 0.30$ and $P = 0.60$. The element selectively transmits one component of the mixture. It has not generated optical energy; half the incident intensity was lost.

A fitted Mueller matrix maps one measured positive input into $P = 1.08$. How should the inference be checked?

▼ Work through the solution

That output violates positivity for this convention. Check dark subtraction, intensity normalisation, handedness of the circular analyser and conditioning of the preparation and analysis matrices. Collect more independent states with uncertainties and fit within a physical response model. Clipping each output degree to unity separately need not produce one common linear physical Mueller map for the instrument.

Reconstruct a Stokes state numerically from analyser intensities.

TU Delft's optical course covers polarisation states and their measurement.

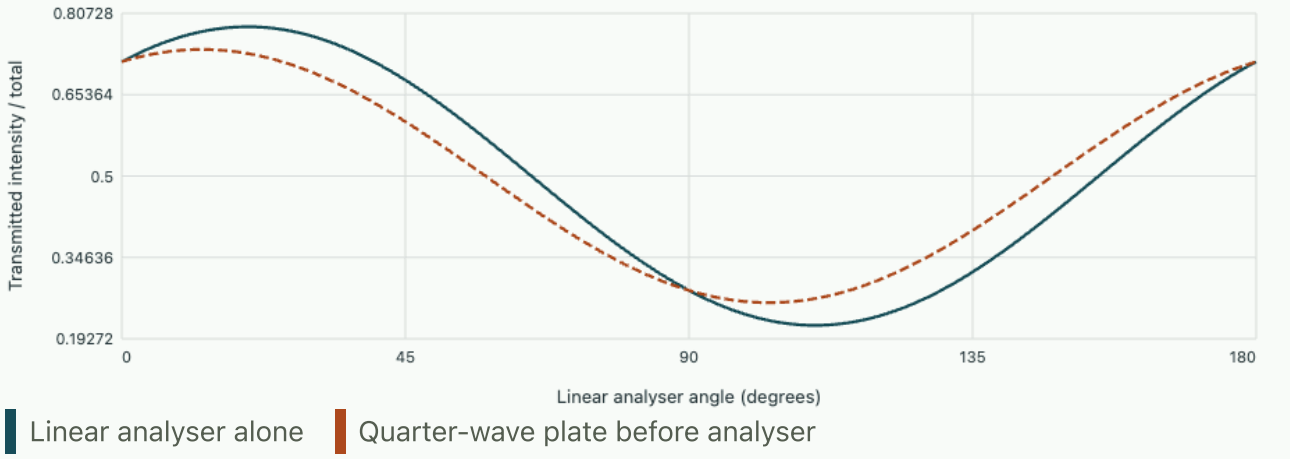
Open this experiment online

Try the calculation

Polarised fraction P **0.6**

Azimuth ψ (degrees) **20**

Ellipticity χ (degrees) **10**



Measured complementary pairs (horizontal, diagonal, circular): 0.716 / 0.284 · 0.6812 / 0.3188 · 0.6026 / 0.3974. Reconstructed $(S_1,S_2,S_3)/S_0$ = (0.4319, 0.3624, 0.2052); degree 0.6.

Optical fields, instruments and observations · 6. Photon counting, bunching and the time resolution of a coincidence test

An intensity trace can be steady while individual detector clicks show correlations. Derive the observable of a two-detector Hanbury Brown–Twiss measurement, quantify how background and bin width change it, and decide whether a dip near zero delay belongs to the source or to a detector.

Derive the coincidence baseline from two count streams

Split light on a beamsplitter and send the two outputs to separate photon-counting detectors. Write their mean count rates as R_1, R_2. For stationary independent streams observed for time T, the expected number of pairs in a delay bin of width Δt near zero is $R_1 R_2 T \Delta t$ when the bin is short relative to T. Each click in arm one defines a short opportunity of duration Δt in arm two. At long delays near the edges of a finite record, the available pair exposure is reduced; explicitly accounting for overlap time avoids a spurious triangular envelope.

The normalised second-order correlation $g^{(2)}(\tau)$ is the measured coincidence density divided by that independent baseline. In quantum notation it uses a normally ordered intensity product, so one photon cannot register as two distinct photons at the opposite outputs of an ideal splitter. For one occupied mode at equal time,

$$g^{(2)}(0) = \langle n(n-1)\rangle / \langle n\rangle^2.$$

A number state with exactly one photon has zero; a coherent field with Poisson photon statistics has one; one ideal thermal mode gives two. A low mean count rate alone implies none of these cases: attenuation lowers the count rate but independent loss ideally cancels from the normalised ratio.

For $R_1 = 40\,\mathrm{kcounts/s}$, $R_2 = 30\,\mathrm{kcounts/s}$, $T = 100\,\mathrm{s}$, and $\Delta t = 1.0\,\mathrm{ns}$, the expected uncorrelated baseline is 120 counts per bin. A measured 60 counts gives a raw correlation of 0.50. A simple Poisson estimate of the numerator uncertainty is $\sqrt{60}/120 = 0.065$ before baseline, dead-time and other uncertainties. Shrinking the bin can increase temporal contrast but reduces the counts per bin, so reporting a very small $g^{(2)}(0)$ without the actual counts or uncertainty conceals an important limitation.

Calculate the effect of uncorrelated backgrounds

In arm i, let $R_i = S_i + B_i$ with signal fraction $\rho_i = S_i/R_i$. Assume both backgrounds are Poissonian, independent of each other and of the source, and that counting response is linear. Every coincidence involving a background contributes only its independent baseline, while true signal–signal pairs carry the nontrivial source correlation. Expanding the product gives

$$g^{(2)}_{\mathrm{obs}}(\tau) = 1 + \rho_1 \rho_2 [g^{(2)}_{\mathrm{signal}}(\tau) - 1].$$

For equal signal fractions 0.80, an ideal single-emitter zero-delay dip reaches only 0.36 in an otherwise perfect instrument. The same background reduces a thermal single-mode peak from two to 1.64. A correction using ρ_i estimated at another alignment or pump power can overcorrect. Correlated background, detector crosstalk and blinking invalidate the independent-background relation or add longer-delay structures. Measure backgrounds with the relevant source and optical settings whenever possible.

A simple stationary emitter with one exponential recovery time gives

$$g_s^{(2)}(\tau) = 1 - e^{-|\tau|/\tau_c}.$$

This shape can arise in a two-state incoherent pumping model because detecting an emitted photon resets the system to its ground state and the excited population recovers. A thermal bunching example has $g_s^{(2)} = 1 + e^{-|\tau|/\tau_c}$ under a corresponding exponential intensity-correlation assumption. The time constant used in the latter is the intensity-correlation time; it need not equal a field-amplitude coherence time. Multi-level emitters or multimode light require a different model.

Average the correlation over the time bin

A rectangular zero-centred bin reports its mean correlation, not the instantaneous value. For the exponential departure from unity the zero-bin attenuation is $A = (2\tau_c/\Delta t)[1 - e^{-\Delta t/(2\tau_c)}]$. It tends to one for a narrow bin and approximately $2\tau_c/\Delta t$ for a very wide one. An ideal anticorrelation with $\tau_c = 2\,\mathrm{ns}$, $\Delta t = 4\,\mathrm{ns}$ and signal fractions 0.80 in each arm yields

$$g_{\mathrm{bin}}(0) = 1 - 0.64(1 - e^{-1}) = 0.595.$$

Inferring two emitters solely because the observed value exceeds one half would ignore known attenuation by both background and timing.

Electronic timing jitter gives an additional convolution with the probability distribution of the difference of the two timing errors. For independent Gaussian jitters their variances add. The measured bin width and this jitter serve distinct roles: rebinning time-tagged data cannot undo detector jitter. Extremely large bins can make single-emitter and Poisson observations nearly indistinguishable even with many total counts. A calibrated timing response and a physically justified line model can constrain an underlying time scale, but extremely narrow intrinsic features can remain unresolved.

Compare photon statistics with detector statistics

A single detector with a recovery dead time after a click produces an artificial deficit of very short same-detector intervals even for coherent light. Using separate arms avoids that particular single-detector autocorrelation artifact. Each detector's saturation can still alter measured rates, and an electronic dead time shared between channels can create a false coincidence dip. Conversely optical or electronic crosstalk can create a false peak. Send a stable coherent reference through the complete

acquisition chain at comparable rates to test the baseline, and vary attenuation or cable delays to separate features tied to absolute electronics timing from those tied to optical delay.

The live graph displays an ideal recovery dip, its background-attenuated and bin-averaged counterpart, and the independent baseline. Predict the zero-bin counts by multiplying the plotted normalised value by the expected independent counts for a chosen acquisition. Increase bin width and lower signal fraction separately. Both reduce contrast, but only the finite bin averages over neighbouring delays. Measuring data with several bins from the same stored time tags tests that dependence without changing the source.

Open this experiment online

Try the calculation

Recovery time (ns) **2**

Bin width (ns) **4**

Signal fraction in each arm **0.8**



The raw zero-bin correlation is 0.595. At 40 and 30 kcounts/s for 100 s, the expected zero-bin count is 285.8 against an independent baseline of 480.0.

Read pulsed coincidences as areas of a peak train

A pulsed source with repetition period much longer than its emission decay produces coincidence peaks separated by that period. The central peak counts two detections attributed to the same pulse; neighbouring peaks can estimate the independent probability from different pulses. Integrate comparable windows around those peaks rather than comparing single-bin heights if jitter spreads clicks. A normalised central area of 25 counts against a mean adjacent area of 100 gives raw $g^{(2)}(0) \simeq 0.25$ under identical per-pulse brightness and efficiency. The result can coexist with obvious timing spread in every individual peak.

Slow blinking or pulse-to-pulse brightness correlation makes neighbouring peaks poor independent references. Plot several side-peak areas versus separation and compare with the mean at delays longer than the brightness correlation time. A decaying envelope in those areas contains source information that one baseline value conceals. At high detection probability, detector recovery across successive pulses can also suppress one or more nearby peaks. Vary repetition period or attenuation while keeping the binning procedure fixed to distinguish that recovery from intrinsic photon statistics.

Tutorial on coincidence data

Two equally bright independent ideal single-photon emitters contribute to one unresolved beam. Ignoring background and jitter, what zero-delay correlation follows?

