APP NOTE

Two Types of Instruments, One Optical Model

Why pairing spectroscopic ellipsometry with spectrophotometry produces better results than either technique on its own

Plus: how these techniques can be used to reveal hidden layers in your phone

Introduction

Any modern optical device, whether it's the touchscreen on your phone, a solar cell, or an anti-reflection coating on glass, is a complex engineered system. Many optical devices consist of several thin layers stacked according to their behavior when interacting with light. For example, a bandpass filter is a stack of material engineered to transmit a narrow band of wavelengths while reflecting the rest.

Before a stack of thin films is fabricated, the engineer designs the device as a simulation — and that simulation is only as good as the optical constants fed into it. Refractive index and extinction coefficient, n and k, describe how a material bends light and how strongly it absorbs it, wavelength by wavelength. These constants determine how much light reflects at every interface in a coating and which wavelengths of light can make it through a filter. And they often cannot be borrowed from a handbook, because n and k depend on how a film was deposited, not on what it is made of. The optical constants of each film in each optical device must be measured and then those results must be fit to a model to determine the device's optical properties. In this application note, we will show that the best and most reliable models are built by using spectroscopic ellipsometry (SE) in combination with spectrophotometry.

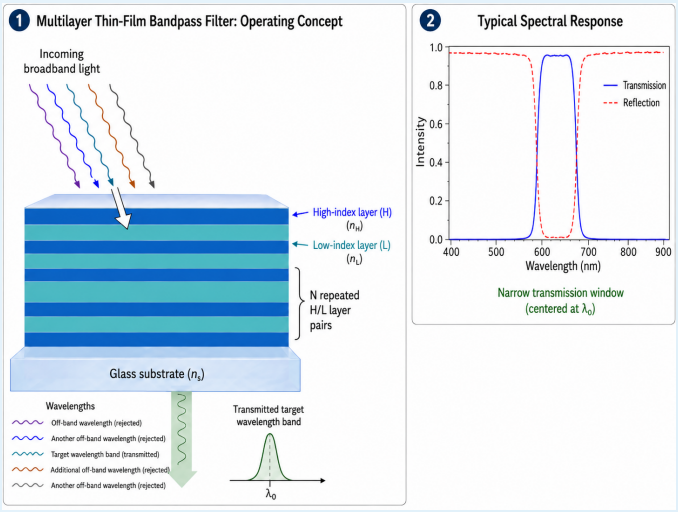


Figure 1: Diagram of bandpass filter showing alternating stacks of material, each reflecting or transmitting different amounts of incoming light (left) and reflectance spectrum (top right). Each boundary between the high-index and low-index material reflects only a small fraction of incoming light. On its own that's a poor mirror. But make every layer a quarter-wavelength thick optically, and the reflections coming back from all those boundaries arrive in phase with each other, so they add constructively. Stack up twenty or thirty pairs and those weak reflections sum to better than 99% reflectivity — better than most metal mirrors, and with almost no absorption loss.

The Problem with Trusting a Single Measurement

Refractive index (n) and extinction coefficient (k) are crucial constants for optical device engineers; n defines how much light slows down inside the material and k defines how strongly the material absorbs light. A designer who guesses at these values or borrows a literature value measured on someone else's film deposited by a different process, designs a filter that misses its target wavelength. The unglamorous reality is that materials as-deposited rarely match handbook values, because n and k depend on density, stoichiometry, crystallinity, and deposition conditions. In order to get trustworthy values of n and k, they must be measured.

However, n and k are not quantities an instrument reads off directly. No detector reports them. What an instrument records is how light behaved: how its polarization changed on reflection, how its phase shifted, or how much of it was transmitted through or reflected by the sample. Getting from that raw signal to optical constants requires building a mathematical model of the sample and adjusting the model until it reproduces the measurement.

This carries a consequence worth stating plainly: When an optical model matches the data, the match is not proof the model is correct. It proves only that the model is capable of producing that data.

Take a simple case: one thin film on a substrate. Its behavior is governed largely by optical thickness (n•d) the physical thickness (d) multiplied by the refractive index (n). Raise the refractive index slightly, reduce the thickness proportionally, and the result is nearly unchanged. Two physically different structures that are reproducing measurement curves to nearly the same accuracy. If the only data in hand is a single transmittance curve, nothing in it reliably separates them — both answers fit almost equally well, and the modeling software will report whichever one it happened to reach. Add more variables and the problem compounds, because a sufficiently flexible model can be driven to fit almost any dataset, including one describing a structure that does not physically exist.

For the engineer relying on those numbers, the failure is not academic. Optical constants are fed into designs that get made into real devices. A set of constants that reproduces the measurement while misstating the physical properties of the film produces a design that behaves correctly in simulation and incorrectly in hardware.

The remedy is not a better fitting algorithm. It is more independent information — and specifically, information of a different kind

Quick Review of Optics Topics

What happens to light when it interacts with a surface?

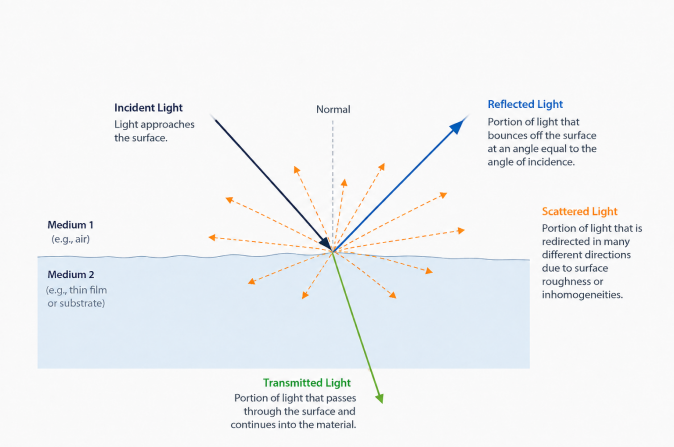


Figure 2: Portions of the incoming light may be reflected from the surface, transmitted into the material, or scattered in multiple directions due to surface roughness or other inhomogeneities.

What is polarization?

