

Probe-Aware Small-Data Inverse Design and Experimental Benchmarking of Free-Space Second-Harmonic Generation in Au/ITO Hybrid Nanopillars

Po-Jui Chiang

pjchiang@nku.edu.tw

National Kaohsiung University of Science and Technology

Yin-Chi Liao

National Kaohsiung University of Science and Technology

Sheng-Xin Pan

National Kaohsiung University of Science and Technology

Shuo-Hsien Wang

National Kaohsiung University of Science and Technology

Research Article

Keywords: Au/ITO hybrid nanopillars, nonlinear nanophotonics, free-space second-harmonic generation, inverse design, exact-face finite elements, probe-offset sensitivity, experimental benchmarking.

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11239594/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Probe-Aware Small-Data Inverse Design and Experimental Benchmarking of Free-Space Second-Harmonic Generation in Au/ITO Hybrid Nanopillars

Po-Jui Chiang^{a,*}, Yin-Chi Liao^a, Sheng-Xin Pan^a, Shuo-Hsien Wang^a

^aNational Kaohsiung University of Science and Technology, Department of Electronic Engineering, Kaohsiung, Taiwan

Abstract. Small-data inverse design of nonlinear nanophotonic structures is difficult to trust, because surrogate labels, near-boundary field probes, and the nonlinear source model can each bias the result. We report a probe-aware workflow in which a surrogate trained on 35 samples only proposes candidate basins, and all five promoted candidates are evaluated by direct simulation without pooling inherited labels. Applied to Au nanopillars on an indium tin oxide (ITO) film, the workflow promotes a 200×70 nm pillar with a $1.25 \times$ larger total proxy and a $2.03 \times$ larger Au-ITO interface proxy than a 140×70 nm reference at a $0.92 \mu\text{m}$ pump. Representative free-space measurements give a corresponding SHG peak ratio of about 2.10 near 460 nm. Probe-offset tests identify which near-boundary metrics are diagnostic rather than label grade, and an independent exact-face, driven- 2ω finite-element cross-check closes the normalized energy balance to within 0.51% and preserves the ranking across 42 unit-norm relative interface-source cases. Because the pump is well separated from the estimated $1.40 \mu\text{m}$ ITO epsilon-near-zero wavelength, the design trend is assigned to localized plasmonic concentration and Au-ITO interface fields rather than to direct ENZ excitation. The combined evidence supports a normalized-relative geometry ranking, not a calibrated susceptibility, absolute conversion efficiency, minimum-order claim, or direct ENZ-enhancement claim.

Keywords: Au/ITO hybrid nanopillars, nonlinear nanophotonics, free-space second-harmonic generation, inverse design, exact-face finite elements, probe-offset sensitivity, experimental benchmarking.

*Po-Jui Chiang, pjchiang@nkust.edu.tw

1 Introduction

Nonlinear nanophotonic structures use optical resonances, field confinement, and engineered symmetry breaking to enhance weak nonlinear processes in subwavelength volumes. Plasmonic metasurfaces are especially attractive for second-harmonic generation (SHG), because metal nanostructures concentrate the fundamental field at surfaces, corners, gaps, and material discontinuities where inversion symmetry is naturally broken.¹⁻³ Recent reviews and demonstrations show that resonant antennas, coupled modes, lattice resonances, and double-resonance designs can strongly increase nonlinear signals.^{4,5} For early-stage design, however, the same strong localization that makes plasmonic SHG appealing also makes the computed objective sensitive to geometry, discretization, and where the near field is sampled.

ITO is often used as an epsilon-near-zero (ENZ) conducting oxide. Near its ENZ wavelength, the longitudinal field and phase response can be strongly reshaped, and ITO can exhibit unusually large optical nonlinearity.⁶⁻⁸ That mechanism is not invoked here: the $0.92\ \mu\text{m}$ pump is far from the estimated $1.40\ \mu\text{m}$ ENZ wavelength of the modeled ITO. We therefore describe the structure throughout as an Au/ITO hybrid and interpret the reported trend as localized plasmonic field concentration and contact-field redistribution. ENZ-region bulk nonlinearities remain relevant motivation for future wavelength-tuned studies, but they are outside the physical claim made here.

A second challenge is that the nonlinear source itself is difficult to define robustly in a small-data workflow. Surface SHG in centrosymmetric media is commonly described by effective surface nonlinear polarization terms.⁹⁻¹² In a numerical grid, however, the strongest fields often lie exactly where material assignment, staircasing, and dispersive boundary conditions are most delicate. A nominally simple change in nanopillar diameter, height, interface overlap, or off-boundary probe distance can therefore move the apparent nonlinear proxy. For this reason, a useful inverse-design target should not only be large; it should also reveal which part of the response is stable enough to use as a label and which part is merely diagnostic.

Machine-learning inverse design has become a standard tool for nanophotonic structures.¹³⁻²⁰ For nonlinear metasurfaces, however, the cost of generating labels is a bottleneck. Large datasets are hard to obtain, full 2ω source-and-radiation calculations are more expensive than linear screening, and a trained model can easily learn discretization artifacts rather than a physically transferable response. Adjoint methods are powerful for differentiable electromagnetic objectives,^{21,22} but a fully coupled nonlinear objective with calibrated material coefficients and far-field radiation is more involved than early screening often permits.

This motivates an intermediate regime between purely heuristic design and fully calibrated

nonlinear-device simulation. The workflow used here computes normalized linear finite-difference time-domain (FDTD) fields, post-processes them into a Rudnick-Stern-type relative surface-SHG proxy, separates shell and Au-ITO interface channels, trains a small surrogate for basin discovery, and promotes candidates only after direct FDTD validation. The surrogate is therefore used as a search aid, not as the final evidence for a device claim. The final evidence chain is deliberately conservative: direct labels, resolution gates, probe-offset sensitivity, and a representative free-space SHG benchmark must all point in the same direction before the promoted geometry is discussed as a useful candidate. In the experimental branch, fabricated arrays are measured in a free-space microscope geometry with monitored femtosecond excitation, filtered 2ω collection near 460 nm, and auxiliary linear reflection/scattering checks for alignment.

The present manuscript asks three specific questions. First, can a small inherited surrogate and a limited C2-A extension identify a stronger Au-ITO nanopillar basin without requiring a large new nonlinear dataset? Second, once such a basin is found, is the improvement carried by the total proxy, the absolute Au-ITO interface proxy, or by a fragile interface-fraction metric? Third, does a fabricated free-space SHG comparison follow the same ranking trend as the simulated channel proxy? These questions are intentionally narrower than an absolute conversion-efficiency claim, but they are directly relevant to how small-data nonlinear nanophotonic design can be made auditable.

