

Spectral BRDF Prediction of Solar Panels Using a Five-Parameter Model and ExtraTrees

Haibin Tian^{1,2}, Zhou Wang³, Dan Li^{2,4}, Hongxing Cai^{2,4}, Ziwei Zhou³, Xuguo Zhang³, Zhenwen Liu^{1*}

¹School of Mathematics and Statistics, Changchun University of Science and Technology, Changchun, China

²School of Physics, Changchun University of Science and Technology, Changchun, China

³Beijing Institute of Space Mechanics & Electricity, Beijing, China

⁴Jilin Provincial Key Laboratory of Spectral Detection Science and Technology for Higher Education, Changchun, China

Email: *liuzw@cust.edu.cn

How to cite this paper: Tian, H.B., Wang, Z., Li, D., Cai, H.X., Zhou, Z.W., Zhang, X.G. and Liu, Z.W. (2026) Spectral BRDF Prediction of Solar Panels Using a Five-Parameter Model and ExtraTrees. *Journal of Applied Mathematics and Physics*, **14**, 4049-4065.

<https://doi.org/10.4236/jamp.2026.1410196>

Received: September 6, 2026

Accepted: October 17, 2026

Published: October 20, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Sparse angular measurements limit the prediction of solar-panel spectral bidirectional reflectance distribution functions (BRDFs), particularly near localized high-reflectance features. This study combines a wavelength-wise five-parameter semi-empirical model with two ExtraTrees regressors that predict a log-ratio residual and a direct log-spectrum. Their predictions are fused in log space, with reconstruction scale and fusion weight selected exclusively within each outer training partition. Measurements of a small solar-panel model containing solar cells under a glass covering comprise 120 angular configurations and 342 wavelengths from 500.08 to 759.85 nm. All configurations, including the highest-reflectance configuration, enter the primary evaluation; physically equivalent normal-view geometries remain together during splitting. Across 100 grouped random splits, the nested fusion model achieves mean global R^2 of 0.7932, RMSE of 0.002421 sr^{-1} , cosine-weighted normalized squared error of 4.2545%, and spectral angle of 2.1984° . Four incidence-angle holdouts yield mean R^2 of 0.8378 and RMSE of 0.002176 sr^{-1} . The high-reflectance configuration has mean held-out RMSE of 0.008614 sr^{-1} over its 15 test occurrences in the random protocol, indicating substantial local sensitivity despite improved aggregate amplitude prediction. These results support internal prediction over the sampled geometry and spectral range, while limited angular coverage, nominal white-reference calibration, and unavailable repeat measurements constrain physical interpretation and uncertainty assessment.

Keywords

Solar Panel, Spectral BRDF, Five-Parameter Model, ExtraTrees, Nested Validation

1. Introduction

The bidirectional reflectance distribution function (BRDF) relates reflected radiance to incident irradiance and specifies their directional dependence [1]. It is used in optical-scattering characterization, target radiometry, and scene simulation. Rough-surface, microfacet, anisotropic, and comprehensive reflection models describe different aspects of surface appearance [2]-[5]. Diffuse-reflection models, alternative angular parameterizations, measured reflectance representations, and comparative studies extend these descriptions to more complex materials [6]-[10]. Their usefulness under sparse measurement depends on both the suitability of the model and the geometry of the available observations.

Multiband optical observations support material characterization, while statistical five-parameter models provide compact descriptions of measured BRDFs [11] [12]. Solar-cell and photovoltaic-module measurements show directional responses that depend on the surface and its covering; periodic structure, out-of-plane scattering, structured glass, and wave-optics effects have been studied in this context [13]-[17]. Such findings motivate geometry-dependent spectral prediction, but do not establish the detailed optical structure of an uncharacterized specimen. Here, the specimen is a small model fitted with actual solar cells and covered by glass. Its cell technology, glass coating, and fine surface structure have not been independently characterized.

Automated BRDF measurement and learned reflectance representations offer additional routes to material and appearance characterization [18]-[20]. ExtraTrees, random forests, boosting, and multi-output regression provide nonlinear prediction tools [21]-[24], while principal components, Gaussian processes, and high-dimensional calibration methods offer alternative representations of correlated outputs [25]-[28]. This study uses two multi-output ExtraTrees regressors to learn complete spectra from angular coordinates. A positive semi-empirical anchor supports multiplicative residual learning, and log-domain fusion combines that representation with a direct spectral predictor.

Evaluation must distinguish model selection from performance assessment [29]-[31] and respect dependence among observations [32]-[34]. We therefore reevaluate the model using training-only nested selection, physically grouped random splits, and held-out incidence-angle layers. The primary analysis retains all 120 measured configurations. A paired 119-configuration sensitivity analysis and explicit held-out prediction of the high-reflectance configuration assess the effect of the exclusion used in the earlier evaluation. The revised analysis tests aggregate amplitude recovery, spectral shape, and local failure modes without interpreting split variability as measurement repeatability.

2. Methods

2.1. Specimen, Measurement Data, and Prediction Task

The specimen is a small solar-panel model fitted with actual solar cells beneath an external glass covering. Measurements were made using the indoor spectral BRDF

platform at Changchun University of Science and Technology. The acquisition documentation describes a solar simulator and reflecting mirror, a sample elevation platform, azimuth and zenith positioning structures, and fiber-coupled collection using a 74-UV optic and an Ocean QE65 Pro spectrometer. The documented procedure includes at least 15 min of warm-up, laser-assisted alignment, and integration-time adjustment to avoid saturation. The listed spectrometer range is 199.36 - 994.96 nm, with an optical resolution specification of 0.47 nm FWHM. These are instrument-document specifications; they are not an uncertainty estimate for this measurement batch.