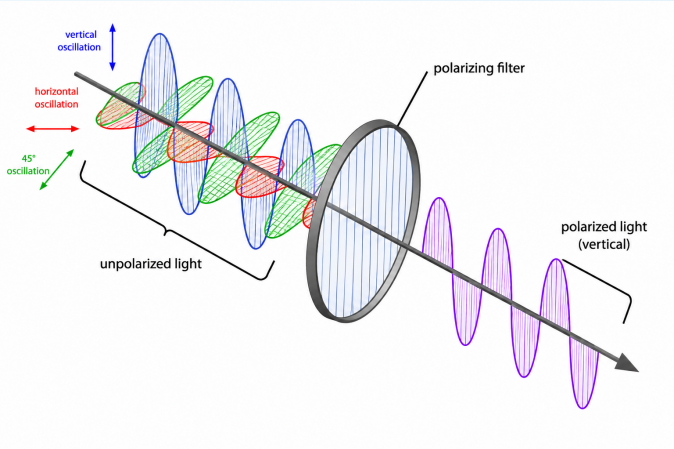


Figure 3: The polarization of light defines the direction of vibration of each wave's electric field. Unpolarized light is light that is oscillating in a random mix of all directions. A polarizing filter allows light that is polarized in a specific direction to pass through while blocking all the rest.

How do surfaces polarize reflected light?

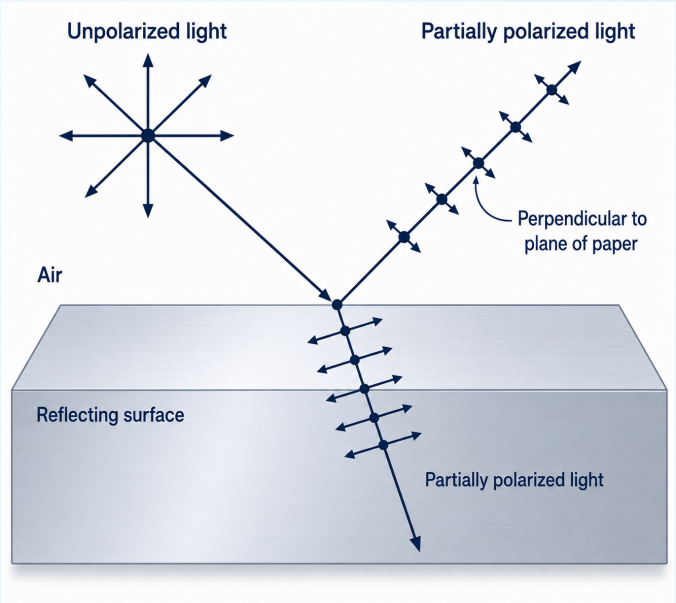


Figure 4: When unpolarized light reflects off a surface at oblique incidence, it becomes partially polarized. Reflection favors the s component of the light, the component that is parallel to the reflection surface (perpendicular to the plane of incidence).

What is a phase shift?

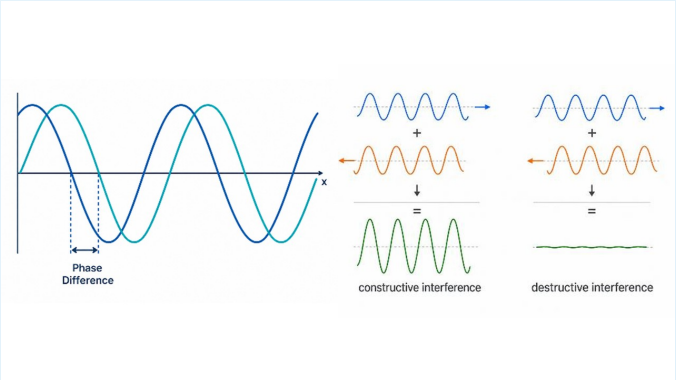


Figure 5: A phase shift means the wave has been nudged forward or backwards along its wave cycle, so its peaks arrive slightly earlier or later than they otherwise would — and when two waves are compared, a phase shift between them means their peaks no longer line up. This will result in interference which can change the amplitude of resulting waves.

How do surfaces change the phase of reflected light?

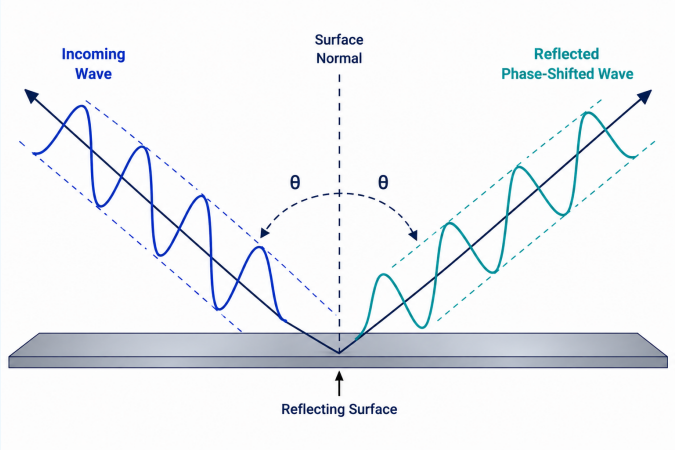


Figure 6: When light reflects at a boundary it doesn't simply bounce — the interaction itself nudges the wave along in its cycle, and the size of that nudge depends on the two materials meeting at the boundary, the angle the light arrives at, and which direction the wave is oscillating.

What is spectroscopic ellipsometry?

Spectroscopic ellipsometry (SE) shines a polarized light beam on a sample at a given angle and records changes in phase and polarization of the reflected beam. The instrument repeats this measurement across multiple wavelengths and angles of incidence.



Figure 7: Image of the ellipsometer we use at Covalent for wavelengths from 193 to 1690 nm (deep UV through most of the near IR range), J.A. Woollam RC2-DI.

An ellipsometer does not measure how much light a sample returns, it just measures changes in phase and changes in polarization. Two components of the incoming wave, the component oscillating in the plane of incidence (p) and the one across it (s) are reflected by the surface differently and the beam comes back with p and s no longer in their original relationship. Two quantities describe the difference: Ψ and Δ.

Ψ records how the amplitude ratio of the polarization changed: which component survived the reflection better, and by how much. Δ is defined by the change in phase difference between the p and s polarized components. The important take away is that both of these variables are defined by ratios describing p relative to s.

Because Ψ and Δ are derived from the ratio of two reflection coefficients measured from the same incident beam, any common absolute-intensity term such as source output, detector response or optical throughput cancels out, leaving SE sensitive only to the relative amplitude and phase difference between p- and s-polarized light. The ratio structure means these values are inherently self-normalizing and extremely precise; the phase information (Δ) makes it exceptionally sensitive to film thickness and to the real index n, down to films only nanometers thick.



Figure 8: Image of the ellipsometer we use at Covalent for wavelengths from 1.7 to 30 microns (near IR into the far IR), J.A. Woollam IR-VASE Mark II.