The main result is a 200×70 nm Au nanopillar on ITO, denoted C2-A5, that is promoted by direct FDTD checks. At r_{110} and $0.92 \mu\text{m}$, it gives a $1.25 \times$ larger total proxy and a $2.03 \times$ larger non-fractional Au-ITO interface proxy than the 140×70 nm reference. In the representative experiment, fabricated 200×70 and 140×70 arrays are pumped at $0.92 \mu\text{m}$ and collected near 460 nm; the resulting $200 \times 70 / 140 \times 70$ SHG peak ratio is about 2.10, close to the simulated interface-proxy ratio. At the same time, the interface fraction is not stable: it changes substantially with resolution




Fig 1 Workflow overview. Project C supplies 35 inherited training samples for basin proposal only; five C2-A candidates are evaluated with direct labels and are not pooled into the inherited surrogate. The validation branch applies the mesh/probe audit and the stated 0.08 dex, 5×10^{-4} algebraic-residual, and 0.03 energy-residual gates before a normalized-relative claim is retained.

and probe definition, and is therefore rejected as a final optimization label. The contribution is thus not a record nonlinear metasurface, but an experimentally benchmarked, probe-aware validation chain for early Au/ITO nonlinear-nanophotonic design.

Figure 1 is the logical map for the rest of the paper. The upper screening row corresponds to the low-cost stage: generate normalized FDTD fields, post-process them into shell/interface SHG-proxy targets, train the small surrogate, and propose candidates. The lower validation row is the claim filter: selected candidates are rerun by direct FDTD, checked by resolution and probe-offset diagnostics, and then compared with a representative free-space SHG measurement. The current result follows the retained-basin branch of this workflow rather than the absolute-device branch: C2-A5 becomes the strongest r110 total and absolute-interface design, the representative SHG peak ratio follows the simulated interface-proxy trend, and the r130 $0.91 \mu\text{m}$ maximum is kept as a fixed-probe diagnostic because it fails the stricter probe-offset robustness test.

2 Device and Simulation Model

The primary screening device is a square-periodic Au nanopillar on a Drude ITO film over glass, with normal-incidence excitation from the air side. The principal candidate geometries discussed below are a promoted 200×70 nm Au pillar and the retained 140×70 nm and 220×50 nm diagnostic references, where the first number is diameter and the second is height. The ITO film thickness is 150 nm. The Au pillar is intentionally overlapped into ITO by 2 nm so that the voxelized model cannot form an accidental one-cell air gap at the metal-oxide contact. The nominal interface probes are placed 3 nm above and below the contact plane, and the sensitivity study varies this offset to 2 and 4 nm.

The material model uses a single-pole Drude form

$$\varepsilon_m(\omega) = \varepsilon_{\infty,m} - \frac{\omega_{p,m}^2}{\omega(\omega + i\gamma_m)}, \quad (1)$$

with Au parameters $\varepsilon_{\infty} = 9.5$, $\omega_p = 8.95$ eV, and $\gamma = 0.069$ eV from a Drude-only noble-metal fit.²³ The representative ITO model uses $\varepsilon_{\infty} = 3.9$, $\omega_p = 1.78$ eV, and $\gamma = 0.155$ eV, giving an ENZ estimate near $1.40 \mu\text{m}$. Accordingly, the $0.92 \mu\text{m}$ pump used below is not claimed to excite the ITO ENZ resonance itself; it is chosen because the Au-pillar/ITO unit cell shows a local plasmonic near-field feature and proxy maximum around $0.90\text{--}0.94 \mu\text{m}$. In this operating band, the simulated enhancement is attributed mainly to plasmonic field concentration and the Au-ITO surface/interface channel, while calibrated ENZ-region bulk nonlinearities remain outside the present model. All simulations use MEEP FDTD,²⁴ z -directed perfectly matched layers (PMLs),²⁵ periodic lateral

Table 1 Unit-cell parameter ranges and fixed validation settings.

Parameter	Value or range	Role
Au nanorod length L	100–300 nm	original 35-point LHS variable
Au nanorod width W	50–150 nm	original 35-point LHS variable
Au nanorod height H	30–80 nm	original 35-point LHS variable
Period P	400–800 nm	original 35-point LHS variable
Polarization angle θ	0–90 deg	original 35-point LHS variable
ITO film thickness	150 nm	C2 Au/ITO platform
Au-ITO overlap	2 nm	contact-stability setting
Interface probe offset	2, 3, 4 nm	sensitivity diagnostic
Shell offset	40 nm	near-surface monitor distance

boundaries, and a normally incident Gaussian source. The incident polarization is

$$\mathbf{E}_{\text{inc}}(t) = E_0 s(t)(\cos \theta \hat{\mathbf{x}} + \sin \theta \hat{\mathbf{y}}). \quad (2)$$

Across material interfaces the linear fields satisfy the standard tangential and normal boundary conditions

$$\hat{\mathbf{n}} \times (\mathbf{E}_2 - \mathbf{E}_1) = 0, \quad \hat{\mathbf{n}} \cdot (\mathbf{D}_2 - \mathbf{D}_1) = \sigma_f, \quad (3)$$

which motivates off-boundary field probes rather than sampling exactly on a dispersive-material discontinuity. Figure 2(a) shows the retained 140×70 nm reference geometry, and Figure 2(b) shows the 220×50 nm diagnostic geometry. In both panels, the shell monitors sample near-metal field concentration, while the paired Au-side and ITO-side probes define the interface channel used later for probe-offset sensitivity tests.

3 Surface-SHG Proxy Formulation

The nonlinear target is computed in a linear-plus-post-processing framework. Linear FDTD supplies complex fundamental fields on shell planes and Au-ITO interface probes. These fields are




Fig 2 Representative three-dimensional Au/ITO hybrid unit-cell architecture. (a) 140×70 nm and (b) 220×50 nm Au nanopillars on a 150 nm Drude ITO film. The figure marks the incident wave, polarization direction, Au-side and ITO-side interface probes, and shell monitor surfaces used for proxy extraction; the same model is used for the promoted 200×70 nm candidate.

then converted to a relative surface-SHG proxy using a local surface tensor. This is physically motivated by surface second-harmonic generation in centrosymmetric media,^{9–12} but it is not a full 2ω reradiation calculation.

For a monitor plane Γ_q , the field is decomposed into normal and tangential components,

$$E_{\perp} = \mathbf{E} \cdot \hat{\mathbf{n}}, \quad \mathbf{E}_{\parallel} = \mathbf{E} - E_{\perp} \hat{\mathbf{n}}. \quad (4)$$

Using normalized Rudnick-Stern-type weights, the local nonlinear surface polarization is approxi-

124 mated as

$$P_{\perp}^{(2)}(2\omega) = \chi_{\perp\perp\perp} E_{\perp}^2 + \chi_{\perp\parallel\parallel} (E_{\parallel,1}^2 + E_{\parallel,2}^2), \quad (5)$$

$$P_{\parallel,j}^{(2)}(2\omega) = 2\chi_{\parallel\perp\parallel} E_{\perp} E_{\parallel,j}, \quad j = 1, 2. \quad (6)$$

125 The shell-channel proxy at frequency index k is

$$J_k^{\text{shell}} = \sum_{q \in \mathcal{S}} \int_{\Gamma_q} \left(|P_{\perp}^{(2)}|^2 + |P_{\parallel,1}^{(2)}|^2 + |P_{\parallel,2}^{(2)}|^2 \right) dS, \quad (7)$$

126 and the Au-ITO interface channel is

$$J_k^{\text{int}} = \sum_{q \in \mathcal{I}} \int_{\Gamma_q \cap A_{\text{Au}}} \left(|P_{\perp}^{(2)}|^2 + |P_{\parallel,1}^{(2)}|^2 + |P_{\parallel,2}^{(2)}|^2 \right) dS. \quad (8)$$

127 The reported total target is

$$J_k = \log_{10} \left[\epsilon + J_k^{\text{shell}} + J_k^{\text{int}} \right], \quad (9)$$

128 where ϵ prevents numerical singularities. The corresponding diagnostic interface fraction is

$$f_k^{\text{int}} = \frac{J_k^{\text{int}}}{J_k^{\text{shell}} + J_k^{\text{int}}}. \quad (10)$$

129 Equations (7)–(10) define relative labels only. They omit coherent radiation, phase-matched far-field

130 summation, and bulk ENZ $\chi^{(2)}$ or higher-order nonlinearities.

