Available Online: 10 October 2026 DOI: 10.54254/2977-3903/2026.37571

Physics-validated defect-radius inversion in a square-lattice photonic crystal

Qishuo Wang

School of Information Engineering, Tianjin University of Commerce, Tianjin, China

2069928001@qq.com

Abstract. Defect-parameter inversion requires numerical error, mode-localization constraints, and the conditioning of the inverse map to be considered separately. A two-dimensional square-lattice photonic crystal containing a point defect is studied using TM eigenmodes calculated with MIT Photonic Bands. Supercell-size and grid-resolution tests confirm frequency convergence, while the central energy fraction, inverse participation ratio, and effective mode area restrict the reliable calibration interval to $0.02 \leq r_d/a \leq 0.14$, yielding 25 calibration samples. A PCHIP interpolant is compared with seven regression models. Nested leave-one-out cross-validation gives the lowest MAE for PCHIP ($6.42 \times 10^{-5}\ a$), while BP-MLP performs best among the machine-learning models ($1.47 \times 10^{-4}\ a$). Thirteen independent off-grid simulations confirm the accuracy of the leading nonlinear methods, although paired statistical tests do not distinguish BP-MLP, GPR, and SVR-RBF conclusively. Noise-propagation analysis further shows that practical inversion error is governed mainly by local inverse sensitivity, particularly near $r_d/a = 0.02$, where the forward frequency–radius slope is small. The resulting framework combines numerical convergence, physical mode screening, model comparison, and uncertainty analysis in a reproducible workflow for photonic-crystal parameter retrieval.

Keywords: photonic crystal, defect-radius inversion, MPB, BP neural network, mode localization

1. Introduction

Photonic crystals control electromagnetic waves through a periodic modulation of permittivity. The ideas of a photonic band gap and photon localization were introduced independently by Yablonovitch and John in 1987 [1, 2]. A point defect breaks the translational symmetry of an otherwise periodic crystal and can create a discrete state inside the gap. Such states underpin photonic-crystal cavities, filters, sensors, and lasers [3-5]. In a forward calculation, the geometry is known and the spectrum is computed. Device inspection and fabrication diagnosis often pose the opposite problem: a spectral feature is measured, while the underlying dimension has to be inferred.

Machine learning is now widely used in photonic inverse design [6-8], although a small regression error does not by itself establish a sound physical result. A frequency inside a nominal gap may belong to a folded, extended supercell band rather than to a localized defect state. Coarse plane-wave resolution can shift the band edges, and a supercell that is too small allows periodic images of the defect to interact. With few samples, a

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

single random train-test split can also give an optimistic estimate of performance. There is a simpler issue as well: for a smooth, strictly monotone one-dimensional relation, a conventional monotone interpolant may work as well as, or better than, a flexible learning model. These concerns are addressed here before any regression model is fitted.

The aim is not to propose a new learning algorithm. Instead, we put the full chain of evidence in one place: physical mode selection, numerical convergence, inverse modelling, independent validation, and uncertainty analysis. The first TM gap is established and the adequacy of a 7 × 7 defect supercell is checked. Localized modes are screened using the fraction of energy near the defect, the Inverse Participation Ratio (IPR), and the effective mode area. This leaves a 25-point calibration set and a separate set of 13 off-grid simulations. PCHIP is then compared with seven regression models under nested leave-one-out cross-validation, with additional attention paid to BP-MLP. Bootstrap resampling, frequency-noise propagation, BP initialization tests, learning curves, and local inverse sensitivity are used to show where the reported accuracy is reliable and where it is not.

2. Physical model and numerical verification

2.1. Two-dimensional square lattice

We consider an infinite two-dimensional square lattice of dielectric rods in air. The lattice constant a is the unit of length. Each bulk rod has relative permittivity 12 and radius $0.2a$. The relevant TM solutions have a nonzero out-of-plane electric component. MPB solves the magnetic-field form of the frequency-domain Maxwell eigenproblem in a plane-wave basis [9]. For a periodic, nonmagnetic medium, the eigenproblem is

$$\nabla \times [\frac{1}{\varepsilon(r)} \nabla \times H(r)] = (\frac{\omega}{c})^2 H(r) \tag{1}$$

Frequencies are reported as $f_n = \omega a/(2\pi c) = a/\lambda$. The bulk bands are sampled along Γ-X-M-Γ. A point defect is formed by changing the radius of the rod at the center of a square supercell to r_d; all other rods retain the bulk radius of $0.2a$.

Table 1. Physical model and final numerical settings

Item	Final setting	Reason
Lattice and polarization	2D square lattice; TM (E_z)	Target structure
Material	ε_{rod} = 12; air background	Lossless, nondispersive model
Bulk rod radius	r/a = 0.20	Reference photonic crystal
Bulk k path	Γ-X-M-Γ	Boundary of the first Brillouin zone
Bulk resolution	40 pixels/a	Converged band edges
Defect supercell	7 × 7	< 0.043% from 9 × 9
Defect resolution	40 pixels/a	< 0.024% from resolution 48
Tracked mode	TM band 49 at M = (0.5, 0.5, 0)	Unique isolated in-gap branch; frequency and field continuity checked
Calibration interval	$0.02 \le r_d/a \le 0.14$	Strong-localization screening
Training / validation	25 / 13 samples	On-grid / independent off-grid

The final settings are summarized in Table 1. The corner of the supercell Brillouin zone, M = (0.5, 0.5, 0), is used as a fixed operational sampling point so that every radius, resolution, and supercell calculation has

identical Bloch boundary conditions; the reported value is therefore the M-point supercell eigenfrequency rather than a full defect-band average. The mode is not identified from the integer label alone. At all 25 calibration radii, band 49 is the only eigenstate inside the bulk gap and varies smoothly and monotonically. Band 48 stays below 0.279868 and band 50 stays at or above 0.420482, leaving minimum separations of 0.013640 and 0.046356 from band 49. Representative field maps retain the same defect-centered symmetry and pass the localization test through r_d/a = 0.14. Together with the 7 × 7-to-9 × 9 convergence results, these checks provide no evidence of a crossing or mode switch within the calibration interval. A full k-path calculation would nevertheless be required to quantify any residual defect-band dispersion.

