Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Fabrication-constrained tandem inverse design of photonic crystal fiber SPR sensors with uncertainty-aware Pareto screening

Md Shahanur Islam Shagor

smshagor.ru@gmail.com

Voronezh State University of Forestry and Technologies https://orcid.org/0009-0003-7730-3202

Research Article

Keywords: photonic crystal fiber, surface plasmon resonance, inverse design, tandem neural network, conformal prediction, out-of-distribution detection, fabrication constraints, multi-objective optimization

Posted Date: September 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10905069/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Fabrication-constrained tandem inverse design of photonic crystal fiber SPR sensors with uncertainty-aware Pareto screening

Md Shahanur Islam Shagor

Department of Computer Science and Microelectronics Engineering,Voronezh State University of Forestry and Technologies, Voronezh, Russia

Corresponding author: shagormd.s.i@iunp.edu.rs

Abstract

Inverse design of photonic crystal fiber surface plasmon resonance (PCF-SPR) sensors is intrinsically non-unique: different geometries can produce similar spectral responses, while a numerically attractive solution may violate fabrication bounds or lie outside the region represented by the training data. This paper develops a fabrication-constrained tandem inverse-design methodology that treats these failure modes explicitly. A forward surrogate maps four physical geometry variables and analyte refractive index to wavelength sensitivity, figure of merit (FOM), and resonance wavelength. A conditioned inverse generator then produces a population of candidate geometries for a requested sensing target. Training is regularized by differentiable fabrication penalties, and post-generation screening combines conformal residual intervals, Mahalanobis out-of-distribution (OOD) distance, projection distance, and non-dominated Pareto ranking. The validation protocol keeps all refractive-index samples belonging to the same base geometry in one data split and recomputes sensitivity and FOM from complete fixed-geometry refractive-index sweeps rather than treating these relational quantities as independent row-wise observations. In a software-consistency benchmark comprising 100 synthetic geometries and five refractive-index points per geometry, the fixed-geometry analysis produced near-linear resonance shifts by construction. A three-seed ablation showed that fabrication regularization increased the mean raw generated-geometry validity from 79.0% to 88.3%, while mean inverse target-satisfaction R^2 changed from 0.830 to 0.786, exposing an explicit feasibility-accuracy trade-off. Baseline analysis further showed that resonance wavelength was easy to model in the synthetic generator whereas local sensitivity was not, motivating geometry-sweep-aware target definitions. The reported numerical values validate the computational methodology and its diagnostics; they are not presented as physical sensor-performance claims. The framework provides a mathematically traceable route from target specification to risk-aware candidate selection and independent physics-ready verification.

Keywords: photonic crystal fiber; surface plasmon resonance; inverse design; tandem neural network; conformal prediction; out-of-distribution detection; fabrication constraints; multi-objective optimization

1 Introduction

Photonic crystal fibers (PCFs) provide unusual control over optical confinement because their microstructured cladding can be engineered over a wide geometric range (Russell 2003). When a plasmonic metal is brought into the evanescent-field region of a guided mode, the phase-matching interaction between the guided optical mode and a surface plasmon polariton can produce a pronounced wavelength-dependent loss feature. Tracking the position and shape of this resonance is the basis of many surface plasmon resonance (SPR) sensing schemes (Homola 2008). PCF-SPR designs therefore combine two degrees of freedom that are individually powerful but jointly difficult to optimize: the modal structure of the fiber and the dispersive response of the plasmonic interface.

The conventional workflow is forward and iterative. A geometry is proposed, an electromagnetic solver is executed, the resonance curve is extracted, sensing metrics are calculated, and the geometry is modified. This cycle can be effective for a small number of parameters, but its cost increases rapidly when pitch, air-hole diameter, plasmonic thickness, channel size, analyte refractive index, material model, and wavelength sampling are considered together. Recent PCF-SPR studies have therefore used regression, explainable machine learning, and multi-objective optimization to reduce the number of expensive simulations required during design exploration (Ahammed and Hossain 2024; Khatun and Islam 2025; Li et al. 2025; Romeiro et al. 2026; Khatun and Islam 2026). Similar trends are visible across nanophotonics, where learned surrogates and inverse models are

increasingly used to replace or accelerate repeated electromagnetic solves (Peurifoy et al. 2018; Molesky et al. 2018; Gencal et al. 2026).

A forward surrogate alone does not solve the inverse problem. If the mapping from geometry to response is many-to-one, the formal inverse is not a single-valued function. This is a practical difficulty rather than a purely mathematical curiosity: a network trained to reconstruct one geometry for a response that can be produced by many geometries is forced to average incompatible solutions. Tandem learning addresses this difficulty by training a forward model first and then optimizing the inverse model through the frozen forward model, so that the generated geometry is judged by whether its predicted response matches the requested target rather than by whether it reproduces an arbitrary training geometry (Liu et al. 2018).

For sensor design, however, non-uniqueness is only one of several risks. An inverse model can return geometries outside the fabrication envelope. A surrogate can be confident far from the support of the training data. A candidate can match one metric while becoming unacceptable on another. Most importantly, sensitivity and FOM are not primitive row-wise observables when they are obtained from refractive-index sweeps: they are derived from how the resonance changes while the physical geometry remains fixed. Randomly splitting rows from the same geometry across training and validation partitions can therefore leak nearly identical geometric information into both sets and inflate apparent generalization.

This work formulates these issues as one connected design problem. The proposed CyberPhotonics-SPR methodology combines geometry-grouped refractive-index handling, a conditioned tandem inverse generator, differentiable fabrication regularization, conformal residual calibration, Mahalanobis OOD screening, Pareto ranking, and an independent fixed-geometry validation path. The contribution is not a claim that machine learning itself is new in PCF-SPR design. Rather, the contribution is the mathematical and computational integration of inverse generation, feasibility, uncertainty, distribution shift, relational sensing metrics, and evidence-aware verification within one reproducible workflow.

The paper has four specific objectives. First, it states the sensor and inverse-design quantities in physical units and defines which quantities are local, which are sweep-derived, and which are only ranking aids. Second, it derives the loss and screening functions used to generate and rank designs. Third, it evaluates failure modes with grouped baselines and a multi-seed fabrication-constraint ablation. Fourth, it separates software-consistency evidence from electromagnetic and experimental evidence. This separation is essential because a synthetic generator can be useful for validating code paths without establishing the sensitivity of a real PCF-SPR device.

Table 1 Selected recent work and the methodological gap addressed here.