The angular coordinates are defined in **Figure 1**. The exported variables x , z , and y are in degrees, with:

$$\theta_i = z, \quad \phi_i = 0^\circ, \quad \theta_o = 90^\circ - y, \quad \phi_o = x. \quad (1)$$

Here, zenith angles are measured from the surface normal; x denotes observation azimuth and y is the complement of observation zenith. The grid is $x = 30^\circ, 60^\circ, 90^\circ, 120^\circ, 150^\circ$; $z = 30^\circ, 40^\circ, 50^\circ, 60^\circ$; and $y = 15^\circ, 30^\circ, 45^\circ, 60^\circ, 75^\circ, 90^\circ$. It contains 120 labeled configurations. At $y = 90^\circ$, observation is normal to the surface and azimuth is undefined physically. The five x labels at each fixed z are therefore placed in one physical group for splitting, giving 104 physical geometry groups. All measured spectra are retained; no averaging is performed. These equivalent labels are not treated as a documented repeatability experiment.




Figure 1. Measurement geometry and angular-coordinate definitions for the glass-covered solar-panel model. The recorded variables map to the BRDF angles as $\theta_i = z$, $\theta_r = 90^\circ - y$, $\phi_i = 0^\circ$, and $\phi_r = x$. Direction vectors are defined away from the surface.

The reference was a Labsphere Spectralon SRS-99 diffuse-reflectance standard with nominal reflectance of 99% [35]. The individual reference's wavelength-dependent calibration spectrum was unavailable. The exported dataset uses an integration-time-normalized sample/reference ratio and the nominal Lambertian-reference conversion:

$$f_r(\lambda) = \frac{0.99}{\pi} \frac{S(\lambda)/t_s}{W(\lambda)/t_w}.$$

Here, S and W denote the exported sample and white-reference signals, and t denotes integration time. Sample integration times were 800,000 or 2,000,000 μs ; the reference integration time was 100,000 μs . This conversion reproduces the BRDF column throughout the exported dataset. The nominal 0.99 factor is applied once. Manufacturer information describes the standard family but does not replace the unavailable calibration of the individual standard. Wavelength-dependent departures from 0.99, reference non-Lambertian behavior, and instrumental effects are not corrected by this nominal conversion. Although the acquisition procedure specifies background handling, its implementation cannot be independently reconstructed from the exported signals alone.

The analyzed dataset contains 41,040 records marked `quality_flag = ok`, with 342 wavelength samples from 500.08 to 759.85 nm at every configuration. Adjacent exported wavelengths differ by approximately 0.74 - 0.78 nm; this sampling interval is distinct from optical resolution. No additional spectral resampling or smoothing is applied in the reevaluation. Each configuration supplies one complete output vector:

$$\mathbf{y}(x, z, y) = [f_r(\lambda_1), f_r(\lambda_2), \dots, f_r(\lambda_{342})]^T. \quad (2)$$

No dedicated repeated acquisitions at the same configuration are available. Consequently, measurement repeatability, a traceable combined uncertainty budget, and uncertainty attributable to positioning or source drift cannot be estimated from these data. The retained high-reflectance configuration is $(x, z, y) = (30^\circ, 30^\circ, 60^\circ)$, corresponding to incident zenith 30° , observation zenith 30° , and observation azimuth 30° . No configuration is removed because of its reflectance magnitude.

2.2. Five-Parameter Semi-Empirical Anchor

At each wavelength, the five-parameter model [12] is fitted independently:

$$f_z = \frac{k_b D(\alpha; k_r) F(\gamma; a, b) G(\theta_i, \theta_o, \phi)}{\cos \theta_o} + k_d. \quad (3)$$

The parameters k_b and k_d control directional and background components, k_r controls the distribution shape, and a and b control an empirical Fresnel-like term. The half-vector inclination is α , the angle between incident direction and half-vector is γ , and the relative azimuth is $\phi = \varphi_i - \varphi_o$. Thus:

$$\cos^2 \gamma = \frac{\cos \theta_i \cos \theta_o + \sin \theta_i \sin \theta_o \cos \phi + 1}{2}, \quad (4)$$

$$\cos \alpha = \frac{\cos \theta_i + \cos \theta_o}{2 \cos \gamma}. \quad (5)$$

The distribution and empirical reflection terms are:

$$D(\alpha; k_r) = \frac{k_r^2 \cos \alpha}{1 + (k_r^2 - 1) \cos^2 \alpha}, \quad (6)$$

$$F(\gamma; a, b) = \exp \left[b(1 - \cos \gamma)^a \right]. \quad (7)$$

The statistical masking term G uses the retained implementation with constants $\sigma_{\text{sph}} = 0.0136$, $\sigma_r = 0.0141$, $u = 9$, and $v = 1$. Its full computational definition is supplied with the reproducibility code. Bounds are $k_b \in [0, 50]$, $k_d \in [0, 5]$, $k_r \in [0.05, 50]$, $a \in [0.05, 25]$, and $b \in [-120, 0]$. Bounded least squares uses the soft-L1 loss, six initializations, and at most 400 function evaluations per initialization. Residuals are weighted by $\sqrt{\max(\cos \theta_o, 0.05)}$ and scaled by the training-data standard deviation at each wavelength. The previous wavelength's solution is one candidate initialization when available. All fits, including this initialization sequence, use only the applicable training partition.