2.2. Bulk band gap and resolution convergence

At resolution 40, the first complete TM gap has lower and upper edges of 0.280874659 and 0.417137799. Its midgap-normalized width is 39.043%, as shown by the shaded region in Figure 1. The gap is used to identify candidate frequencies, not as proof of localization: an in-gap state can still be too extended to serve as a defect marker.

$$\Delta_{rel} = \frac{2(f_{upper} - f_{lower})}{f_{upper} + f_{lower}} \times 100\% \tag{2}$$

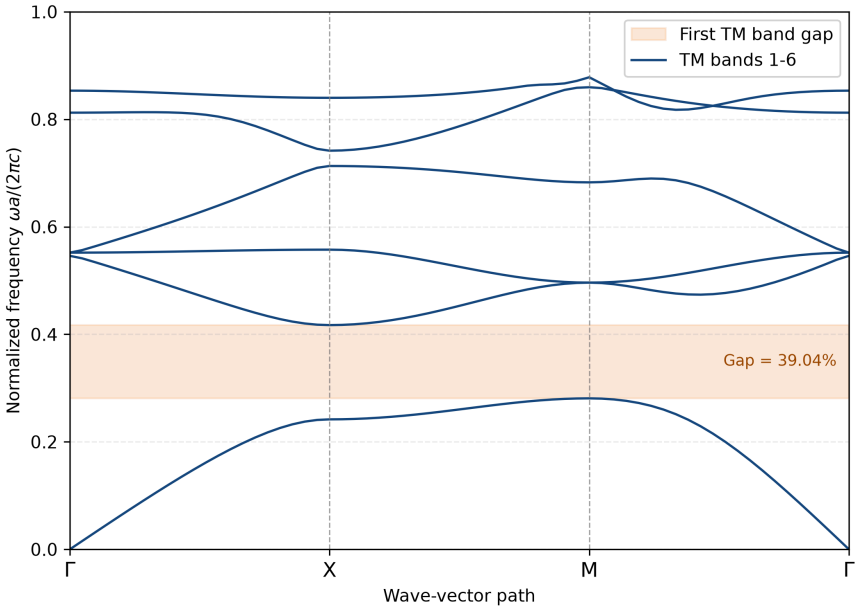


Figure 1. TM band structure of the complete square lattice at resolution 40. The shaded region marks the first TM band gap (0.280875-0.417138)

The bulk calculation was repeated at resolutions 20, 24, 32, and 40. Taking resolution 40 as the reference, the lower-edge errors at 20, 24, and 32 were approximately 0.237%, 0.080%, and 0.026%; the upper edge changed even less. Resolution 40 was retained for the final calculations so that the data set would be internally consistent.

2.3. Convergence with supercell size and defect resolution

Supercell convergence was tested at r_d/a = 0.07 and 0.14 using 5 × 5, 7 × 7, and 9 × 9 cells. At r_d/a = 0.07, the corresponding frequencies were 0.351315, 0.350595, and 0.350522. At r_d/a = 0.14, they were 0.293624, 0.292800, and 0.292675. Relative to the 9 × 9 reference, the 7 × 7 errors were 0.0208% and 0.0427%. The 7 ×

7 cell was therefore adopted as a practical compromise between accuracy and cost. A separate resolution test at r_d/a = 0.02, 0.07, and 0.14 gave changes of about 0.0086%, 0.0128%, and 0.0232% between resolutions 40 and 48. These differences are small enough to support resolution 40 for the radius sweep. Figure 2 collects the bulk-resolution, supercell-size, and defect-resolution convergence results.

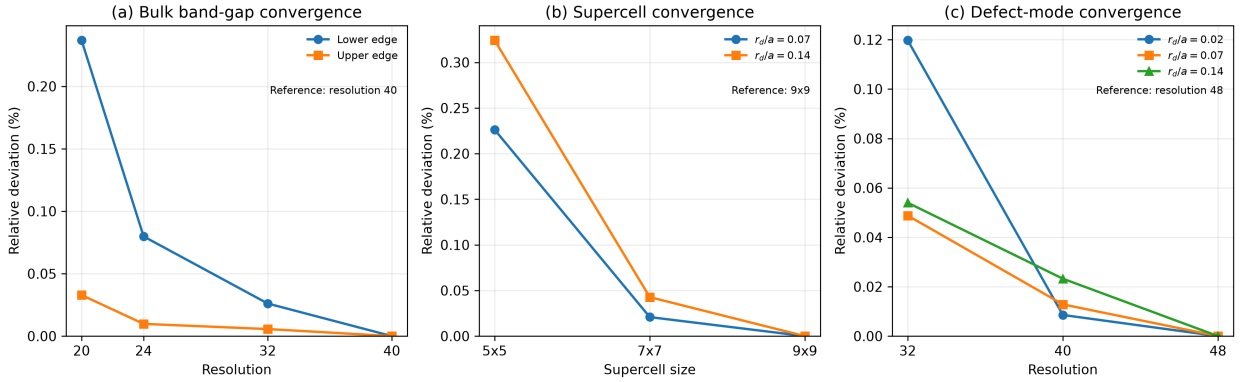


Figure 2. Numerical convergence summary: (a) bulk band edges versus resolution; (b) defect frequency versus supercell size; and (c) defect-frequency convergence at representative radii