131 **Algorithm 1. Surface-SHG proxy calculation**

132 **Require:** Normalized FDTD run, shell planes $\mathcal{S}$, interface probes $\mathcal{I}$, frequencies ω_k

```

133 Ensure: Total proxy vector  $\mathbf{J}$  and channel diagnostics
134   1: for each frequency index  $k$  do
135     2:   for each shell or interface monitor plane  $\Gamma_q$  do
136       3:     Load complex DFT fields  $(E_x, E_y, E_z)$ .
137       4:     Assign local normal vector  $\hat{\mathbf{n}}$  and decompose  $E_\perp, \mathbf{E}_\parallel$ .
138       5:     Compute  $P_\perp^{(2)}$  and  $P_{\parallel,j}^{(2)}$  using (5)–(6).
139       6:     Apply the plane mask: shell-box mask for sidewalls, circular footprint mask for interface
140         probes.
141       7:     Accumulate shell or interface integral by numerical quadrature over the monitor grid.
142     8:   end for
143   9:   Store  $J_k = \log_{10}[\epsilon + J_k^{\text{shell}} + J_k^{\text{int}}]$  and  $f_k^{\text{int}}$ .
144 10: end for

```

145 **4 Small-Data Surrogate and Inverse Design**

146 The small-data inverse-design stage uses a 35-sample Latin-hypercube dataset over the five geo-
 147 metric and polarization variables listed in Table 1. Table 2 makes the transfer boundary explicit.
 148 Project C is the original Au-nanorod screening branch and supplies only the standardized training
 149 coordinates, three relative proxy targets, and the KRR basin-proposal policy. C2 is a separate
 150 Au/ITO nanopillar validation branch comprising five proposed geometries. Its labels are generated
 151 by new direct FDTD runs; they are neither appended to Project C nor used to retrain the reported
 152 cross-validation model. Consequently, Project C errors quantify inherited basin-search performance,
 153 whereas every C2-A ranking claim rests on deterministic reruns, channel decomposition, and probe-
 154 offset diagnostics. For the surrogate stage, the input vector is standardized as $\mathbf{z}_i$, and the target

Table 2 Data provenance and transfer boundary. No C2-A response label is pooled into the inherited Project C surrogate.

Branch	Size		Origin and labels		Role in this work
Project C	35 points	LHS	Original branch;	Au-nanorod five geometric/polarization inputs and three normalized proxy targets	Train and cross-validate KRR for basin proposal only
C2-A	5 candidates		New Au/ITO nanopillar direct-FDTD runs with shell/interface decomposition		Deterministic ranking and promotion; not used to retrain Project C
Exact-face audit	2 geometries, 42 source cases each		Independent HCurl pump and driven- 2ω solves	p2 mixed- and driven-	Test normalized-relative ranking, energy closure, and radiative/absorptive partition

155 vector is $\mathbf{y}_i = [J_0, J_1, J_2]$. The radial-basis-function KRR model is

$$K_{ij} = \exp\left(-\frac{\|\mathbf{z}_i - \mathbf{z}_j\|^2}{2\ell^2}\right), \quad \mathbf{A} = (K + \lambda I)^{-1}Y. \quad (11)$$

156 The prediction at a new design is $\hat{\mathbf{y}}(\mathbf{z}) = \mathbf{k}(\mathbf{z})^T \mathbf{A}$. Five-fold cross-validation gives a KRR mean
157 absolute error of $0.493 \pm 0.123 \log_{10}$ units, compared with $0.557 \pm 0.150 \log_{10}$ units for degree-1
158 ridge regression and $0.930 \pm 0.311 \log_{10}$ units for degree-2 ridge regression. Figure 3(a) shows
159 that KRR is the most competitive inherited-dataset regressor by MAE/RMSE, but these metrics
160 are not used as C2 Au-ITO label uncertainties. Figure 3(b) shows the corresponding predicted-
161 versus-FDTD scatter: the spread is acceptable for basin discovery but too large for direct device
162 ranking without deterministic validation. Algorithm 2 therefore summarizes the validation-gated
163 candidate-promotion loop used after surrogate ranking.

Algorithm 2. Validation-gated inverse-design loop

165 **Require:** Initial dataset $\mathcal{D}_0$, candidate budget B , direct-FDTD gate, probe offsets Δz

166 **Ensure:** Accepted label-grade candidate set or diagnostic rejection




Fig 3 Inherited-dataset surrogate performance. (a) Five-fold mean absolute error (MAE) and root-mean-square error (RMSE) for the linear ridge and KRR models. (b) KRR predicted-versus-FDTD proxy values for the stored train/test split, used as a visual check rather than a standalone claim of model superiority. The MAE/RMSE statistics motivate basin discovery but are not C2 nanopillar label uncertainties.

- 1: Train KRR and ridge baselines on $\mathcal{D}_0$.
- 2: Rank unsimulated or locally refined designs by surrogate score.
- 3: Select the top B candidates and run direct normalized FDTD.
- 4: Compute shell and interface SHG proxy channels using Algorithm 1.
- 5: **if** flux sanity and proxy extraction pass **then**
- 6: Promote the candidate to label-grade analysis.
- 7: Repeat the interface-probe calculation at $\Delta z = 2, 3, 4$ nm.
- 8: **if** total proxy is stable and the physical interpretation does not rely on unstable interface fraction **then**
- 9: Retain the candidate as an absolute-response or diagnostic design.
- 10: **else**
- 11: Reject the candidate as a probe-definition artifact.

179 12: **end if**

180 13: **else**

181 14: Reject or rerun the simulation.

182 15: **end if**

183 The probe audit can also be incorporated directly into a future acquisition function instead of
184 being applied only after ranking. For design $\mathbf{x}$, resolution r , and probe offset δ , we propose

$$\begin{aligned} \mathcal{A}_{\text{robust}}(\mathbf{x}) = & \text{median}_{r,\delta}[\log_{10} J(\mathbf{x}; r, \delta)] - \lambda_{\delta} \left(\max_{\delta} \log_{10} J - \min_{\delta} \log_{10} J \right) \\ & - \lambda_r \left(\max_r \log_{10} J - \min_r \log_{10} J \right), \end{aligned} \quad (12)$$

185 subject to the existing 0.08 dex spread gate and numerical-residual gates. The first term rewards
186 a large typical response, while the two range penalties suppress designs whose rank depends on
187 a particular sampling plane or mesh. A conservative alternative is to replace the median by the
188 worst-case response. This formulation converts the present probe-offset diagnostic into an active
189 robustness criterion without treating the unstable interface fraction as a target.

