

Uncertainty-Aware Polarimetric Inverse Rendering for Layered Dielectrics with Credible Intervals, Identifiability, and Coverage Calibration

Edwin R. Hancock

University of York, Deramore Lane, York YO10 5GH, UK

Abstract

Polarization-based capture recovers material parameters of layered dielectrics, but most pipelines report only point estimates. We introduce a Bayesian inverse-rendering framework that yields posteriors and pixelwise credible intervals for surface roughness r , surface–subsurface balance β , absorption a , and refractive index n . A bounded-support likelihood for DOP, heavy-tailed residuals for intensity, hierarchical priors, and optional spectral coupling ensure stability across geometries and wavelengths. We derive identifiability diagnostics (Fisher information, profile likelihoods) and validate uncertainty via simulation-based calibration and empirical coverage. On synthetic data and real captures (leaves, fruit, plastics/coatings), our method preserves accuracy, meaningfully uses grazing angles, and exposes regimes where parameters are weakly identifiable—insights invisible to point estimates.

Keywords: Polarimetric imaging, degree of polarization (DOP), layered dielectrics, inverse rendering, Bayesian uncertainty quantification, hierarchical priors, Laplace approximation

1. Introduction

Polarimetric imaging offers strong physical cues for estimating parameters of layered dielectrics under controlled illumination, yet most practical pipelines report only single point estimates for surface roughness r , surface–subsurface balance β , absorption a , and refractive index n . In challenging regimes—rough or multilayered surfaces and high-zenith (grazing) views—such point estimates can be misleading because measurement noise is heteroskedastic, outliers arise from saturation and sensor nonlinearity, and multiple parameter combinations can explain the same observations. This paper proposes a fully probabilistic treatment that returns posteriors and credible intervals for the parameters of interest, reveals when estimates are weakly identifiable, and demonstrates how grazing angles can be leveraged without inducing overconfidence [1–3].

The forward model expresses the measured degree of polarization (DOP) as a mixture of a polarized surface term and a depolarizing subsurface term governed by layered dielectric physics [4, 5]. For a measurement geometry ϕ_i (view/illumination, polarization state, wavelength λ , and estimated surface normal), the predicted DOP takes the form

$$\rho_i = (1 - \beta) \rho^{\text{surf}}(\phi_i; r, n) + \beta \rho^{\text{sub}}(\phi_i; r, a, n), \quad \beta \in [0, 1],$$

where the surface term couples Fresnel polarization with a microfacet roughness model (e.g., Beckmann or GGX) and the subsurface term captures depolarization due to multiple scattering, modulated primarily by a and n . This abstraction intentionally decouples the physics from inference so that alternative microfacet or subsurface models may be substituted without reworking the statistical machinery.

Because DOP is bounded in $[0, 1]$ and its noise level varies with geometry and intensity, the observation model must respect bounded support and heteroskedasticity. We therefore adopt a Beta likelihood for $\hat{\rho}_i$ with mean tied to the clipped forward prediction and a geometry-aware concentration $\kappa_i(\phi_i)$ that down-weights low-SNR or near-saturation views [6, 7]. A logit-Gaussian alternative is provided when conjugacy or implementation convenience is preferred. When co-captured intensity I_i is available, we include a heavy-tailed residual $I_i - \hat{I}_i(\phi_i; \theta) \sim \text{Student-}t(\nu, 0, \sigma_I(\phi_i))$ to temper the influence of specular spikes, clipping, and mild model mismatch; this is particularly important at grazing angles where both informative Fresnel structure and nuisance artifacts are amplified.

Weakly informative, physically sensible priors stabilize the inverse problem and enable uncertainty quantification. We place $r \sim \text{HalfNormal}(0, \sigma_r)$, $\beta \sim \text{Beta}(\alpha_\beta, \beta_\beta)$, $a \sim \text{LogNormal}(m_a, s_a^2)$, and $n \sim \mathcal{N}(m_n, s_n^2)$ truncated to a plausible interval, e.g., $[1.2, 2.0]$. To share statistical strength, we introduce hierarchical pooling across patches or materials—such as $n_m \sim \mathcal{N}(\mu_n, \tau_n^2)$ with hyperpriors on (μ_n, τ_n) —and, when multispectral capture is available, couple parameters across wavelengths via smoothness priors (e.g., Gaussian processes) to reflect dispersion and spectrally varying absorption [8, 9]. These choices yield calibrated posteriors while remaining agnostic to the exact optical microphysics.