2.4. Quantitative assessment of mode localization

The complex electric field is reconstructed from the real and imaginary MPB output arrays. In a lossless, nondispersive dielectric, the weight $\varepsilon|E_z|^2$ is proportional to the time-averaged electric energy density; the omitted constant cancels from all ratios. The metrics are invariant to an arbitrary common field amplitude, so no absolute modal normalization is required. Numerically, the weight is normalized over the full supercell after multiplication by the grid-cell area. For the 7 × 7 cell at resolution 40, the grid is 280 × 280 and the area element is included explicitly in both IPR integrals. The reported normalized IPR is the physical IPR multiplied by the supercell area. The central energy fraction, physical IPR, and effective mode area are defined by

$$\eta_R = \frac{\int_{|r-r_0|\le R} W(r)\,dA}{\int_{A_{cell}} W(r)\,dA}, \quad IPR = \frac{\int_{A_{cell}} W^2(r)\,dA}{[\int_{A_{cell}} W(r)\,dA]^2}, \quad A_{eff} = \frac{1}{IPR} \tag{3}$$

For this structure, strong localization is defined operationally by a primary threshold of $\eta_{1a} \ge 0.60$, supported by a defect-centered field maximum, an isolated continuous in-gap branch, and consistent IPR and effective-area trends. The central fraction is 0.632, 0.696, and 0.636 at r_d/a = 0.02, 0.07, and 0.14, respectively. It drops to 0.594 at 0.145 and 0.538 at 0.150; hence 0.14, rather than 0.145 or 0.150, is the upper boundary of the strong-localization interval. It falls further to 0.370 at 0.160 and 0.118 at 0.180, while the effective mode area grows from $0.447a^2$ at 0.145 to $1.319a^2$ and $5.850a^2$. At r_d/a = 0.20, the field maximum is displaced from the defect and η_{1a} is only 0.0599. Figures 3 and 4 show the accompanying field-pattern change and quantitative localization trends. We therefore use $0.02 \le r_d/a \le 0.14$ for calibration, classify 0.145-0.160 as a transition, and treat $r_d/a \ge 0.17$ as weakly localized or bulk-like. These thresholds are operational definitions for the present structure, not universal constants.

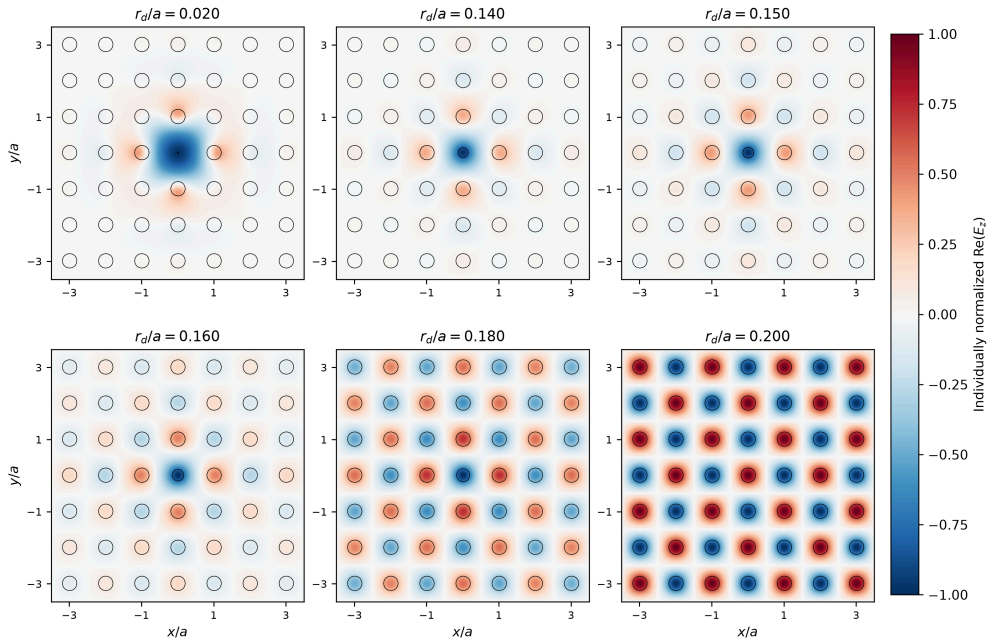


Figure 3. Evolution of E_z from a defect-centered localized state to an extended bulk-like pattern as r_d increases. Each panel is normalized separately; color should therefore be read as spatial structure, not as absolute amplitude across panels