190 **5 Results: Candidate Alignment and Channel Decomposition**

191 *5.1 Candidate Identification at r110 Own-Peak Points*

192 Candidate alignment is organized as an absolute-channel landscape rather than as an interface-
193 fraction optimization. A five-point C2-A screen identified a promoted 200×70 nm candidate,
194 hereafter C2-A5, which was then checked at r110 over 0.88–0.94 μm . At the nominal 3 nm
195 interface-probe offset, C2-A5 peaks at 0.92 μm with total proxy 1.0344, shell proxy 0.8148, and
196 absolute Au-ITO interface proxy 0.2195. A follow-up r130 wavelength gate over 0.90–0.94 μm

keeps this nominal $0.92\ \mu\text{m}$ point stable, with total proxy 1.0799 and a $+0.019\ \log_{10}$ change relative to r110. The same fixed-probe r130 sweep also resolves a larger neighboring $0.91\ \mu\text{m}$ value, with total proxy 2.6225, shell proxy 1.5435, and interface proxy 1.0790. That $0.91\ \mu\text{m}$ value is reported as a refined fixed-probe diagnostic row, not as a label-grade amplitude, because the probe-offset test below shows a $0.292\ \log_{10}$ total-proxy span at this wavelength. The previous 140×70 r110 reference at $0.90\ \mu\text{m}$ gives total proxy 0.8288 and interface proxy 0.1080, while the 220×50 diagnostic reference at $0.92\ \mu\text{m}$ gives total proxy 0.1963 and interface proxy 0.0240.

Thus C2-A5 has the largest r110 relative proxy in the current label set: its total proxy is $1.25\times$ higher than 140×70 , and its interface proxy is $2.03\times$ higher. The interface fraction tells a different, and less robust, story. In the r70 screen, C2-A5 had interface fraction 48.3%, but at r110 the same nominal 3 nm definition gives 21.2%. This drop reinforces the main diagnostic rule of the paper: exact interface fraction should not be used as a final optimization label. The robust claim is therefore narrower than simply maximizing a scalar proxy: direct FDTD exposes an r110 relative total-versus-interface channel landscape in which C2-A5 is the retained nondominated point, while higher-resolution fixed-probe extrema must still be checked for probe-offset stability before being promoted to label-grade amplitudes.

Figure 4(a) separates the two absolute contributions that make up the channel proxy at the r110 local-peak rows. This panel should therefore be read as an absolute shell-plus-interface comparison, not as a percentage plot. The C2-A5 200×70 nm candidate is stronger than the 140×70 nm reference because both terms remain favorable: the shell term stays high, and the absolute Au-ITO interface term increases from 0.1080 to 0.2195. The 220×50 nm diagnostic geometry has a useful local resonance for convergence testing, but Figure 4(a) shows that both its shell and interface terms are much smaller than those of C2-A5 at r110.

Figure 4(b) recasts the same r110 rows as a design landscape. The horizontal axis is the total proxy, and the vertical axis is the absolute interface proxy, so the upper-right direction is the desired direction when the goal is to improve total response without losing the interface channel. In this view, C2-A5 is the nondominated retained point: it has both the largest total proxy and the largest absolute Au-ITO interface proxy among the plotted candidates. The marker area in Figure 4(b) records nominal interface fraction only as a diagnostic. It is not the optimization target, because the r70-to-r110 drop in C2-A5 interface fraction already shows that fractional near-boundary quantities can be strongly resolution- and probe-definition-sensitive. The absolute-channel coordinates are more useful for ranking than a fraction, but they are not assumed to be automatically probe-invariant; the r130 C2-A5 check below explicitly tests that boundary.

Figure 5(a) then tests whether the C2-A5 channel ranking is a single-wavelength artifact. In this panel, line style encodes resolution tier: solid curves are r110 rows, and dashed curves are r130 wavelength-gate rows at the nominal 3 nm probe definition. Color and marker shape encode channel: circles show the shell proxy, and squares show the absolute Au-ITO interface proxy. With this convention, the C2-A5 r110 shell and interface curves both peak at the nominal $0.92\ \mu\text{m}$ point, while the r130 rerun resolves a larger neighboring $0.91\ \mu\text{m}$ fixed-probe point. This shift means the refined sweep identifies a $0.91\text{--}0.92\ \mu\text{m}$ spectral neighborhood, but it does not by itself promote the $0.91\ \mu\text{m}$ amplitude to a label-grade optimum. Importantly, the fixed-probe $0.91\ \mu\text{m}$ point is not produced by the interface channel alone: both shell and interface proxies increase at that wavelength. Figure 5(b) shows the opposite role for the 220×50 diagnostic geometry: its r110 and r130 spectra preserve a $0.92\ \mu\text{m}$ feature, but the absolute magnitudes remain far below C2-A5, so it is retained as a resolution diagnostic rather than as a winning geometry.




Fig 4 Channel decomposition at r110 local peaks. (a) Shell and absolute Au-ITO interface proxy contributions for the promoted C2-A5 200×70 nm candidate, the 140×70 reference, and the 220×50 diagnostic geometry. (b) The same r110 rows plotted as a total-proxy versus absolute-interface-proxy channel landscape. Marker area indicates nominal interface fraction, which remains diagnostic rather than a standalone optimization label.

Table 3 Candidate alignment and diagnostic rows. C2-A5 is the promoted r110 label-grade winner at 0.92 μm . The r130 wavelength gate keeps that nominal point within +0.019 $\log_{10}$ units and resolves a larger fixed-probe local maximum at 0.91 μm , but that r130 0.91 μm amplitude is retained as a diagnostic row because the probe-offset span is 0.292 $\log_{10}$ units. The 220×50 r130 row is retained as a resolution-gate diagnostic around its 0.92 μm peak.

Geometry	r px/ μm	λ μm	Total proxy	Shell proxy	Interface proxy	f^{int} %
C2-A5 200×70	110	0.92	1.0344	0.8148	0.2195	21.2
C2-A5 200×70	130	0.91	2.6225	1.5435	1.0790	41.1
C2-A5 200×70	130	0.92	1.0799	0.6601	0.4198	38.9
140×70	110	0.90	0.8288	0.7207	0.1080	13.0
220×50	110	0.92	0.1963	0.1724	0.0240	12.2
220×50	130	0.92	0.2419	0.1901	0.0518	21.4

5.2 Field-Distribution Evidence

Figure 6(a) visualizes the promoted C2-A5 winner using central MEEP DFT field slices rather than the proxy monitor planes used for numerical post-processing. Its sidewall and contact-region enhancement is broad, consistent with the larger absolute shell and interface proxies in Figure 4(a). Figure 6(b) shows that the 140×70 reference has a similar qualitative hotspot pattern but weaker absolute interface response. Figure 6(c) shows the 220×50 diagnostic geometry: the field remains




Fig 5 Frequency-resolved shell/interface channel decomposition. Line style encodes resolution: solid curves are r110 rows, and dashed curves are r130 wavelength-gate rows at the nominal 3 nm probe definition. Color and marker shape encode channel: circle markers are shell proxies, and square markers are absolute Au-ITO interface proxies. (a) C2-A5 200×70 nm over the 0.88–0.94 μm local basin. The r130 check keeps the 0.92 μm nominal pump point close to its r110 value but resolves a larger 0.91 μm fixed-probe local maximum that is treated as diagnostic until it passes probe-offset robustness. (b) The 220×50 diagnostic geometry over the matched spectral window.

resonant near the contact region, but the absolute proxy levels are lower, matching its diagnostic rather than winning role in Figure 4(b). Across all three panels, the maps are visual evidence for hotspot morphology, not a replacement for the quantitative proxy bars.