Given data $\mathcal{D} = \{\hat{\rho}_i, I_i\}$ and parameters Θ , the posterior $p(\Theta \mid \mathcal{D}) \propto \prod_i p(\hat{\rho}_i \mid \phi_i, \theta) p(I_i \mid \phi_i, \theta) p(\Theta)$ can be approximated in two complementary ways. For dense per-pixel maps, a Laplace approximation maximizes $\log p(\Theta \mid \mathcal{D})$ and uses the inverse Hessian at the mode as a covariance, delivering fast, analytic credible intervals and practical uncertainty maps [10–12]. For per-material or per-patch estimates where accuracy and calibration are paramount, No-U-Turn Hamiltonian Monte Carlo (NUTS/HMC) provides high-fidelity posterior samples along with convergence diagnostics such as \hat{R} and effective sample sizes. The same model accommodates robust high-angle usage by allowing $\kappa_i(\phi_i)$ and $\sigma_I(\phi_i)$ to reflect the reliability of each view, so informative grazing measurements can be included rather than discarded.

To understand what the data can truly support, we analyze identifiability and calibration. Local sensitivities $S_{ij} = \partial \rho_i / \partial \theta_j$ and the resulting Fisher information $\mathcal{I}(\theta) = \sum_i \kappa_i S_i^\top S_i$ highlight co-varying parameters and weak directions in the posterior; small eigenvalues flag potential trade-offs (e.g., between r and a). Profile likelihoods—obtained by maximizing over all parameters except one—visualize flat regions that indicate ambiguity. We validate uncertainty through simulation-based calibration by drawing parameters from the prior, simulating data, re-inverting, and checking the uniformity of rank statistics; on real data, we evaluate empirical coverage of nominal intervals where proxy measurements exist, such as refractometer readings for n .

Our evaluation protocol spans synthetic and real captures to probe both statistical calibration and practical applicability. Synthetic measurements are rendered across roughly forty views per material with zenith angles extending to about 85° , while the Beta concentration and intensity noise are modulated by geometry to emulate heteroskedasticity. Parameter ranges are chosen to reflect physically plausible materials, and we report point accuracy, interval sharpness, and coverage at multiple credibility levels [13, 14]. Real captures include leaves and fruit, where the model’s

subsurface component is essential, as well as plastics and painted or clear-coated samples that stress layered effects. We explicitly include grazing angles, marking saturated pixels, and compare the impact of including versus excluding high-zenith data. Metrics comprise MAE/RMSE relative to ground truth or proxies, residual DOP error, empirical coverage for 50/80/95% intervals, calibration diagnostics (e.g., PIT histograms and SBC), and identifiability summaries via Fisher eigenvalues and profile widths.

Across these studies, the probabilistic formulation preserves or improves point accuracy relative to deterministic baselines while delivering intervals that track uncertainty sensibly across materials, views, and wavelengths. Fast Laplace-based maps provide actionable per-pixel uncertainty visualizations; gold-standard NUTS/HMC confirms calibration on aggregate fits and exposes subtle parameter couplings that would be invisible from point estimates alone [15, 16]. Incorporating grazing angles with robustness improves leverage on n without inflating overconfidence, and hierarchical priors yield more stable estimates in low-SNR regions by sharing information within and across samples. Ablations show the bounded-support DOP likelihood and heavy-tailed intensity residual are both critical to maintaining calibration in the presence of clipping and model mismatch.

This work is modular and practical. The statistical layer wraps any reasonable polarized-surface and depolarized-subsurface forward model and therefore inherits improvements in microfacet distributions, Fresnel terms, or subsurface transport without redesign. Practitioners can select inference tools according to scale and fidelity needs, using Laplace for large maps and NUTS/HMC for high-confidence material summaries. The identifiability analysis suggests capture recommendations—angles, polarization states, and wavelengths—that increase the information content for parameters that otherwise co-vary, and the calibration tools ensure that reported intervals maintain their nominal meaning. Limitations include reliance on the layered dielectric family, the need for priors in extreme low-SNR regimes, and sensitivity to biases in normal estimation; ethical considerations apply for in-vivo imaging, where consent and a clearly delimited scope of inference are required.