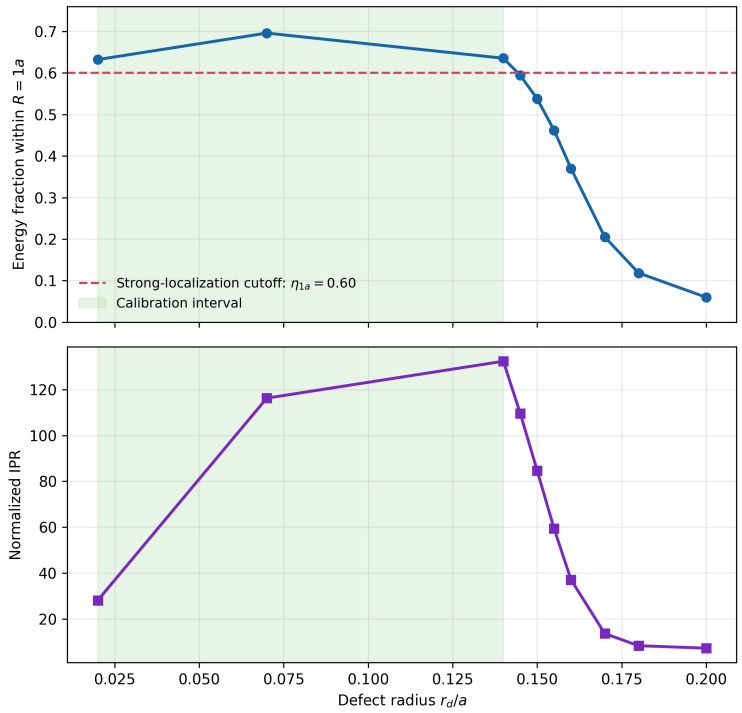


Figure 4. Localization measures as a function of defect radius, including the strongly localized, transition, and weakly localized ranges

3. Defect-radius inversion

3.1. Calibration and independent validation sets

The final calibration set contains 25 MPB simulations with r_d/a running from 0.020 to 0.140 in steps of 0.005. The tracked frequency of TM band 49 decreases strictly from 0.374126 to 0.292710. The missing-rod case $r_d = 0$ and the point $r_d/a = 0.01$ were not used for training. The former is physically a distinct endpoint, while the frequency difference between the latter and its neighbor is already comparable with the numerical discretization scale. Independent validation uses 13 radii that were not involved in model selection: 0.0225, 0.0325, ..., 0.1325, and 0.1375. Figure 5 shows the resulting strictly monotone calibration curve.

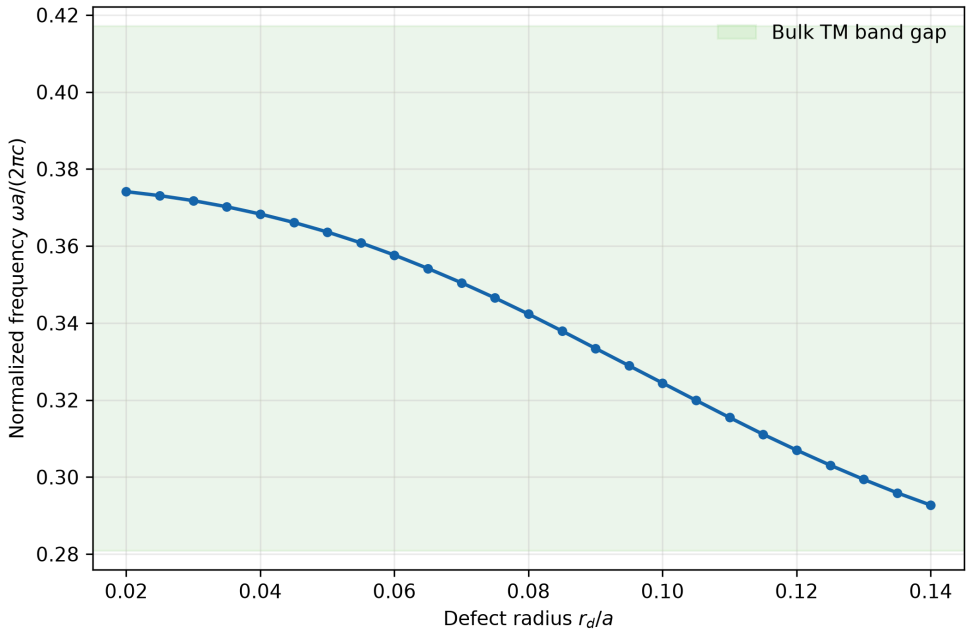


Figure 5. Radius-frequency relation for the final 7 × 7, resolution-40 data in the strongly localized interval. Strict monotonicity makes the inverse map single-valued

3.2. Comparison models and BP configuration

Normalized frequency is the sole input and r_d/a is the target. The comparison includes a Back-Propagation Multilayer Perceptron (BP-MLP) [10], RBF Support-Vector Regression (SVR) [11], Gaussian-Process Regression (GPR) [12], random forest [13], monotone Piecewise Cubic Hermite Interpolation (PCHIP) [14], second- and third-order polynomial regression, and linear regression. PCHIP is included deliberately because the present mapping is one-dimensional, smooth, and monotone, making a shape-preserving deterministic interpolant a relevant baseline. The machine-learning models were implemented in scikit-learn [15].

The BP network has one scalar input, one hidden layer, and one scalar output. Input and target scaling are refitted within each training fold to avoid leakage. Hidden-layer width, activation, and L2 regularization are selected in the inner validation loop. The most frequently selected setting uses five tanh units and $\alpha = 10^{-5}$, optimized with L-BFGS. Random seed 42 is used for the reproducible final fit. A small network is sufficient for this smooth one-dimensional relation; a deep architecture would be difficult to justify with 25 samples.