However, the ratio structure also sets limits. Because Ψ and Δ compare p against s, they carry no information about the absolute quantity of light reflected— if the whole reflected beam is dimmed, both values stay the same. Ellipsometry therefore has no way to notice that light has gone missing due to scattering; it cannot check its own books. And where a film absorbs only weakly, that faint absorption barely alters the reflected polarization, so very small values of k remain poorly determined. Ellipsometry can establish that a film is nearly transparent. It cannot say precisely how nearly. To build a truly accurate model, we must account for all the light. This gap calls for a measurement that reports light on an absolute scale. This is where spectrophotometry becomes crucial.



Figure 9: Image of the Fourier Transformed Infrared Spectrometer (FTIR) we use at Covalent for wavelengths from 2.5–25 μm (mid IR to far IR). Thermo Fisher Nicolet 6700

What is Spectrophotometry?

Spectrophotometry (UV-Vis and FTIR) records the absolute intensity of light transmitted, reflected, or absorbed by a sample at normal incidence, expressed as a percentage of incident intensity at each wavelength.



Figure 10: Image of the UV-Vis-NIR instrument we use at Covalent for wavelengths 190–2500 nm (UV through the complete NIR band). Perkin Elmer Lambda 1050

Spectrophotometry's strength is direct sensitivity to weak absorption accumulated over the full optical path. Because the signal integrates absorption through the entire sample thickness, transmittance picks up subtle absorption features that reflection-based techniques like ellipsometry can miss. It also returns an absolute magnitude — a real percentage of light transmitted — which ellipsometry, as a ratio and phase measurement, does not directly provide.

The limitation is that transmittance is degenerate on its own. A single % transmittance curve can be reproduced by more than one combination of thickness and index, so the fit isn't unique — many thickness/index pairs can generate the same curve. That's because transmittance carries no phase information which is typically needed to break index-thickness degeneracy.

Scattering: Another Problem

A perfectly clear film transmits light in a straight line. Introduce scattering — from surface roughness, grain boundaries, voids, inclusions — and some of that light is redirected into other angles. Enough of it and the sample looks hazy.

This creates a trap for conventional analysis. When a spectrophotometer records less light arriving at the detector than went in, the usual interpretation is that the sample absorbed the difference. That interpretation is only valid if nothing was scattered. If light was instead deflected out of the detector's path, it gets counted as absorption that never happened — and the extracted absorption values come out inflated. The measurement looks perfectly reasonable, however, it is wrong.

The way out is to stop assuming and start accounting. Every photon entering the sample ends up in one of four places: transmitted straight through, transmitted but deflected off its original path, reflected straight back the way it came (specular reflection), or reflected back scattered in many directions (diffuse reflection). Whatever cannot be found in those four channels was genuinely absorbed. So, the measurement strategy is to close the books on all four. An integrating sphere makes this possible.

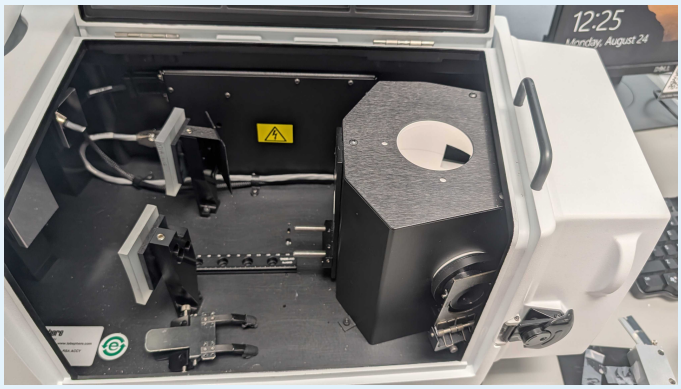


Figure 11: Image of the integrating sphere we use a Covalent. The sphere is contained within the black compartment on the right side of the image. This instrument allows us to account for all light.

The integrating sphere's interior collects light arriving from every direction and delivers it to a single detector, so it can capture deflected light that a conventional detector would miss. Positioning the sphere on the far side of the sample measures transmission; moving it to the near side measures reflection. In each geometry four quantities are recorded:

1. **Total transmittance** — all light emerging forward, deflected or not

2. **Specular transmittance** — only the light that passed through undeviated

3. **Total reflectance** — all light returning, in any direction

4. **Specular reflectance** — only the mirror-direction reflection

The difference between total and specular in each geometry is the diffuse, scattered portion. Four numbers, and the light budget is complete.

Combining Ellipsometry with Spectrophotometry to Account for All Light

Those four spectrophotometry measurements still cannot be interpreted alone. Extracting absorption and scattering from them requires fitting a radiative transfer model that correctly accounts for the ordinary boundary reflections occurring at every interface, the same way a clean piece of glass reflects a little light with no scattering involved. Those boundary reflections depend on the refractive index. Ellipsometry has vastly better sensitivity to index than the intensity measurements, so using the ellipsometry data in conjunction with the spectrophotometry data is what makes the rest of the accounting trustworthy.

Why Fitting Both Techniques Together Wins

There are four distinct reasons a combined SE-plus-transmittance fit outperforms either technique run alone.

1) Complementary sensitivity

SE constrains thickness and refractive index tightly, because the ellipsometric ratio is exquisitely sensitive to both. However, it has comparatively little leverage on weak absorption, on the small-k end of the spectrum, where the extinction coefficient barely perturbs Ψ and Δ. Transmittance is the mirror image: it's directly sensitive to weak absorption accumulated over the full optical path, precisely where SE goes blind. Fit the two together and each technique covers the other's blind spot, so a single model determines all the parameters well across the whole spectrum, rather than half the parameters well and half poorly.

With the light budget closed and the index known, a regression fit separates two quantities that are otherwise entangled: an absorption coefficient and a scattering coefficient. The result is an absorption value no longer contaminated by scattering losses, and a quantitative scattering measurement in its own right.

A third parameter describes the direction of the scattering rather than its magnitude: the anisotropy parameter. It runs from −1 to +1, where positive values indicate light scattered predominantly forward, negative values predominantly backward, and zero an even split. It isn't measured directly — it emerges from the fit to the integrating-sphere data.

2) Breaking parameter degeneracy

In any single dataset, thickness and index trade off against each other. A slightly thicker film with a slightly lower index can mimic a slightly thinner film with a slightly higher index, at least well enough to fool one measurement. But SE and spectrophotometry are physically different measurement types: reflection-mode polarization at oblique angles versus absolute intensity at normal incidence. A wrong-but-convenient thickness/index pair that happens to satisfy one measurement has no reason to also satisfy the other. Requiring the model to fit both at once collapses the family of "good enough" fits down toward the one combination that's actually physically correct.