Study	Physics/data source	Inverse design	Fabrication constraint	UQ/OOD	Multi-objective
Ahammed and Hossain (2024)	PCF-SPR simulation/ML	Yes	Not central	No	No
Khatun and Islam (2025)	PCF-SPR + ML/XAI	Primarily prediction/optimization	Not explicit	Interpretability	Limited
Li et al. (2025)	FEM + MLP	Optimization	Via search bounds	No	NSGA-II
Romeiro et al. (2026)	Simulation + hybrid ML	Optimization	Geometry ranges	No	Yes
Khatun and Islam (2026)	FEM + ML/XAI	Design optimization	Parameter ranges	Uncertainty analysis	Partial
Present work	Synthetic/software validation; physics-ready interface	Conditioned one-to-many tandem	Differentiable penalties + projection	Conformal + Mahalanobis OOD	Pareto risk ranking

2 Theoretical and computational formulation

The methodology is expressed in terms of a physical geometry vector and a sensing-response vector:

$$\mathbf{g} = [\Lambda,\ d/\Lambda,\ t_{\mathrm{Au}},\ r_c], \qquad \mathbf{m} = [S_\lambda,\ \mathrm{FOM},\ \lambda_{\mathrm{res}}] \tag{1}$$

Here Λ is the PCF pitch, d/Λ is the normalized air-hole diameter, t_Au is the gold-film thickness, r_c is the sensing-channel radius, n_a is the analyte refractive index, S_lambda is wavelength sensitivity, FOM is the wavelength-domain figure of merit, and lambda_res is the resonance wavelength. The forward and inverse relationships are written as

$$\mathbf{m} = F(\mathbf{g}, n_a), \qquad \mathbf{g} = F^{-1}(\mathbf{m}, n_a) \tag{2}$$

The inverse in Eq. (2) is set-valued in practice. The method therefore seeks a family of feasible geometries whose forward responses lie near the requested target rather than a unique geometric reconstruction.

2.1 Material-dispersion relations

The software retains explicit material relations so that physical regularization and future solver coupling use consistent units. Fused-silica dispersion follows the three-term Sellmeier equation reported by Malitson (1965):

$$n^2_{\mathrm{SiO_2}}(\lambda) = 1 + \sum_{j=1}^{3} \frac{B_j \lambda^2}{\lambda^2 - C_j} \tag{3}$$

with B = [0.6961663, 0.4079426, 0.8974794] and C = [0.0684043^2, 0.1162414^2, 9.896161^2] micrometre squared when wavelength is expressed in micrometres. A compact Lorentz-Drude representation is used for the gold dielectric response as a regularization aid, consistent in form with established optical-constant models for noble metals (Rakic et al. 1998):

$$\varepsilon_{\mathrm{Au}}(E) = \varepsilon_\infty - \frac{f_D \omega_p^2}{E(E + i\gamma_D)} + \sum_j \frac{f_j \omega_p^2}{E_j^2 - E^2 - i\gamma_j E} \tag{4}$$

The compact dispersion model is not treated as a substitute for measured optical constants or a high-fidelity electromagnetic model. In the present work it provides physically meaningful spectral regularization and a consistent interface for later COMSOL or other finite-element verification.

2.2 Resonance extraction, FWHM, sensitivity, and FOM

For a sampled confinement-loss spectrum L(lambda), the discrete maximum is first identified. A local three-point quadratic fit reduces wavelength-grid bias. With x = lambda - lambda_i around the discrete peak, the local approximation and vertex are

$$L(x) = ax^2 + bx + c, \qquad \lambda_{\mathrm{res}} = \lambda_i - \frac{b}{2a} \tag{5}$$

The refinement is accepted only when the fitted parabola opens downward and its vertex remains inside the neighbouring samples; otherwise the discrete peak is retained. This prevents sub-grid interpolation from becoming extrapolation.

FWHM is measured relative to the local baseline rather than zero loss. If L_base is the minimum represented loss and L_peak the resonance maximum, the half-height level and width are

$$L_{1/2} = L_{\mathrm{base}} + \frac{L_{\mathrm{peak}} - L_{\mathrm{base}}}{2}, \qquad \mathrm{FWHM} = \lambda_R - \lambda_L \tag{6}$$

Left and right half-height crossings are linearly interpolated between adjacent wavelength samples. For a fixed geometry, wavelength sensitivity and FOM are then

$$S_\lambda \approx \frac{\Delta\lambda_{\mathrm{res}}}{\Delta n_a}, \qquad \mathrm{FOM} = \frac{|S_\lambda|}{\mathrm{FWHM}} \tag{7}$$

A key distinction in this work is that the independent validation layer does not reuse the local finite-difference sensitivity labels used during learning. Instead, it fits a line through the complete fixed-geometry refractive-index sweep and uses the least-squares slope

$$S_{\text{fit}} = \frac{\sum_i (n_i - \bar{n})(\lambda_i - \bar{\lambda})}{\sum_i (n_i - \bar{n})^2} \tag{8}$$

Linearity is summarized with the coefficient of determination

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2} \tag{9}$$

This creates an intentional separation between the learning target and the verification statistic. A design that looks acceptable on individual rows can still fail if the complete refractive-index sweep is nonlinear or internally inconsistent.

Table 2 Design variables and fabrication envelope used by the software.

Variable	Repository field	Unit	Lower bound	Upper bound	Role
Pitch, Λ	pitch_um	µm	0.80	4.00	PCF lattice scale
Air-hole ratio, d/Λ	d_over_lambda	-	0.20	0.90	Normalized hole diameter
Gold thickness, t_Au	metal_thickness_nm	nm	15	80	Plasmonic coating
Channel radius, r_c	channel_radius_um	µm	0.20	1.50	Analyte interaction region