3.3. Nested validation and error measures

Generalization error is estimated by an outer Leave-One-Out Cross-Validation (LOOCV). In each outer split, BP and SVR hyperparameters are chosen by five-fold cross-validation using only the remaining 24 samples. Preprocessing is fitted on the same training data. The held-out point therefore has no role in model selection. The 25 outer predictions are pooled to compute MAE, RMSE, R^2, maximum absolute error, and MAPE.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|r_i - \hat{r}_i|, \quad RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(r_i - \hat{r}_i)^2}, \quad R^2 = 1 - \frac{\sum_{i=1}^{n}(r_i-\hat{r}_i)^2}{\sum_{i=1}^{n}(r_i-\bar{r})^2} \tag{4}$$

Several checks supplement the main comparison: BP stability over 50 random seeds; learning curves with 7, 9, 13, 17, 21, and 25 samples; Monte Carlo frequency-noise propagation; bootstrap intervals for off-grid MAE; paired Wilcoxon signed-rank tests; and local inverse sensitivity obtained from the monotone calibration curve. For the noise test, each of the 13 off-grid frequencies is multiplied by one plus a zero-mean Gaussian perturbation. Relative standard deviations of 0, 0.05%, 0.10%, 0.20%, and 0.50% are evaluated with 2,000 realizations per level using seed 20260905. Perturbed frequencies are not clipped: out-of-domain cases remain in the error calculation and their fraction is reported, while the practical prediction tool flags them as outside the calibrated range.

4. Results and discussion

4.1. Nested leave-one-out cross-validation

Table 2. Nested-LOOCV performance

Model	MAE (a)	RMSE (a)	R^2	Max. error (a)	MAPE (%)
PCHIP	0.0000642	0.0001247	0.999988	0.0005357	0.191
BP-MLP	0.0001469	0.0002475	0.999953	0.0009822	0.375
GPR	0.0001554	0.0002624	0.999947	0.0010037	0.449
SVR-RBF	0.0002063	0.0003883	0.999884	0.0015766	0.600
Polynomial-3	0.0012385	0.0015807	0.998078	0.0050435	2.771
Random forest	0.0028752	0.0032865	0.991692	0.0077900	5.234
Polynomial-2	0.0031174	0.0037604	0.989123	0.0098294	6.331
Linear	0.0040100	0.0050210	0.980607	0.0141737	9.153

Table 2 and Figure 6 summarize the nested-LOOCV results. PCHIP ranks first, as expected for a densely sampled, noise-free, monotone curve. BP-MLP ranks second overall and first among the machine-learning models. This result does not establish that a neural network is generally superior; it shows that a compact BP model can recover the present inverse relation accurately and offers a natural route to a future multi-input formulation. Polynomial and tree-based alternatives are less reliable near the interval boundaries.

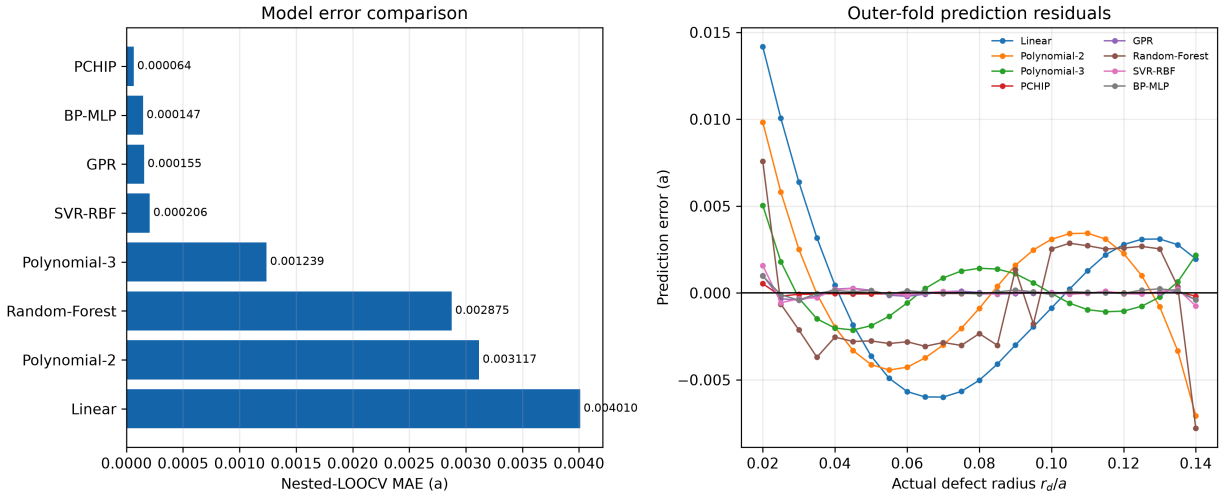


Figure 6. Nested-LOOCV comparison. Hyperparameters for BP-MLP and SVR were selected independently within each outer fold

4.2. Independent off-grid validation

Table 3. Independent off-grid validation results

Model	MAE (a)	RMSE (a)	R^2	Max. error (a)	MAPE (%)
PCHIP	0.00000823	0.00001342	0.99999987	0.00003195	0.0163
GPR	0.00005544	0.00008794	0.99999430	0.00021459	0.1207
BP-MLP	0.00006461	0.00007896	0.99999540	0.00015698	0.1037
SVR-RBF	0.00007785	0.00011822	0.99998969	0.00030145	0.2103
Polynomial-3	0.00097449	0.00117097	0.99898853	0.00244595	2.0785
Random forest	0.00135615	0.00147156	0.99840258	0.00188000	2.0532
Polynomial-2	0.00266950	0.00306974	0.99304873	0.00633802	5.1371
Linear	0.00360914	0.00445894	0.98533353	0.01075949	7.8515