3) Absolute scale paired with a built-in consistency check

SE measures a ratio (Ψ and Δ) while transmittance measures an absolute quantity, a real percentage of light transmitted. Because those are fundamentally different kinds of quantities, requiring one model to reproduce both simultaneously is a genuine test of that model rather than just another curve to fit. If the model matches Ψ and Δ well but predicts the wrong % transmittance, that's a signal that something about the model — the layer structure, the dispersion relation, an assumption about the interfaces — is wrong. Single-technique fitting has no equivalent check available; a good fit to one ratio-based dataset says nothing about whether the underlying physical picture is actually correct.

4) Accounting for light that goes missing

If a purely absorbing model — one built only from the k values that fit SE — predicts a certain transmittance, and the measured transmittance comes in lower than that prediction, the shortfall can't be absorption; absorption is already accounted for in the model. The deficit points instead to a loss mechanism the model doesn't yet include: scattering from surface roughness or bulk haze. Telling genuine absorption apart from scattering loss requires both datasets, because SE alone can't distinguish the two and transmittance alone can't isolate which one is responsible for a given loss.

Example From Literature

In a comprehensive paper by James Hilfiker from J.A. Woolam et al, the concept of improving optical models by pairing ellipsometry data with spectrophotometry values is laid out beautifully[2]. Hilfiker explains that a "uniqueness test" can be applied to a fitted parameter to judge whether or not the parameter is well determined by the data. The mean squared error (MSE) can be plotted against a fit parameter, for example, the thickness of a film. In figure 9, the red-highlighted "Uniqueness Range" marks how far you can push thickness away from the best-fit value before MSE degrades by 10%. It's found by taking the two thickness values (one above, one below the minimum) where the curve crosses 1.10, and reporting the interval between them — roughly 55.7 to 57 nm in this example.

If the uniqueness range is narrow, thickness is tightly constrained by the measurement — the optical constants can't compensate for an incorrect thickness, so the fit is unique. If the curve were shallow and flat instead of a sharp parabola, a wide range of thicknesses would all fit almost equally well, meaning thickness is correlated with the optical constants and the solution isn't uniquely determined — a real risk for thin absorbing films where thickness, n, and k trade off against each other.

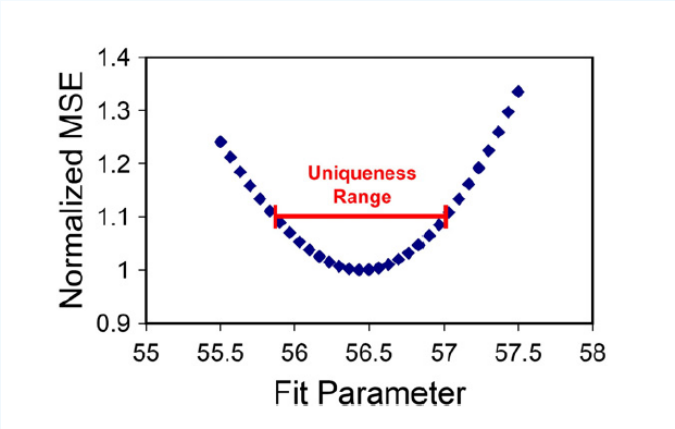


Figure 12: Illustration of the Uniqueness Test, a method for assessing whether a fitted model parameter is genuinely determined by the measurement rather than an artifact of the regression. Any parameter in the model can be tested this way; in the example shown, the tested parameter is film thickness. The parameter under test is fixed at a series of values spanning a chosen range; at each fixed value, every other parameter in the model is re-optimized to find the best possible match to the experimental data, and the resulting mean squared error (MSE) is recorded. MSE values are normalized to the global minimum (best overall fit = 1.0) and plotted against the tested parameter, producing the curve shown. A sharp, well-defined minimum indicates the parameter is tightly constrained by the data — no other value can be compensated for by adjusting the remaining parameters without noticeably worsening the fit. The Uniqueness Range (red) marks the interval between the two parameter values, above and below the minimum, where normalized MSE first exceeds 1.10 (10% away from the best fit); a narrower range reflects greater confidence in the reported value.
Adapted from Hilfiker et al.[2]

Hilfiker demonstrates the uniqueness test using SE on its own, which resulted in a broad, flat parabola with many possible solutions for thickness. However, when paired with either transmittance or reflectance intensity measurements, the parabola becomes much sharper with a clear minimum error.

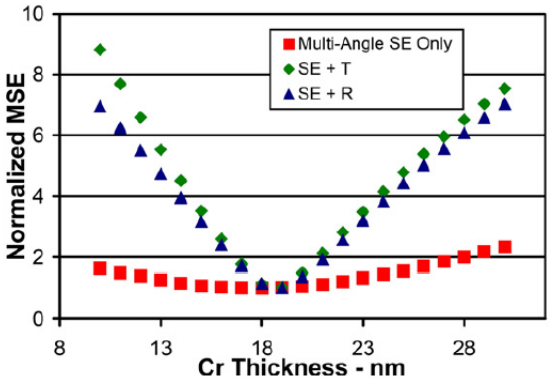


Figure 13: Application of the Uniqueness Test to a case where the measured data alone fails to uniquely determine film thickness. Multi-angle spectroscopic ellipsometry (SE) data from a thin absorbing film was analyzed as described above; rather than the sharp minimum obtained when the test is well constrained, the resulting normalized MSE curve from SE only is broad and shallow across a wide thickness range, indicating that no single thickness can be distinguished from the SE data alone — many thickness values, each paired with a correspondingly adjusted set of optical constants, fit the data almost equally well. Supplementing the SE measurement with a second, independent data type — transmitted intensity or reflected intensity, each collected simultaneously and regressed jointly with the SE data — sharpens the MSE minimum considerably and narrows the Uniqueness Range, showing that the added measurement resolves the thickness/optical-constant correlation that SE alone could not.
Adapted from Hilfiker et al.[2]

Case 1: What is hiding inside your phone's touch screen?

The functionality of the touch screen on your phone comes down to a thin film of Indium Tin Oxide (ITO) coated on glass. ITO is a fascinating material because it's both highly electrically conductive and transparent to visible light. Most conductors (copper, aluminum, gold) are opaque, even shiny! Most transparent materials (glass, plastics) are electrical insulators. ITO does both jobs at once. That unique combination makes it indispensable wherever electronics and optics need to coexist. The conductivity of that ITO layer lets the screen sense the tiny change in capacitance created by your fingertip, pinpointing exactly where you touched, while its transparency lets you see the display's image passing through it, undisturbed, from underneath.



Figure 14: Image of a touch screen on a phone. ITO plays a crucial role in making the touch screen both transparent and conductive. SE and spectrophotometry can be used to validate engineering decisions for ITO films.