3 Fabrication-constrained tandem inverse design

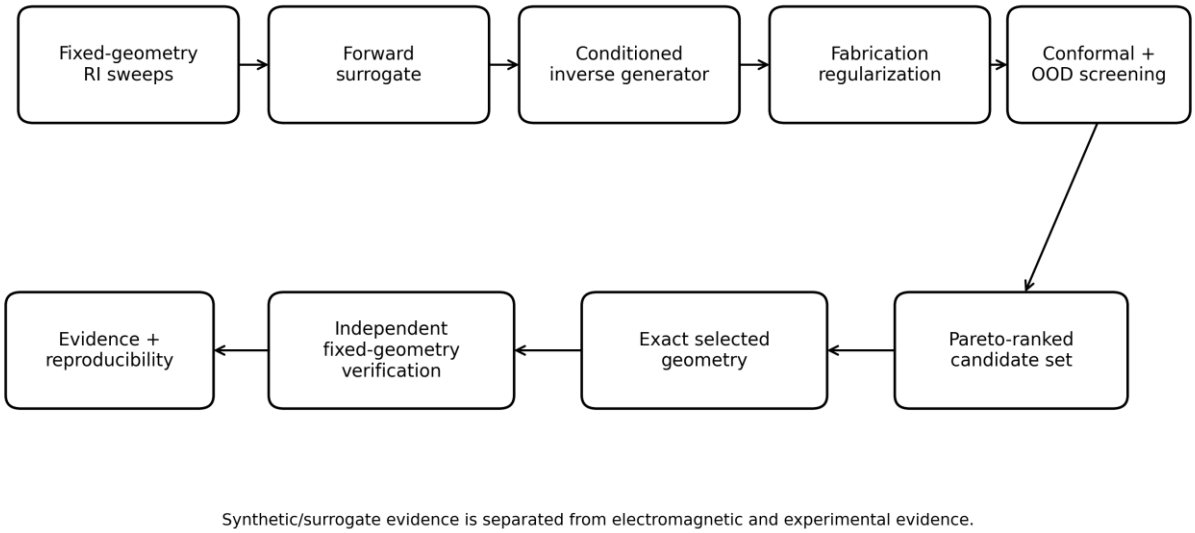


Fig. 1 End-to-end workflow. Geometry-grouped refractive-index data train a forward surrogate and a conditioned inverse generator. Candidate geometries are regularized for fabrication, screened for uncertainty and distribution shift, Pareto ranked, and then passed to an independent fixed-geometry verification path. The evidence gate keeps software-only validation distinct from electromagnetic or experimental evidence.

3.1 Leakage-resistant grouped data handling

Every row contains four geometry variables, one analyte refractive index, and three model targets. Rows are standardized before training, but the standardizers are fitted only on the training partition. The important constraint is that all refractive-index points belonging to one geometry remain in the same partition. Let G(g) denote the stable group identifier of a geometry. Then the split contract is

$$G(\mathbf{g}_i) = G(\mathbf{g}_j) \Rightarrow \text{split}(i) = \text{split}(j) \tag{23}$$

The release implementation uses a group-aware shuffle split. This is stricter than an ordinary random row split because a five-point refractive-index sweep cannot be divided between training and validation. The same geometry would otherwise appear in both partitions with only n_a changed, creating a particularly severe form of leakage for resonance trends and any derived sensitivity target.

3.2 Standardization and forward surrogate

Geometry, analyte condition, and sensing metrics are standardized independently:

$$\tilde{x} = \frac{x - \mu_x}{\sigma_x} \tag{10}$$

The forward network receives five standardized inputs - four geometry variables and analyte refractive index - and predicts three standardized metrics. The release architecture is a fully connected multilayer perceptron with hidden widths 128, 128, and 64. Its training objective is mean-squared error in standardized metric space:

$$\mathcal{L}_f = \frac{1}{N} \sum_{i=1}^{N} \| f_\theta(\tilde{\mathbf{g}}_i, \tilde{n}_i) - \tilde{\mathbf{m}}_i \|_2^2 \tag{11}$$

The forward model is trained first. During inverse training its parameters are frozen, so the inverse generator cannot improve its objective by changing the surrogate itself. This separation is the defining tandem mechanism used to address non-unique inverse mappings (Liu et al. 2018).

3.3 Conditioned one-to-many inverse generator

The inverse generator receives the requested metric vector, analyte refractive index, and a latent vector. Candidate k is generated as

$$\mathbf{g}^{(k)} = G_\phi(\mathbf{m}^\star, n_a, \mathbf{z}^{(k)}), \qquad \mathbf{z}^{(k)} \sim \mathcal{N}(\mathbf{0}, \sigma_z^2 \mathbf{I}) \tag{12}$$

The release uses a four-dimensional latent vector. Sampling the latent space produces a candidate population rather than forcing the network to collapse the inverse mapping to one geometry. The resulting candidate is evaluated through the frozen forward surrogate. The inverse objective is the target mismatch plus fabrication-related regularizers:

$$\mathcal{L}_{\text{inv}} = \| f_\theta(G_\phi(\mathbf{m}^\star, n_a, \mathbf{z}), n_a) - \mathbf{m}^\star \|_2^2 + \alpha \mathcal{L}_{\text{ov}} + \beta \mathcal{L}_{\text{bd}} + \gamma \mathcal{L}_{\text{disp}} \tag{13}$$

The term L_disp is optional. In the release configuration used for the ablation reported below, the dispersion weight is zero; therefore the ablation is described as fabrication-regularized rather than as a physics-informed neural network experiment. This wording matters: no Maxwell-equation residual or electromagnetic boundary-condition residual is minimized during the reported tandem training.

3.4 Differentiable fabrication regularization

The air-hole diameter is d = Λ(d/Λ). A normalized overlap penalty activates only when the diameter reaches or exceeds the pitch:

$$\mathcal{L}_{\text{ov}} = \text{mean}\left[\text{ReLU}\left(\frac{d - \Lambda}{\max(\Lambda, \varepsilon)}\right)^2\right] \tag{14}$$

For a generic physical variable x with lower and upper limits l and u, the boundary penalty is

$$\mathcal{L}_{\text{bd}}(x) = \text{mean}\left[\text{ReLU}\left(\frac{l - x}{u - l}\right)^2 + \text{ReLU}\left(\frac{x - u}{u - l}\right)^2\right] \tag{15}$$

The total boundary penalty is summed over pitch, normalized air-hole diameter, gold thickness, and channel radius. The framework records both raw inverse outputs and projected geometries. This distinction prevents post-hoc clipping from hiding an inverse model that frequently generates invalid designs.

4 Uncertainty calibration, OOD analysis, and Pareto screening

4.1 Residual conformal intervals

A point prediction alone is not enough for selecting a design for expensive numerical verification. The framework therefore computes absolute residuals on a held-out group-aware partition and uses the finite-sample conformal quantile. For n calibration residuals, the targetwise half-width is

$$q_{1-\alpha} = \text{Quantile}_{\lceil (n+1)(1-\alpha) \rceil / n}(|y_i - \hat{y}_i|) \tag{16}$$

and the corresponding interval is

$$C_\alpha(x) = [\hat{y}(x) - q_{1-\alpha},\ \hat{y}(x) + q_{1-\alpha}] \tag{17}$$

Conformal prediction provides a distribution-free reliability statement under exchangeability assumptions (Vovk et al. 2005). In this implementation the interval half-width is augmented by ensemble disagreement when multiple forward members are available. Importantly, the later ranking score is not interpreted as a calibrated probability of design success.