Table 3 and Figure 7 compare the models on 13 independently simulated off-grid radii. All four smooth nonlinear methods reproduce these points closely. The largest BP error is 1.57×10^{-4} a, and no BP pointwise relative error exceeds 0.28%. PCHIP reaches an MAE of 8.23×10^{-6} a because every validation radius lies between closely spaced, noise-free calibration points. This value is an interpolation error within the same numerical model, not an estimate of experimental or absolute inversion accuracy. Converting the observed supercell and resolution shifts with the local inverse sensitivity gives equivalent radius shifts of approximately 5.9×10^{-5} to 2.1×10^{-4} a, which exceed the PCHIP off-grid MAE. The interpolation result must therefore be reported together with the numerical-convergence scale, physical domain, and assumed frequency noise.

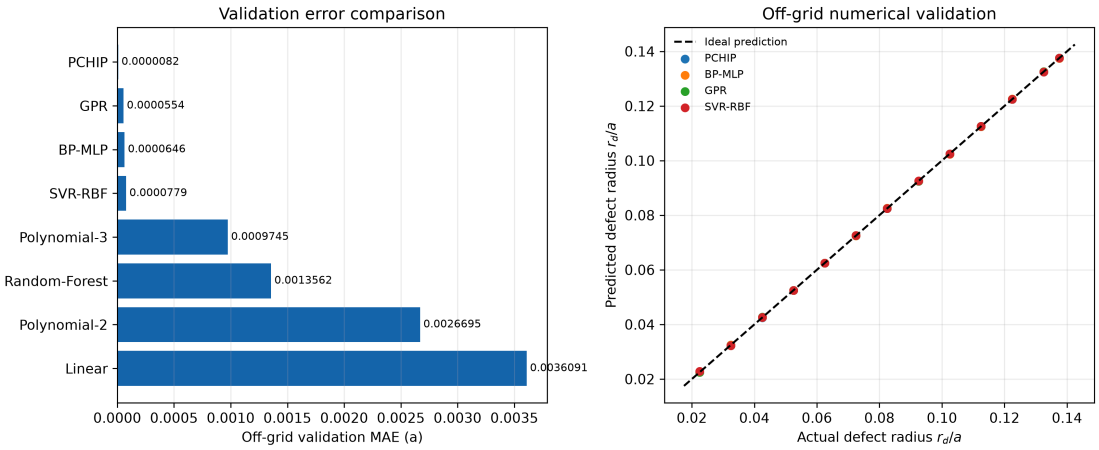


Figure 7. Predictions and absolute errors for the 13 independently simulated off-grid radii

4.3. Statistical comparison and BP stability

Percentile 95% confidence intervals were obtained from 20,000 paired bootstrap resamples of the 13 off-grid cases. They are [3.01 × 10^{-6}, 1.45 × 10^{-5}] a for PCHIP, [2.25 × 10^{-5}, 9.60 × 10^{-5}] a for GPR, [4.12 × 10^{-5}, 9.01 × 10^{-5}] a for BP, and [3.54 × 10^{-5}, 1.31 × 10^{-4}] a for SVR. Pointwise absolute errors were also compared with two-sided paired Wilcoxon signed-rank tests. Holm correction was applied across the three prespecified BP contrasts. The adjusted p value is 0.00732 for BP versus PCHIP and 1.000 for both BP versus GPR and BP versus SVR. Thus PCHIP has lower error on these noise-free interpolation points, whereas the present data do not distinguish BP, GPR, and SVR. Because only 13 paired cases are available, these significance results are treated as exploratory rather than definitive.

BP was also refitted from 50 random initializations. The mean, standard deviation, and median off-grid MAE were 1.12 × 10^{-4} a, 6.54 × 10^{-5} a, and 9.44 × 10^{-5} a. The 95th percentile was 1.75 × 10^{-4} a, and the observed range was 5.58 × 10^{-5} a to 3.87 × 10^{-4} a. Seed 42 gave 6.46 × 10^{-5} a. Reporting the distribution is important here: it avoids presenting a favorable run as typical and shows that initialization variability is not negligible relative to the noise-free interpolation error.

4.4. Noise robustness and inverse conditioning

Figure 8 shows that radius error grows almost linearly with the imposed Gaussian frequency-noise level. At 0.05% relative standard deviation, the mean MAE of the four leading methods lies between 2.51 × 10^{-4} a and 2.58 × 10^{-4} a. At 0.10%, it rises to 4.88 × 10^{-4}-4.99 × 10^{-4} a, and at 0.50% to 2.46 × 10^{-3}-2.53 × 10^{-3} a. Once input noise is present, most of PCHIP's advantage on exact data disappears. At the 0.50% level, 4.48% of perturbed frequencies fall outside the calibration domain; these extrapolated cases are included in the reported errors but must be flagged in practical use.