The sample we used for this study was a thin film on ITO deposited directly onto fused silica, a clean, well-behaved reference substrate for validating the optical properties of the ITO film. The bare substrate was characterized first and independently (fit MSE = 5.573), so the substrate's optical constants are known rather than assumed when the ITO film model is built on top of it. Getting the substrate right is a prerequisite for getting the film right, a step that's easy to skip and quietly corrupts results.

We collected results on four different tools:

1. An ellipsometer with a light source that sweeps through the deep UV through most of the near IR range. (JA Woollam RC2-DI: 193–1690 nm)
2. An ellipsometer with a light source that sweeps from the near IR to the far IR. (JA Woollam IR-VASE Mark II: 1.7–30 µm)
3. A spectrophotometer with a light source that sweeps through UV through the complete NIR band. (Perkin Elmer Lambda 1050: 190–2500 nm)
4. A spectrophotometer with a light source that sweeps through the mid IR, reaching into the far IR at its upper end. (ThermoFisher Nicolet 6700 FTIR: 2.5–25 µm)

The light spectrum range of each tool is slightly different, and one picks up where the last left off. Using the experimental measurements from all four tools together, we created a model for Ψ and Δ at multiple angles of incidence (55°, 65°, 75°) as well transmittance (T) from incident light across all wavelengths from UV to IR.

All four datasets were regressed simultaneously against a single substrate model, yielding a combined fit of MSE = 5.573. Characterizing the bare substrate independently, before any ITO film is added, fixes its optical constants so they can be carried forward as known values, rather than free parameters, when fitting the coated sample.

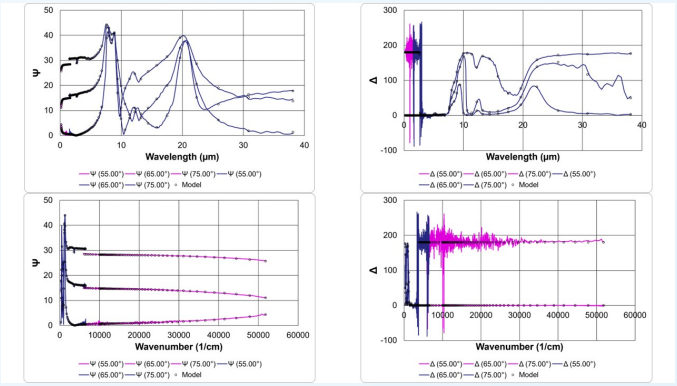


Figure 15: Model fits for Ψ and Δ on the bare fused silica substrate, measured across two variable-angle spectroscopic ellipsometry instruments covering deep UV through near IR (magenta) and near IR through far IR (blue). The graphs in the top row show Ψ and Δ as a function of wavelength, the graphs on the bottom row show Ψ and Δ as a function of wavenumber.

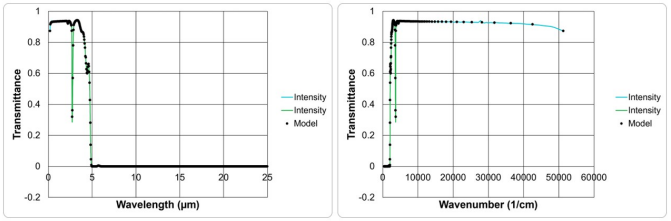


Figure 16: Model fits for transmittance intensity on the bare fused silica substrate from two spectrophotometers covering UV through near IR and mid IR (blue) through far IR (green). The graph on the left shows transmittance as a function of wavelength and the graph on the right shows transmittance as a function of wave number.

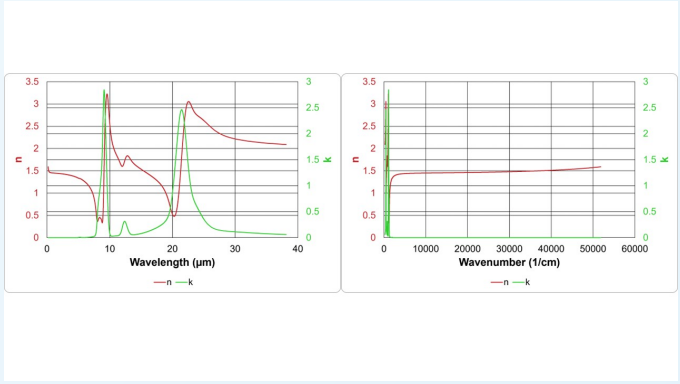


Figure 17: Optical constants (n, k) of the bare fused silica substrate, derived from the combined variable-angle SE and transmittance fit shown in the preceding slide. Across the visible and near-IR, n is smooth and near 1.45, with negligible absorption. This is consistent with published values for high-purity fused silica. Two distinct absorption features appear further into the IR, near 9 μm and 21–22 μm, corresponding to the well-known Si–O–Si bending and stretching vibrational bands of fused silica. Each shows up not just as a dip in absorption, but with the smooth, correspondingly-shaped shift in the refractive index around it that a genuine absorption band should produce — rather than appearing as an isolated bump with no matching effect on n, which would suggest noise or a fitting artifact. That correspondence serves as a validation check that the model is describing real material behavior. These substrate constants are fixed and carried forward, unmodified, into the ITO-coated sample's model.

With known properties of the bare substrate in hand, we then performed the same routine to model the ITO film deposited on the silica substrate.

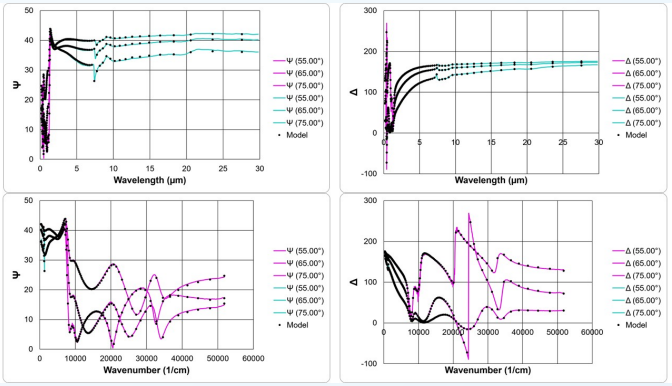


Figure 18: Model fits for Ψ and Δ from ITO thin film deposited on fused silica, measured across two variable-angle spectroscopic ellipsometry instruments covering deep UV through near IR (magenta) and near IR through far IR (blue). The models in the top row show Ψ and Δ as a function of wavelength, the models on the bottom row show Ψ and Δ as a function of wave number.