4.2 Mahalanobis distribution-shift screening

A generated design can satisfy a neural surrogate while being far from the standardized geometry-condition region represented in training. To detect this condition, the framework constructs u = [g, n_a], estimates the training mean and a regularized covariance matrix, and computes the Mahalanobis distance

$$D_M(\mathbf{u}) = \sqrt{(\mathbf{u} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{u} - \boldsymbol{\mu})} \tag{18}$$

The distance is normalized by a high quantile of distances observed in the training reference set. A normalized value above one marks a candidate as outside the calibrated design domain. This use follows the broader principle of Mahalanobis-based OOD detection introduced for deep models by Lee et al. (2018), but here it is applied to standardized physical design variables and operating condition rather than classifier features.

4.3 Multi-objective candidate ranking

For each target, the candidate pool is evaluated with normalized target errors, fabrication-projection distance, and OOD distance. The minimization objective vector is

$$\mathbf{J} = [e_S,\ e_{\text{FOM}},\ e_\lambda,\ d_{\text{proj}},\ d_{\text{OOD}}] \tag{19}$$

Non-dominated ranks are assigned by Pareto dominance:

$$\mathbf{J}^{(a)} \prec \mathbf{J}^{(b)}: \quad J_j^{(a)} \le J_j^{(b)} \text{ for all } j, \text{ and } J_k^{(a)} < J_k^{(b)} \text{ for at least one } k \tag{20}$$

Within the resulting fronts, the implementation uses a composite ordering aid

$$\text{score} = \bar{e} + 0.25 d_{\text{proj}} + 0.20 \max(d_{\text{OOD}} - 1, 0) + 0.10(1 - c) \tag{21}$$

where e-bar is the mean normalized target error, d_proj is the normalized projection distance, d_OOD is the normalized Mahalanobis score, and c is the fraction of target metrics covered by the calibrated intervals. The

constants are ranking weights, not empirically calibrated utilities. Their purpose is to make the selection policy explicit and inspectable. The use of non-dominated sorting follows the general multi-objective principle exemplified by NSGA-II (Deb et al. 2002), although the repository computes Pareto ranks directly for the generated candidate pool rather than running an evolutionary optimizer.

5 Computational experiment

The numerical study in this manuscript is deliberately a software-consistency experiment. It tests whether the grouped data path, metric extraction, baselines, tandem optimization, and fabrication constraints behave as intended. It does not replace electromagnetic simulation. The synthetic generator used in the preparation benchmark contains a strong approximately linear refractive-index term; consequently, the resulting sensitivity magnitude is expected by construction and must not be interpreted as a predicted physical sensitivity of a fabricated sensor.

5.1 Synthetic protocol

Table 3 Software-consistency benchmark protocol.

Item	Value
Base geometries	100
Analyte refractive indices	1.33, 1.35, 1.37, 1.39, 1.41
Total rows	500
Wavelength samples per spectrum	256
Wavelength interval	350-950 nm
Primary random seed	7
Grouped split	By exact four-variable geometry identity
Tandem ablation seeds	7, 17, 29
Ablation epochs	50 forward + 50 inverse

For each base geometry, all five analyte conditions form a fixed-geometry sweep. The scientific validation path extracts resonance, FWHM, a global least-squares sensitivity, FOM, and RI-linearity statistics from the sweep. Bootstrap confidence intervals for mean sweep statistics are produced with deterministic resampling.

5.2 Baselines and ablation design

The repository baseline is ridge regression evaluated with the same grouped protocol. To understand whether the synthetic targets are linearly or nonlinearly learnable, the preparation study also evaluated Random Forest, Extremely Randomized Trees, and a scikit-learn multilayer perceptron under the same grouped split. The extension models are included as diagnostic baselines rather than claimed as tuned state-of-the-art competitors. Ridge regression follows Hoerl and Kennard (1970), Random Forest follows Breiman (2001), and Extra Trees follows Geurts et al. (2006).

The tandem ablation compares two settings across seeds 7, 17, and 29. In the no-penalty setting, the inverse generator minimizes target mismatch only. In the fabrication-regularized setting, the air-hole overlap and physical-boundary terms are enabled. The forward network is identical within each seed, so differences between the two inverse variants isolate the effect of fabrication regularization. Dispersion regularization is disabled in both variants.

5.3 Evaluation metrics

Model quality is summarized with R^2, RMSE, and MAE:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_i (y_i - \hat{y}_i)^2}, \qquad \mathrm{MAE} = \frac{1}{N}\sum_i |y_i - \hat{y}_i| \tag{22}$$

Because the three targets have different physical units and very different difficulty in the synthetic generator, target-wise R^2 is reported alongside the uniform-average score. Generated-geometry validity requires finite values, all four variables inside their fabrication intervals, and $d < \Lambda$ before post-processing projection.

6 Results

6.1 Fixed-geometry software-consistency checks

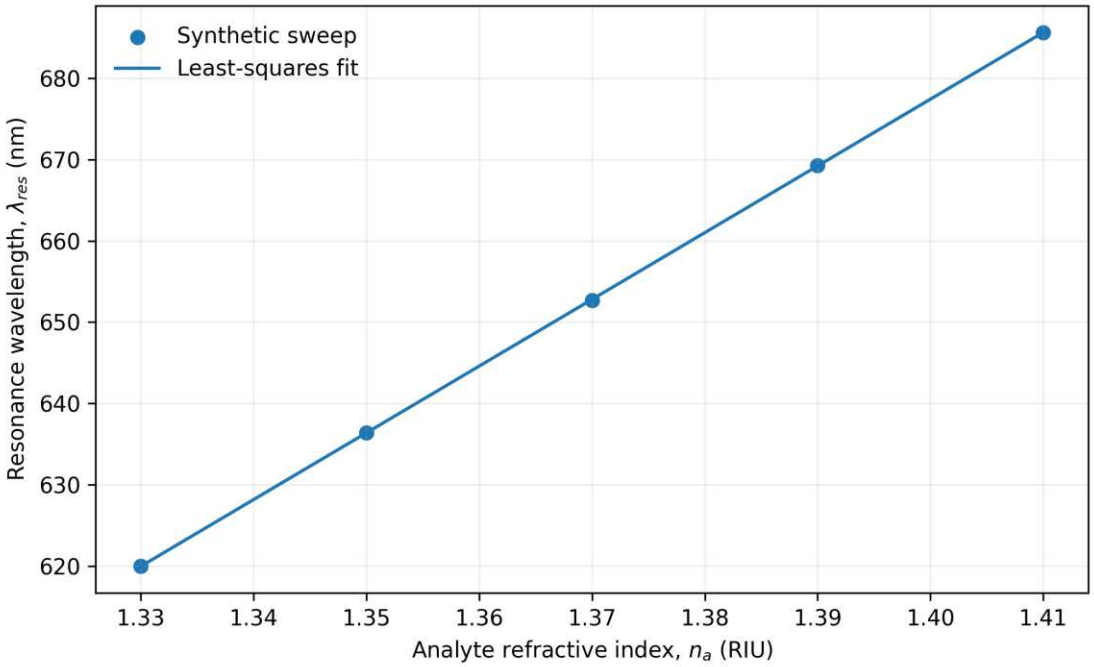


Fig. 2 Illustrative fixed-geometry synthetic refractive-index sweep and least-squares resonance fit. The near-linear slope demonstrates the validation calculation, not the performance of a physical sensor.