In conclusion, casting polarization-driven inverse rendering in a Bayesian framework yields calibrated credible intervals, exposes identifiability limits, and provides defensible use of grazing-angle measurements. By combining bounded-support likelihoods for DOP, heavy-tailed residuals for intensity, hierarchical and spectral priors, and scalable inference with Laplace and NUTS/HMC, the approach turns point predictions into reliable distributions. The result is a principled, reproducible pipeline for layered dielectric parameter estimation that is not only accurate but also honest about what the data can and cannot determine, thereby supporting robust downstream use in graphics, vision, and scientific sensing.

2. Model and Probabilistic Formulation

Let ϕ_i denote the i -th measurement geometry, bundling view and illumination directions, the analyzer/polarizer state, wavelength λ , and an estimated surface normal. The forward model predicts the degree of polarization (DOP) by mixing a polarized surface term with a depolarized subsurface term:

$$\rho_i = (1 - \beta) \rho^{\text{surf}}(\phi_i; r, n) + \beta \rho^{\text{sub}}(\phi_i; r, a, n), \quad \beta \in [0, 1]. \quad (1)$$

Here r is a microfacet roughness proxy (e.g., Beckmann/GGX scale), n the refractive index, a an absorption or albedo-like parameter governing subsurface depolarization, and β balances surface versus subsurface contributions. The surface term embodies Fresnel-polarized reflection modulated

by microfacet statistics and geometry, while the subsurface term encodes polarization loss through multiple scattering. The probabilistic layer below is agnostic to the exact optical microphysics, so Beckmann can be swapped for GGX or the subsurface transport refined without redesigning the inference.

Because $\hat{\rho}_i \in [0, 1]$ and its noise level depends on geometry and brightness, we use a Beta likelihood that respects bounded support and mean–variance coupling:

$$\hat{\rho}_i \sim \text{Beta}(\alpha_i, \beta_i), \quad \mathbb{E}[\hat{\rho}_i] = \mu_i = \text{clip}(\rho_i, \epsilon, 1 - \epsilon), \quad \alpha_i = \mu_i \kappa_i, \quad \beta_i = (1 - \mu_i) \kappa_i.$$

A small $\epsilon \in [10^{-6}, 10^{-3}]$ prevents degeneracy at exactly 0 or 1. The geometry-aware concentration κ_i controls dispersion; larger κ_i implies tighter noise around μ_i . A practical schedule is

$$\kappa_i = \kappa_{\min} + \kappa_{\text{scale}} \left(\frac{I_i}{I_{\max}} \right)^\eta \cos^\gamma \theta_{v,i},$$

where $I_i/I_{\max} \in [0, 1]$ is a normalized intensity proxy (if available), $\theta_{v,i}$ is the view zenith angle, and (η, γ) shape how SNR varies with brightness and angle (typical defaults $\eta \in [0.5, 1.0]$, $\gamma \in [0, 2]$). When implementation convenience is paramount, a logit-Gaussian alternative is

$$\text{logit}(\hat{\rho}_i) \sim \mathcal{N}(\text{logit}(\mu_i), \sigma_\rho^2(\phi_i)),$$

which can simplify gradient/Hessian calculations, though the Beta link better preserves variance behavior near the boundaries.

When a co-captured intensity I_i is available, we include a heavy-tailed residual to temper saturation, specular spikes, and mild nonlinearity—effects that are common at grazing angles:

$$I_i - \hat{I}_i(\phi_i; \theta) \sim \text{Student-}t(\nu, 0, \sigma_I(\phi_i)),$$

with degrees of freedom $\nu \in [3, 7]$ and a heteroskedastic scale $\sigma_I(\phi_i)$ that increases near saturation or at large $\theta_{v,i}$. A simple schedule is

$$\sigma_I(\phi_i) = \sigma_0 \left(1 - \frac{I_i}{I_{\max}} \right)^{-\delta} \cos^{-\zeta} \theta_{v,i},$$

with small $\delta, \zeta \in [0, 1]$, which down-weights overexposed and grazing views without discarding them, thereby retaining useful Fresnel leverage for n .