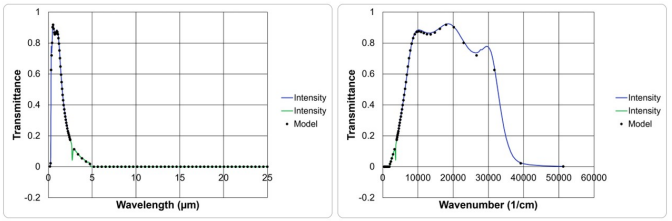


Figure 19: Model fits for transmittance intensity on ITO thin film deposited on fused silica from two spectrophotometers covering UV through near IR and mid IR (blue) through far IR (green). The graph on the left shows transmittance as a function of wavelength and the graph on the right shows transmittance as a function of wavenumber.

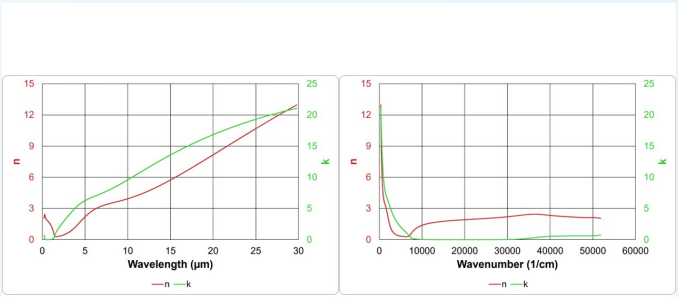


Figure 20: Optical constants (n, k) of the ITO thin film, derived from the combined variable-angle SE and transmittance fit on fused silica (MSE = 7.982), with the substrate's own optical constants fixed from the independent reference measurement rather than treated as free parameters. Across the visible range, both n and k stay low, reflecting the transparency window that allows the film to sit over a display without degrading the image beneath it. Moving into the mid-to-far IR, both rise substantially and k rises from near zero to above 20 at a wavelength of 30 μm. This is due to the free-carrier absorption from the same doped charge carriers that are responsible for the film's electrical conductivity.

An interpretation to this data that needs to be drawn out is that one continuous model spans two behaviors that look, at first glance, like they belong to different materials. Through the visible range, both n and k stay modest — ITO barely absorbs light there, which is the entire reason it can sit directly over a display without degrading what you see through it. That transparency window is what makes ITO usable as a touchscreen material in the first place. But follow the same curve into the mid-to-far infrared and both n and k climb steeply, k in particular rising from near zero to over a dozen by 30 µm. That rise isn't a separate phenomenon or a modeling quirk — it's free-carrier absorption, the optical signature of the same doped charge carriers responsible for ITO's electrical conductivity. Those carriers are essentially transparent to visible light but increasingly absorptive as wavelength grows, so the property that makes ITO conductive enough to sense a touch is the same property that eventually announces itself in the infrared. Capturing both regimes in a single, physically consistent model — rather than fitting the visible and infrared behavior as two disconnected problems — is what lets this transition read as one continuous optical response, and it's a direct, physically grounded confirmation that the deposited film is behaving exactly the way a functioning transparent conductor should.

Case 2: How does the facial recognition camera ksee only what it needs to?

Many phones now unlock with a look instead of a touch, using a dedicated infrared camera that projects and reads back a pattern of IR light on your face.



Figure 21: Facial recognition camera on a phone.

That camera only works if it can tell its own illumination apart from everything else in the scene from light sources including sunlight or indoor lighting. The component that makes that separation possible is a bandpass filter: a stack of thin films placed over the IR camera's sensor that transmits only a narrow slice of the spectrum matching the illuminator's wavelength and rejects essentially everything else. Where the touch screen's ITO film must do two conflicting jobs at once in a single layer, this filter does its job by stacking many layers, each an incremental step toward one very specific optical outcome.

The filter measured here is a 28-layer stack of alternating amorphous silicon (a-Si) and silicon dioxide (SiO_2) deposited on a glass substrate, capped with about 1.7 nm of surface roughness. a-Si is the high-index layer (n ≈ 3.7–3.9 across this wavelength range) and SiO_2 is the low-index layer (n ≈ 1.45–1.57); pairing a very high-index material with a very low-index one is what lets a comparatively compact stack (the entire stack of 28 layers is 4 µm total!) do the reflective work that would otherwise take many more layers. Thickness is not uniform through the stack: comparatively thin layers (22–65 nm and 124–162 nm, alternating) make up roughly the outer half nearest the surface, thickening sharply in the interior to 330–370 nm pairs, then thinning again to 156 nm and 55 nm right at the substrate. That thickness pattern isn't arbitrary; grading the layers this way, rather than repeating a fixed pair, is what gives the filter a clean, ripple-free passband with sharp edges.

Layer	Material	Thickness (nm)	Layer	Material	Thickness (nm)
Roughness	—	1.68	14	a-Si	61.77
28	a-Si	22.34	13	SiO2	147.62
27	SiO2	141.04	12	a-Si	28.93
26	a-Si	61.71	11	SiO2	161.24
25	SiO2	153.87	10	a-Si	128.32
24	a-Si	51.91	9	SiO2	364.04
23	SiO2	132.47	8	a-Si	133.69
22	a-Si	57.14	7	SiO2	331.12
21	SiO2	154.42	6	a-Si	127.17
20	a-Si	62.26	5	SiO2	369.33
19	SiO2	146.89	4	a-Si	130.64
18	a-Si	55.43	3	SiO2	351.60
17	SiO2	124.09	2	a-Si	156.50
16	a-Si	63.49	1	SiO2	55.11
15	SiO2	161.55			

Table 1: SE-modeled layer structure of the short-wave infrared (SWIR) bandpass filter, from the outermost layer (28, nearest the surface) to the innermost (1, at the substrate). Total stack thickness: 3937.4 nm.

We measured this filter the same way as the touch screen film earlier in this document: multi-angle spectroscopic ellipsometry combined with spectrophotometry, regressed together against one optical model rather than treated as separate problems.

1. Variable-angle ellipsometry at seven angles of incidence (45°, 50°, 55°, 60°, 65°, 70°, and 75°) on the J.A. Woollam RC2-DI (193–1690 nm), recording Ψ and Δ at each angle.
2. Transmittance intensity from 190–2500 nm on the Perkin Elmer Lambda 1050.

All seven angles and the transmittance data were fit simultaneously to a single 28-layer optical model, converging to a combined MSE of 16.754. That is higher than the single- and few-layer fits earlier in this case study, which is expected since over 28 individual layers, even minute mismatch between model and reality on each layer will compound. The value of MSE is highly sample-specific and a somewhat larger MSE here still represents a tightly constrained result given the complexity of the model.