2.3. Geometry-Informed Correction

A unit direction vector and the corresponding specular direction are:

$$\boldsymbol{\omega}(\theta, \phi) = [\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta]^T, \quad (8)$$

$$\boldsymbol{\omega}_{\text{spec}} = 2(\mathbf{n} \cdot \boldsymbol{\omega}_i) \mathbf{n} - \boldsymbol{\omega}_i, \quad \mathbf{n} = [0, 0, 1]^T. \quad (9)$$

The specular angular deviation is:

$$\Delta_{\text{spec}} = \arccos(\boldsymbol{\omega}_o \cdot \boldsymbol{\omega}_{\text{spec}}). \quad (10)$$

For an in-plane reference direction φ_g , define:

$$\begin{bmatrix} q_x \\ q_y \end{bmatrix} = \begin{bmatrix} \sin \theta_o \cos(\phi_o - \phi_g) - \sin \theta_i \cos(\phi_i - \phi_g) \\ \sin \theta_o \sin(\phi_o - \phi_g) - \sin \theta_i \sin(\phi_i - \phi_g) \end{bmatrix}. \quad (11)$$

The retained correction functions are:

$$g_{\text{spec}} = \exp \left[-\frac{1}{2} \left(\frac{\Delta_{\text{spec}}}{\sigma_{\text{spec}}} \right)^2 \right], \quad \sigma_{\text{spec}} = 3^\circ, \quad (12)$$

$$g_{\phi_g} = \exp \left[-\frac{1}{2} \left(\frac{q_x^2}{\sigma_x^2} + \frac{q_y^2}{\sigma_y^2} \right) \right], \quad \sigma_x = \sigma_y = 0.10. \quad (13)$$

The implementation retains $\varphi_g = 0^\circ$ and 90° . Because $\sigma_x = \sigma_y$, these two basis functions are identical under rotation. They provide a redundant, regularized correction, rather than identifying two independent structural axes or measured grid-line directions. Their coefficients are estimated from training residuals, with $\mathbf{g} = [g_{\text{spec}}, g_0, g_{90}]^T$ and training design matrix $\mathbf{B}$:

$$\mathbf{c}(\lambda) = (\mathbf{B}^T \mathbf{B} + 10^{-7} \mathbf{I})^{-1} \mathbf{B}^T [\mathbf{y}(\lambda) - \mathbf{f}_z(\lambda)]. \quad (14)$$

The resulting anchor is:

$$f_{\text{phys}}(\lambda) = \max \{ 10^{-10}, f_z(\lambda) + \mathbf{g}^T \mathbf{c}(\lambda) \}. \quad (15)$$

The subscript *phys* is retained as a model label. Positivity and explicit geometric terms do not establish energy conservation, reciprocity, or a uniquely identified optical mechanism.

2.4. ExtraTrees Branches and Training-Only Fusion Selection

Both regressors use $[x, z, y]$ as input and predict a 342-component vector. Input standardization is fitted on training configurations. Each regressor has 900 trees, minimum leaf size 2, and `max_features = 1.0`. Tree seeds equal the applicable fit seed plus 17 and plus 29. The same architecture and fitting settings are used throughout; the server executes 16 independent outer jobs concurrently with one computational thread per job.

The ratio branch learns the logarithmic residual:

$$\mathbf{r}_{\text{ratio}} = \log(\mathbf{y} / \mathbf{f}_{\text{phys}}). \quad (16)$$

Its reconstruction uses a scale η :

$$\hat{\mathbf{y}}_{\text{ratio}} = \mathbf{f}_{\text{phys}} \odot \exp(\eta \hat{\mathbf{r}}_{\text{ratio}}). \quad (17)$$

The direct branch learns:

$$\ell = \log \mathbf{y}, \quad (18)$$

and its anchored reconstruction is:

$$\hat{\mathbf{y}}_{\text{direct}} = \exp[(1 - \eta) \log \mathbf{f}_{\text{phys}} + \eta \ell]. \quad (19)$$

The two reconstructions are fused using ratio-branch weight w :

$$\hat{\mathbf{y}} = \exp[w \log \hat{\mathbf{y}}_{\text{ratio}} + (1 - w) \log \hat{\mathbf{y}}_{\text{direct}}]. \quad (20)$$




Figure 2. Nested evaluation workflow. Only inner-validation predictions select the reconstruction scale and fusion weight. Outer test targets are supplied to scoring after model fitting and selection.

All divisions, logarithms, and exponentials are elementwise; positive values are clipped below at 10^{-10} . We use η for reconstruction scale to distinguish it from the half-vector angle α . The two branches share η , preserving the structure of the earlier 1.05/1.05 formulation. For every outer split, three inner folds evaluate $\eta \in \{0.95, 1.00, 1.05, 1.10\}$ and $w \in \{0, 0.25, 0.50, 0.65, 0.75, 1.00\}$. The pair minimizing pooled inner-validation RMSE in linear BRDF units is selected. Pooling

sums squared errors over all validation wavelengths and configurations before taking the square root. Ties favor η nearest 1, then w nearest 0.5, followed by ascending η and w .