▼ **Work through the solution**

Each emitter cannot supply two simultaneous signal photons, while cross-emitter pairs remain possible. For N identical independent sources $g^{(2)}(0) = 1 - 1/N$. Two therefore give one half. If their intensities differ, the reduction is the sum of squared fractional intensities, so equal brightness is a physical assumption. Background and finite resolution can mimic a change of emitter number.

If both arms are attenuated by a factor two and the signal-to-background ratio stays fixed, how should the expected coincidence counts and normalised correlation change?

▼ **Work through the solution**

The independent baseline $R_1 R_2 T \Delta t$ falls by a factor four and correlated counts scale the same way for independent linear losses. Normalised $g^{(2)}$ stays the same but counting uncertainties grow. If detector dark counts stay fixed while the signal falls, signal fractions also fall and the normalised contrast decreases.

A measured zero-delay peak gives $g^{(2)} = 1.40$ for thermal light with both signal fractions 0.80 and negligible timing loss. What effective number of equally populated independent thermal modes would explain it?

▼ **Work through the solution**

For M independent equal modes the excess is $1/M$. Correcting background gives $g_s^{(2)} - 1 = 0.40/0.64 = 0.625$, so $M_{\mathrm{eff}} = 1.60$. The noninteger value can represent unequally weighted modes or unresolved averaging. It is an effective contrast parameter unless the mode distribution has been established independently.

Hanbury Brown and Twiss analyse intensity fluctuations and correlations.

Connect source operation above threshold with the separate measurement of photon statistics.

Optical fields, instruments and observations · 7. Build a time lens from dispersion and a quadratic temporal phase

A temporal imaging system can stretch the time separation of two optical features without requiring them to arrive more slowly at the source. Derive the imaging condition from delay and frequency, calculate the sign of time magnification, and estimate the finite temporal aperture that limits which features can be recovered.

Derive two elementary maps in delay and frequency

Describe a narrowband wave packet by its delay t relative to a reference group and angular-frequency offset $\Omega = \omega - \omega_0$. With Fourier convention $A(t) \propto \int \tilde{A}(\Omega) e^{-i\Omega t} d\Omega$, a dispersive element contributes spectral phase $D\Omega^2/2$. Its derivative is group delay $D\Omega$. A narrow local time-frequency feature obeys $t' = t + D\Omega,\ \Omega' = \Omega$. This is a shear in time-frequency space, directly analogous in algebra to propagation of a paraxial ray, but D has units of time squared and the second coordinate is angular frequency.

A time lens applies temporal phase $+Kt^2/2$ to the envelope. With this sign convention the instantaneous angular-frequency offset is minus the derivative of envelope phase, so $\Omega' = \Omega - Kt$ and $t' = t$. The coefficient K has units of inverse time squared. Producing this modulation requires a physical phase-modulating or nonlinear optical process synchronised to the signal. A passive time-invariant piece of glass cannot implement an arbitrary quadratic temporal phase on an unknown incident time profile; it acts on spectrum instead.

Derive the imaging condition and track an individual feature

Place dispersion D_1 before the time lens and D_2 after it. Start with (t_0, Ω_0). At the lens $t_\ell = t_0 + D_1\Omega_0$; just after it $\Omega_\ell = \Omega_0 - Kt_\ell$. The output time is

$$t_{\text{out}} = (1 - KD_2)t_0 + (D_1 + D_2 - KD_1D_2)\Omega_0.$$

Sharp imaging requires the output time of an input feature to be independent of its initial frequency spread, so set $K = D_1^{-1} + D_2^{-1}$. Then $t_{\text{out}} = Mt_0$ with magnification $M = -D_2/D_1$. This is the finite-conjugate time-lens equation. For positive pre- and post-dispersion the magnification is negative: a later input feature appears earlier relative to the image reference. This is temporal inversion about a chosen reference delay, not propagation of information backward in laboratory time. The whole apparatus has a positive baseline transit delay common to the processed waveform.

With $D_1 = 100\,\text{ps}^2$, $D_2 = 400\,\text{ps}^2$, the required lens curvature is $0.0125\,\text{ps}^{-2}$ and $M = -4$. Two pulses initially separated by $8.0\,\text{ps}$ emerge $32.0\,\text{ps}$ apart in reversed relative order. A detector having a response width of $10\,\text{ps}$ may now resolve the separation, provided that the optical imaging system itself supplied adequate temporal resolution. Ideal energy-conserving magnification rescales the intensity envelope by $1/|M|$ as it stretches time, apart from losses and a possible output phase. Multiplying every time separation by four without reducing its peak intensity would create optical energy.

If the curvature is 2% below nominal,

$$K = 0.01225\,\mathrm{ps}^{-2}.$$

Then the coefficient of Ω_0 is $500 - 0.01225(40000) = 10\,\mathrm{ps}^2$, and $1 - KD_2 = -3.90$. An input spectral spread $0.20\,\mathrm{rad/ps}$ contributes an output delay spread of order $2.0\,\mathrm{ps}$ before other aberrations. A calibration pulse with a different spectral width can thus appear sharper even at the same imaging setting. Measure the output dependence on carrier detuning of a narrow feature to test the frequency coefficient directly.

Recover imaging from the wave integral and locate an aperture limit

The temporal propagation kernel of a quadratic spectral phase, up to constants and a uniform phase, is $\exp[-i(t-t')^2/(2D)]$ under the adopted transform convention. In the two-propagation integral multiply the first kernel, the lens phase $e^{iKt_\ell^2/2}$, and the second kernel. The coefficient of t_ℓ^2 is $(K - D_1^{-1} - D_2^{-1})/2$. At the imaging condition it vanishes. The remaining linear phase in t_ℓ is $t_\ell(t_0/D_1 + t_{\mathrm{out}}/D_2)$. Integrating over an unlimited lens time produces a delta condition $t_{\mathrm{out}} = -D_2 t_0/D_1$, recovering the ray-style map directly from wave interference. Quadratic input and output phase factors remain, even where the intensity envelope is correctly imaged.

A finite flat temporal aperture of total width T_a replaces the delta condition by a sinc amplitude. For a point feature its first zero in output delay is at $2\pi|D_2|/T_a$; referring it back to the input divides by $|M|$, giving $2\pi|D_1|/T_a$. With $T_a = 100\,\mathrm{ps}$, $D_1 = 100\,\mathrm{ps}^2$ this input first-zero scale is $6.28\,\mathrm{ps}$. A much larger time magnification cannot restore features the lens failed to collect over its temporal aperture. The criterion is a first zero of a rectangular-aperture field response; an intensity width or a Gaussian gate would use a different numerical factor.

Construct and calibrate an approximate phase profile

An electro-optic modulator driven sinusoidally can approximate a quadratic phase near an extremum. Expanding $\Phi_0\cos(\omega_m t)$ gives a constant, $-\Phi_0\omega_m^2 t^2/2$, then a quartic term $\Phi_0\omega_m^4 t^4/24$. Choose bias phase and sign to obtain the required K. The ratio of quartic to quadratic magnitude is approximately $(\omega_m t)^2/12$, so only part of the modulation period is a good lens. Widening the accepted temporal window improves nominal aperture resolution but can increase aberration in the wings. An independent phase calibration and a scan over input delays expose the tradeoff.

Distinguish magnification calibration from the image of a point

Launch two independently timed reference pulses flanking the centre of the field of view. If their input times are -5.0 and $+5.0\,\mathrm{ps}$, an ideal $M = -4$ instrument places their output times at $+20$ and $-20\,\mathrm{ps}$ relative to a common delay. Their $40\,\mathrm{ps}$ separation calibrates temporal magnification without

requiring knowledge of an absolute cable or optical path delay. The widths of the two output features probe the point-spread response and detector resolution. A single pulse can demonstrate broadening but cannot by itself distinguish image magnification from an imperfectly focused, unmagnified pulse.

Vary K around the predicted imaging value and use two thin pulses with different carrier detunings at the same nominal time. The spread of their output positions is proportional to the frequency coefficient $D_1 + D_2 - KD_1D_2$, which changes sign through best focus. Once that coefficient vanishes, repeat with inputs at different reference delays to measure M. This sequence separates a focus error from a scale error. A medium with large cubic dispersion can make the zero coefficient depend on carrier detuning; sampling a third detuning diagnoses the failure of the purely quadratic matrix model.

The live phase-space calculation tracks several narrow spectral components across the device for selected input times. At the correct lens strength their output lines have equal output time at fixed input time. Detuning the lens gives them different output arrival times. Changing signs of D_2 distinguishes an inverted image from one that preserves time ordering.

Open this experiment online

Try the calculation

First dispersion (ps²) **100**

Second dispersion (ps²) **400**

Lens K (10^{-3} ps^{-2}) **12.5**



Temporal map coefficients A = -4.00, B = 0.00 ps². The image is focused at M = -4.00.

Tutorial on an imaging system for time

A signal traverses $D_1 = +100\,\mathrm{ps}^2$ and $D_2 = -200\,\mathrm{ps}^2$. Which lens curvature and magnification form an image?

▼ **Work through the solution**

The lens condition gives

$$K = 0.010 - 0.005 = 0.005\,\mathrm{ps}^{-2},$$

with $M = +2$. The output keeps the ordering of features while doubling their separation. The signs refer to group-delay dispersion and temporal lens phase under the chapter's Fourier convention.

The previous positive-dispersion system images two pulses well near the centre of its time window but distorts them near the edge. Does this demonstrate a wrong constant D_2?

▼ **Work through the solution**

A uniform error of the imaging condition gives residual sensitivity to input frequency throughout the window. Edge-specific distortion suggests finite gate clipping, nonquadratic temporal phase or higher-order dispersion acting on different parts of the signal. Translate the same reference pulse across the lens timing and independently sweep its carrier detuning to separate field-of-view and chromatic errors.

If the input contains a short feature at $t_0 = 0$, can its optical phase always be inferred from a measured intensity image?

▼ **Work through the solution**

The ideal system maps complex envelopes with known quadratic phase factors, but an intensity-only detector measures the squared magnitude. It cannot generally recover unknown input phase from that one image. A coherent reference or additional phase-diverse measurements can supply missing information. The time-lens condition guarantees a mapping of the field under ideal propagation, not complete observability of every input degree of freedom.

