Inverse design of open nanophotonic cavities through differentiable mode expansions

Author list and affiliations in preparation

Working draft — September 5, 2026

Abstract

Modal solvers describe photonic-crystal resonances far more economically than general-purpose field solvers, and differentiating them has opened a fast route to inverse design. The existing differentiable formulations, however, solve a real-frequency eigenproblem on a periodic cell, where the radiative linewidth is not part of the spectrum. We remove that restriction. A finite, one-dimensional nanobeam cavity is written as an outgoing complex-frequency pole of a nonlinear operator $\mathbf{A}(\widetilde{\omega}, \boldsymbol{\theta})\boldsymbol{u} = 0$ assembled from the guided and radiation-continuum spectrum of a reference waveguide and a complete Zernike–Legendre basis inside each etched hole. The outgoing condition is carried by a Sommerfeld contour in longitudinal momentum rather than by an absorbing box, so the imaginary part of the pole—and hence Q—is generated by the continuum itself. Implicit differentiation of the pole through the left and right null vectors returns the complete geometry gradient at the cost of one additional linear solve, independently of the number of geometry variables. We exercise the identical operator protocol on two materially distinct devices: an air-clad silicon Quan–Lončar nanobeam (15 holes per side, 45 variables) and a five-layer SiO_2/Si_3N_4/anthracene/PVA cavity with an explicitly subtracted guided feed pole (42 holes per side, 126 variables). Independent finite-difference time-domain (FDTD) evaluation of frozen endpoint geometries gives quality-factor gains of 3.8407× and 3.0666×. The central finding is a dissociation: along the second device's trajectory the surrogate moves from under-predicting FDTD Q (ratio 0.472) to over-predicting it by more than an order of magnitude (ratio 12.31)—a drift of 27.5× in the ratio of the two observables—while the rank ordering of designs within a fixed FDTD protocol is preserved (Spearman 0.911, $n = 23$). Ascent directions transfer; absolute linewidths do not. Both gains are quoted with the exact ringdown conditions under which they were measured, including the endpoint fits obtained over less than one photon lifetime, and are labelled provisional pending a predeclared matched convergence panel.

1 Introduction

Nanophotonic cavity design is an unusually sharp test of an optimization method. The target mode must be confined to a wavelength-scale volume while its weak coupling to a continuum of radiation channels sets a quality factor that can move by orders of magnitude under subwavelength boundary displacements. The quantity being optimized is therefore a small imaginary part riding on a large real one: at $Q = 10^4$ the loss rate is only 5×10^{-5} of the resonant frequency, so a discretization error that is entirely harmless for the resonance wavelength can dominate—and reverse—the Q gradient.

Two families of solver are used for such problems. General-purpose Maxwell solvers optimize arbitrary structures with adjoint methods [1, 2, 3], but every iteration pays for a discretized domain and an absorbing boundary, and the gradient inherits whatever the absorber does to the outgoing field. Modal solvers compress the same physics into a small structured basis; making them differentiable has consequently become an important complementary direction. Guided-mode expansion in particular has been cast in a fully differentiable form and applied to photonic-crystal slab inverse design [4, 5, 6].

That construction solves a real-frequency eigenproblem on a periodic cell. The radiative linewidth is recovered afterwards, perturbatively, from the coupling of a guided mode to the light cone. This is well suited to periodic slabs and is not available for a finite cavity whose entire figure of merit is the linewidth: a closed or truncated modal basis can return a numerically precise and physically meaningless Q, because a handful of guided or absorber modes can reproduce the near field while getting the tiny continuum coupling completely wrong.

The quasinormal-mode (QNM) literature supplies the correct object [7, 8, 9, 10]. A resonance of an open system is a pole of the scattering problem at complex frequency $\widetilde{\omega}$ with $\operatorname{Im}\widetilde{\omega} < 0$; its normalization is not the closed-cavity energy integral, and substituting the latter into an outgoing problem generally gives the wrong imaginary-frequency derivative even when the real shift looks plausible. What has been missing is a formulation in which the QNM of a *finite, patterned* device is defined by a small nonlinear operator that is (i) built from an analytically known reference spectrum including its continuum, (ii) outgoing by contour choice rather than by numerical absorption, and (iii) differentiable in closed form with respect to every geometric degree of freedom.