The five-parameter fits, correction coefficients, input scaler, and both regressors are fitted afresh inside each inner training fold. After selection, the complete pipeline is refitted on the outer training data and evaluated on the untouched outer test data. **Figure 2** summarizes the separation. The historical values $\eta = 1.05$ and $w = 0.65$ are retained only as a labeled comparison. Their original selection history does not establish independence from this dataset. The candidate grid was frozen before the new outer results were generated; it was not expanded or reselected using those results.

2.5. Comparison Methods

Eight variants use identical outer partitions: the five-parameter baseline; the geometry-corrected anchor; the ratio branch with $w = 1$; the anchored direct branch with $w = 0$; equal-weight fusion with $w = 0.5$; nested fusion with η and w selected jointly; historical fixed fusion with $\eta = 1.05$ and $w = 0.65$; and pure ExtraTrees predicting the direct log-spectrum. For the ratio, anchored-direct, and equal-fusion variants, η is independently selected from the same inner folds with their respective w fixed. Pure ExtraTrees is obtained at $\eta = 1$ and $w = 0$, where the anchor cancels exactly. It uses the same fitted direct-log regressor, so it supplies a controlled anchor-free comparison. The earlier capped-ratio variant is omitted from the revised comparison set, which instead emphasizes explicit branch controls and the historical fixed configuration.

2.6. Evaluation Design and Metrics

The primary evaluation comprises 100 random splits with seeds 42 - 141 and a test fraction of 0.15 applied to the 104 physical groups. Each split holds out 16 physical groups, corresponding to 16 - 32 labeled configurations because normal-view groups contain five spectra. The observed test counts are 16, 20, 24, 28, and 32 in 51, 36, 11, 1, and 1 splits, respectively; training counts range from 104 to 88. Inner validation uses three shuffled physical-group folds. All wavelengths of a configuration and all physically equivalent normal-view labels remain on the same side at both levels. These random splits assess within-grid interpolation and can place neighboring geometries in training and test sets.

For angularly blocked evaluation, each incidence zenith $z = 30^\circ, 40^\circ, 50^\circ$, or 60° is held out in turn: 90 configurations train and 30 test. Inner validation leaves out one of the three remaining incidence layers at a time. The 40° and 50° tests interpolate between training incidence levels; the 30° and 60° tests lie outside the incidence range of their respective training folds. The azimuth and observation-elevation grids remain shared, so this is an incidence-block test, not a test of an entirely new angular domain.

Two additional designs examine the high-reflectance configuration. First,

paired restricted-data splits remove it from whichever side contains it in each full-data partition, then repeat inner selection and outer fitting on the remaining 119 configurations. Separately, full-data predictions are rescored after omitting only this configuration from the test metric; that diagnostic makes no refit. Second, its predictions are reported for all 15 occasions on which it is an outer test case in the primary random protocol. A supplementary leave-one-configuration-out run trains on all other 119 spectra and tests only this configuration, using the same nested selection rule. This larger training set is identified explicitly and is not substituted for the matched-protocol result.

For N test wavelength-configuration pairs, global R^2 and RMSE are:

$$R^2 = 1 - \frac{\sum_{j=1}^N (y_j - \hat{y}_j)^2}{\sum_{j=1}^N (y_j - \bar{y})^2}, \quad (21)$$

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{j=1}^N (y_j - \hat{y}_j)^2}. \quad (22)$$

The cosine-weighted normalized squared error is:

$$\text{CW-NSE} = 100\% \times \frac{\sum_{j=1}^N [(\hat{y}_j - y_j) \cos \theta_{o,j}]^2}{\sum_{j=1}^N [y_j \cos \theta_{o,j}]^2}. \quad (23)$$

Equation (23) is a normalized squared pointwise error with effective $\cos^2$ weighting. It contains neither a square root nor solid-angle quadrature and must not be interpreted as hemispherical reflected-energy error. Its numerical definition is preserved while its name and interpretation are corrected. Spectral angle mapper (SAM) compares the complete spectra of individual configurations [36] [37]:

$$\text{SAM}(\mathbf{y}, \hat{\mathbf{y}}) = \arccos \left(\frac{\mathbf{y}^T \hat{\mathbf{y}}}{\|\mathbf{y}\|_2 \|\hat{\mathbf{y}}\|_2} \right). \quad (24)$$

SAM is converted to degrees and averaged across test configurations. Global metrics aggregate all test wavelength-configuration pairs within a split; reported protocol means then weight splits equally. Signed relative mean bias is 100 times mean (prediction-measurement) divided by mean (measurement). All raw spectra retain equal weight in the pointwise metrics, including the retained normal-view labels. Split distributions are descriptive sensitivity summaries: repeated splits overlap, and neither their spread nor their count represents independent measurement replication. All 205 outer jobs use the existing specimen dataset, making this a retrospective internal reevaluation rather than validation on newly acquired independent data.

3. Results


3.1. Full-Dataset Within-Grid Interpolation

Table 1 reports all eight variants using the complete 120-configuration dataset.

The primary nested fusion model achieves mean global R^2 of 0.7932, RMSE of 0.002421 sr^{-1} , CW-NSE of 4.2545%, and SAM of 2.1984°. Relative to the five-parameter baseline, the mean amplitude-sensitive errors decrease substantially. Pure ExtraTrees has lower SAM, showing that improvements in amplitude and spectral direction are distinct objectives.

Table 1. Mean test metrics across 100 grouped random splits of all 120 configurations.