Kolner and Nazarathy's original paper derives temporal imaging with a time lens.

Compare selection by a finite spatial Fourier aperture with the temporal acceptance window.

Optical fields, instruments and observations · 8. Recover an optical phase from one off-axis intensity frame

A camera measures intensity and discards field phase. A known tilted reference wave can shift a cross term away from the intensity background so the missing complex field becomes accessible. Derive how much camera bandwidth that costs, and convert the recovered phase into an optical-path measurement with a stated ambiguity.

Expand the recorded intensity and identify four terms

Let $O(x)$ be the unknown coherent object field on the detector and $R(x) = A_r e^{i2\pi f_c x}$ a known reference plane wave. They must have sufficient mutual coherence at the relative path delay, and their polarisation projections on the analyser must overlap. A square-law camera records $I = |O|^2 + |A_r|^2 + OA_r^* e^{-i2\pi f_c x} + O^* A_r e^{i2\pi f_c x}$. The first two terms are a real background and the last two are complex conjugates making real interference fringes. A local phase shift in O changes the fringe position without requiring the camera itself to follow the optical frequency.

Suppose the complex object spectrum has support approximately $|f| \le B$. The spectrum of $|O|^2$ is its autocorrelation and can extend to $|f| = 2B$. The cross terms occupy intervals centred at $\pm f_c$ with half-width B. In an ideal one-dimensional separation, $f_c - B > 2B$ keeps a sideband clear of the central intensity spectrum. A camera with pitch p has Nyquist frequency $f_N = 1/(2p)$, requiring $f_c + B < f_N$. Both inequalities hold simultaneously only if $B < f_N/4$ under these conservative support assumptions. Some fields have much narrower background or sparse spectra, but a designed full-bandwidth instrument cannot assume that without evidence.

Filter the negative-frequency sideband containing OA_r^*, shift it to zero frequency, and inverse transform. Dividing by the known complex reference amplitude recovers the object field up to an overall calibration phase. Filtering the opposite sideband gives the conjugate field. A reference tilt calibration error leaves a linear phase ramp after reconstruction. Finite detector field of view multiplies the hologram by a window and spreads its Fourier peaks. Use a tapered window or leave guard bands at the cost of spatial resolution. The filter radius can suppress noise while also discarding high object frequencies; it is a physical resolution choice as well as a plotting parameter.

Choose a carrier that a real pixel lattice can record

Take a camera pitch $p = 6.5\,\mu\mathrm{m}$, giving $f_N = 76.9$ cycles per millimetre. If the field bandwidth is $B = 12$ cycles per millimetre, the conservative carrier interval is $36 < f_c < 64.9$ cycles per millimetre. A choice $f_c = 48$ cycles per millimetre places the selected sideband between 36 and 60, clear of the central interval extending to 24. The fringe period is $20.8\,\mu\mathrm{m}$, or 3.21 pixels. At vacuum wavelength $532\,\mathrm{nm}$ and propagation close to normal in air the tilt angle is $\arcsin(\lambda f_c) \approx 1.46^\circ$. Raising the angle indefinitely would eventually alias the sideband back into the usable spectrum.

Each camera pixel integrates across its physical sensitive area. A full-width square pixel attenuates a sinusoidal fringe amplitude by $\mathrm{sinc}(\pi f_c p)$, with $\mathrm{sinc}\, x = \sin x / x$. At the example carrier this factor is about 0.847. Frequencies on either side of the band receive different attenuation, so a very broad sideband can have a frequency-dependent instrumental response even inside Nyquist. A flat-field

intensity correction fixes pixel-to-pixel gain, but does not undo this finite-area transfer. The sensor's usable dynamic range and read noise also constrain the reference amplitude.

Convert a reconstructed phase into a path change

Let the reconstruction relative to a reference frame provide $O_2O_1^* = |O_1O_2|e^{i\Delta\phi}$. The angle $\arg(O_2O_1^*)$ gives phase change modulo 2π, so $\Delta\mathrm{OPL} = \lambda\Delta\phi/(2\pi)$ modulo one vacuum wavelength. For transmission through a uniform nonabsorbing layer replacing ambient index n_a, an added thickness d gives optical-path change $(n - n_a)d$. For a normally reflecting surface displacement h the outgoing and return paths both change, giving $h = \lambda\Delta\phi/(4\pi)$ in air. Using the transmission conversion on reflection doubles inferred height.

A reflection measurement at $532\,\mathrm{nm}$ gives a local phase change $0.80\,\mathrm{rad}$, implying a height change $33.9\,\mathrm{nm}$ modulo $266\,\mathrm{nm}$ under the ideal normal-incidence geometry. In transmission through a layer of contrast $n - n_a = 0.10$, the same phase corresponds to $0.677\,\mu\mathrm{m}$ of thickness modulo $5.32\,\mu\mathrm{m}$. A single phase image cannot simultaneously determine thickness and refractive index without independent information. Phase unwrapping uses spatial or temporal continuity to choose integer cycles; at dark amplitude nodes or steep unresolved steps those choices can be ambiguous.

Background · How can a second wavelength extend the unambiguous path range?

Reconstruct the phase of the same optical path change d at two vacuum wavelengths. Assume the path change is independent of wavelength and both reconstructions share a calibrated zero. Each phase is known modulo 2π. Their difference increases more slowly with path than either phase alone.

$$\Delta\phi = 2\pi d\left(\frac{1}{\lambda_1} - \frac{1}{\lambda_2}\right) \pmod{2\pi},$$

$$\Lambda = \frac{\lambda_1\lambda_2}{|\lambda_2 - \lambda_1|}.$$

Wavelengths of 630 and 640 nm give a synthetic path period of 40.32 µm. For normal reflection the physical surface displacement is half the round-trip path change. A 0.05 rad uncertainty in the difference phase corresponds to about 0.32 µm of path uncertainty at this synthetic period. The long period helps identify a fringe order, while a single-wavelength phase can refine the displacement once that order is known.

Measure noise and reference aberration separately

The fringe visibility for uniform signal and reference intensities is $2\sqrt{I_oI_r}/(I_o + I_r)$ when coherence and polarisation overlap are perfect. Equal intensities maximise visibility. Increasing reference power can raise heterodyne signal above camera read noise, but once reference shot noise dominates,

continually increasing it does not improve sensitivity proportionally and can saturate pixels. A fixed total photon budget and a fixed object photon budget are different optimisation problems. Background subtraction should use a reference-only frame, an object-only frame when feasible, and a blank complex phase calibration acquired with the same optical geometry.

A wavefront aberration in the reference multiplies the recovered object by a spatially varying phase. Dividing the complex field by a blank-field reconstruction can remove a stable aberration, but division near very low blank amplitude amplifies noise. A phase map through masked pixels should disclose where phase was interpolated and where the signal actually supported it. Reconstruct at several sideband filter radii and carrier angles to check whether a fine feature moves or changes shape with the processing. A real specimen feature should remain in specimen coordinates under those controlled changes within the valid resolution.

Decide whether a thin specimen can be measured without unwrapping

If surface steps in reflection are known a priori to lie within a single interval of width $\lambda/2$, the wrapped phase can be mapped uniquely into that chosen range up to a common reference height. For $532\,\mathrm{nm}$ a known $60\,\mathrm{nm}$ relief is safely smaller than the $266\,\mathrm{nm}$ height period. An isolated $300\,\mathrm{nm}$ step can instead appear as $34\,\mathrm{nm}$ if there is no continuity information across its edge. A smoothness algorithm can select precisely the wrong branch at a real discontinuity.

Two wavelengths offer a different reference scale. For nearby wavelengths λ_1, λ_2, subtract their interferometric phases for the same optical path to obtain a synthetic optical-path period $\lambda_1\lambda_2/|\lambda_1 - \lambda_2|$. Reflection height has half that period. With 532 and $540\,\mathrm{nm}$, the height period is about $18.0\,\mu\mathrm{m}$, much larger than for either colour alone. The corresponding height sensitivity to a given phase noise is worse, so the synthetic phase often chooses the coarse branch while one individual wavelength supplies fine precision. This assumes both measurements interrogate the same position and calibrated path without a significant wavelength-dependent sample response.

Tutorial on carrier and phase measurements

The same $6.5\,\mu\mathrm{m}$ camera is asked to recover $B = 22$ cycles per millimetre with full separation from an object intensity spectrum extending to $2B$. Is there any carrier satisfying both conditions?

▼ **Work through the solution**

The carrier would have to exceed 66 and be below $76.9 - 22 = 54.9$ cycles per millimetre. No such carrier exists. Change magnification to reduce object bandwidth on the detector, use smaller pixels, or exploit multiple phase-shifted frames under an appropriate measurement model. Merely tuning the tilt cannot create detector bandwidth.

A reconstruction of a flat mirror shows an apparent phase ramp of 2π across $5\,\mathrm{mm}$ after demodulation. What carrier correction would remove a pure calibration ramp?

▼ **Work through the solution**

A phase change of one cycle over $5\,\mathrm{mm}$ corresponds to 0.20 cycles per millimetre. Adjust the assumed carrier by that magnitude with sign selected from the observed phase slope, then repeat on the reference mirror. Correcting a sample ramp without a blank could accidentally remove a real tilt of the sample.

A wrapped phase map appears very smooth even though the object amplitude almost vanishes on one side of an image. Does smoothness certify accuracy?

▼ **Work through the solution**

Near zero complex amplitude the phase responds very strongly to noise. Smoothing or interpolation can produce a visually regular map without adding phase information. Examine the sideband signal-to-noise ratio, repeat at independent exposures, and state an amplitude-based confidence mask. An optical-path conversion is quantitative only in the region that supplied a reliably measured complex field.

Cuche, Marquet and Depeursinge demonstrate Fourier separation of off-axis holographic orders.