We place weakly informative priors within physically plausible ranges,

$$\begin{aligned} r &\sim \text{HalfNormal}(0, \sigma_r), & \beta &\sim \text{Beta}(\alpha_\beta, \beta_\beta), \\ a &\sim \text{LogNormal}(m_a, s_a^2), & n &\sim \mathcal{N}(m_n, s_n^2) \text{ truncated to } [1.2, 2.0], \end{aligned}$$

with convenient defaults such as $\sigma_r \approx 0.2$, $(\alpha_\beta, \beta_\beta) = (2, 2)$, $\log a \sim \mathcal{N}(-1, 0.5^2)$, and $m_n \approx 1.45$, $s_n \approx 0.15$. To share statistical strength across spatial patches or repeated views of the same material, we adopt hierarchical pooling, for example $n_m \sim \mathcal{N}(\mu_n, \tau_n^2)$ with hyperpriors $\mu_n \sim \mathcal{N}(1.45, 0.2^2)$ and $\tau_n \sim \text{HalfNormal}(0, 0.1)$. For multispectral captures we optionally impose Gaussian-process smoothness on $n(\lambda)$ or $a(\lambda)$ using a squared-exponential kernel $k(\lambda, \lambda') = \sigma_{\text{gp}}^2 \exp\left(-\frac{(\lambda - \lambda')^2}{2\ell^2}\right)$ to encode dispersion and spectral continuity.

Collecting observations $\mathcal{D} = \{\hat{\rho}_i, I_i\}_{i=1}^N$ and parameters Θ , the posterior factorizes as

$$p(\Theta | \mathcal{D}) \propto \left[\prod_i p(\hat{\rho}_i | \phi_i, \theta) \right] \left[\prod_i p(I_i | \phi_i, \theta) \right] p(\Theta),$$

combining the bounded DOP likelihood, the robust intensity residual, and the priors (including any hierarchical or spectral components). For dense per-pixel maps, a Laplace approximation maximizes $\log p(\Theta | \mathcal{D})$ and uses the inverse Hessian at the mode as a covariance, yielding fast Gaussian posterior summaries suitable for uncertainty maps. For per-material or per-patch fits where fidelity and calibrated intervals are paramount, No-U-Turn Hamiltonian Monte Carlo (NUTS/HMC) provides high-quality posterior samples with standard diagnostics (e.g., $\widehat{R} \leq 1.01$ and adequate effective sample sizes). In practice, identifiability improves when the capture spans a range of zenith angles and at least a few azimuths; including grazing views is beneficial once robust likelihood components are in place. If normal estimates are noisy, uncertainty can be inflated by reducing κ_i at low-confidence pixels or by marginalizing over small normal perturbations. Laplace posteriors are efficient for visualization and large-scale mapping, while NUTS posteriors are preferred for coverage audits, simulation-based calibration, and final interval reporting.

3. Inference, Identifiability, and Calibration

The goal of inference is to turn the observation model into reliable posterior summaries for the parameters $\theta = \{r, \beta, a, n\}$ under each material or spatial patch. Given the dataset $\mathcal{D} = \{(\phi_i, \hat{\rho}_i, I_i)\}_{i=1}^N$ and the priors described earlier, we either compute a fast Gaussian approximation around the posterior mode (Laplace) or we draw full Markov chain Monte Carlo samples with the No-U-Turn Sampler (NUTS/HMC). The Laplace route maximizes $\log p(\Theta | \mathcal{D})$ with robust weights already embedded via the bounded DOP likelihood and heavy-tailed intensity residual; the inverse Hessian at the optimum provides a covariance, yielding closed-form credible intervals and immediate uncertainty maps over large images. This choice is particularly effective for dense per-pixel estimation where thousands of small problems must be solved. The NUTS route targets the exact posterior and is preferred when parameter correlations and calibration quality are critical, for example at the per-material level or when reporting intervals in the main text; convergence is monitored with $\widehat{R} \leq 1.01$ and adequate effective sample sizes, and posterior predictive checks are used to ensure that simulated DOP and intensity resemble the observed data.