5.3 Resolution Hierarchy and r130 Wavelength Gate

The validation policy uses three levels. First, low and moderate resolutions identify a candidate basin. Second, r110 own-peak checks establish the main candidate comparison, including the promoted C2-A5 result. Third, the promoted C2-A5 candidate receives an r130 wavelength gate over 0.90–0.94 μm, and the 220×50 diagnostic candidate receives an r130 label gate at 0.91, 0.92, and 0.93 μm. For C2-A5, the nominal 0.92 μm row changes from 1.0344 at r110 to 1.0799 at r130, a +0.019 log₁₀ shift that is inside the ±0.08 log₁₀ amplitude gate. At the same nominal wavelength and 3 nm probe definition, however, Table 3 shows that interface fraction jumps from 21.2% at r110




Fig 6 Near-field visual evidence from central DFT field slices. Columns are treated as panels: (a) promoted C2-A5 200×70 nm candidate, (b) 140×70 reference, and (c) 220×50 diagnostic geometry. Within each column, the rows show $\log_{10}(|E|^2/I_{bg})$ in central x - z , central y - z , and ITO-side x - y sections. The overlays mark the Au pillar, ITO film, Au-ITO contact plane, and incident polarization direction.

to 38.9% at r130. This same-wavelength contrast is a useful internal control: the total proxy can pass the resolution amplitude gate while the fractional interface metric still moves substantially, reinforcing why interface fraction is retained only as a localization diagnostic. The r130 spectrum also resolves a larger neighboring $0.91\ \mu\text{m}$ fixed-probe point with total proxy 2.6225 and interface fraction 41.1%. This point is evidence for a nearby spectral feature, but it is not label-grade by amplitude because the later probe-offset check gives a $0.292\ \log_{10}$ total-proxy span, far above the $0.08\ \log_{10}$ gate. The 220×50 r130 check gives a peak at $0.92\ \mu\text{m}$ with total proxy 0.2419 and interface fraction 21.4%, while neighboring wavelengths give 0.1980 and 0.1522. This confirms that the 220×50 feature survives a higher-resolution local wavelength check, but it does not approach the r110 label-set response of C2-A5.

The convergence ladder in Figure 7(a) and Table 4 is retained as a cautionary control for the 220×50 diagnostic candidate. Figure 7(a) shows that the r70, r90, r110, and r130 spectra all place the local peak at $0.92\ \mu\text{m}$, so the resonance position is stable under refinement. Figure 7(b), however, shows that the amplitude still changes with resolution: at $0.92\ \mu\text{m}$, r70 is lower than r110 by 0.459 $\log_{10}$ units, r90 is lower by 0.165 $\log_{10}$ units, and r130 is higher than r110 by 0.091 $\log_{10}$ units. This is close to, but slightly outside, the $0.08\ \log_{10}$ amplitude gate used for the label check. Therefore, the paper treats the resonance location as robust but keeps the absolute amplitude and interface fraction in a resolution-qualified category.

6 Probe-Offset Sensitivity and Diagnostic Limits

Probe-offset sensitivity is the key diagnostic that motivates the paper's current claim boundary. Figure 8(a) compares 2, 3, and 4 nm off-boundary interface probes for the 140×70 and 220×50 candidates, each at its own r110 local peak. For 140×70 , the total proxy changes only from 0.790




Fig 7 Resolution ladder for the 220×50 diagnostic candidate. (a) r70, r90, r110, and r130 total-proxy spectra around the 0.92 μm local peak. (b) Successive-resolution changes at 0.92 μm relative to the r110 value and the $\pm 0.08 \log_{10}$ amplitude gate. The peak position is stable, while the amplitude remains resolution-qualified.

Table 4 Resolution hierarchy for the 220×50 diagnostic candidate at 0.92 μm . The r60 fine-scan row is a coarse frequency-scan reference; the r70–r130 rows define the resolution ladder and label-gate comparison.

Run	r px/ μm	Total proxy	$\log_{10}$ total proxy	f^{int} %	Δ vs r110 $\log_{10}$
r60 fine scan	60	0.3489	-0.4573	25.2	+0.250
r70 ladder	70	0.0682	-1.1664	8.1	-0.459
r90 ladder	90	0.1342	-0.8721	14.6	-0.165
r110 label	110	0.1963	-0.7070	12.2	0.000
r130 gate	130	0.2419	-0.6164	21.4	+0.091

to 0.903 across the offset range, corresponding to a 0.058 $\log_{10}$ span. For 220×50 , the total proxy changes from 0.186 to 0.220, corresponding to a 0.074 $\log_{10}$ span. These total-proxy changes are moderate for an early-stage label.

Figure 8(b) shows why the interface fraction is much less stable. For 140×70 , it changes from 8.8% to 20.2%; for 220×50 , it changes from 7.2% to 21.7%. The direction of the change is also tied to which side of the Au-ITO boundary is sampled. This confirms that exact interface fraction should not be treated as an optimization label in the current formulation. For these two r110 diagnostic

rows, the total proxy is comparatively more stable than interface fraction and is therefore more useful for early ranking. The paper does not assume that total proxy is always probe-invariant; the C2-A5 r130 check below is included precisely to test where that assumption breaks.

The promoted C2-A5 basin was then subjected to the same 2, 3, and 4 nm off-boundary probe test at r130. Figure 9(a) shows the total-proxy sensitivity for the two wavelength rows most relevant to the refined basin, 0.91 and 0.92 μm . At 0.91 μm , the total proxy spans 1.969–3.861, or 0.292 $\log_{10}$ units; at 0.92 μm , it spans 0.828–1.553, or 0.273 $\log_{10}$ units. Both spans are more than three times larger than the $\pm 0.08 \log_{10}$ amplitude gate used for resolution promotion. Figure 9(b) shows the corresponding interface-fraction sensitivity, with spans of 21.6–60.0% at 0.91 μm and 20.3–57.5% at 0.92 μm . The shell proxy is unchanged because the shell monitors and geometry are unchanged, but the interface channel increases strongly as the off-boundary probe moves farther into the ITO-side sampling region. Therefore, the r130 0.91 μm maximum is a useful fixed-probe warning about a nearby high-field feature, but it is not promoted to a label-grade amplitude. The retained label-grade statement remains the r110 0.92 μm candidate ranking and its representative experimental SHG ratio; the r130 C2-A5 rows define the next robustness problem rather than closing it.