Model	R^2	RMSE (sr^{-1})	CW-NSE (%)	SAM (°)
Five-parameter	-0.0389	0.005568	23.0735	5.8621
Geometry-corrected	-0.0022	0.005405	21.6780	5.6710
Ratio branch (tuned)	0.6568	0.003021	9.1964	3.0197
Direct anchored (tuned)	0.6571	0.003195	8.9503	1.2002
Equal fusion (tuned)	0.7904	0.002440	4.4555	1.7276
Nested fusion (primary)	0.7932	0.002421	4.2545	2.1984
Historical fixed fusion	0.8063	0.002315	3.8223	2.0881
Pure ExtraTrees	0.6732	0.003113	8.4699	1.1300



100 outer splits; all 120 configurations eligible, including the high-reflectance configuration.

Figure 3. Distributions across the 100 full-data grouped random splits. Boxes show the interquartile range and median; whiskers extend to observations within 1.5 interquartile ranges and points show the remaining observations. “Nested” denotes training-only fusion selection; “Fixed” denotes the historical configuration. These distributions describe partition sensitivity, not measurement uncertainty.

The historical fixed configuration has slightly better mean R^2 , RMSE, and CW-

NSE than nested fusion in these random splits. This comparison does not establish an independently selected optimum because those fixed values arose during the earlier development on the same dataset. The primary result remains the training-only selection procedure. **Figure 3** shows the observed split distributions. Nested-fusion R^2 has median 0.8259, interquartile range 0.7414 - 0.8593, and range 0.4227 - 0.9511. Its RMSE ranges from 0.000925 to 0.005187 sr^{-1} , indicating substantial dependence on which geometries are withheld.

3.2. Fusion and Selection Sensitivity

Geometry correction alone raises mean R^2 from -0.0389 to -0.0022 . The larger gain accompanies ExtraTrees prediction: ratio and anchored-direct branches achieve mean R^2 of 0.6568 and 0.6571, respectively, while equal fusion and nested fusion achieve 0.7904 and 0.7932. The pure direct-log ExtraTrees baseline achieves 0.6732. Thus, combining the two representations improves mean amplitude metrics over either branch alone in this evaluation, although nested fusion does not minimize SAM or outperform every comparison on every metric.

Across the 100 full-data random splits, inner validation selects $\eta = 0.95, 1.00, 1.05$, and 1.10 in 1, 4, 13, and 82 cases. Selected w values are 0.50, 0.65, 0.75, and 1.00 in 16, 30, 45, and 9 cases. Frequent selection of the upper η boundary indicates sensitivity to the bounded search; it does not show that 1.10 is a global optimum. The frozen search is retained to preserve the interpretation of the outer evaluation.

3.3. Incidence-Angle Holdouts

Table 2 reports nested fusion for the four angular blocks. The four-fold mean R^2 is 0.8378, RMSE is 0.002176 sr^{-1} , CW-NSE is 3.5922%, and SAM is 1.6736° . The 60° block has the largest RMSE and SAM. **Figure 4** compares the amplitude metrics with the five-parameter and pure ExtraTrees baselines; their four-fold mean R^2 values are -0.0038 and 0.7239, respectively.

Table 2. Nested fusion with each incidence zenith layer held out. Each fold uses 90 training and 30 test configurations.

Held-out z	R^2	RMSE (sr^{-1})	CW-NSE (%)	SAM ($^\circ$)
30°	0.8510	0.002404	2.8639	0.9954
40°	0.8727	0.002037	2.4360	1.1257
50°	0.8776	0.001574	2.0852	1.0910
60°	0.7498	0.002687	6.9837	3.4822
Four-fold mean	0.8378	0.002176	3.5922	1.6736

Inner validation selects $w = 0.5$ in every blocked fold, and nested and equal-fusion predictions coincide in these four tests. The historical fixed fusion gives mean R^2 of 0.8269 and RMSE of 0.002252 sr^{-1} . The higher blocked mean R^2 than the random-split mean does not imply that blocked prediction is generally easier:

the test composition, target variance, and training geometry differ. The result supports the specific incidence-layer transfers evaluated here, including two boundary-layer tests, rather than unrestricted angular extrapolation.




Figure 4. Incidence-angle holdout results for the five-parameter model, pure ExtraTrees, and nested fusion. Each bar represents one outer fold; the remaining incidence levels alone supply training and inner validation.

3.4. Sensitivity to the High-Reflectance Configuration

Table 3 separates the primary full-data result, a scoring-only diagnostic, and the paired 119-configuration refit. Removing the high-reflectance test spectrum from scoring lowers mean RMSE to 0.002296 sr^{-1} . Removing it from the dataset and repeating selection and fitting lowers mean RMSE further to 0.002268 sr^{-1} , with SAM of 1.2468° . These are different operations and cannot be treated as interchangeable estimates. Changes in R^2 are not necessarily monotonic because removal also changes the test-set variance used in its denominator. In the restricted-data refits, equal fusion has mean R^2 of 0.8065 and RMSE of 0.002205 sr^{-1} , again precluding a claim that free weight selection wins universally.

Table 3. Nested-fusion sensitivity analysis, averaged over 100 paired random partitions.