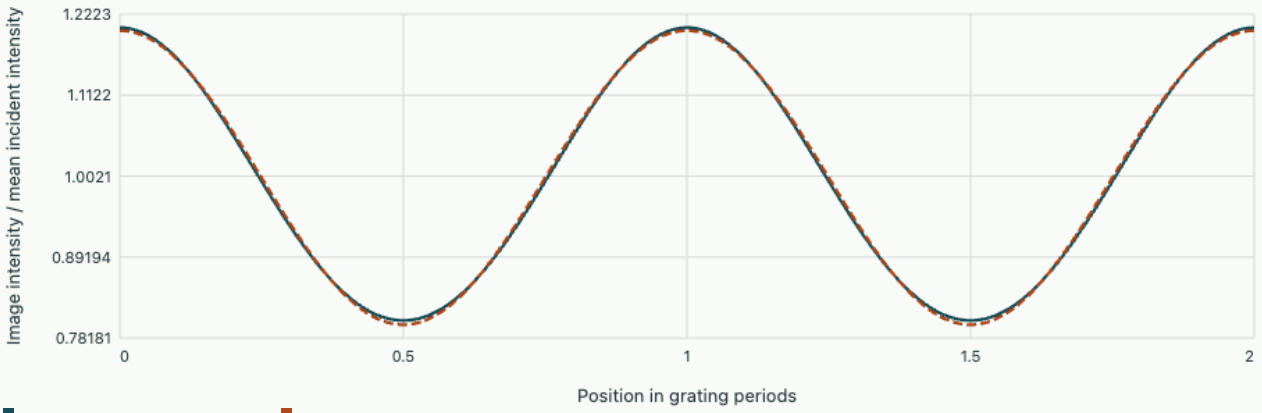
Vary a Fourier aperture and observe the resolution cost of filtering.

Open this experiment online

Try the calculation

Specimen phase amplitude ε (rad) **0.1**

Central order shift ϕ (degrees) **90**



All grating orders First-order phase contrast

Zero-order amplitude 0.9975. Peak-to-trough intensity difference 0.3983; linear prediction 0.4.

Optical fields, instruments and observations · 9. Resolve two nearby wavelengths by watching fringes disappear

A Michelson interferometer shows very fine bright and dark fringes as one mirror moves. Two narrow spectral lines produce a much slower variation of the fringe contrast. Use the mirror travel between contrast minima to determine the separation of the lines.

Combine two independent optical frequencies

Let d be mirror displacement from equal paths. The optical path difference of the returning beams is $2d$ and the phase at vacuum wavenumber $\sigma = 1/\lambda$ is $4\pi\sigma d$. Assume equal arm intensities for each individual line and an ideal detector that averages quickly oscillating optical fields and the beat between the distinct source frequencies. The measured intensities from the two spectral components add.

If the second line carries a fraction r of the first line's intensity, the modulated part of the signal divided by the total mean becomes

$$s = [\cos(4\pi\sigma_1 d) + r\cos(4\pi\sigma_2 d)]/(1+r).$$

Writing the two cosines as real parts separates a rapidly rotating phase from the complex envelope. Taking the magnitude produces the normalised visibility,

$$V(d) = \frac{\sqrt{1 + r^2 + 2r\cos(4\pi\Delta\sigma d)}}{1+r},$$

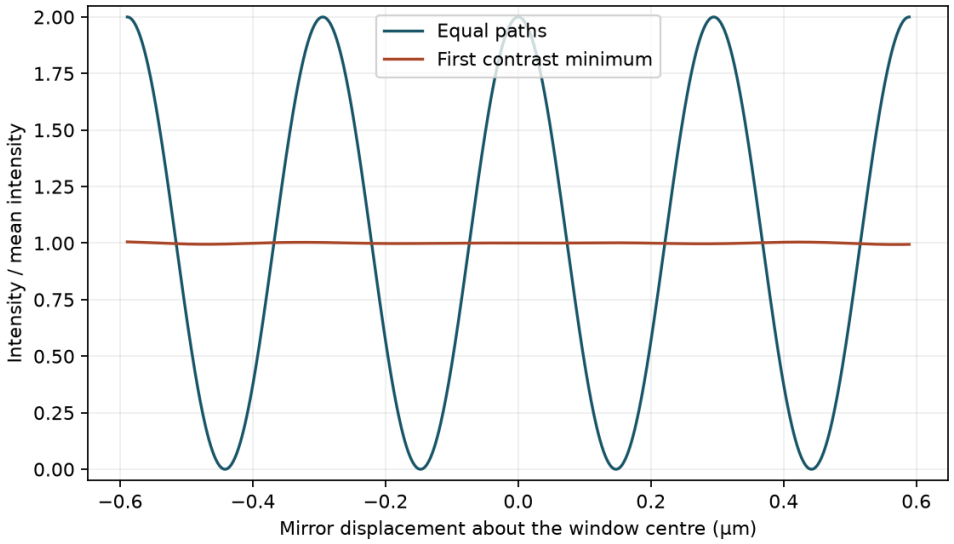
$$\Delta\sigma = |\sigma_1 - \sigma_2|.$$

This is the visibility obtained by scanning enough fine fringes to estimate local maxima and minima whilst the envelope changes negligibly. For equally strong lines it simplifies to $|\cos(2\pi\Delta\sigma d)|$. The first loss of contrast occurs at $d = 1/(4\Delta\sigma)$ and successive minima are spaced by $1/(2\Delta\sigma)$. The factor of two relative to optical path difference comes directly from reflection at the moving mirror.

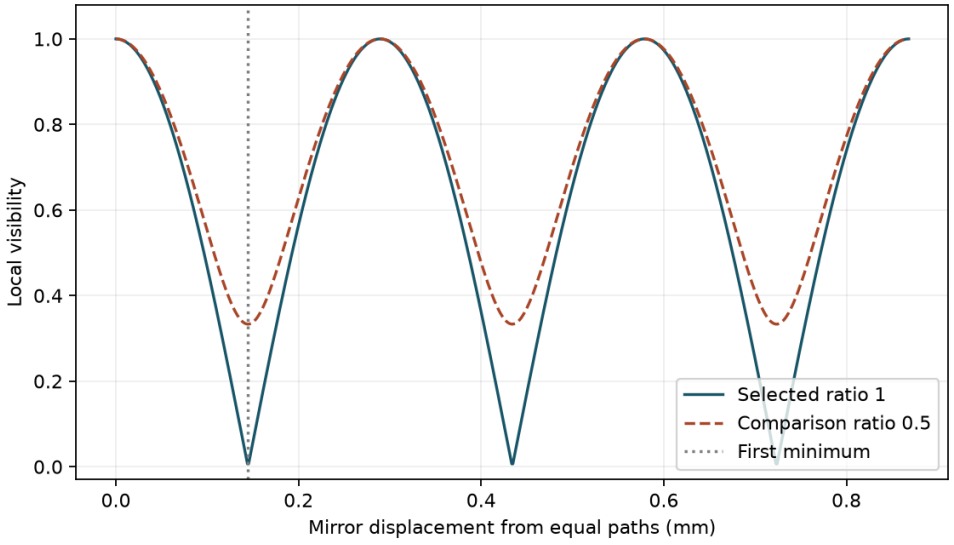
Compare mirror travel with a single optical fringe

Choose vacuum wavelengths 589.0 and 589.6 nm. Their wavenumber separation is $\Delta\sigma = 0.6/(589.0 \times 589.6)$ nm⁻¹, about 1.73×10^3 m⁻¹. The first visibility minimum lies at mirror displacement 0.145 mm and contrast revives at 0.289 mm. An individual mirror fringe spans only about 0.295 µm. Thus hundreds of carrier fringes fit inside a single envelope lobe. A coarse plot of raw intensity over a millimetre can alias those fast oscillations.

Suppose the second line carries half the intensity of the first. At the same minimum the two modulation vectors point in opposite directions. The surviving magnitude relative to their summed intensities is $|1 - 0.5|/(1 + 0.5) = 1/3$. An incomplete minimum can therefore arise even when both components are individually sharp. A finite width of either line or imperfect overlap of the arms in a real interferometer can also reduce the observed contrast. Independent measurements near equal path help identify such instrumental loss.



Each curve shows a short scan around the position indicated in its legend. Only the window near equal paths gives full contrast when the lines have equal intensity.



The separation between minima determines wavenumber splitting even when unequal line intensities prevent complete cancellation.

Exercise · infer a line separation from successive minima

A mirror moves 0.400 mm between two consecutive contrast minima. Estimate the wavelength separation for lines centred near 600 nm.

▼ **Work through the solution**

The wavenumber difference is $\Delta\sigma = 1/(2 \times 0.400\,\mathrm{mm}) = 1250$ m⁻¹. The narrow-doublet approximation $|\Delta\lambda| \simeq \lambda^2 \Delta\sigma$ gives 0.450 nm. Using the optical path difference as though it were mirror displacement would introduce a factor of two.

Exercise · infer the weaker line fraction

An ideal instrument observes maximum visibility one and minimum visibility 0.20. How strong is the weaker line relative to the stronger one? Can visibility label which line has the longer wavelength?

▼ **Work through the solution**

Calling the weaker-to-stronger ratio $r \leq 1$ gives $(1 - r)/(1 + r) = 0.20$, hence $r = 2/3$. The visibility magnitude is symmetric when the two lines are exchanged. Assigning the stronger intensity to a particular wavelength requires additional spectral information.

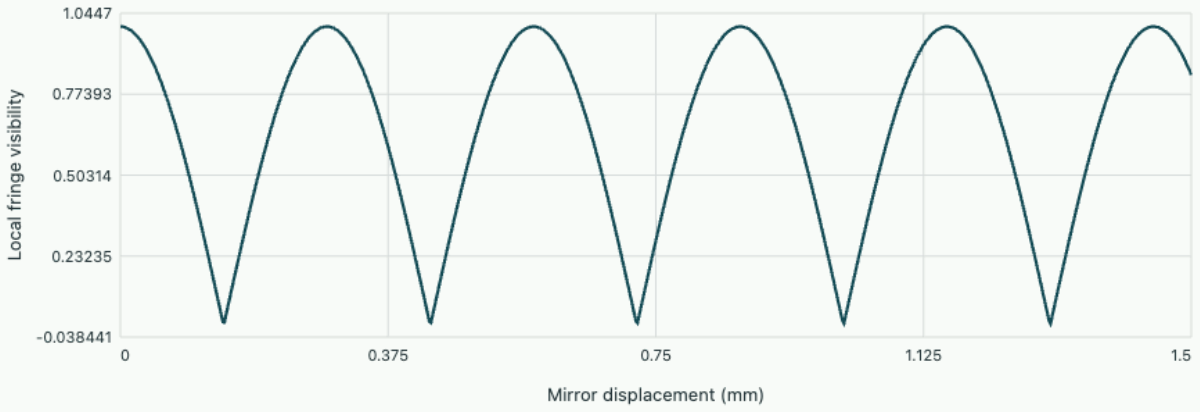
Try it · Doublet

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Wavelength splitting (nm)	**0.6**
Second / first line intensity	**1**



Incoherent spectral doublet

First contrast minimum 0.1447 mm; minimum visibility 0. Wavelength of first line is 589 nm.