Contributions. This paper reports the following.

(i) **An outgoing, differentiable mode expansion for finite 1D cavities** (Section 2). The cavity is the root of $\det \mathbf{A}(\widetilde{\omega}, \boldsymbol{\theta}) = 0$ for an operator assembled from the exact guided-plus-continuum Green tensor of the unpatterned beam and a complete Zernike–Legendre basis inside each through-hole. Openness is enforced by a Sommerfeld contour in longitudinal momentum, analytically continued to the lower half plane along an explicit homotopy.

(ii) **A closed-form geometry gradient** (Section 3). Implicit differentiation at a simple pole gives $\partial\widetilde{\omega}/\partial\theta_j$ from the left and right null vectors and the analytic derivatives of the hole form factors, so a full 126-variable gradient costs one pole solve plus one reverse-mode contraction rather than 126 forward solves.

(iii) **Operator-level portability across two devices** (Sections 5 and 6). The same protocol runs on a homogeneous air-clad silicon beam and on a five-layer stack whose reference Green tensor requires Fresnel recursion, Dyson strip resummation, and explicit subtraction of the guided feed pole before the longitudinal contour integral is numerically usable.

(iv) **A quantified dissociation between direction transfer and magnitude transfer** (Section 7). This is the paper's most portable result and the one with the clearest consequence for how reduced open-system models should be used.

We are deliberate about what is *not* claimed. The two independently measured Q gains are historical endpoints obtained under stated FDTD settings whose ringdown windows were, in three of four cases, shorter than the fitted photon lifetime (Section 4). They are reported with that condition attached and remain provisional against the predeclared convergence panel of Table 2; no analytical Q is anywhere equated to, or substituted for, an FDTD Q.

2 The cavity as an outgoing pole

2.1 Conventions and the object being computed

We use $\exp(-\mathrm{i}\omega t)$ and nonmagnetic, nondispersive media throughout. A passive outgoing resonance is a pole $\widetilde{\omega} = \omega_r + \mathrm{i}\omega_i$ with $\omega_i < 0$, and

$$Q = -\frac{\operatorname{Re}\widetilde{\omega}}{2\operatorname{Im}\widetilde{\omega}} = -\frac{\omega_r}{2\omega_i}. \tag{1}$$

The optimization target is $\log Q$, for which

$$\delta \log Q = \frac{\delta\omega_r}{\omega_r} - \frac{\delta\omega_i}{\omega_i}, \tag{2}$$

so that multiplicative gains are additive and the two parts of the pole enter on an equal relative footing. Equation (2) also states the difficulty of the problem plainly: since $|\omega_i|/\omega_r = 1/(2Q)$, the second term is the one that must be resolved, and it is the smaller of the two by four orders of magnitude at $Q \sim 10^4$.

2.2 Exact parameterized dielectric

The unpatterned beam of width W and thickness T defines the reference permittivity $\epsilon_0(y,z)$. Hole j is the elliptical through-cylinder

$$H_j(\boldsymbol{\theta}) = \left\{ (x,y,z) : \tfrac{(x-x_j)^2}{r_{x,j}^2} + \tfrac{y^2}{r_{y,j}^2} \le 1,\ |z| \le \tfrac{T}{2} \right\}, \tag{3}$$

so that $\epsilon_{\boldsymbol{\theta}} = \epsilon_0 + \Delta\epsilon_{\boldsymbol{\theta}}$ with $\Delta\epsilon_{\boldsymbol{\theta}} = (\epsilon_c - \epsilon_b)\sum_j \mathbf{1}_{H_j}$. The geometry vector $\boldsymbol{\theta}$ collects the positive-side centres x_j and the two semi-axes $r_{x,j}, r_{y,j}$ of every hole; the mirror image is generated exactly. This parameterization is not a discretized permittivity map: the Fourier factor of a hole is available in closed form,