7 Free-Space SHG Experiment and Proxy Comparison

Figure 10 shows the free-space optical geometry used for the SHG benchmark and auxiliary linear alignment. A tunable femtosecond laser in the 0.88–0.94 μm range is attenuated, monitored by a beam sampler and power meter, polarization controlled, expanded, and redirected by a 45° steering mirror through a 50 $\times$, NA 0.65 focusing objective. The sample is mounted on an XYZ piezo stage. The transmitted second-harmonic signal near 460 nm is collected by a separate




Fig 8 Probe-offset contrast at r110 own-peak points. (a) Total proxy relative to the 3 nm baseline. (b) Interface fraction versus off-boundary probe offset. Total proxy is comparatively stable, while interface fraction is strongly probe-definition dependent.

objective below the sample and routed through dichroic, short-pass, and 460 nm band-pass filters before spectrometer or photon-counting/imaging detection. A separate dashed branch denotes the CCD/photodiode alignment channel. No beam splitter is shared between the first-row pump path and the transmission-side SHG path.

With this measurement geometry fixed, Figure 11(a) summarizes the theory-to-experiment comparison used here as a free-space consistency check. The simulation predicts a $1.25\times$ increase in the total surface-SHG proxy for the promoted 200×70 nm candidate relative to the 140×70 nm reference, but a larger $2.03\times$ increase in the absolute Au-ITO interface proxy. The representative measured SHG peak ratio is 2.10, which is much closer to the simulated interface-channel ratio than to the total-proxy ratio. This agreement supports the interpretation that the measured free-space SHG contrast is dominated by the contact/interface hotspot channel, but it should not be overread as




Fig 9 C2-A5 r130 probe-offset sensitivity at the refined local-basin wavelengths. (a) Total proxy normalized to the 3 nm baseline for 0.91 and 0.92 μm . (b) Interface fraction versus off-boundary probe offset. The shell term is fixed by the unchanged shell monitors, while the interface channel and total proxy remain sensitive to the near-boundary sampling definition; the 0.91 μm fixed-probe maximum is therefore diagnostic rather than label-grade.

a calibrated conversion-efficiency prediction because the proxy omits absolute nonlinear coefficients and coherent 2ω radiation.

Figure 11(b) shows representative SHG spectra under 0.92 μm femtosecond excitation. Both arrays emit near the expected 460 nm second-harmonic wavelength, while the 200×70 nm array has the larger peak above the background. Figure 11(c) shows that the SHG intensity follows an approximately quadratic pump-power dependence over the displayed mW-scale range, consistent with a second-order response; the inset repeats the key comparison between the simulated $2.03 \times$ interface-proxy ratio and the measured $2.10 \times$ SHG peak ratio. Because this figure is a representative device-level check rather than a repeated-array statistics set, the slope uncertainty, fit residuals, and device-to-device spread are not used as quantitative labels in the present inverse-design workflow.

Figure 11(d) provides fabrication and measurement context. Top-view SEM images verify the two nominal lateral geometries, AFM maps give heights of 78 ± 6 nm for the 140×70 nm reference and 84 ± 6 nm for the 200×70 nm promoted array, and the cross-section SEM confirms Au

334 nanopillars on a 150 nm ITO film over glass. The optical micrograph indicates the array and pump
335 hotspot used for free-space collection. These images establish that the measured SHG comparison
336 was made on fabricated Au-ITO nanopillar arrays matching the simulation geometry at the level
337 needed for a representative theory-experiment check.

Free-space SHG measurement geometry

Pump delivery and transmission-side SHG collection are shown on physically separate paths





Fig 10 Free-space SHG and linear optical measurement geometry. The solid red path denotes pump delivery through a 45° steering mirror and focusing objective. The solid blue path is the physically separate transmission-side SHG collection path through a collection objective, dichroic, short-pass, and 460 nm band-pass filters. The dashed red branch denotes auxiliary CCD/photodiode alignment.



Pump: 0.92 μm , femtosecond excitation • SHG collection: 460 $\pm$ 10 nm • Objective: 50 $\times$, high-NA • Sample: Au nanopillar on 150-nm ITO / glass
 Note: Spectrum in (b) is representative of typical devices; individual spectra may vary slightly.

Fig 11 Free-space SHG experiment and proxy comparison. (a) Simulated total and absolute-interface proxy ratios for the promoted 200×70 nm candidate relative to the 140×70 nm reference, with the measured SHG peak ratio shown as a representative comparison. (b) Representative SHG spectra under 0.92 μm femtosecond excitation and 460±10 nm collection. (c) Pump-power dependence showing an approximately quadratic log-log trend, with an inset comparing simulated and measured ratios. (d) Sample verification images from top-view SEM, cross-section SEM, AFM, and optical microscopy. The panel is used as a representative free-space consistency check, not as a statistical device-to-device variation study.

8 Exact-Face Driven- 2ω Cross-Check

To test whether the FDTD proxy ranking survives an explicit second-harmonic field solve, we independently evaluated the 140×70 nm Reference and 200×70 nm C2-A5 geometries with a frequency-domain finite-element formulation. Periodic p2 HCurl pump fields drive split-domain p2 HCurl fields at 2ω through independent order-1 HDivSurface mortar spaces on the exact Au-air and Au-ITO faces. The normalized mixed system is

$$\mathbf{A}_{2\omega} \mathbf{u} + \mathbf{B} \boldsymbol{\lambda} = 0, \quad \mathbf{B}^T \mathbf{u} = \mathbf{g}, \quad (13)$$

where $\mathbf{g}$ contains unit-norm relative complex coefficients for the two interface-source families. This is a driven, undepleted-pump calculation: it includes propagation, material absorption, and upward/downward radiation at 2ω , but it does not include nonlinear back-action from the generated field onto the pump.

The audit spans 42 matched relative-source cases: Au-ITO/Au-air amplitude ratios 0, 0.25, 0.5, 1, 2, 4, and the Au-ITO-only endpoint, with eight relative phases for each nonzero finite ratio. These coefficients deliberately test ranking sensitivity and are not a calibrated complex $\chi_s^{(2)}$. Upward and downward fluxes are integrated on exact material-side faces, absorption is integrated in Au and ITO volumes, and weak-form source work provides the energy reference. Table 5 shows that every predefined numerical gate passes without relaxation.

Absorption remains the dominant normalized energy channel for both geometries, while C2-A5 has a larger radiative fraction over the ensemble (Figure 12). Across all 42 matched source cases, the C2-A5/Reference normalized source-response ratio ranges from 11.732 to 11.996, with median 11.784. Boundary-excluded mean-field changes are also positive in every reported region: the

Table 5 Exact-face numerical verification over the matched 42-case ensemble.

Quantity	Reference	C2-A5	Gate
Pump algebraic residual	4.585×10^{-4}	2.338×10^{-4}	$\leq 5 \times 10^{-4}$
2ω system residual	1.388×10^{-4}	4.920×10^{-5}	$\leq 5 \times 10^{-4}$
Maximum energy residual	0.00467	0.00504	≤ 0.03
Surface-quadrature change	5.45×10^{-14}	5.49×10^{-14}	≤ 0.002
Exact-face area error (μm^2)	1.72×10^{-15}	1.72×10^{-15}	$\leq 10^{-10}$

median changes are 0.616 dex in the Au core, 0.917 dex in ITO beneath the pillar, 0.956 dex above the cap, and 0.932 dex in the lateral air shell (Figure 13). Because the relative source coefficients are unit normalized rather than material calibrated, these values establish only the stability and direction of the geometry ranking. They are not predictions of the measured $2.10\times$ ratio, absolute SH power, or conversion efficiency.