Download the corresponding editable notebook. The browser recalculates the reduced equations directly; the notebook exposes the full calculation and its checks.

Open editable notebook Download .ipynb

Optical fields, instruments and observations · 10. Where a thick lens focuses and which aperture an instrument actually sees

The distance from a lens mount to a sharply focused sensor is often shorter than the focal length printed on the barrel. Two principal planes, together with the image of the physical aperture stop, reconcile these measurements and determine which rays an instrument accepts.

Compose refraction and translation with an explicit sign convention

Light travels left to right. Write a paraxial ray as (y, u) with $u = n\theta$. A positive radius has its centre of curvature to the right of the surface. Small-angle Snell refraction at a spherical surface gives $u_2 = u_1 - \Phi_i y$ with $\Phi_i = (n_2 - n_1)/R_i$. Translation by d in material of index n gives $y_2 = y_1 + (d/n)u_1$. Heights and slopes can therefore be propagated in the same two-component vector even when refractive index changes.

Multiply refraction at surface 1, internal translation through thickness d, and refraction at surface 2, writing the last transformation on the left. Let $\ell = d/n$. The elements of the vertex-to-vertex matrix in air are

$$M = \begin{pmatrix} 1 & 0 \\ -\Phi_2 & 1 \end{pmatrix} \begin{pmatrix} 1 & \ell \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -\Phi_1 & 1 \end{pmatrix} = \begin{pmatrix} 1 - \ell\Phi_1 & \ell \\ -\Phi & 1 - \ell\Phi_2 \end{pmatrix}, \qquad \Phi = \Phi_1 + \Phi_2 - \ell\Phi_1\Phi_2.$$

Its determinant is one. In the thin limit the surface powers add. Treating widely separated surfaces as one zero-thickness refractor discards the product term and generally changes both the power and where distances must be measured.

Factor the matrix and locate the cardinal points

For any air-to-air matrix with entries A, B, C, D and nonzero $\Phi = -C$, factor it into an air translation l_1 to the first principal plane, a thin lens of power Φ, and an air translation l_2 from the second principal plane to the back vertex. Matrix multiplication fixes $l_1 = (1 - D)/\Phi$ and $l_2 = (1 - A)/\Phi$. The effective focal length is $f = 1/\Phi$. Hence the back focal distance from the back vertex is $f - l_2 = A/\Phi$.

For a symmetric biconvex lens with $R_1 = +50$ mm, $R_2 = -50$ mm, $n = 1.50$ and $d = 10$ mm, each surface contributes 10 m⁻¹ and $\Phi = 19.333$ m⁻¹. The effective focal length is 51.72 mm. Each principal plane is 3.45 mm inside its adjacent vertex and the back focal distance is 48.28 mm. An object 200 mm left of the first principal plane gives an image 69.77 mm right of the second plane with lateral magnification -0.349.

Recall · What is fixed when the medium changes?

The expressions $f = 1/\Phi$ and unit angular magnification at nodal points here use equal unit indices in the incident and image media. With unequal external indices the front and back focal distances differ. An afocal system has $\Phi = 0$ and is handled directly by the ray matrix.

Identify the entrance and exit pupils

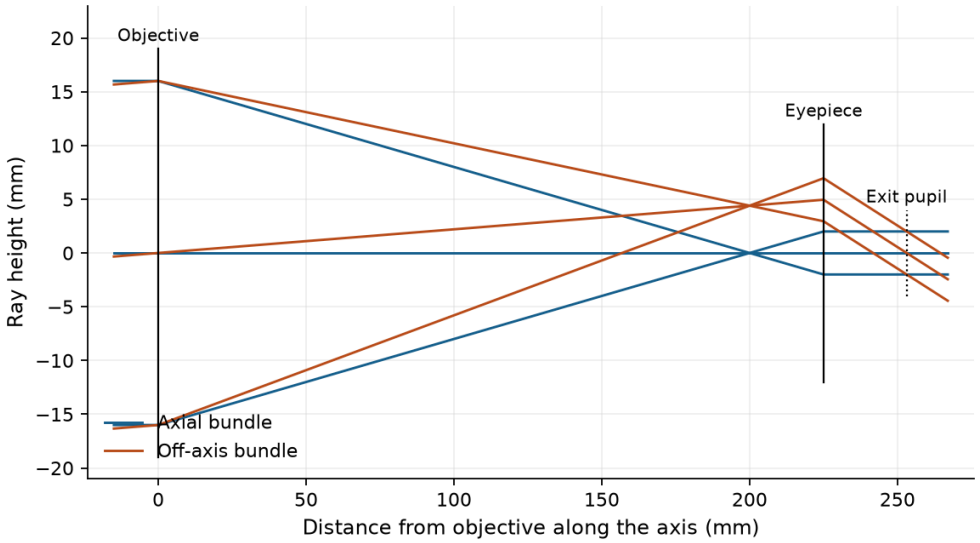
An axial point emits a bundle of possible rays. The element that first rejects an increasing cone angle in object space is the aperture stop. The entrance pupil is its image formed by the optics before it, viewed from the object. The exit pupil is its image formed by later optics, viewed from the image. A field stop constrains which off-axis object points are visible. It can be physically different from the aperture stop.

Put a 6 mm diameter iris 25 mm behind an air thin lens of focal length 50 mm. To find what an observer in front sees, trace rays backward from the iris. The thin-lens equation with propagation now to the left gives image distance -50 mm relative to the lens in that reversed convention. The stop therefore appears as a virtual 12 mm diameter opening 50 mm behind the lens. An axial object 150 mm in front of the lens sees a half-angle near $6/200 = 0.030$ rad. A sufficiently small lens rim would become the actual stop instead, invalidating this calculation.

Place the eye and relate field to throughput

An afocal two-positive-lens telescope separates its objective and eyepiece by $f_o + f_e$. A distant angle θ makes an intermediate height $f_o\theta$. The eyepiece then sends that point into an approximately parallel bundle at angle $-f_o\theta/f_e$. If the objective rim is the stop, its image through the eyepiece is the exit pupil. The eyepiece images an object at $f_o + f_e$ to distance $f_e(f_o + f_e)/f_o$ behind it. The exit-pupil diameter is $D_o f_e/f_o$.

With $f_o = 200$ mm, $f_e = 25$ mm and a 32 mm objective, angular magnification is -8, the exit pupil is 4 mm across and lies 28.1 mm behind the eyepiece. Moving a camera away from that plane can vignette off-axis bundles even with a sharp central image. A field diaphragm at the intermediate image separately clips the angular field.



Axial and off-axis paraxial bundles through the example telescope. Both fill the same 4 mm exit pupil at the indicated plane, while the eyepiece redirects the two bundles at different angles.

Try it · Telescope

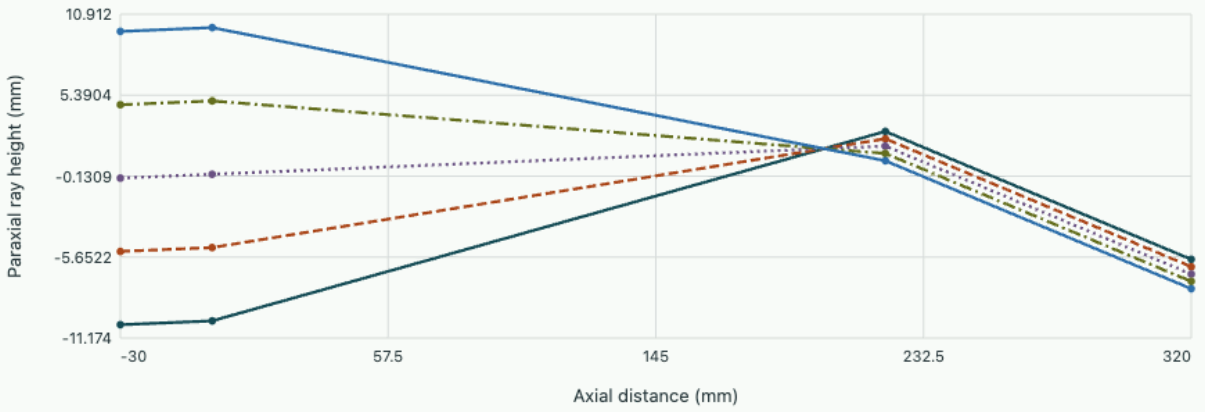
Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Eyepiece focal length (mm) 20

Field angle (degrees) 0.5



Afocal spacing 220 mm; angular magnification −10. A 20 mm objective gives an exit pupil 2 mm in diameter, 22 mm after the eyepiece.

Tutorial on lens-card and aperture measurements

1 · Focus a sensor from the vertex

▼ **Work through the solution**

In the thick-lens example the image for the 200 mm object is 69.77 mm from the second principal plane, hence 66.32 mm beyond the back lens vertex. A ruler referenced to the vertex must use the second value.

2 · Test whether the lens rim could be the stop

▼ **Work through the solution**

In the iris example an axial object at 150 mm and half-angle 0.030 rad uses a beam radius of about 4.5 mm at the lens. A 7 mm clear lens diameter clips it before the 6 mm iris can define the full bundle. Repeat the ray check for each potentially limiting aperture.

---

3 · Specify a telescope field stop

▼ **Work through the solution**

An intermediate focal-plane field stop of 8 mm diameter with the 200 mm objective gives an approximate full angular field $2(4/200) = 0.040$ rad or 2.29°. Halving its diameter halves this field. In a system without vignetting it leaves the axial exit-pupil diameter unchanged.

For further study see Hecht, Optics, 5th edition, chapters 5–6 on geometrical optics, cardinal planes and instruments.