$$\widehat{\mathbf{1}}_{H_j}(\boldsymbol{k}) = e^{-\mathrm{i}k_x x_j}\left[2\pi r_{x,j} r_{y,j}\,\tfrac{J_1(\rho_j)}{\rho_j}\right]\left[\tfrac{2\sin(k_z T/2)}{k_z}\right], \tag{4}$$

with $\rho_j^2 = (k_x r_{x,j})^2 + (k_y r_{y,j})^2$, and so are its derivatives, e.g. $\partial_{x_j}\widehat{\mathbf{1}}_{H_j} = -\mathrm{i}k_x\widehat{\mathbf{1}}_{H_j}$. Every geometry derivative used later is therefore analytic rather than a finite difference of a remeshed structure.

2.3 Outgoing volume-integral pole equation

With $\mathcal{L}_0(\omega) = \nabla\times\nabla\times - (\omega^2/c^2)\epsilon_0$ and $\mathbf{G}^+_{\mathrm{wg}}(\omega) = \mathcal{L}_0(\omega)^{-1}$ taken on the outgoing sheet, the source-free Maxwell problem is equivalent to the homogeneous Lippmann–Schwinger equation $\mathbf{E} = (\omega^2/c^2)\,\mathbf{G}^+_{\mathrm{wg}}\Delta\epsilon_{\boldsymbol{\theta}}\mathbf{E}$. Restricting the unknown to the finite hole support gives the exact pole condition

$$D(\omega,\boldsymbol{\theta}) = \det\nolimits_{\mathrm{reg}}\left[I - \tfrac{\omega^2}{c^2}\mathbf{G}^+_{\mathrm{wg}}(\omega)\,\Delta\epsilon_{\boldsymbol{\theta}}\right] = 0. \tag{5}$$

The regularized determinant is written deliberately. The electric dyadic Green tensor carries a singular local term; an unregularized collocation determinant of the same kernel is not the Fredholm determinant of a trace-class operator and does not have (5) as its root. We keep the longitudinal projector in Fourier space, where it is an algebraic term rather than a $1/R^3$ singularity requiring a depolarization guess (Supplementary Section S4).

2.4 Reference Green tensor: guided poles and continuum

The unpatterned beam is invariant along x, so

$$\mathbf{G}_{\mathrm{wg}}(x-x';\omega) = \int_{\mathcal{C}_\beta}\frac{\mathrm{d}\beta}{2\pi}\,e^{\mathrm{i}\beta(x-x')}\,\widetilde{\mathbf{G}}_{\mathrm{wg}}(\beta;\omega). \tag{6}$$