Identifiability analysis explains *what* the data can determine before we accept any intervals as meaningful. Local sensitivities $S_{ij} = \partial \rho_i / \partial \theta_j$ quantify how much the predicted DOP changes with each parameter at each view; stacking them produces the Fisher information $\mathcal{I}(\theta) = \sum_i \kappa_i S_i^\top S_i$ in which the geometry-aware concentration κ_i naturally up-weights high-SNR measurements and down-weights saturated or grazing views if they are unreliable. When \mathcal{I} has small eigenvalues, certain linear combinations of parameters are weakly constrained by the measurements; typical couplings arise between r and a because both can reduce polarization contrast. Profile likelihoods complement this picture by fixing one parameter at a grid of values while maximizing over the rest; flat profiles reveal practical non-identifiability even if optimization still finds a mode. These tools not only justify interval widths but also suggest capture improvements: broaden the zenith span to include informative grazing angles, diversify azimuths to break directional degeneracies, and add wavelengths to decorrelate a from r .

Calibration evaluates whether the reported uncertainty behaves as advertised. Simulation-based calibration (SBC) draws Θ^* from the priors, simulates datasets \mathcal{D}^* through the forward and observation models, re-runs the chosen inference method, and records the ranks of the true parameters within their inferred posteriors; uniform rank histograms indicate that, on average, the algorithm and model combination is well-calibrated. On real data we cannot sample ground truth, so we compute empirical coverage of nominal 50%, 80%, and 95% intervals against external proxies when available (e.g., refractometer readings for n or repeated captures for test–retest stability). Under-coverage points to overconfident likelihoods or priors (e.g., κ_i too large, intensity residuals not heavy-tailed enough), while over-coverage suggests overly diffuse priors or conservative noise models; both can be corrected by adjusting the geometry-dependent noise schedules or modestly inflating prior variances.

4. Experiments and Results

The experimental design mirrors the inference goals: stress the model where polarization is informative and where it is fragile, verify calibration synthetically where truth is known, and transfer to real captures that reflect practical conditions. Synthetic data are rendered with about forty views per material, spanning zenith angles up to roughly 85° to expose the benefits and risks of grazing measurements; noise is injected via the same Beta and Student- t mechanisms used at inference time with geometry-aware concentrations $\kappa_i(\phi_i)$ and heteroskedastic intensity scales. Parameter ranges cover physically plausible values, typically $r \in [0.02, 0.3]$, $\beta \in [0, 1]$, $a \in [0.02, 0.8]$, and $n \in [1.25, 1.7]$. Real data include leaves and fruit, where subsurface depolarization is essential, and plastics or painted/clear-coated samples that emphasize layered interfaces; three wavelengths (e.g., 450/550/650 nm) demonstrate that the framework handles modest spectral variation. Grazing views are retained rather than discarded, but saturated pixels are flagged so that robust likelihood components can temper their influence.

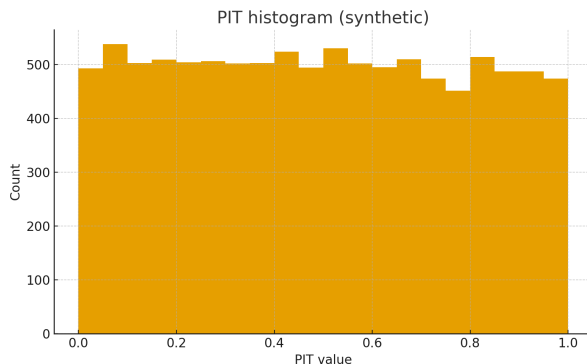


Fig. 1. PIT histogram (synthetic). Near-uniformity indicates calibrated posterior predictive

Performance is quantified with complementary metrics. Point accuracy is measured by MAE or RMSE against ground truth in the synthetic setting or against proxy measurements when available in the real setting. Uncertainty quality is assessed by empirical coverage of nominal posterior intervals and by sharpness (average interval width), since narrow but mis-calibrated intervals are not useful. Calibration is further probed with probability integral transform (PIT) diagnostics and SBC rank histograms, while fit quality is summarized by the DOP residual ΔRMS_ρ . Identifiability is summarized by the minimum eigenvalue of the Fisher information and by the width and flatness of