Optical fields, instruments and observations · 11. From measured wavefront error to coma, astigmatism and focus

A point source can become a comet-shaped patch or two perpendicular line foci. Describing these observations by a single blur diameter loses the distinction between aberrations and suggests the wrong correction. The gradient of a pupil wavefront predicts the ray displacement at the detector.

Derive the image displacement associated with a pupil error

In a circular pupil of radius a write the normalised Cartesian coordinates $u = x/a$, $v = y/a$. A phase plate or optical error contributes a path difference $W(u, v)$ in units of length. Under monochromatic, paraxial conditions a pupil wave has the extra phase kW. Its local wave-vector tilt is $k^{-1}\nabla_\perp(kW) = \nabla_\perp W$. A focal-length f lens converts a small tilt into a displacement in its focal plane. With phase convention e^{ikW},

$$(\delta X, \delta Y) = \frac{f}{a}(\partial_u W, \partial_v W).$$

A constant piston gives no displacement. A linear tilt gives the same displacement to every ray. A quadratic rotationally symmetric term gives a radial displacement proportional to radius and can be cancelled by shifting the detector. The relation gives a geometrical spot construction. Near the diffraction limit, integrate the complex pupil amplitude to obtain the point-spread function.

Distinguish removable defocus from a fourth-order pupil error

On axis the pupil error of a centred rotationally symmetric lens has even powers of $\rho = \sqrt{u^2 + v^2}$. After choosing a focus the next independent term is $W = S\rho^4$. Its radial slope is $4S\rho^3$. Moving the detector can add a slope proportional to ρ and sharpen some zones while leaving others out of focus. Closing the aperture reduces the outer-radius contribution quickly, but increases the diffraction width on the sensor.

A lens whose full physical pupil is trimmed to half the radius still has the same physical curvature errors. For a fourth-power physical-height term, the edge value of W drops by 2^4 and its edge slope by 2^3. These scalings are measurements that can separate an aberrated lens from a diffraction-limited one.

Derive two off-axis spot signatures

Let the field point lie along the horizontal direction and let h be a dimensionless field coordinate. The leading coma term allowed by symmetry has the form $W_C = Ch\,u(u^2 + v^2)$. Differentiating gives

$$\delta X_C = (f/a)Ch(3u^2 + v^2),$$
$$\delta Y_C = (f/a)Ch(2uv).$$

At opposite edges of the horizontal diameter the horizontal error has the same sign. A round pupil therefore maps to an asymmetric fan with a brighter concentration of rays on one side. Reversing the field direction reverses the fan. With $f = 100$ mm, $a = 5$ mm and $Ch = 0.20$ µm, marginal horizontal rays shift by 12 µm.

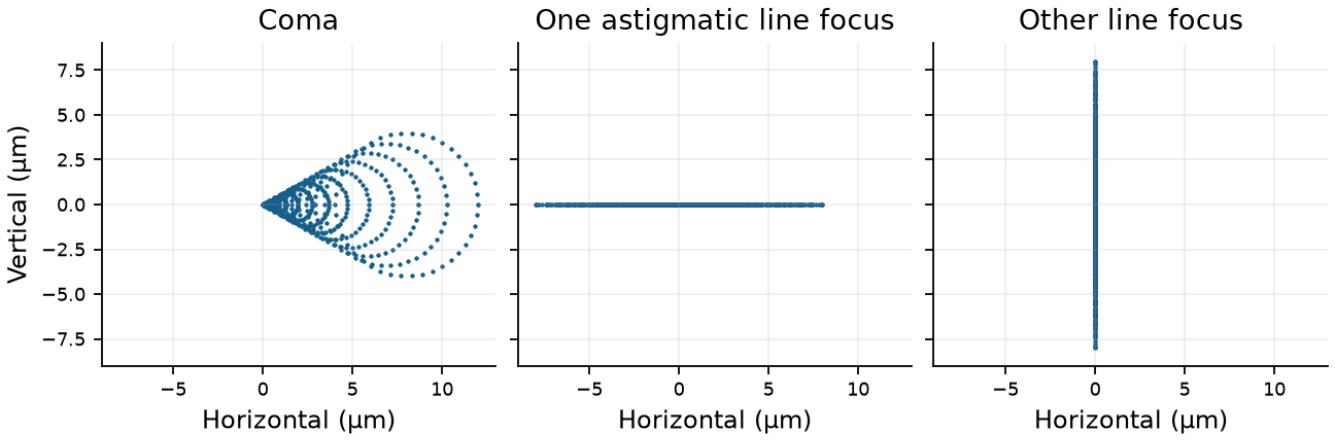
After subtraction of a field-dependent common focus error an astigmatic term can be written $W_A = Ah^2(u^2 - v^2)$. Horizontal slopes vary as $+2Ah^2u$, vertical slopes as $-2Ah^2v$. Add a small detector defocus with radial slope of adjustable sign. It can bring either meridional family to its own line focus, but cannot make both vanish together. A longitudinal detector shift Δz changes a marginal ray position by about $a\Delta z/f$. Equating its magnitude to $2f|A|h^2/a$ gives separation of the two line foci

$$|\Delta z_{\text{pair}}| \simeq 4f^2|A|h^2/a^2.$$

For $|A|h^2 = 0.10$ µm and the same geometry this gives 0.160 mm. Translating the detector between the two foci rotates the long axis of the spot by roughly 90°. Reversing field leaves this term even, giving an independent distinction from coma.

Recall · Separate field curvature and distortion

A common quadratic pupil term varying with the square of field shifts the best image surface. A field-dependent pupil tilt shifts the image position without blurring a single field point in this approximation. Measure multiple field positions and focus settings before attributing all deviations to coma.



Ray-intercept diagrams from the pupil gradients in the worked example. Coma produces an asymmetric fan. Astigmatism produces two perpendicular line foci at distinct detector positions. Diffraction would broaden each geometrical spot.

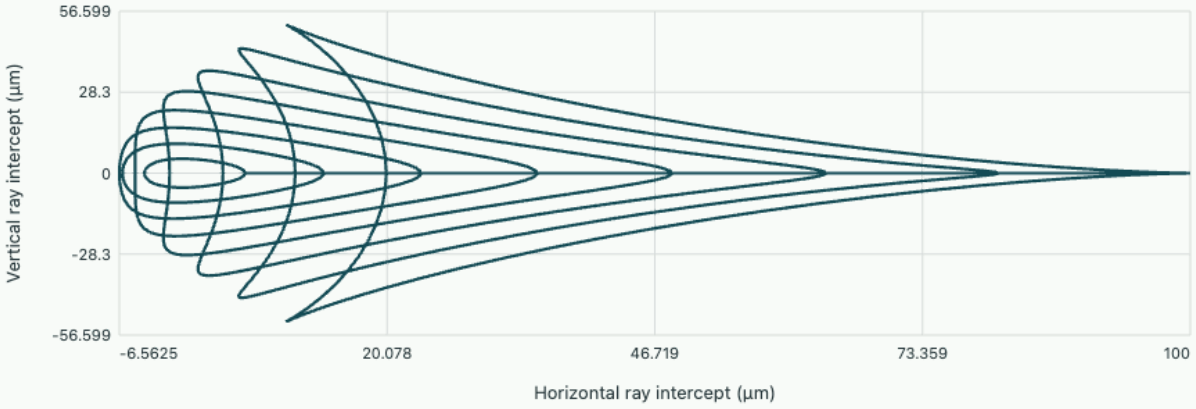
Try it · Aberration

Change the inputs and compare the recalculated curves with the preceding derivation and baseline figure. The axes give the units; the numerical readout states the model assumptions and any limiting cases. Reset the inputs to return to the reference setting.

Open this live calculation

Try the calculation

Coma coefficient (µm)	1
Astigmatism coefficient (µm)	1
Defocus coefficient (µm)	0



Pupil rings mapped by wavefront gradient

W = C u(u²+v²) + A(u²−v²) + D(u²+v²), with f/a=20. D=±A gives the two astigmatic line foci when C=0. This geometrical intercept diagram omits diffraction.

Tutorial on recovering an aberration from data

1 · Diagnose an aperture sweep

▼ Work through the solution

A measured edge contribution to blur falls approximately eightfold when diameter is halved while the centre stays fixed. In geometrical optics that agrees with a cubic ray slope from primary spherical aberration. For coma the marginal physical-height dependence is quadratic at fixed field. At small apertures the Airy contribution eventually grows.

2 · Extract an astigmatic path error

▼ **Work through the solution**

A line-focus separation of 0.08 mm at $f = 100$ mm and $a = 5$ mm implies $|A|h^2 = \Delta z\, a^2/(4f^2) = 0.050$ µm. A single image at best focus would not distinguish the signs of the two principal wavefront curvatures.

---

3 · Test a tilted detector as an alternative explanation

▼ **Work through the solution**

Record a grid of pinhole images on both sides of the axis and scan focus. A detector tilt creates a systematic best-focus gradient across the field. Coma changes sign with reversed field, and astigmatism produces two line foci at the same field position. These measurements separate errors that can look similar in a single frame.

For further study see Hecht, Optics, 5th edition, chapters 5–6 on instruments, ray tracing and aberrations.

Optical fields, instruments and observations · 12. Why light can generate its second harmonic and then lose it again

A nonlinear polarisation can oscillate twice during one period of the applied field. Whether the generated wave grows through a crystal depends on how successive source layers arrive in phase. This spatial phase relation makes crystal length and dispersion as important as local nonlinearity.

Generate a new frequency by expanding the material response

In a transparent, nonmagnetic dielectric take one relevant polarisation and a specified tensor projection. For moderate fields write $P = \epsilon_0(\chi^{(1)}E + \chi^{(2)}E^2 + \cdots)$. For $E = E_0\cos(k_1 z - \omega t)$ the quadratic term contains both a steady component and an oscillatory one:

$$P^{(2)} = \tfrac{1}{2}\epsilon_0\chi^{(2)}E_0^2\{1 + \cos(2k_1 z - 2\omega t)\}.$$

In a centrosymmetric bulk medium the electric-dipole response must change sign when the electric field changes sign. This symmetry forces the bulk electric-dipole $\chi^{(2)}$ to vanish. Surfaces, suitable noncentrosymmetric crystals or other coupling mechanisms can supply a nonzero response.