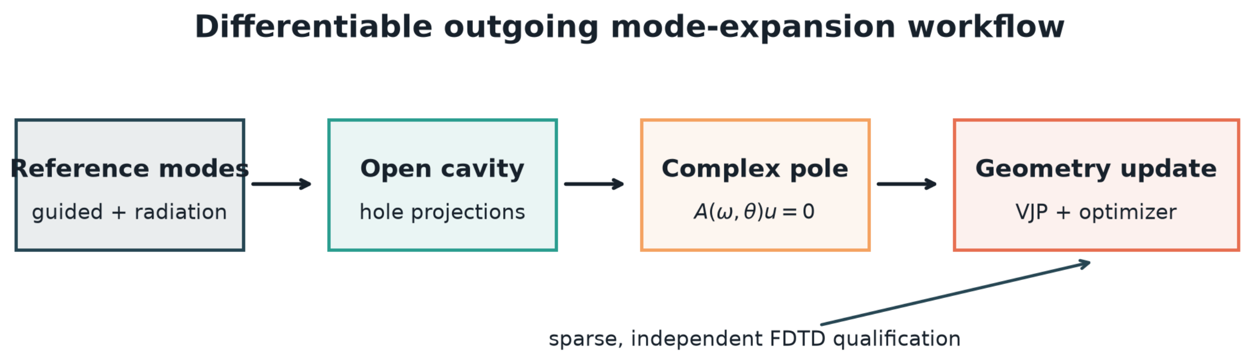

Figure 1: Operator flow. The reference waveguide supplies an exact guided-plus-continuum Green tensor; the finite hole array supplies a complete interior basis; their contraction on an outgoing longitudinal contour defines $\mathbf{A}(\widetilde{\omega}, \boldsymbol{\theta})$, whose root is the cavity pole. The geometry gradient is obtained by implicit differentiation at that root. FDTD is an independent, sparse qualification layer and never supplies an optimizer gradient.

At fixed β the transverse problem has compact support. For a homogeneous cladding the symbol is exact, $\widetilde{\mathbf{G}}_{\mathrm{c}} = P_T/(k^2 - k_c^2) - P_L/k_c^2$ with $P_L = \boldsymbol{k}\boldsymbol{k}^T/k^2$, and the beam cross-section is restored by a Dyson resummation over its own susceptibility χ_b,

$$\widetilde{\mathbf{G}_{\mathrm{wg}}}(\beta,\omega) = \left[I - k_0^2\chi_b\widetilde{\mathbf{G}}_{\mathrm{c}}(\beta,\omega)\right]^{-1}\widetilde{\mathbf{G}}_{\mathrm{c}}(\beta,\omega). \tag{7}$$

Guided modes are therefore not inserted by hand: they are the zeros of $\det[I - k_0^2\chi_b\widetilde{\mathbf{G}}_{\mathrm{c}}]$, i.e. poles of the same resolvent that carries the radiation continuum, and their residues are normalized by the pole derivative rather than by an inconsistent finite-box norm.

For the multilayer device the homogeneous symbol is replaced by the stratified one. Generalized Fresnel reflections are composed recursively through each film, $R_{j-1} = (r_{j-1,j} + R_j e^{2\mathrm{i}k_{zj}d_j})/(1 + r_{j-1,j}R_j e^{2\mathrm{i}k_{zj}d_j})$, the vertical kernel in the source layer is $g(z,z') = -u_L(z_<)u_R(z_>)/W_{\mathrm{r}}$ with Wronskian $W_{\mathrm{r}} = 2\mathrm{i}k_z(e^{-\mathrm{i}k_z d} - r_b r_t e^{\mathrm{i}k_z d})$, and the patterned strip is added by the same Dyson step as (7). Replacing the films by an effective cladding would delete their guided and leaky poles and is not an admissible reference for Q optimization (Supplementary Section S6).

2.5 Outgoing contour and analytic continuation

Openness enters through $\mathcal{C}_\beta$, not through a numerical box. In the scaled variable $u = \beta/k_c(\omega)$ the radiation branch points are pinned at $u = \pm 1$ and the guided poles sit near $u = \pm n_{\mathrm{eff}}/n_c$. The implementation uses the Sommerfeld path

$$u(t) = t - \mathrm{i}h\tanh(t/d)\,\exp\!\left[-(t/L)^8\right], \tag{8}$$

which passes below every positive-u singularity and above its negative-u partner, and reverts to the real axis outside $|t| \sim L$. Holding $u(t)$ fixed while moving ω from the causal upper half plane down to the resonance is an explicit continuation homotopy; a pole or branch point crossing the frozen path is detected as a discontinuity rather than silently changing sheet. Substituting a lower-half-plane frequency into a real-axis quadrature lands on the wrong sheet, and did so in an early revision of this work (Supplementary Section S18).