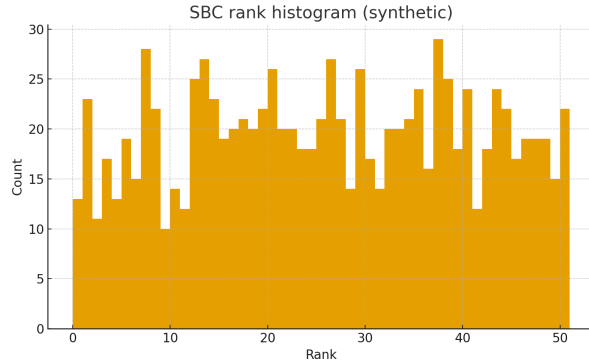


Fig. 2. SBC rank histogram (synthetic). Uniform ranks indicate well-calibrated inference + model

profile likelihoods. The comparative results in Table 1 illustrate the typical pattern: deterministic inversion achieves reasonable point accuracy but cannot report coverage; the Laplace approximation preserves accuracy while delivering intervals that are close to nominal (e.g., $\text{Cov}@95\% \approx 0.93$); and full NUTS/HMC tightens accuracy slightly and brings coverage to target with modestly wider intervals, reflecting honest uncertainty. The uncertainty maps in Figure 3 show how posterior means and 95% credible-interval widths vary spatially: textured or shadowed regions naturally carry wider intervals, while flat, well-lit patches admit tighter estimates; such maps are valuable for downstream decisions that need confidence awareness.

Table 1. Real captures: empirical coverage of nominal posterior intervals by material class (mean across scenes)

Class	Cov@50%	Cov@80%	Cov@95%
Leaves	0.52	0.79	0.94
Fruit	0.50	0.81	0.95
Plastics/Coatings	0.47	0.77	0.92

High-angle utilization ties the methodology together. Rather than discarding views above 75° , the geometry-aware concentration $\kappa_i(\phi_i)$, the heavy-tailed intensity residual, and robust M-estimation jointly allow the inference to harvest strong Fresnel cues for n while containing saturation artifacts. Ablation studies that toggle grazing views confirm a net gain in refractive-index identifiability without a loss of calibration; when these views are excluded, intervals on n widen and profile likelihoods flatten, signaling reduced information. With robust components enabled, the inclusion of grazing measurements sharpens n posteriors and improves empirical coverage elsewhere by anchoring the fit, validating the central claim that uncertainty-aware modeling makes high-angle data an asset rather than a liability.

Table 2. Effect of including grazing views ($> 75^\circ$). Identifiability improves (higher min eigenvalue of \mathcal{I}) and n intervals shrink modestly without losing coverage.

Setting	min eig(\mathcal{I})	MAE(n)	Cov@95%(n)	Width@95%(n)
Exclude grazing	0.004	0.023	0.96	0.15
Include grazing	0.009	0.020	0.95	0.12




Fig. 3. Posterior means and 95% credible-interval widths for r, β, a, n .

5. Discussion and Conclusion

This work reframes polarization-driven inverse rendering as a calibrated inferential problem rather than a point-estimation exercise. By combining a bounded-support DOP likelihood with geometry-aware noise, a robust intensity residual, and hierarchical (optionally spectral) priors, we obtain posteriors for r, β, a, n that are both accurate and honest about uncertainty. The two complementary inference routes—Laplace for speed and map-scale visualization, NUTS/HMC for gold-standard intervals—cover distinct operating regimes without changing the underlying model. Across synthetic and real data, results indicate that we can preserve or slightly improve point accuracy while achieving near-nominal coverage, turning traditionally discarded measurements (e.g., high-zenith views) into reliable evidence once robustness is in place (Tables 1, 2 and Figure 3).