Across 100 fixed-geometry sweeps, the mean fitted sensitivity was 820.26 nm/RIU with a 95% bootstrap interval of 819.97-820.56 nm/RIU. The mean fitted FOM was 14.11 RIU^{-1} with a 95% bootstrap interval of 13.70-14.52 RIU^{-1}. Median linearity R^2 was 0.999993 and the minimum was 0.999955. Mean FWHM was 59.38 nm. The mean difference between the stored local sensitivity and the independent global-fit sensitivity was -0.14 nm/RIU.

These values are useful as internal consistency checks because the synthetic generator itself imposes an approximately 820 nm/RIU refractive-index coefficient. They therefore confirm that resonance extraction, grouping, global slope fitting, and FOM calculation are mutually consistent. They do not establish a PCF-SPR sensitivity. A physical claim would require a specified electromagnetic geometry, mesh, material data, boundary conditions, convergence study, and preferably experimental comparison.

Table 4 Fixed-geometry software-consistency summary (synthetic benchmark).

Quantity	Result	Interpretation
Geometry sweeps	100	Independent fixed-geometry groups
Mean fitted sensitivity	820.26 nm/RIU	Generator-consistency value; not physical performance

95% CI of mean sensitivity	819.97-820.56 nm/RIU	Bootstrap over geometry sweeps
Mean fitted FOM	14.11 RIU^{-1}	Derived from fitted slope and mean FWHM
95% CI of mean FOM	13.70-14.52 RIU^{-1}	Bootstrap over geometry sweeps
Median RI-linearity R^2	0.999993	Expected near-linearity of synthetic protocol
Minimum RI-linearity R^2	0.999955	No sweep failed linearity in this benchmark
Mean FWHM	59.38 nm	Synthetic spectral width

6.2 Grouped forward baselines expose target-definition asymmetry

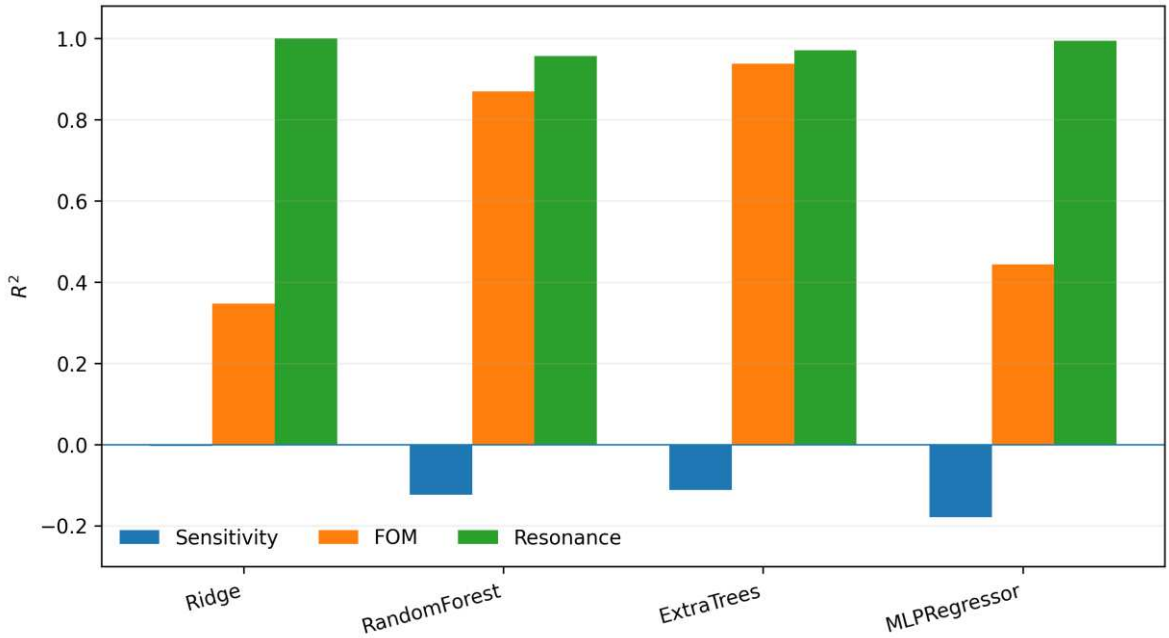


Fig. 3 Target-wise grouped R^2 for four diagnostic forward baselines. Resonance wavelength is easy to learn in the synthetic generator, whereas the row-wise local sensitivity target is poorly modeled by all baselines.

Table 5 Grouped forward-model baselines on the synthetic benchmark.

Model	Mean R^2	Sensitivity R^2	FOM R^2	Resonance R^2	Mean MAE
Ridge	0.448	-0.004	0.347	1.000	2.577
RandomForest	0.568	-0.124	0.870	0.957	5.112
ExtraTrees	0.598	-0.112	0.937	0.970	4.342
MLPRegressor	0.420	-0.179	0.444	0.994	3.627

Ridge regression achieved a mean R^2 of 0.448, but its resonance-wavelength R^2 was 0.99998 with a resonance MAE of 0.153 nm. The same model produced sensitivity R^2 = -0.004 and FOM R^2 = 0.347. Tree ensembles improved the FOM target substantially: Random Forest reached FOM R^2 = 0.870 and Extra Trees reached 0.937. Nevertheless, sensitivity R^2 remained negative for all four baselines. The diagnostic MLP also failed to make the row-wise sensitivity target reliable.

This result reveals a modeling issue that would be hidden by reporting only an average score. Resonance wavelength is a direct row-level property in the synthetic generator, whereas sensitivity is a relational quantity derived from multiple refractive-index points of the same geometry. Learning a local finite-difference sensitivity as though it were an independent row-level target injects derivative noise and endpoint effects. For a physics

campaign, a stronger formulation is to predict primitive spectral quantities such as resonance wavelength and FWHM for each n_a, then derive sensitivity and FOM from the complete predicted sweep. An alternative is to attach one global fitted sensitivity and FOM consistently to all rows of a geometry. The present framework keeps the independent global-fit validation specifically to expose this class of target inconsistency.