Normalized energy partition across the 42-case interface ensemble

Distributions over unit-norm relative complex Au-air/Au-ITO coefficients




Fig 12 Normalized exact-face energy partition for 42 matched unit-norm relative complex Au-air/Au-ITO source-coefficient cases. Boxes span the interquartile range, horizontal lines mark the median, and whiskers show the minimum-to-maximum range. Absorption exceeds the total upward-plus-downward radiative fraction in every case. Median upward, downward, and absorbed fractions are 0.81%, 4.91%, and 94.54% for the Reference and 1.20%, 5.48%, and 93.59% for C2-A5, respectively. The maximum relative energy-balance residuals are 0.00467 and 0.00504, respectively, below the predefined 0.03 gate.

C2-A5 enhancement relative to Reference

Minimum, median, and maximum across the matched 42-case coefficient ensemble




Fig 13 C2-A5 enhancement relative to the Reference in the exact-face driven- 2ω audit. Markers are medians and whiskers show the full range over 42 matched unit-norm relative source cases. All quantities are normalized-relative; no calibrated surface susceptibility or absolute-power inference is made.

9 Discussion

The present work extends the earlier Project C methodology from a generic surface-shell SHG proxy to an Au/ITO setting with explicit interface channels, a representative free-space SHG benchmark, and an independent exact-face driven- 2ω audit. The main methodological lesson is that small-data inverse design can remain useful when the surrogate is restricted to basin proposal and every promoted label is checked by deterministic simulation. The five-fold statistics show only that KRR is competitive in the inherited Project C window; the C2 contribution is a separate five-candidate validation branch and is not a new large surrogate data set.

Compared with adjoint inverse design, the present workflow is less powerful but easier to audit at small scale. Adjoint methods can optimize many degrees of freedom once the objective is differentiable and coupled to the electromagnetic solver. Here, the screening objective is a post-processed surface and interface proxy derived from stored linear fields, followed by a smaller explicit 2ω audit. The probe-offset diagnostic remains central because it reveals when an apparently attractive interface fraction is a sampling-plane artifact; Equation (12) shows how that diagnostic can become an active penalty in a subsequent design loop.

The operating point is not an ENZ experiment. Although ITO can exhibit large nonlinear response near its ENZ region,^{6,8} the $0.92\ \mu\text{m}$ pump is far from the modeled $1.40\ \mu\text{m}$ ENZ wavelength. The calculation does not include calibrated bulk $\chi^{(2)}$, $\chi^{(3)}$, hot-electron dynamics, nonlocal response, or heating. The exact-face calculation does include coherent driven- 2ω propagation, absorption, and radiation, but only for normalized relative surface-source coefficients and under the undepleted-pump approximation.

The magnitude of the reported enhancement is also intentionally modest relative to deliberately

double-resonant plasmonic SHG designs. Reviews of metal-nanostructure SHG emphasize that resonant antennas, gaps, lattice resonances, and double-resonance mechanisms can produce much larger signal changes;⁴ for example, a dolmen-type gold metasurface with coupled Fano and surface-lattice resonances reported a 230-fold simulated SHG-intensity enhancement relative to a single-resonance case.⁵ Against that context, the conservative C2-A5 r110 comparison in Figure 11(a)– $2.03\times$ in the absolute interface proxy and $2.10\times$ in the representative SHG peak ratio—should be read as a within-platform ranking signal, not as a record-efficiency claim. The r130 $0.91\ \mu\text{m}$ fixed-probe maximum sharpens the local proxy landscape but fails the probe-offset robustness gate, so it is used to motivate a stronger robustness objective rather than to extend the headline label.

10 Limitations and Next Steps

The surface-SHG proxy and the exact-face cross-check are not absolute conversion-efficiency models. The proxy sums Au-side and ITO-side contributions incoherently and uses normalized tensor weights; the exact-face solve restores coherent 2ω radiation but still uses unit-norm relative interface coefficients rather than a calibrated complex susceptibility. The r130 $0.91\ \mu\text{m}$ fixed-probe maximum is not label grade because its total proxy varies by $0.292\ \log_{10}$ units over the 2–4 nm probe-offset test. The free-space result is representative rather than statistical, so device-to-device variation, formal power-law uncertainty, and collection efficiency remain to be quantified. A device-level claim would require repeated fabricated arrays, calibrated material coefficients, uncertainty propagation, a robustness-aware objective such as Equation (12), and a coupled treatment if pump depletion or nonlinear back-action becomes appreciable. The present result is therefore deliberately limited to normalized-relative ranking under the undepleted-pump approximation.

11 Conclusion

This work presents a PC-scale, probe-aware inverse-design workflow for Au/ITO hybrid nonlinear nanopillars, benchmarked by a representative free-space SHG comparison and an exact-face driven- 2ω cross-check. The central idea is to separate basin discovery, direct proxy validation, numerical robustness testing, and physical benchmarking. The workflow combines a provenance-bounded surrogate, direct MEEP FDTD labels, probe and mesh diagnostics, p2 mixed-HCurl verification, and filtered 460 nm measurements. It is an auditable normalized-relative screening pipeline rather than a claim of ENZ-mediated enhancement or absolute conversion efficiency.

Applied to Au nanopillars on an ITO film, the updated C2-A validation identifies C2-A5, a 200×70 nm pillar, as the current r110-confirmed winner at $0.92\ \mu\text{m}$. At this label level, C2-A5 gives total proxy 1.0344 and absolute Au-ITO interface proxy 0.2195. These values are $1.25\times$ and $2.03\times$ larger, respectively, than the previous 140×70 nm r110 reference, and the representative free-space SHG measurement gives a peak ratio of about 2.10 under the same nominal pump wavelength, following the simulated interface-proxy trend. The C2-A5 r130 gate over $0.90\text{--}0.94\ \mu\text{m}$ keeps the nominal $0.92\ \mu\text{m}$ point within $+0.019\ \log_{10}$ units of the r110 value and resolves a larger $0.91\ \mu\text{m}$ fixed-probe local maximum with total proxy 2.6225, but the $0.91\ \mu\text{m}$ point is retained as a diagnostic feature because its total proxy changes by $0.292\ \log_{10}$ units under probe-offset variation. Matched central field-slice maps support the presence of broad sidewall and contact-region enhancement in the promoted geometry, but they do not remove this probe-definition qualification. The 220×50 nm geometry remains useful, but in a different role: it is a resolution and probe-definition diagnostic whose $0.92\ \mu\text{m}$ resonance survives an r130 local wavelength check, while its absolute response remains well below C2-A5.