Analysis	R^2	RMSE (sr^{-1})	CW-NSE (%)	SAM ($^\circ$)
Full 120 (primary)	0.7932	0.002421	4.2545	2.1984
Full fit; peak omitted from scoring	0.7909	0.002296	3.9799	2.1817
Restricted 119; retuned and refitted	0.7939	0.002268	3.9459	1.2468

Figure 5 summarizes each configuration's mean RMSE over its actual test occurrences. The high-reflectance configuration has the second-highest mean RMSE among the 120 labeled configurations; the largest occurs at $(x, z, y) = (150^\circ, 60^\circ, 90^\circ)$, at 0.008657 sr^{-1} . Thus, a large absolute error is not unique to the reflectance peak. Higher-magnitude configurations can contribute strongly to squared-error aggregates, but magnitude alone does not identify the physical cause of a local response. This diagnostic is descriptive and retains all configurations in the primary metrics.




Figure 5. Configuration-level error for nested fusion in the full-data random protocol. (a) Mean held-out RMSE ranked across 120 labeled configurations. (b) The same errors against each configuration's mean measured BRDF. The red point is $(x, z, y) = (30^\circ, 30^\circ, 60^\circ)$. Each point averages its own test occurrences; occurrence counts are supplied with the figure data.

3.5. Explicit Held-Out Prediction of the High-Reflectance Configuration

The configuration $(x, z, y) = (30^\circ, 30^\circ, 60^\circ)$ appears in 15 test partitions of the primary random protocol. Its mean RMSE is 0.008614 sr^{-1} , with a range of $0.003295 - 0.015678 \text{ sr}^{-1}$; mean CW-NSE is 8.0652% and mean SAM is 3.1826° (Table 4). Relative mean bias ranges from -41.48% to $+47.18\%$, despite its average being -6.49% . Thus, a modest average signed bias conceals substantial variation between training partitions. Figure 6(a) shows all 15 predictions and their range, which is not a confidence interval.

The supplementary 119-training fit selects $\eta = 1.05$ and $w = 0.65$. Its RMSE is

0.002514 sr^{-1} , compared with 0.023350 sr^{-1} for the five-parameter baseline and 0.021013 sr^{-1} for pure ExtraTrees. Its relative mean bias is -1.94% , but SAM is 4.2392° . **Figure 6(b)** shows that the prediction contains a localized spectral discontinuity near 704 nm that is absent from the measured curve. Wavelength-wise anchor fitting does not enforce spectral continuity, and this fitted artifact is not interpreted as a material feature. The favorable amplitude result with 119 training configurations does not resolve the variability observed under the smaller matched-protocol training sets.

Table 4. Held-out peak-configuration metrics. The first row averages 15 matched-protocol test cases; the second is one supplementary fit with all other 119 configurations for training.

Peak test design	R^2	RMSE (sr^{-1})	CW-NSE (%)	SAM ($^\circ$)	Mean bias (%)
Matched random (15 cases)	-14.4059	0.008614	8.0652	3.1826	-6.49
119-training holdout (1 case)	-0.1032	0.002514	0.5775	4.2392	-1.94




Figure 6. Explicit held-out predictions for $(x, z, y) = (30^\circ, 30^\circ, 60^\circ)$. (a) All 15 primary-protocol test predictions, their pointwise range, their arithmetic mean, and the measured spectrum. (b) A separate leave-one-configuration-out test using all other 119 spectra for training. The two panels use different training-set sizes and therefore answer different questions.

For this single spectrum, mean random-protocol curve R^2 is -14.4059 and the supplementary fit gives -0.1032 . Within-spectrum variation is small relative to the dataset's between-configuration variation, so curve R^2 and global R^2 are not comparable. The negative curve values nevertheless document residual error and are retained alongside RMSE, CW-NSE, SAM, and signed bias. Sparse angular

sampling still prevents determination of the local angular width or gradient of the high-reflectance feature.

4. Discussion

The full-data reevaluation supports improved aggregate BRDF-amplitude prediction from the combined representations relative to the five-parameter model and pure ExtraTrees. It also narrows the interpretation of that improvement. Nested fusion is the primary procedure because its scale and weight are selected without outer test targets; historical fixed fusion is a transparent development-era comparator. Differences between metrics and between partition designs show that a single preferred weight cannot be inferred as a general property of the specimen.

The geometry correction supplies a positive semi-empirical anchor, but its equal-width in-plane functions do not identify anisotropy, and positivity does not enforce physical conservation laws. The glass covering and actual solar cells establish the broad specimen class. Unknown cell technology, glass thickness or coating, and fine geometry prevent attribution of the isolated response to a particular interface, diffraction process, or grid structure. The current angle grid does not resolve that response locally. Prediction on another specimen, illumination spectrum, spectral band, or azimuthal domain has not been demonstrated.

Measurement and statistical uncertainties must also remain separate. The reported BRDF scale assumes nominal 99% reference reflectance and a Lambertian-reference conversion. Manufacturer specifications cannot recover the missing individual calibration spectrum. No repeated-acquisition series or complete batch-level uncertainty budget is available; train/test variation measures sensitivity of the fitted predictor to the sampled data rather than instrument precision. Computational reevaluation cannot replace these missing physical measurements. Repeated angular measurements, an individual reference calibration, and denser local angular sampling would be needed for stronger metrological claims.

The nested design prevents outer test targets from entering the new fitting and selection steps. It does not undo the prior use of the same dataset to develop the architecture and select retained settings. Moreover, the η search frequently reaches its upper boundary, and the wavelength-wise anchor can introduce discontinuities. Future evaluation should freeze any revised search or spectral regularization before testing it on independently acquired data. These limitations favor reporting the present results as retrospective internal evidence with explicit local failure analysis.