6.3 Fabrication regularization creates a measurable feasibility-accuracy trade-off

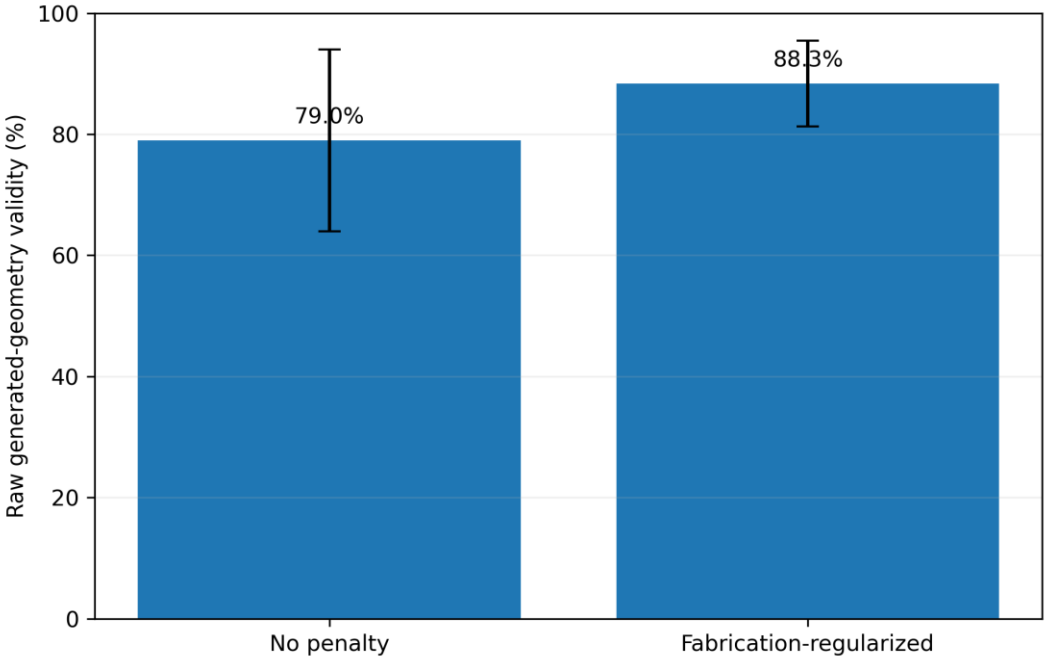


Fig. 4 Three-seed fabrication-constraint ablation. Enabling overlap and boundary penalties increases the mean fraction of raw inverse geometries that are valid before projection.

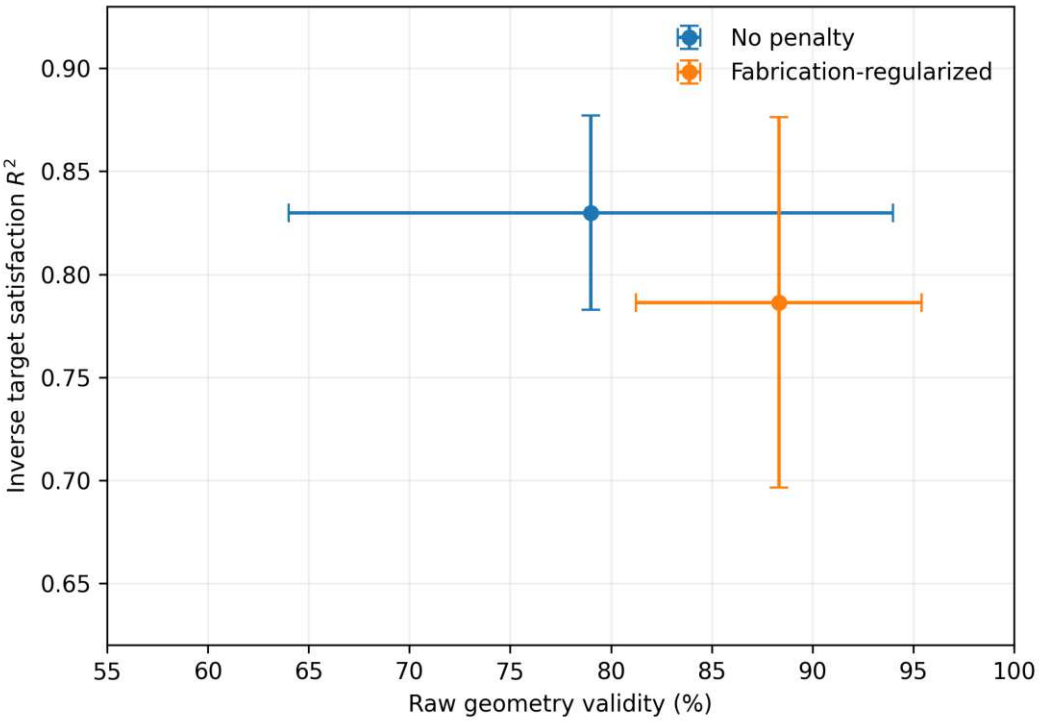


Fig. 5 Mean inverse target-satisfaction R^2 versus raw geometry validity for the same three-seed ablation. Error bars show the standard deviation across seeds and make the feasibility-accuracy trade-off explicit.

Table 6 Three-seed tandem ablation (50 epochs per stage).

Variant	Forward mean R^2	Inverse target R^2	Raw geometry validity
No penalty	0.590 ± 0.057	0.830 ± 0.047	79.0 ± 15.0%
Fabrication-regularized	0.590 ± 0.057	0.786 ± 0.090	88.3 ± 7.1%

The forward network is unchanged by the inverse-loss variant, and its mean R^2 across seeds was 0.590 ± 0.057. Without fabrication penalties, mean inverse target-satisfaction R^2 was 0.830 ± 0.047 and raw geometry validity was 79.0 ± 15.0%. With overlap and boundary penalties, validity increased to 88.3 ± 7.1%, while inverse target-satisfaction R^2 decreased to 0.786 ± 0.090. Thus the constraint term delivered a 9.3 percentage-point improvement in raw feasibility at the cost of a 0.044 reduction in mean target R^2.

This is the expected behavior of a real constraint rather than a free improvement. The unconstrained generator can use geometries outside the intended design envelope to reduce surrogate target error; the constrained generator sacrifices part of that freedom. The result argues against reporting only the post-projection geometry. If every candidate were clipped into bounds before evaluation, the 79.0% versus 88.3% distinction would disappear even though the underlying generators behave differently.

6.4 What the present results establish - and what they do not

The experiments establish five software-level properties. First, geometry-grouped splitting prevents the same base design from appearing in both training and validation. Second, the independent fixed-geometry slope calculation reproduces the synthetic generator trend and therefore validates the metric pipeline. Third, target-wise baselines expose a weakness in the row-level sensitivity definition that an aggregate metric could conceal. Fourth, fabrication penalties measurably change the raw validity distribution. Fifth, the uncertainty and OOD machinery is mathematically defined and can be attached to a candidate population before an expensive solver is called.