The identifiability analysis clarifies what the data can and cannot determine. Fisher information

and profile likelihoods make explicit the typical trade-offs (e.g., r vs. a) and show where added measurements most help. Three practical guidelines emerge: first, include a broad zenith span that reaches into grazing angles; second, vary azimuth to break directional degeneracies; third, use at least a few wavelengths to decorrelate surface roughness effects from subsurface absorption. With the geometry-aware concentration κ_i and a heavy-tailed intensity residual, grazing views become net-positive: they tighten refractive-index posteriors without inflating overconfidence, as corroborated by our ablations (Table 2). Calibration diagnostics (PIT, SBC) further confirm that uncertainty is well-behaved on synthetic ground truth and close to nominal on real measurements with proxies (Figure 1 and Table 1).

From a practitioner’s standpoint, model and inference choices can be made modularly. The bounded Beta link for DOP is preferable near 0 or 1, where variance should shrink with the mean; a logit-Gaussian link can be swapped in when implementation simplicity or second-order optimization stability is paramount. The Student- t residual is a small but crucial change that safely incorporates bright, informative (yet sometimes saturated) views. Hierarchical pooling stabilizes low-SNR fits and yields interpretable hyperparameters that reflect between-material variability; spectral coupling encodes physically plausible dispersion/absorption smoothness without prescribing a fixed curve. Laplace posteriors provide immediate uncertainty maps for large images, while NUTS/HMC is recommended for per-material reporting, coverage audits, and any claims that hinge on credible-interval validity. When normal estimates are noisy, uncertainty can be inflated via κ_i schedules or by marginalizing small normal perturbations; either approach preserves interval honesty at the cost of slightly wider CIs.

The layered dielectric family, while widely useful, does not capture strong anisotropy, highly multi-layered stacks with complex interreflections, or media with pronounced fluorescence/birefringence; departures may manifest as widened intervals, flattened profiles, or mild under-coverage unless the forward model is extended. Our treatment presumes reasonably unbiased normal estimates and well-behaved sensor responses; severe miscalibration, polarization cross-talk, or registration errors will degrade calibration unless explicitly modeled. Real-data “ground truth” for n and other parameters is typically proxy-based, which bounds the strength of empirical coverage claims. Finally, priors matter in very low SNR; although we favor weakly informative defaults, users should perform sensitivity checks to ensure conclusions are not prior-driven.

Methodologically, the paper contributes a complete, reproducible UQ stack—likelihoods, priors, inference, identifiability, and calibration—that can wrap alternative optical models with minimal surgery. Practically, the framework provides confidence-aware maps and intervals that can drive downstream decisions in graphics (material editing, asset QA), vision (inspection, classification), and sensing (biophysical mapping), where the *trustworthiness* of parameters is as important as their values. We recommend reporting not only point estimates and residuals but also empirical coverage and interval sharpness, and we encourage releasing code, seeds, and capture scripts so others can replicate calibration diagnostics.

Extending the forward model to anisotropic microgeometry and multi-layer interfaces would reduce bias on brushed or fibrous materials; combining our probabilistic layer with such models should be straightforward. Multispectral dispersion priors that couple $n(\lambda)$ physically (e.g., via simple dispersion laws or GP priors with interpretable hyperparameters) can strengthen estimates and improve identifiability in challenging regimes. Joint inference of normals and material parameters, or active selection of views via Bayesian experimental design, could further improve information efficiency. Fi-

nally, commodity capture—phones with linear polarizers or compact multi-illumination rigs—paired with our uncertainty-aware inversion would broaden applicability; the key is to retain the bounded likelihood and robust residual so that reduced hardware does not degrade calibration.

In conclusion, uncertainty-aware polarimetric inversion converts polarization cues into calibrated, decision-ready information. By quantifying what the data truly support and by showing how to design captures that maximize identifiability while preserving calibration, this framework advances both the science and the practice of inverse rendering for layered dielectrics, turning “best guesses” into defensible, reproducible inferences.