Add the radiation from successive slices of crystal

At 2ω the wave equation has the driven form $(\partial_z^2 + k_2^2)E_{2\omega} = -\mu_0(2\omega)^2 P^{(2)}_{2\omega}$ in the frequency domain. The linear optical response is already included in k_2. Set $E_{2\omega} = A_2(z)e^{ik_2 z}$ and neglect A_2'' compared with $2ik_2A_2'$. For a plane, monochromatic undepleted pump the source envelope is constant. Apart from an immaterial constant phase, the derivative is proportional to $e^{i\Delta k z}$, with $\Delta k = 2k_1 - k_2$.

$$A_2(L) \propto A_1^2\int_0^L e^{i\Delta k z}\,dz = A_1^2 L\, e^{i\Delta k L/2}\,\mathrm{sinc}(\Delta k L/2).$$

The intensity thus scales as $I_1^2L^2\,\mathrm{sinc}^2(\Delta k L/2)$ under these assumptions. Define $\mathrm{sinc}(x) = \sin x/x$. The coherence length $\pi/|\Delta k|$ gives the first maximum for fixed nonzero mismatch. At twice that length equal contributions from opposite phase regions cancel. A crystal growing longer can therefore reduce the harmonic output.

Predict the coherence length and use the index tensor

For vacuum pump wavelength λ, $k_1 = 2\pi n_1/\lambda$ and $k_2 = 4\pi n_2/\lambda$. In this collinear geometry $L_c = \lambda/[4|n_1 - n_2|]$. Taking $\lambda = 1.0$ µm and an index difference of 0.050 gives $L_c = 5.0$ µm. Birefringent phase matching selects propagation direction and polarisation so the indices felt by the two fields can coincide. The tensor component driving the generated polarisation must also be nonzero for those directions.

Another route periodically reverses the sign of $\chi^{(2)}$ every coherence length. When the source would begin to cancel the previous contribution, reversal restores its sign. For a square-wave reversal of period Λ, the fundamental reciprocal vector $2\pi/\Lambda$ can cancel the mismatch. First-order quasi-phase matching needs $\Lambda = 2L_c$ and gives an ideal first-harmonic coupling magnitude $2/\pi$ times the corresponding uniformly phase-matched value. Finite bandwidth, walk-off and absorption further limit the interaction length.

Recall · Distinguish frequency conversion from spatial harmonics

The second temporal harmonic carries frequency 2ω. A linear diffraction grating can contain twice a spatial frequency while scattering light at the original optical frequency. A spectrometer separates these two mechanisms even if intensity images alone look similar.

Include the pump energy balance

At high conversion the assumed constant source term overestimates the pump. Scale amplitudes so $|a|^2$ and $|b|^2$ are the fractions of incident power at ω and 2ω. Let z now denote distance in nonlinear coupling lengths and δ the correspondingly scaled mismatch. A lossless two-wave model after a phase choice is

$$\frac{da}{dz} = -a^* b\, e^{-i\delta z},$$

$$\frac{db}{dz} = a^2 e^{i\delta z}.$$

Multiply the first equation by a^* and the second by b^*, add complex conjugates and observe the cancellation. Total power is conserved. At $\delta = 0$, $a(0) = 1$, $b(0) = 0$ give real solutions $a = \operatorname{sech} z$ and $b = \tanh z$. This supplies an independent check for integrations with nonzero mismatch, for which power can return to the fundamental.

OneSimulation experiment · Pump depletion and phase mismatch in a lossless crystal

Try phase mismatch, depletion and time-step error

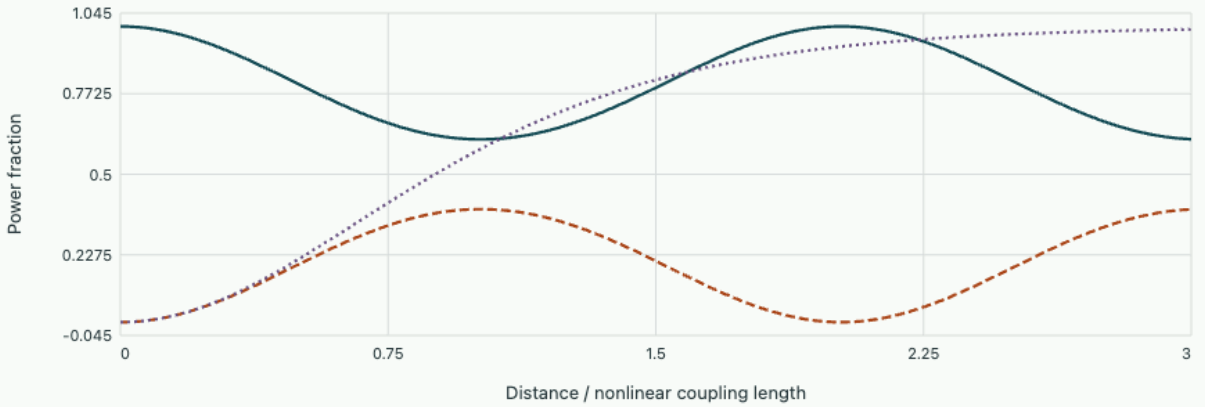
Change the phase mismatch and crystal length to locate a conversion maximum. At zero mismatch the computed second harmonic should follow tanh²(z/Lnl). At nonzero mismatch conversion can reverse; increasing length need not increase the output. Reduce the RK4 step while checking the separate power-balance residual.

The browser integrates the same four real envelope equations as the OneSimulation notebook below. Length and step are measured in nonlinear coupling lengths; mismatch is ΔkLnl. The lossless plane-wave model omits diffraction, absorption and temporal walk-off. Match `mismatch`, `step` and `count` in the notebook's `propagate` call to repeat a selected run.

Open this experiment online

Try the calculation

Phase mismatch Δk·Lnl	2
Crystal length / Lnl	3
RK4 step / Lnl	0.01



Fundamental power · RK4 | Second harmonic · RK4 | Exact second harmonic at Δk = 0

Exit harmonic fraction = 0.381279. Largest |Pω + P2ω − 1| = 3.68e-10. Actual step = 0.01. Normalised, lossless plane-wave model with pump depletion.

Reproduce with OneSimulation

Change the mismatch and step size. Inspect harmonic output and the conserved power separately. At exact phase matching compare the entire returned curve with the analytic hyperbolic-tangent solution.

Run the downloaded notebook with OneSimulation, NumPy and Matplotlib installed. The output below was recorded with the displayed calculation.

```
"""Normalised fundamental and second-harmonic envelopes in a lossless crystal."""
import numpy as np
import matplotlib.pyplot as plt
from onesim.numerics.time import ClassicalRungeKutta4TimeIntegrator

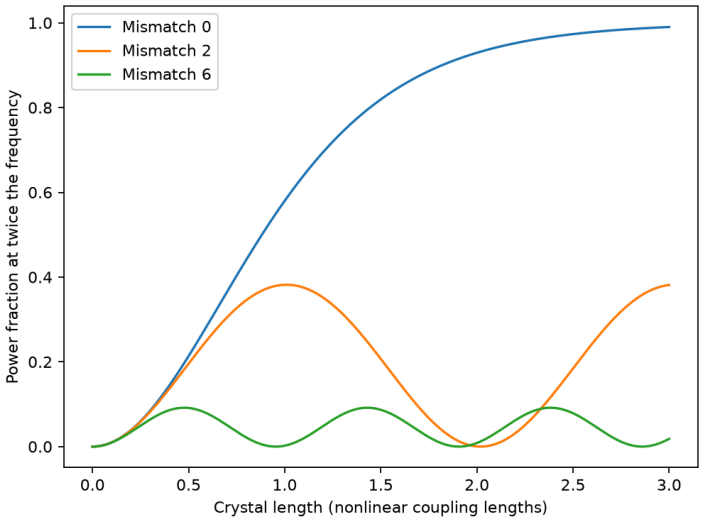
def propagate(mismatch, step, count):
    def rates(state, z):
        a, b = complex(*state[:2]), complex(*state[2:])
        phase = np.exp(1j*mismatch*z)
        da = -a.conjugate()*b/phase
        db = a*a*phase
        return da.real, da.imag, db.real, db.imag
    method = ClassicalRungeKutta4TimeIntegrator()
    state = (1., 0., 0., 0.)
    trajectory = [state]
    for i in range(count):
        state = method.advance(state, rates, step, time=i*step)
```

```
        trajectory.append(state)
    ar, ai, br, bi = np.array(trajectory).T
    return ar**2+ai**2, br**2+bi**2

z = np.linspace(0, 3, 301)
for mismatch in (0., 2., 6.):
    pump, harmonic = propagate(mismatch, .01, 300)
    print(f"Mismatch {mismatch:g}: harmonic fraction at z=3 = {harmonic[-1]:.5f}")
    print(f"Max energy residual = {max(abs(pump+harmonic-1)):.2e}")
    plt.plot(z, harmonic, label=f"Mismatch {mismatch:g}")
    if mismatch == 0:
        assert np.max(abs(harmonic-np.tanh(z)**2)) < 1e-7
plt.xlabel("Crystal length (nonlinear coupling lengths)")
plt.ylabel("Power fraction at twice the frequency")
plt.legend()
plt.tight_layout()
plt.savefig("physics-harmonic-generation.png", dpi=160)
```

Computed output

```
Mismatch 0: harmonic fraction at z=3 = 0.99013
Max energy residual = 5.36e-11
Mismatch 2: harmonic fraction at z=3 = 0.38128
Max energy residual = 3.68e-10
Mismatch 6: harmonic fraction at z=3 = 0.01833
Max energy residual = 6.10e-10
```



Pump depletion and phase mismatch in a lossless crystal.

Download editable notebook · Download Python script

Tutorial on coherence length and power

1 · Infer a dispersion difference

▼ **Work through the solution**

A 1.064 µm pump has its first weak-conversion maximum at length 20 µm. The index difference is $\lambda/(4L_c) = 0.0133$. The next zero is near 40 µm if mismatch and pump intensity remain fixed.