The experiments do not establish an electromagnetic optimum, a physical limit of detection, a fabricated-device sensitivity, or a deployment-speed claim. The repository contains COMSOL orchestration and hardware-facing code paths, but this manuscript does not promote mock or unverified outputs to physical evidence. That evidence boundary is intentional and is part of the proposed validation methodology.

7 Discussion

7.1 Relation to recent PCF-SPR machine-learning work

Recent literature confirms that machine learning for PCF-SPR design is no longer novel by itself. Ahammed and Hossain (2024) demonstrated inverse ML for a PCF-based SPR sensor. Khatun and Islam (2025) combined PCF-SPR optimization with regression and explainable AI. Li et al. (2025) integrated FEM, an MLP surrogate, and NSGA-II for multi-objective sensitivity/FOM optimization. Romeiro et al. (2026) combined simulation and machine learning for geometrical PCF-SPR optimization, while Khatun and Islam (2026) extended ML-XAI design with uncertainty-related analysis. Therefore, a defensible contribution must lie beyond the statement that an ML model predicts or optimizes PCF-SPR metrics.

The present framework differs in the order and coupling of the decisions. It begins with a group contract that treats a refractive-index sweep as one geometric unit. It then uses a latent conditioned inverse generator, not only forward regression, to represent non-unique design solutions. Fabrication penalties are evaluated before projection so that the raw generator can be audited. Candidate reliability is described with both calibrated residual widths and an explicit training-domain distance. Multi-objective ranking incorporates target mismatch, feasibility modification, and OOD risk. Finally, the chosen exact geometry is carried unchanged into a separate fixed-geometry verification path. The novelty is therefore the closed-loop methodology and its assumption-aware evidence discipline, not any single neural-network component.

7.2 Mathematical interpretation of the one-to-many design problem

Suppose two geometries g_1 and g_2 satisfy $\|F(g_1,n_a)-m^*\| < \varepsilon$ and $\|F(g_2,n_a)-m^*\| < \varepsilon$. A direct supervised inverse trained on one arbitrarily selected geometry is penalized for predicting the other even though both are valid solutions. In contrast, the tandem objective only requires $F(G(m^*,n_a,z),n_a)$ to return the target. The latent variable can therefore parameterize multiple points on or near the inverse solution manifold. Fabrication regularization then intersects that response-equivalent set with a feasible design region. OOD screening further restricts attention to the portion of the region for which the surrogate has empirical support. The Pareto stage avoids collapsing these distinct criteria into a single loss during training.

This layered formulation is useful because the criteria have different semantics. Fabrication bounds are hard engineering knowledge; target errors are optimization objectives; conformal intervals quantify calibration residuals; Mahalanobis distance is a support diagnostic; and Pareto rank is a decision rule. Treating them as separate objects makes it possible to inspect why a candidate was rejected rather than hiding all behavior in one neural score.

7.3 Limitations and required physics validation

The most important limitation is the evidence class. The benchmark data are synthetic, and the generator was constructed to produce a strong nearly linear refractive-index dependence. The numerical sensitivity and FOM values in Table 4 therefore cannot be compared with FEM or experimental sensitivity values from the literature as though they were device results. Their only role is to verify the software calculations and reveal modeling behavior.

Second, the release workflow uses a two-way grouped train/validation partition, while conformal calibration is fitted on the validation partition. A definitive predictive-performance paper should use three disjoint geometry groups: training, calibration, and final test. That change would allow empirical conformal coverage to be measured on data not used to estimate the conformal radius. Coverage should also be stratified by OOD status to test whether the Mahalanobis gate meaningfully separates interpolation from extrapolation.

Third, the forward model in the short three-seed study has only moderate average R^2 because the synthetic sensitivity target is poorly conditioned. The baseline evidence suggests reformulating the target representation before claiming predictive superiority. Fourth, the ablation uses only three seeds and a fixed 50-epoch schedule; a final model comparison should add early stopping, best-checkpoint selection, more seeds, and confidence intervals on model-level differences. Fifth, material dispersion in the regularizer is not equivalent to a Maxwell PDE residual, and the present method should not be described as a PINN in the Raissi et al. (2019) sense.

The next validation layer is therefore well defined. The exact Pareto-selected geometry should be solved with a verified finite-element PCF-SPR model using explicit boundary conditions, mesh-convergence checks, wavelength resolution, silica and gold material data, and a fixed analyte-RI sweep. The independent validation code should then recompute resonance, FWHM, sensitivity, FOM, and linearity directly from those spectra. Only those results should be used for physical sensor-performance claims. Experimental measurements, if available later, form a separate evidence tier again.

7.4 Reproducibility and software traceability

The implementation is released as CyberPhotonics-SPR under the Apache-2.0 license. Version 1.0.0rc1 is publicly tagged and distributed with source and wheel artifacts. The package targets Python 3.10-3.13 and uses NumPy, pandas, SciPy, scikit-learn, PyTorch, matplotlib, and PySide6, with optional interfaces for ONNX, COMSOL, explainability, and hardware communication. Reproducibility utilities capture seeds, configuration, hashes, and evidence metadata. These software mechanisms do not prove scientific validity, but they reduce ambiguity about which code and configuration produced a reported artifact.

8 Conclusion

This paper presented a mathematically explicit, fabrication-constrained tandem inverse-design methodology for PCF-SPR sensors. The method treats non-unique inverse mapping, fabrication feasibility, uncertainty, distribution

shift, and multi-objective selection as separate but connected parts of one design loop. Geometry-grouped splitting protects refractive-index sweeps from leakage, a frozen forward surrogate provides the tandem training signal, differentiable penalties guide the inverse generator toward feasible structures, conformal residuals quantify calibration uncertainty, Mahalanobis distance flags extrapolation, and Pareto ranking preserves trade-offs before one exact geometry is passed to independent verification.

The software-consistency experiments show why this separation is useful. The fixed-geometry validation correctly recovers the trend embedded in the synthetic generator; baseline models reveal that resonance is far easier to learn than the row-level local sensitivity target; and the three-seed ablation shows that fabrication regularization improves raw geometric validity from 79.0% to 88.3% while reducing inverse target R^2 from 0.830 to 0.786. The constraint therefore changes the feasible solution set rather than producing a free accuracy gain.

The central outcome is a research workflow rather than a physical sensor-performance record. The next step is to rerun the same exact-geometry pipeline using a verified FEM or COMSOL model on a disjoint train/calibration/test protocol, then compare the selected designs with conventional optimization under identical physics. Because the evidence boundary is explicit, those future results can extend the present methodology without conflating synthetic software validation with numerical-physics or experimental evidence.