The most important methodological result is that interface fraction is not a robust final optimization label in this formulation, and that even total proxy must pass a probe-offset check before a refined hotspot amplitude is promoted. C2-A5 initially appeared strongly interface weighted at r70, with a nominal interface fraction of 48.3%, but the same candidate drops to 21.2% at r110. Together with the r110 2/3/4 nm probe-offset tests and the C2-A5 r130 check, this behavior shows that near-boundary quantities can be strongly affected by resolution and sampling-plane definitions. The safer result is a probe-qualified relative-channel landscape: total proxy and non-fractional interface proxy can identify stronger r110 candidates, while interface fraction and fixed-probe r130 extrema remain diagnostics until they satisfy the same robustness criteria.

The broader contribution is therefore not a record-efficiency metasurface, but a reproducible validation chain for early Au/ITO nonlinear-nanophotonic design. Its transferable element is the sequence of provenance control, direct reruns, mesh/probe gates, exact-face energy closure, and experimental ranking checks. The next stage is to add repeated-array statistics and calibrated interface coefficients so that the normalized-relative workflow can progress toward device-level power and efficiency predictions.

Supporting Information

Supporting Information, source files, simulation scripts, and processed representative measurement data are available from the corresponding author and will be deposited with the public repository associated with this manuscript.

Disclosures

The authors declare that there are no financial interests, commercial affiliations, or other potential conflicts of interest that could have influenced the objectivity of this research or the writing of this paper.

Code, Data, and Materials Availability

The simulation and post-processing scripts, candidate manifests, processed CSV/JSON summaries, figure-generation scripts, and processed representative SHG spectra/images supporting this study will be made available in a public repository at the time of submission or publication. Large raw MEEP field-output files are not included because of their size, but the input scripts and parameters required to regenerate them are provided; raw field outputs and unprocessed experimental files can be made available from the corresponding author upon reasonable request.

Acknowledgments

The authors acknowledge National Kaohsiung University of Science and Technology for institutional support. This research received no external funding. OpenAI Codex was used during manuscript preparation for language editing, LaTeX reformatting, and consistency checking; all scientific statements, data, and conclusions were reviewed by the authors.

References

- 1 M. Kauranen and A. V. Zayats, “Nonlinear plasmonics,” *Nature Photonics* **6**(11), 737–748 (2012).

- 2 J. Butet, P.-F. Brevet, and O. J. F. Martin, “Optical second harmonic generation in plasmonic nanostructures: From fundamental principles to advanced applications,” *ACS Nano* **9**(11), 10545–10562 (2015).
- 3 A. Krasnok, M. Tymchenko, and A. Alù, “Nonlinear metasurfaces: A paradigm shift in nonlinear optics,” *Materials Today* **21**(1), 8–21 (2018).
- 4 C.-C. Zhang, J.-Y. Zhang, J.-R. Feng, *et al.*, “Plasmon-enhanced second harmonic generation of metal nanostructures,” *Nanoscale* **16**(12), 5960–5975 (2024).
- 5 X. Sun, L. Gui, H. Xie, *et al.*, “230-fold enhancement of second-harmonic generation by coupled double resonances in a dolmen-type gold metasurface,” *Journal of Applied Physics* **135**(19), 193102 (2024).
- 6 M. Z. Alam, I. De Leon, and R. W. Boyd, “Large optical nonlinearity of indium tin oxide in its epsilon-near-zero region,” *Science* **352**(6287), 795–797 (2016).
- 7 I. Liberal and N. Engheta, “Near-zero refractive index photonics,” *Nature Photonics* **11**, 149–158 (2017).
- 8 N. Kinsey, C. DeVault, A. Boltasseva, *et al.*, “Near-zero-index materials for photonics,” *Nature Reviews Materials* **4**(12), 742–760 (2019).
- 9 N. Bloembergen and P. S. Pershan, “Light waves at the boundary of nonlinear media,” *Physical Review* **128**(2), 606–622 (1962).
- 10 J. Rudnick and E. A. Stern, “Second-harmonic radiation from metal surfaces,” *Physical Review B* **4**(12), 4274–4290 (1971).
- 11 J. E. Sipe, D. J. Moss, and H. M. van Driel, “Phenomenological theory of optical second- and

third-harmonic generation from cubic centrosymmetric crystals,” *Physical Review B* **35**(3), 1129–1141 (1987).

12 J. I. Dadap, J. Shan, and T. F. Heinz, “Theory of optical second-harmonic generation from a sphere of centrosymmetric material: Small-particle limit,” *Journal of the Optical Society of America B* **21**(7), 1328–1347 (2004).

13 J. Peurifoy, Y. Shen, L. Jing, *et al.*, “Nanophotonic particle simulation and inverse design using artificial neural networks,” *Science Advances* **4**(6), eaar4206 (2018).

14 D. Liu, Y. Tan, E. Khoram, *et al.*, “Training deep neural networks for the inverse design of nanophotonic structures,” *ACS Photonics* **5**(4), 1365–1369 (2018).

15 S. An, C. Fowler, B. Zheng, *et al.*, “A deep learning approach for objective-driven all-dielectric metasurface design,” *ACS Photonics* **6**(12), 3196–3207 (2019).

16 W. Ma, Z. Liu, Z. A. Kudyshev, *et al.*, “Deep learning for the design of photonic structures,” *Nature Photonics* **15**(2), 77–90 (2021).

17 S. Molesky, Z. Lin, A. Y. Piggott, *et al.*, “Inverse design in nanophotonics,” *Nature Photonics* **12**(11), 659–670 (2018).

18 C. Kang, C. Park, M. Lee, *et al.*, “Large-scale photonic inverse design: Computational challenges and breakthroughs,” *Nanophotonics* **13**(20), 3765–3792 (2024).

19 J. Kim, J.-Y. Kim, J. Kim, *et al.*, “Inverse design of nanophotonic devices enabled by optimization algorithms and deep learning: Recent achievements and future prospects,” *Nanophotonics* **14**(2), 121–151 (2025).

20 H. Zhang, L. Kang, S. D. Campbell, *et al.*, “Data driven approaches in nanophotonics: A review of AI-enabled metadevices,” *Nanoscale* **17**(41), 23788–23803 (2025).

- 509 21 C. M. Lalau-Keraly, S. Bhargava, O. D. Miller, *et al.*, “Adjoint shape optimization applied to
510 electromagnetic design,” *Optics Express* **21**(18), 21693–21701 (2013).
- 511 22 N. Li, J. Zhang, D. N. Neshev, *et al.*, “Inverse design of nonlinear metasurfaces for sum
512 frequency generation,” *Nanophotonics* **13**(18), 3363–3372 (2024).
- 513 23 A. D. Rakić, A. B. Djurišić, J. M. Elazar, *et al.*, “Optical properties of metallic films for
514 vertical-cavity optoelectronic devices,” *Applied Optics* **37**(22), 5271–5283 (1998).
- 515 24 A. F. Oskooi, D. Roundy, M. Ibanescu, *et al.*, “MEEP: A flexible free-software package
516 for electromagnetic simulations by the FDTD method,” *Computer Physics Communications*
517 **181**(3), 687–702 (2010).
- 518 25 J.-P. Bérenger, “A perfectly matched layer for the absorption of electromagnetic waves,”
519 *Journal of Computational Physics* **114**(2), 185–200 (1994).