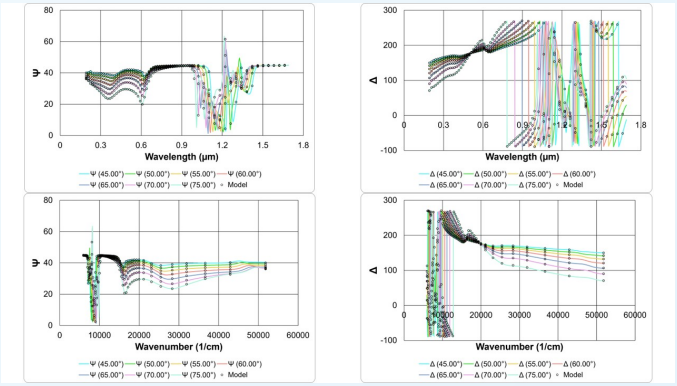


Figure 22: Model fits for Ψ and Δ on the SWIR bandpass filter, measured at seven angles of incidence (45°–75°) on the J.A. Woollam RC2-DI across 193–1690 nm.

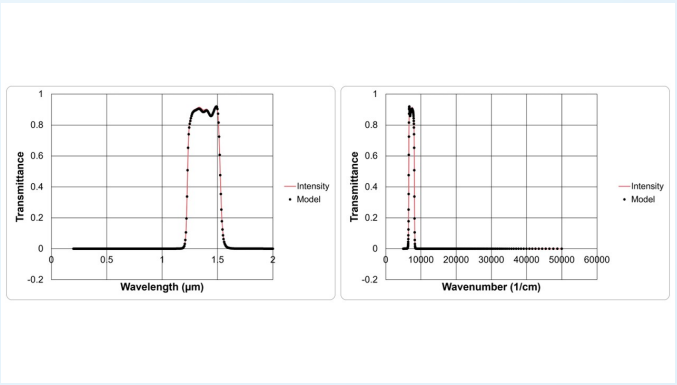


Figure 23: Model fit for transmittance intensity on the SWIR bandpass filter from 190–2500 nm, showing the measured and modeled passband.

The result is a clean, single bandpass. Transmittance stays under 0.1% through the visible range and short-wave near-IR out to about 1150 nm, then rises sharply to a broad passband centered near 1380 nm and roughly 300 nm wide at half-maximum, peaking at 91.5% transmittance at 1335 nm, before falling back under 1% by 1600 nm and staying low out past 2000 nm. That is exactly the behavior a dedicated IR camera needs: everything the sensor sees outside its illuminator's band — sunlight, indoor lighting, the visible scene itself — is suppressed by roughly three orders of magnitude, while the wavelength the camera is actually built to use gets through with less than a tenth of its light lost.

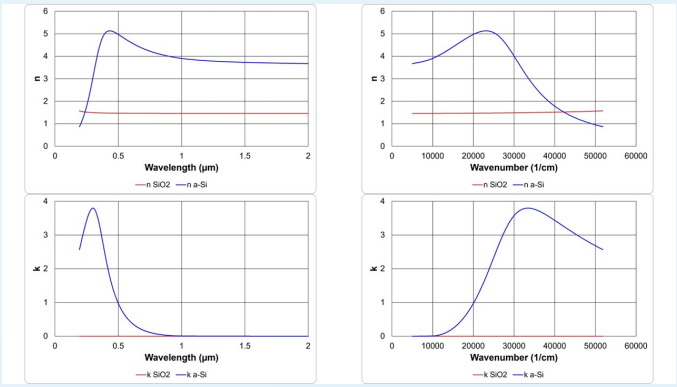


Figure 24: Optical constants (n, k) of the a-Si and SiO_2 films, derived from the combined SE and transmittance fit.

Across this range, SiO_2 behaves close to its handbook value ($n \approx 1.45–1.57$, $k \approx 0$) — an inert, well-understood low-index spacer. a-Si is the film worth watching: its refractive index is not a fixed material constant but a direct product of how the film was deposited, and it is also the single largest lever on where the passband sits. Simulations run across a range of plausible a-Si dispersion curves show that a 1% change in a-Si's index shifts the passband center by roughly 7.3 nm — on a filter whose entire job is admitting one narrow band and rejecting the rest, that is the difference between a camera that works and one that does not. Because a-Si's index is deposition-dependent rather than a fixed material property, and because SE is essentially the only technique sensitive enough to measure it precisely on the finished film, SE becomes not just a characterization tool here but a natural process-control checkpoint for this kind of filter.

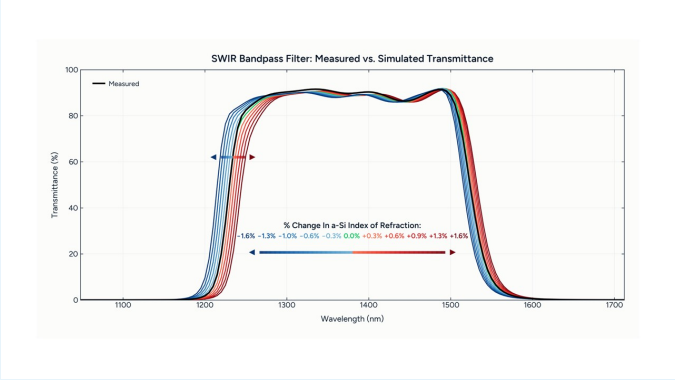


Figure 25: Simulated shift in the SWIR filter's transmittance passband as a function of a-Si refractive index, showing the passband's sensitivity to a deposition-dependent film property.

Checking the optical model against the physical device

Optical modeling normally must make its case entirely on optical evidence, cross-checking ellipsometry against spectrophotometry the way this case study has done throughout. For this filter, we had the opportunity to go one step further with a completely independent, non-optical measurement: a focused-ion-beam (FIB) cross-section of the actual device, imaged by transmission electron microscopy (TEM) and mapped by energy-dispersive X-ray spectroscopy (EDS).

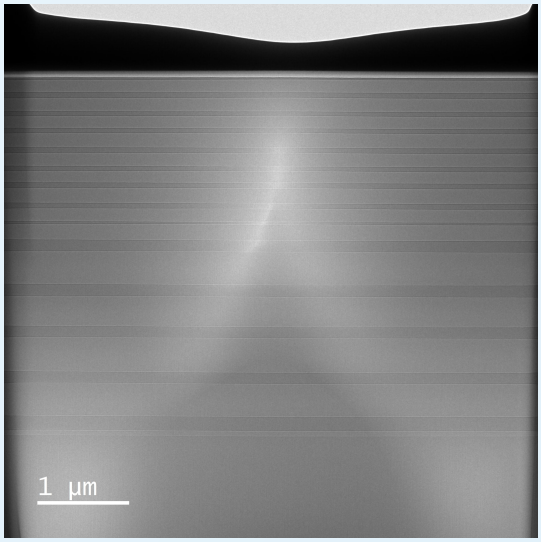


Figure 26: Low-magnification TEM cross-section (7000X) of the SWIR bandpass filter, showing the alternating a-Si/SiO_2 multilayer stack.