Declarations

Funding: No external funding was received for this study.

Competing interests: The author declares no competing interests.

Ethics approval: Not applicable. The study concerns software and computational methodology and does not involve human participants or animals.

Consent to participate: Not applicable.

Consent for publication: Not applicable.

Author contribution: Md Shahanur Islam Shagor conceived the methodology, developed the software framework, performed the computational analysis, curated the reproducibility artifacts, prepared the figures, and wrote and reviewed the manuscript.

Data and code availability: The source code, documentation, and reproducibility workflow are publicly available in the CyberPhotonics-SPR repository (version 1.0.0rc1): https://github.com/smshagor-dev/CyberPhotonics-SPR . The numerical results reported here are software-consistency benchmark values derived from the synthetic protocol described in Section 5; no confidential or human-subject data were used. A stable archival DOI should be added to the final accepted manuscript when the manuscript-specific release is deposited.

Use of generative AI: During manuscript preparation, the author used OpenAI ChatGPT for drafting support, language structuring, and document formatting. The author independently reviewed the complete manuscript, verified the technical claims, calculations, references, and source-to-claim mappings, and takes full responsibility for the final content. No generative-AI images are used in this manuscript.

References

1. Ahammed T, Hossain MF (2024) Inverse design of a PCF-based SPR sensor using machine learning. In: 2024 IEEE International Conference on Computing, Applications and Systems (COMPAS). https://doi.org/10.1109/COMPAS60761.2024.10797137
2. Breiman L (2001) Random forests. Mach Learn 45:5-32. https://doi.org/10.1023/A:1010933404324
3. Chen N, Chang M, Lu X, Zhou J, Zhang X (2019) Photonic crystal fiber plasmonic sensor based on dual optofluidic channel. Sensors 19:5150. https://doi.org/10.3390/s19235150
4. Deb K, Pratap A, Agarwal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans Evol Comput 6:182-197. https://doi.org/10.1109/4235.996017
5. Gal Y, Ghahramani Z (2016) Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: Proceedings of the 33rd International Conference on Machine Learning, PMLR 48:1050-1059

6. Gencal H, Aksoy A, Yigit E, Aydemir U, Demirtas M et al (2026) A novel deep learning approach for accurate and efficient design of LNOI power splitters. Opt Quant Electron 58:239. https://doi.org/10.1007/s11082-026-08835-y
7. Geurts P, Ernst D, Wehenkel L (2006) Extremely randomized trees. Mach Learn 63:3-42. https://doi.org/10.1007/s10994-006-6226-1
8. Hoerl AE, Kennard RW (1970) Ridge regression: biased estimation for nonorthogonal problems. Technometrics 12:55-67. https://doi.org/10.1080/00401706.1970.10488634
9. Homola J (2008) Surface plasmon resonance sensors for detection of chemical and biological species. Chem Rev 108:462-493. https://doi.org/10.1021/cr068107d
10. Khatun MR, Islam MS (2025) Design optimization of high-sensitivity PCF-SPR biosensor using machine learning and explainable AI. PLOS ONE 20:e0330944. https://doi.org/10.1371/journal.pone.0330944
11. Khatun MR, Islam MS (2026) ML-XAI assisted design of PCF-SPR biosensors for refractive index-based analyte detection. Optik 350:172761. https://doi.org/10.1016/j.ijleo.2026.172761
12. Lee K, Lee K, Lee H, Shin J (2018) A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In: Advances in Neural Information Processing Systems 31
13. Li P, Wang H, Tong Z, Zhang W, Ma J (2025) Machine learning-assisted multi-objective optimization of photonic crystal fiber SPR sensors. Photonics Nanostruct Fundam Appl 67:101478. https://doi.org/10.1016/j.photonics.2025.101478
14. Liu D, Tan Y, Khoram E, Yu Z (2018) Training deep neural networks for the inverse design of nanophotonic structures. ACS Photonics 5:1365-1369. https://doi.org/10.1021/acsphotonics.7b01377
15. Malitson IH (1965) Interspecimen comparison of the refractive index of fused silica. J Opt Soc Am 55:1205-1209. https://doi.org/10.1364/JOSA.55.001205
16. Molesky S, Lin Z, Piggott AY, Jin W, Vuckovic J, Rodriguez AW (2018) Inverse design in nanophotonics. Nat Photonics 12:659-670. https://doi.org/10.1038/s41566-018-0246-9
17. Pan H, Pan F, Zhanga A, Cao C, Xue F et al (2022) Wide refractive index detection range surface plasmon resonance sensor based on D-shaped photonic crystal fiber. Opt Quant Electron 54:393. https://doi.org/10.1007/s11082-022-03805-6
18. Peurifoy J, Shen Y, Jing L, Yang Y, Cano-Renteria F, DeLacy BG, Joannopoulos JD, Tegmark M, Soljacic M (2018) Nanophotonic particle simulation and inverse design using artificial neural networks. Sci Adv 4:eaar4206. https://doi.org/10.1126/sciadv.aar4206
19. Raissi M, Perdikaris P, Karniadakis GE (2019) Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J Comput Phys 378:686-707. https://doi.org/10.1016/j.jcp.2018.10.045
20. Rakic AD, Djurisic AB, Elazar JM, Majewski ML (1998) Optical properties of metallic films for vertical-cavity optoelectronic devices. Appl Opt 37:5271-5283. https://doi.org/10.1364/AO.37.005271
21. Romeiro AF, Silva AO, Costa JCWA, Giraldi MTR, Guerreiro A, Santos JL (2026) Shaping photonic crystal fibers: geometric optimization for SPR sensor performance. Opt Quant Electron 58:105. https://doi.org/10.1007/s11082-025-08665-4
22. Russell P (2003) Photonic crystal fibers. Science 299:358-362. https://doi.org/10.1126/science.1079280
23. Shagor MSI (2026) CyberPhotonics-SPR, version 1.0.0rc1. GitHub repository. https://github.com/smshagor-dev/CyberPhotonics-SPR
24. Tsakyridis A, Moralis-Pegios M, Giamougiannis G, Kirtas M, Passalis N, Tefas A, Pleros NN (2024) Photonic neural networks and optics-informed deep learning fundamentals. APL Photonics 9:011102. https://doi.org/10.1063/5.0169810
25. Unal HS, Durgun AC (2025) A physics-aware neural network for effective refractive index prediction of photonic waveguides. Opt Quant Electron 57:107. https://doi.org/10.1007/s11082-024-08009-8
26. Vovk V, Gammerman A, Shafer G (2005) Algorithmic Learning in a Random World. Springer, New York. https://doi.org/10.1007/b106715