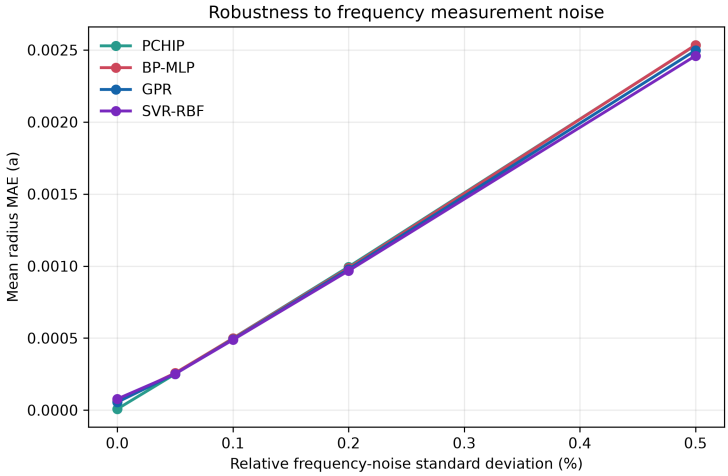


Figure 8. Mean radius MAE under relative frequency noise. PCHIP, BP-MLP, GPR, and SVR-RBF show nearly the same rate of error growth

Figure 9 relates this behavior to local inverse sensitivity. At r_d/a = 0.02, $|df/dr|$ = 0.1795 and $|dr/df|$ = 5.571, so a 0.1% relative frequency uncertainty corresponds to an estimated radius standard deviation of about 0.00208a. At r_d/a = 0.10, the estimate is only 0.000360 a. The small-radius, high-frequency end of the interval is therefore intrinsically ill-conditioned. Changing the regression algorithm cannot remove that physical limitation.

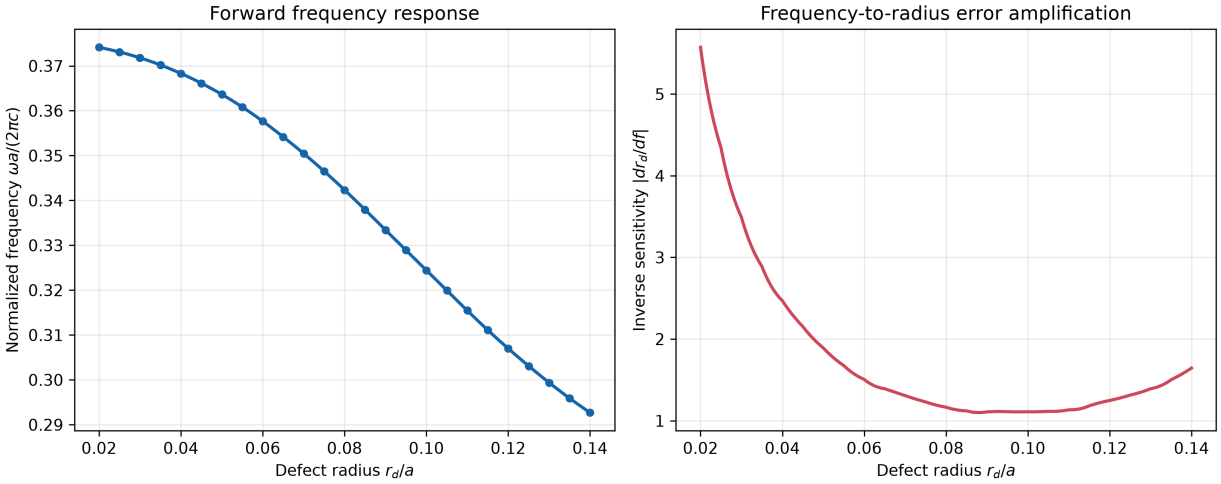


Figure 9. Forward frequency response and local inverse sensitivity across the calibration interval

4.5. Sample size and practical use

The learning curve changes most when the calibration set grows from 7 to 13 points. The median BP MAE falls from 1.95×10^{-4} a to 1.17×10^{-4} a and reaches 1.04×10^{-4} a with all 25 points. Improvement beyond 21 points is modest and partly obscured by variation among initializations. PCHIP continues to benefit from shorter local spacing and reaches 8.23×10^{-6} a with 25 points. Figure 10 shows that twenty-five samples are adequate for this one-parameter proof of concept, but they do not support broad claims about the high-dimensional generalization of neural networks.

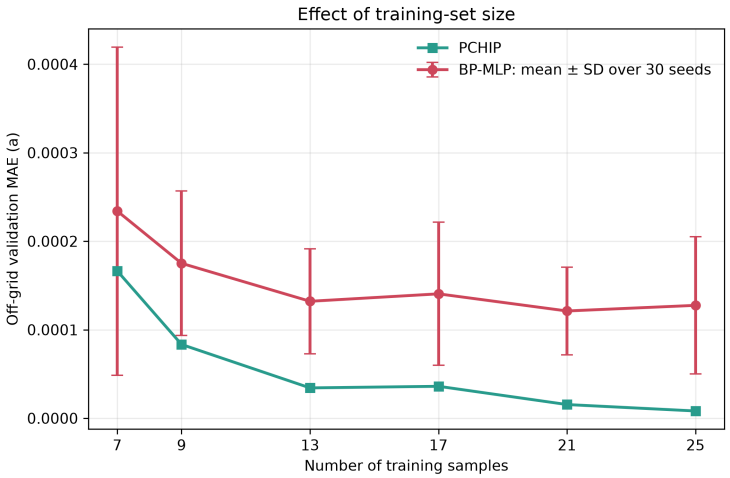


Figure 10. Off-grid learning curves for BP-MLP and PCHIP. The BP statistics include variability caused by random initialization

In use, the measured frequency should first be checked against the strongly localized calibration range, 0.292710-0.374126. Predictions from the leading models can then be reported together with their spread and a noise-propagated interval. For f = 0.348526, BP, GPR, SVR, and PCHIP return 0.072527 a, 0.072569 a, 0.072537 a, and 0.072501 a. With 0.1% frequency noise, the BP 95% interval is approximately [0.071656, 0.073382] a. This interval is much wider than the error on exact off-grid data and should determine the number of meaningful digits in the reported radius.