5. Conclusion

A five-parameter-anchored ExtraTrees fusion model was evaluated using all 120 measured configurations of a glass-covered solar-panel model. Training-only nested selection gives mean R^2 of 0.7932 and RMSE of 0.002421 sr^{-1} over 100 grouped random splits; four incidence-angle holdouts give mean R^2 of 0.8378 and RMSE of 0.002176 sr^{-1} . The corrected CW-NSE metric is a normalized squared

pointwise error and carries no reflected-energy interpretation. Explicit prediction of the high-reflectance configuration reveals substantial partition sensitivity, with mean matched-protocol RMSE of 0.008614 sr^{-1} . The results support aggregate amplitude prediction over the evaluated geometry and wavelength range, while calibration limitations, missing repeat measurements, local angular undersampling, and spectral fitting artifacts constrain broader physical and predictive claims.

Reproducibility Materials

The accompanying materials contain the input dataset, retained model implementation, frozen experiment protocol and source hashes, all outer partitions, inner candidate scores, selected parameters, test predictions, metric tables, and independent numerical-audit results. The complete run contains 100 full-data random jobs, 100 paired restricted-data jobs, four incidence-block jobs, and one peak-only job. The archived protocol records software versions and random seeds. No new physical measurements were acquired for this reevaluation.

Funding

CAST Innovation Fund (KYC-XZ-XM-2025-021, Intelligent Recognition of Space-Target States Based on Snapshot Spectral Characteristics).

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- [1] Nicodemus, F.E., Richmond, J.C., Hsia, J.J., *et al.* (1977) Geometrical Considerations and Nomenclature for Reflectance. National Bureau of Standards.
- [2] Torrance, K.E. and Sparrow, E.M. (1967) Theory for Off-Specular Reflection from Roughened Surfaces. *Journal of the Optical Society of America*, **57**, 1105-1114. <https://doi.org/10.1364/josa.57.001105>
- [3] Cook, R.L. and Torrance, K.E. (1982) A Reflectance Model for Computer Graphics. *ACM Transactions on Graphics*, **1**, 7-24. <https://doi.org/10.1145/357290.357293>
- [4] Ward, G.J. (1992) Measuring and Modeling Anisotropic Reflection. *ACM SIGGRAPH Computer Graphics*, **26**, 265-272. <https://doi.org/10.1145/142920.134078>
- [5] He, X.D., Torrance, K.E., Sillion, F.X. and Greenberg, D.P. (1991) A Comprehensive Physical Model for Light Reflection. *ACM SIGGRAPH Computer Graphics*, **25**, 175-186. <https://doi.org/10.1145/127719.122738>
- [6] Oren, M. and Nayar, S.K. (1994) Generalization of Lambert's Reflectance Model. *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, New York, 24-29 July 1994, 239-246. <https://doi.org/10.1145/192161.192213>
- [7] Rusinkiewicz, S.M. (1998) A New Change of Variables for Efficient BRDF Representation. In: Drettakis, G. and Max, N., Eds., *Eurographics*, Springer, 11-22. https://doi.org/10.1007/978-3-7091-6453-2_2
- [8] Matusik, W., Pfister, H., Brand, M. and McMillan, L. (2003) A Data-Driven Reflectance Model. *ACM Transactions on Graphics*, **22**, 759-769.

- <https://doi.org/10.1145/882262.882343>
- [9] Ngan, A., Durand, F. and Matusik, W. (2005) Experimental Analysis of BRDF Models. In: Bala, K. and Dutré, P., Eds., *Eurographics Symposium on Rendering*, Eurographics Association, 117-126.
 - [10] Guarnera, D., Guarnera, G.C., Ghosh, A., Denk, C. and Glencross, M. (2016) BRDF Representation and Acquisition. *Computer Graphics Forum*, **35**, 625-650. <https://doi.org/10.1111/cgf.12867>
 - [11] Hall, D. (2010) Surface Material Characterization from Multi-Band Optical Observations. *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, 14-17 September 2010, 61-75.
 - [12] Zhang, H., Wu, Z., Cao, Y., et al. (2010) Measurement and Statistical Modeling of BRDF of Various Samples. *Optica Applicata*, **40**, 197-208.
 - [13] Hou, Q., Wang, Z., Su, J. and Tan, F. (2019) Measurement of Equivalent BRDF on the Surface of Solar Panel with Periodic Structure. *Coatings*, **9**, Article 193. <https://doi.org/10.3390/coatings9030193>
 - [14] Small, T.V., Butler, S.D. and Marciniak, M.A. (2021) Solar Cell BRDF Measurement and Modeling with Out-of-Plane Data. *Optics Express*, **29**, 35501-35515. <https://doi.org/10.1364/oe.440190>
 - [15] Babin, M., Thorsteinsson, S., Jakobsen, M.L. and Spataru, S.V. (2022) Glare Potential Evaluation of Structured PV Glass Based on Goniorelectrometry. *IEEE Journal of Photovoltaics*, **12**, 1314-1318. <https://doi.org/10.1109/jphotov.2022.3189779>
 - [16] Compean, M., Small, T., Hyde, M. and Marciniak, M. (2023) Wave Optics Approach to Solar Cell BRDF Modeling with Experimental Results. *Optics Express*, **31**, 26289-26300. <https://doi.org/10.1364/oe.494284>
 - [17] Bucher, C., Bohren, A., Hess, D., et al. (2023) Two-Dimensional Representation of the Bidirectional Reflectance Distribution Function of Photovoltaic Modules. *40th European Photovoltaic Solar Energy Conference and Exhibition*, Lisbon, 18-22 September 2023.
 - [18] Yao, D., Sun, Y., He, L., Wu, H., Lin, G., Wang, J., et al. (2026) Automated BRDF Measurement for Aerospace Materials and 1D-CNN-Based Estimation of Mixed-Material Composition. *Sensors*, **26**, Article 1560. <https://doi.org/10.3390/s26051560>
 - [19] Deschaintre, V., Aittala, M., Durand, F., Drettakis, G. and Bousseau, A. (2018) Single-image SVBRDF Capture with a Rendering-Aware Deep Network. *ACM Transactions on Graphics*, **37**, 1-15. <https://doi.org/10.1145/3197517.3201378>
 - [20] Nam, G., Lee, J.H., Gutierrez, D. and Kim, M.H. (2018) Practical SVBRDF Acquisition of 3D Objects with Unstructured Flash Photography. *ACM Transactions on Graphics*, **37**, 1-12. <https://doi.org/10.1145/3272127.3275017>
 - [21] Geurts, P., Ernst, D. and Wehenkel, L. (2006) Extremely Randomized Trees. *Machine Learning*, **63**, 3-42. <https://doi.org/10.1007/s10994-006-6226-1>
 - [22] Breiman, L. (2001) Random Forests. *Machine Learning*, **45**, 5-32. <https://doi.org/10.1023/a:1010933404324>
 - [23] Chen, T. and Guestrin, C. (2016) XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, 13-17 August 2016, 785-794. <https://doi.org/10.1145/2939672.2939785>
 - [24] Borchani, H., Varando, G., Bielza, C. and Larrañaga, P. (2015) A Survey on Multi-output Regression. *WIREs Data Mining and Knowledge Discovery*, **5**, 216-233. <https://doi.org/10.1002/widm.1157>