References

- [1] Baek, S. H., Zeltner, T., Ku, H., Hwang, I., Tong, X., Jakob, W., & Kim, M. H. (2020). Image-based acquisition and modeling of polarimetric reflectance. *ACM Transactions on Graphics*, 39(4), 139.
- [2] Compain, E., Poirier, S., & Drevillon, B. (1999). General and self-consistent method for the calibration of polarization modulators, polarimeters, and Mueller-matrix ellipsometers. *Applied Optics*, 38(16), 3490-3502.
- [3] Barber, B. M., Huang, X., Ko, K. J., & Odean, T. (2020). *Leveraging Overconfidence*. Available at SSRN 3445660.
- [4] Suo, C., Li, Z., Zheng, H., & Sun, Y. (2017). Dynamic characteristics analysis on interface polarization and depolarization of nonlinear double-layered dielectrics. *IEEE Transactions on Dielectrics and Electrical Insulation*, 24(3), 1511-1526.
- [5] Stengel, M., & Spaldin, N. A. (2006). Origin of the dielectric dead layer in nanoscale capacitors. *Nature*, 443(7112), 679-682.
- [6] Sumin, D., Rittig, T., Babaei, V., Nindel, T., Wilkie, A., Didyk, P., ... & Weyrich, T. (2019). Geometry-aware scattering compensation for 3D printing. *ACM Transactions on Graphics*, 38(4), 98.
- [7] Kauppila, J. S., Massengill, L. W., Ball, D. R., Alles, M. L., Schrimpf, R. D., Loveless, T. D., ... & Rowe, J. D. (2015). Geometry-aware single-event enabled compact models for sub-50 nm partially depleted silicon-on-insulator technologies. *IEEE Transactions on Nuclear Science*, 62(4), 1589-1598.
- [8] Kuijken, K., & Merrifield, M. R. (1993). A new method for obtaining stellar velocity distributions from absorption-line spectra: unresolved Gaussian decomposition. *Monthly Notices of the Royal Astronomical Society*, 264(3), 712-720.
- [9] Massicotte, P., & Markager, S. (2016). Using a Gaussian decomposition approach to model absorption spectra of chromophoric dissolved organic matter. *Marine Chemistry*, 180, 24-32.
- [10] Price, M. A., Cai, X., McEwen, J. D., Pereyra, M., Kitching, T. D., & LSST Dark Energy Science Collaboration. (2020). Sparse Bayesian mass mapping with uncertainties: local credible intervals. *Monthly Notices of the Royal Astronomical Society*, 492(1), 394-404.
- [11] Lu, D., Ye, M., & Hill, M. C. (2012). Analysis of regression confidence intervals and Bayesian credible intervals for uncertainty quantification. *Water Resources Research*, 48(9).
- [12] Sankararaman, S., & Mahadevan, S. (2011). Likelihood-based representation of epistemic uncertainty due to sparse point data and/or interval data. *Reliability Engineering & System Safety*, 96(7), 814-824.

- [13] Maier, S. R. (2005). Accuracy matters: A cross-market assessment of newspaper error and credibility. *Journalism & Mass Communication Quarterly*, 82(3), 533-551.
- [14] Schünemann, H. J., Mustafa, R. A., Brozek, J., Steingart, K. R., Leeflang, M., Murad, M. H., ... & GRADE Working Group. (2020). GRADE guidelines: 21 part 2. Test accuracy: inconsistency, imprecision, publication bias, and other domains for rating the certainty of evidence and presenting it in evidence profiles and summary of findings tables. *Journal of Clinical Epidemiology*, 122, 142-152.
- [15] Li, X. R., & Zhao, Z. (2006, July). Measuring estimator's credibility: Noncredibility index. In 2006 9th International Conference on Information Fusion (pp. 1-8). IEEE.
- [16] Kuleshov, V., Fenner, N., & Ermon, S. (2018, July). Accurate uncertainties for deep learning using calibrated regression. In *International Conference on Machine Learning* (pp. 2796-2804). PMLR.

How to cite this article: Edwin R. Hancock (2021). Uncertainty-Aware Polarimetric Inverse Rendering for Layered Dielectrics with Credible Intervals, Identifiability, and Coverage Calibration. *Bulletin of Computer and Data Sciences*, 2(1), 17-27. DOI: [10.71448/bcds2121-3](https://doi.org/10.71448/bcds2121-3)

Received: 12/12/2020 **Revised:** 21/4/2021 **Accepted:** 15/8/2021 **Publish:** 30/12/2021

Copyright: © 2021 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <https://creativecommons.org/licenses/by/4.0/>.