4.6. Limitations

The model is two-dimensional, lossless, and nondispersive, and only one TM mode is considered. Finite slab thickness, radiation loss, fabrication disorder, and material dispersion are outside the present calculation. Radius is also the only unknown, so the inverse relation remains single-valued. A multi-parameter problem may be spectrally non-unique and could require several modal features, field information, or a probabilistic output. The off-grid validation points are independent of model fitting, but they were generated by the same solver under the same physical assumptions; experimental validation is still needed. Supercell convergence was demonstrated at representative radii, not uniformly over every possible defect. Finally, the localization ranges are operational definitions for this structure and would have to be reassessed for another lattice or observable.

5. Conclusions

We have tested a complete route from a TM defect-mode frequency to the radius of a point defect in a two-dimensional square-lattice photonic crystal. The first TM band gap is 0.280875-0.417138, and a 7 × 7 supercell at resolution 40 gives converged defect frequencies for the cases examined. Quantitative field measures restrict the reliable calibration range to $0.02 \leq r_d/a \leq 0.14$; beyond this range, the tracked state gradually becomes weakly localized and bulk-like. Within the calibrated interval, several nonlinear inverse models are accurate. PCHIP is the best baseline for noise-free one-dimensional interpolation. BP-MLP is the strongest machine-learning model in nested LOOCV and is statistically comparable with GPR and SVR on the independent off-grid set. In practical retrieval, frequency noise and local inverse sensitivity matter more than

small differences in model ranking. The main outcome is therefore the traceable workflow itself: mode screening, convergence testing, model comparison, independent validation, and uncertainty analysis. The same framework can be extended to photonic-crystal inverse problems with more than one unknown parameter.

Data and code availability

The accompanying research archive contains the MPB control files and outputs, convergence tables, the 25-point calibration set, the 13-point off-grid validation set, localization metrics, Python analysis scripts, fitted models, and publication figures in PNG and PDF formats. The final modelling environment used NumPy 2.5.2, pandas 3.0.5, SciPy 1.18.1, scikit-learn 1.9.0, and joblib 1.6.0.

References

[1] Yablonovitch, E. (1987). Inhibited spontaneous emission in solid-state physics and electronics. *Physical Review Letters, 58*(20), 2059-2062. https://doi.org/10.1103/PhysRevLett.58.2059

[2] John, S. (1987). Strong localization of photons in certain disordered dielectric superlattices. *Physical Review Letters, 58*(23), 2486-2489. https://doi.org/10.1103/PhysRevLett.58.2486

[3] Joannopoulos, J. D., Johnson, S. G., Winn, J. N., & Meade, R. D. (2008). *Photonic crystals: Molding the flow of light* (2nd ed.). Princeton University Press.

[4] Painter, O., Lee, R. K., Scherer, A., Yariv, A., O'Brien, J. D., Dapkus, P. D., & Kim, I. (1999). Two-dimensional photonic band-gap defect mode laser. *Science, 284*(5421), 1819-1821. https://doi.org/10.1126/science.284.5421.1819

[5] Akahane, Y., Asano, T., Song, B.-S., & Noda, S. (2003). High-Q photonic nanocavity in a two-dimensional photonic crystal. *Nature, 425*(6961), 944-947. https://doi.org/10.1038/nature02063

[6] Molesky, S., Lin, Z., Piggott, A. Y., Jin, W., Vučković, J., & Rodriguez, A. W. (2018). Inverse design in nanophotonics. *Nature Photonics, 12*(11), 659–670. https://doi.org/10.1038/s41566-018-0246-9

[7] Peurifoy, J., Shen, Y., Jing, L., Yang, Y., Cano-Renteria, F., DeLacy, B. G., Joannopoulos, J. D., Tegmark, M., & Soljačić, M. (2018). Nanophotonic particle simulation and inverse design using artificial neural networks. *Science Advances, 4*(6), eaar4206. https://doi.org/10.1126/sciadv.aar4206

[8] Ma, W., Liu, Z., Kudyshev, Z. A., Boltasseva, A., Cai, W., & Liu, Y. (2021). Deep learning for the design of photonic structures. *Nature Photonics, 15*(2), 77–90. https://doi.org/10.1038/s41566-020-0685-y

[9] Johnson, S. G., & Joannopoulos, J. D. (2001). Block-iterative frequency-domain methods for Maxwell's equations in a plane-wave basis. *Optics Express, 8*(3), 173–190. https://doi.org/10.1364/OE.8.000173

[10] Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature, 323*(6088), 533–536. https://doi.org/10.1038/323533a0

[11] Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning, 20*, 273–297. https://doi.org/10.1007/BF00994018

[12] Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. MIT Press.

[13] Breiman, L. (2001). Random forests. *Machine Learning, 45*, 5–32. https://doi.org/10.1023/A:1010933404324

[14] Fritsch, F. N., & Carlson, R. E. (1980). Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis, 17*(2), 238–246. https://doi.org/10.1137/0717021

[15] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*(85), 2825–2830.