- [25] Jolliffe, I.T. and Cadima, J. (2016) Principal Component Analysis: A Review and Recent Developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **374**, Article 20150202. <https://doi.org/10.1098/rsta.2015.0202>
- [26] Rasmussen, C.E. and Williams, C.K.I. (2005) Gaussian Processes for Machine Learning. The MIT Press. <https://doi.org/10.7551/mitpress/3206.001.0001>
- [27] Higdon, D., Gattiker, J., Williams, B. and Rightley, M. (2008) Computer Model Calibration Using High-Dimensional Output. *Journal of the American Statistical Association*, **103**, 570-583. <https://doi.org/10.1198/016214507000000888>
- [28] Camps-Valls, G., Verrelst, J., Munoz-Mari, J., Laparra, V., Mateo-Jimenez, F. and Gomez-Dans, J. (2016) A Survey on Gaussian Processes for Earth-Observation Data Analysis: A Comprehensive Investigation. *IEEE Geoscience and Remote Sensing Magazine*, **4**, 58-78. <https://doi.org/10.1109/mgrs.2015.2510084>
- [29] Stone, M. (1974) Cross-Validatory Choice and Assessment of Statistical Predictions. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, **36**, 111-133. <https://doi.org/10.1111/j.2517-6161.1974.tb00994.x>
- [30] Varma, S. and Simon, R. (2006) Bias in Error Estimation When Using Cross-Validation for Model Selection. *BMC Bioinformatics*, **7**, Article No. 91. <https://doi.org/10.1186/1471-2105-7-91>
- [31] Cawley, G.C. and Talbot, N.L.C. (2010) On Over-Fitting in Model Selection and Subsequent Selection Bias in Performance Evaluation. *Journal of Machine Learning Research*, **11**, 2079-2107.
- [32] Roberts, D.R., Bahn, V., Ciuti, S., Boyce, M.S., Elith, J., Guillerá-Arroita, G., et al. (2017) Cross-Validation Strategies for Data with Temporal, Spatial, Hierarchical, or Phylogenetic Structure. *Ecography*, **40**, 913-929. <https://doi.org/10.1111/ecog.02881>
- [33] Vabalas, A., Gowen, E., Poliakoff, E. and Casson, A.J. (2019) Machine Learning Algorithm Validation with a Limited Sample Size. *PLOS ONE*, **14**, e0224365. <https://doi.org/10.1371/journal.pone.0224365>
- [34] Kapoor, S. and Narayanan, A. (2023) Leakage and the Reproducibility Crisis in Machine-Learning-Based Science. *Patterns*, **4**, Article 100804. <https://doi.org/10.1016/j.patter.2023.100804>
- [35] LABSPHERE (2019) Spectralon Diffuse Reflectance Standards. PB-13058, Revision 2, Manufacturer Datasheet.
- [36] Kruse, F.A., Lefkoff, A.B., Boardman, J.W., Heidebrecht, K.B., Shapiro, A.T., Barloon, P.J., et al. (1993) The Spectral Image Processing System (SIPS)—Interactive Visualization and Analysis of Imaging Spectrometer Data. *Remote Sensing of Environment*, **44**, 145-163. [https://doi.org/10.1016/0034-4257\(93\)90013-n](https://doi.org/10.1016/0034-4257(93)90013-n)
- [37] Du, Y., Chang, C., Ren, H., Chang, C., Jensen, J.O. and D'Amico, F.M. (2004) New Hyperspectral Discrimination Measure for Spectral Characterization. *Optical Engineering*, **43**, 1777-1786. <https://doi.org/10.1117/1.1766301>