Higher-magnification images with layer boundaries measured directly on the micrograph let those thicknesses be compared, layer by layer, against the SE-derived model above.

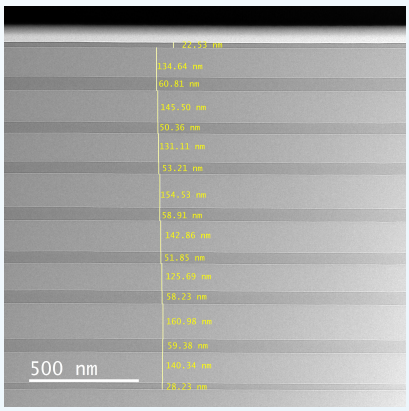


Figure 27: TEM cross-section (22,500X) with measured thicknesses annotated for the outermost 17 layers of the stack (layers 12–28, nearest the surface).

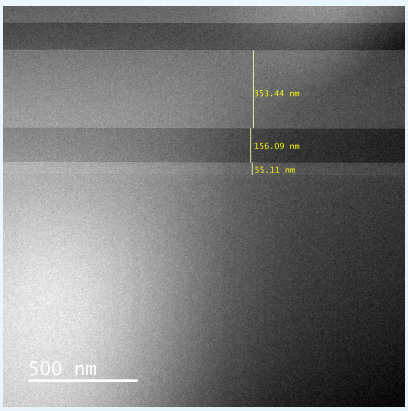


Figure 28: TEM cross-section (22,500X) at the substrate interface, with measured thicknesses annotated for the innermost 3 layers (layers 1–3).

Layer	Material	SE model (nm)	TEM measured (nm)	Difference
28	a-Si	22.34	22.53	0.9%
27	SiO2	141.04	134.64	4.5%
26	a-Si	61.71	60.81	1.5%
25	SiO2	153.87	145.50	5.4%
24	a-Si	51.91	50.36	3.0%
23	SiO2	132.47	131.11	1.0%
22	a-Si	57.14	53.21	6.9%
21	SiO2	154.42	154.53	0.1%
20	a-Si	62.26	58.91	5.4%
19	SiO2	146.89	142.86	2.7%
18	a-Si	55.43	51.85	6.5%
17	SiO2	124.09	125.69	1.3%
16	a-Si	63.49	58.23	8.3%
15	SiO2	161.55	160.98	0.4%
14	a-Si	61.77	59.38	3.9%
13	SiO2	147.62	140.34	4.9%
12	a-Si	28.93	28.23	2.4%
1	SiO2	55.11	55.11	0.0%
2	a-Si	156.50	156.09	0.3%
3	SiO2	351.60	353.44	0.5%
Mean / median / max diff				**3.0% / 2.7% / 8.3%**

Table 2: SE-modeled vs. TEM-measured thickness for 20 of the filter's 28 layers (71% of the stack), spanning the outermost 17 and innermost 3 layers. Mean absolute difference: 3.0%.

Across these 20 layers, the TEM measurements land within an average of 3.0% of the SE-derived thicknesses (median 2.7%, worst case 8.3%, several layers agreeing to within a few tenths of a percent). That is a strong result: TEM thickness measurement carries its own uncertainty from image calibration and boundary placement, typically on the order of a few percent, so agreement at this level means the optical model is not just internally self-consistent, it is describing the physical device that was actually built.



Figure 29: STEM-EDS elemental map of the same cross-section (Si: blue, O: yellow, C: magenta, W: red), confirming the alternating a-Si/SiO_2 layer identity. The tungsten/carbon band at the top is the FIB protective cap deposited during sample preparation, not a device layer.

This is the same argument this case study has been making about pairing spectrophotometry with ellipsometry, carried one step further. A model that only has to satisfy one kind of measurement can miss the mark in physical reality without anyone noticing. Here, the SE-and-transmittance model was checked against a measurement of a genuinely different kind, a physical image of the actual cross-section, not another optical fit, and it held up. For a filter whose entire function depends on 28 layers each landing within a few percent of its design thickness, that is a level of confirmation an optical model can never fully provide on its own.

Covalent's SE and Spectrophotometry Services

Most labs run a single measurement technique and report whatever numbers it produces. When ellipsometry alone can't uniquely determine thickness and optical constants for an absorbing film (as this case study has shown, repeatedly) that limitation typically goes unmentioned in the reported result. The differentiator here isn't a fancier instrument; it's fitting one physically consistent model to measurements from multiple instruments simultaneously and being able to explain exactly why that combined answer deserves more confidence than a single-technique number would. The model has been tested against data it wasn't built to fit in isolation, and the correlation that would otherwise leave thickness and optical constants ambiguous has been resolved rather than assumed away.

For the customer, that translates into two concrete advantages. First, optical constants that actually predict how the device behaves instead of curve-fit values that happen to match one dataset but fail the moment they're checked against an independent measurement of the same sample. Second, an honest accounting of scattering losses instead of absorption numbers quietly inflated by them: without validating the absorption model against a second, independent measurement first, scattering losses get silently folded into k, making a material look more absorbing than it really is.

This isn't a one-off research exercise — it's a standard, growing service line at Covalent, applicable anywhere a thin or multilayer film needs verified optical characterization: transparent conductors like ITO, optical coatings, and multilayer stacks such as bandpass filters. And it's backed by more than methodology — Covalent has developed custom software specifically for this class of joint, multi-instrument analysis, turning combined SE-and-intensity fitting (and the scattering measurement built on top of it) into a repeatable, deliverable service rather than a bespoke exercise rebuilt from scratch for every sample. Furthermore, we've partnered with J.A. Woolam, the company that manufactures our ellipsometers and wrote the book on optical modeling to bring our clients the best results possible.



"Covalent has been a great partner to the J.A. Woollam Company, helping make expert ellipsometry services accessible to the broader materials and optics community. The installation of an IR-VASE adds powerful infrared capability to their lab, expanding the spectral range and application space they can serve."

James Hilfiker · J.A. Woollam Company

Citations

Hilfiker, J. N., Singh, N., Tiwald, T., Convey, D., Smith, S. M., Baker, J. H., & Tompkins, H. G. (2008). Survey of methods to characterize thin absorbing films with Spectroscopic Ellipsometry. *Thin Solid Films*, 516(22), 7979–7989. https://doi.org/10.1016/j.tsf.2008.04.060

Covalent

CONTACT US

Looking for help with your Material Analysis needs?

Contact Us