---

2 · Separate weak conversion and depletion

▼ **Work through the solution**

Doubling length under perfect phase matching multiplies weak harmonic power by four. For a depleted pump at $z = 1$ and 2 the fractions are $\tanh^2(1) = 0.580$ and $\tanh^2(2) = 0.929$. Multiplying 0.580 by four would exceed the available power.

---

3 · Tune a grating with a measured mismatch

▼ **Work through the solution**

For $|\Delta k| = 6.28 \times 10^5$ m⁻¹ first-order quasi-phase matching needs period approximately 10.0 µm and sign reversal every 5.0 µm. A residual mismatch of 10^3 m⁻¹ in a 4 mm sample yields a weak-conversion intensity factor $\operatorname{sinc}^2(2) = 0.207$ relative to exact matching.

For further study see Hecht, Optics, 5th edition, chapter 13 on nonlinear optics and optical frequency conversion.

Optical fields, instruments and observations · 13. How a material changes the speed, attenuation and direction of light

A detector behind a sample loses light both when matter absorbs energy and when light is scattered outside its acceptance. A microscopic oscillator connects refraction to absorption. A radiating dipole then shows how a transparent particle can also reduce the measured direct beam.

Start with the electromagnetic wave equation in a medium

For a nonmagnetic, neutral, homogeneous dielectric with no free currents, Maxwell's curl equations give $\nabla \times \nabla \times \mathbf{E} = -\mu_0 \partial_t^2 (\epsilon_0 \mathbf{E} + \mathbf{P})$. A transverse plane wave in an isotropic material has $\mathbf{k} \cdot \mathbf{E} = 0$. Substitution for time dependence $e^{-i\omega t}$ gives $k^2 = (\omega^2/c^2)(1 + \chi)$. The local linear response is $\mathbf{P} = \epsilon_0 \chi(\omega) \mathbf{E}$. If $k = (\omega/c)(n + i\kappa)$, the real index controls phase accumulation while positive κ attenuates intensity as

$$I(z) = I(0) e^{-\alpha z},$$

$$\alpha = 2\omega\kappa/c.$$

Two samples of the same material with identical polished interfaces can separate bulk attenuation from a nearly constant surface loss. If increasing thickness by 5.0 mm reduces transmitted intensity by a factor of 0.80, the apparent bulk coefficient is $-\ln(0.80)/0.005 = 44.6$ m⁻¹. Multiple internal reflections, stray light and different surface treatments would spoil this cancellation.

Derive the dispersive and absorptive parts of the response

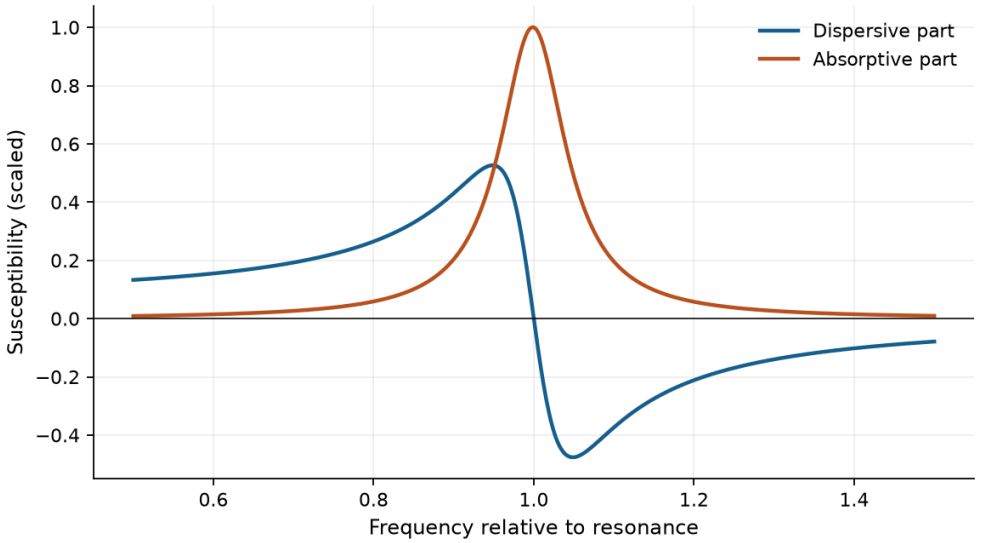
Let a bound charge of magnitude e and mass m have a restoring frequency ω_0 and damping rate Γ. A charge $-e$ obeys $m(\ddot{x} + \Gamma \dot{x} + \omega_0^2 x) = -eE$. The polarisation is the number density N times the dipole $-ex$. Solving the driven equation gives

$$\chi(\omega) = \frac{Ne^2/(\epsilon_0 m)}{\omega_0^2 - \omega^2 - i\Gamma\omega}.$$

Below resonance the induced dipole tends to follow the applied field. Near resonance the imaginary part peaks and the real part evolves rapidly through a sign change. In a transparent region a narrow pulse has group index $n_g = n + \omega\, dn/d\omega$ from differentiating $k(\omega)$. Near strong absorption the pulse reshapes. A time-of-peak measurement alone may then give an unreliable travel time. The model links the two pieces of material response through a causal equation of motion.

Recall · Sign convention for a lossy material

With $e^{i(kz-\omega t)}$ a positive imaginary part of the wave number makes the amplitude decay. A time convention with the opposite sign reverses how the imaginary susceptibility is written while the positive measured attenuation stays the same.



A single passive resonance with damping rate one tenth of its undamped angular frequency. The dispersive phase response reverses sign near the frequency where absorption peaks.

Try it · Em dispersion

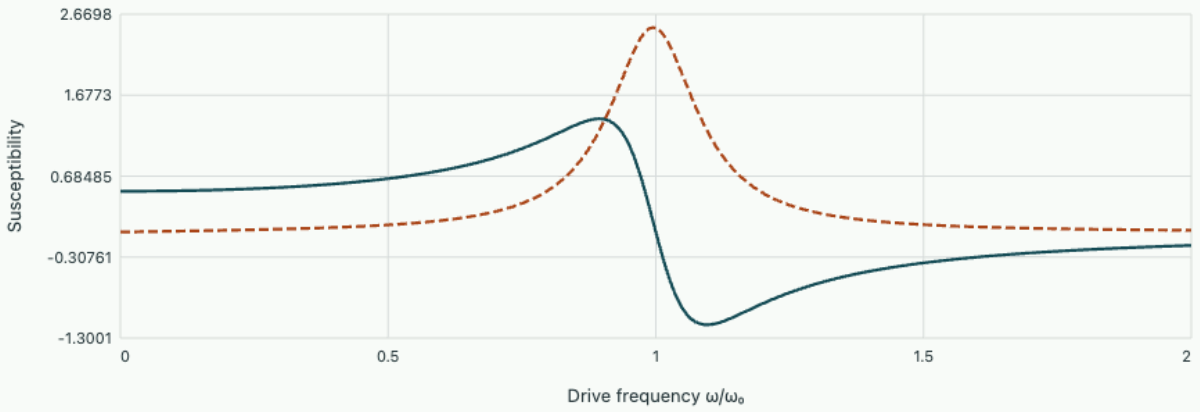
Explore the same physical relation with adjustable inputs. This independent parameter study has its own displayed reference settings; compare its assumptions and units before matching the static example.

Open this live calculation

Try the calculation

Damping γ/ω₀ **0.2**

Oscillator strength ω_p^2/ω_0^2 **0.5**



At ω = ω₀ the susceptibility is purely imaginary with Im χ = 2.5. At zero frequency the real susceptibility is 0.5.

Give a lossless particle a radiative channel

For a particle much smaller than the wavelength in vacuum use its electrostatic dipole amplitude $p = \epsilon_0 \alpha_V E$, with α_V in cubic metres. For a sphere of radius a and relative permittivity ϵ_r the electrostatic boundary conditions give $\alpha_V = 4\pi a^3 (\epsilon_r - 1)/(\epsilon_r + 2)$. Its transverse far field is proportional to $k^2 p \sin\theta/(4\pi\epsilon_0 r)$. Squaring and integrating $\sin^2\theta$ over solid angle gives

$$\sigma_{\mathrm{sca}} = \frac{k^4}{6\pi}|\alpha_V|^2 = \frac{8\pi}{3} k^4 a^6 \left|\frac{\epsilon_r - 1}{\epsilon_r + 2}\right|^2 .$$

Here the scattering angle is measured from the driven dipole direction. For a dilute independent suspension the collimated intensity approximately obeys $I = I_0 e^{-N\sigma_{\mathrm{ext}} z}$ with $\sigma_{\mathrm{ext}} = \sigma_{\mathrm{abs}} + \sigma_{\mathrm{sca}}$. The static polarizability gives the leading scattering power. For a quantitatively consistent forward

amplitude include the small radiation-reaction term in the dipole response. If ka ceases to be small, higher multipoles contribute and the Rayleigh scaling ceases to hold.

For equally sized dilute scatterers with weakly dispersive contrast the cross-section ratio at 450 and 650 nm is $(650/450)^4 = 4.35$. Doubling radius at fixed particle number gives a factor of 64 when Rayleigh conditions hold at both radii. At fixed total particle volume, the number falls by eight on doubling the radius, so the attenuation instead increases by a factor of eight.

Tutorial on measuring optical constants

1 · Recover the imaginary index

▼ Work through the solution

At 600 nm the example attenuation 44.6 m⁻¹ corresponds to $\kappa = \alpha\lambda_0/(4\pi) = 2.13 \times 10^{-6}$ if attenuation is purely absorptive and scattering is negligible. Angularly resolved or integrating-sphere measurements test that last condition.

---

2 · Separate phase and pulse travel times

▼ Work through the solution

A transparent 2 cm slab has $n = 1.50$ and $n_g = 1.56$. The phase and group times are 100 ps and 104 ps respectively using $c = 3.0 \times 10^8$ m/s. A narrowband phase measurement alone does not determine the slope of the index.

---

3 · Test whether attenuation is molecular absorption

▼ Work through the solution

Measure transmitted collimated power and total scattered power with matched geometries. A wavelength-dependent loss proportional to λ^{-4} and increasing strongly with particle size supports Rayleigh scattering. Material absorption may also vary with wavelength, so independently test size and angular dependence.

For further study see Hecht, Optics, 5th edition, chapters 2–4 on electromagnetic waves and the propagation of light.