2.6 Interior basis and assembly

Inside hole j we use normalized hole coordinates $x - x_j = r_{x,j}\rho\cos\theta$, $y = r_{y,j}\rho\sin\theta$, and expand in a complete Zernike–Legendre vector basis,

$$\boldsymbol{\phi}_{j,nm\sigma\ell s} = \frac{Z_{nm}^{\sigma}(\rho,\theta)}{\sqrt{r_{x,j}r_{y,j}N_{nm}}}\sqrt{\tfrac{2\ell+1}{T}}\,P_\ell(2z/T)\,\boldsymbol{e}_s, \tag{9}$$

with $R_n^m(\rho) = (-1)^{(n-m)/2}\rho^m P_{(n-m)/2}^{(m,0)}(1-2\rho^2)$ evaluated in Jacobi form for stability. Because the hole spans the full beam thickness and shares the waveguide's Legendre functions in z, the thickness projection is exact. Denoting by $C_j(\beta)$ the projection of this basis into the cross-section basis at longitudinal momentum β, the projected outgoing Green matrix is

$$\mathbf{G}_H^+(\omega) = \int_{\mathcal{C}_+}\frac{\mathrm{d}\beta}{2\pi}\,C(-\beta)^T\,\widetilde{\mathbf{G}_{\mathrm{wg}}}(\beta,\omega)\,C(\beta), \tag{10}$$

and the finite operator whose root is the cavity is

$$\mathbf{A}(\omega,\boldsymbol{\theta}) = I - k_0^2(\epsilon_c - \epsilon_b)\,\mathbf{G}_H^+(\omega,\boldsymbol{\theta}). \tag{11}$$

The left factor in (10) is the transpose $C(-\beta)^T$, not the Hermitian conjugate $C(\beta)^\dagger$. The two coincide on the real-frequency boundary for real basis functions, but only the transpose is analytic in ω; using the conjugate destroys the continuation into the lower half plane. Reciprocity then predicts a complex-symmetric $\mathbf{A}$, which is a stringent numerical check rather than an assumption (Section 4).

Truncating (9) at radial order N_{xy} and vertical order N_z and the contour at quadrature order N_β yields a sequence $\mathbf{A}^{(N)}$ whose roots $\widetilde{\omega}_N$ converge to the cavity pole; N_β is a continuum quadrature order, not a finite number of physical radiation modes.

3 Geometry gradient of the pole

Let $\boldsymbol{u}$ and $\boldsymbol{v}$ be right and left null vectors of the assembled operator at an isolated simple pole,

$$\mathbf{A}(\widetilde{\omega},\boldsymbol{\theta})\,\boldsymbol{u}=0, \qquad \boldsymbol{v}^T\mathbf{A}(\widetilde{\omega},\boldsymbol{\theta})=0. \tag{12}$$

Differentiating (12) with respect to a geometry coordinate θ_j and contracting on the left with $\boldsymbol{v}$ eliminates the unknown field derivative, giving

$$\boxed{\frac{\partial\widetilde{\omega}}{\partial\theta_j}=-\frac{\boldsymbol{v}^T\big(\partial_{\theta_j}\mathbf{A}\big)\boldsymbol{u}}{\boldsymbol{v}^T\big(\partial_{\omega}\mathbf{A}\big)\boldsymbol{u}}} \tag{13}$$

The denominator is the QNM normalization for this operator. It is generated by the frequency dependence of $\mathbf{A}$ itself—including that of the reference Green tensor and of the contour—and any arbitrary rescaling of either null vector cancels between numerator and denominator. This is the safe interface: a closed-cavity energy normalization substituted into an outgoing problem generally returns a wrong $\partial\omega_i/\partial\theta_j$ even when $\partial\omega_r/\partial\theta_j$ looks correct, and it is precisely ω_i that carries Q through (2).

Three properties make (13) practical.

Cost. The denominator and both null vectors are computed once per accepted geometry. Each additional coordinate costs one contraction of the already-assembled $\partial_{\theta_j}\mathbf{A}$, whose ingredients are the analytic form-factor derivatives of (4). Reverse-mode accumulation over the operator primitives therefore returns the entire gradient—45 coordinates for the first device, 126 for the second—at a cost set by the pole solve, not by the number of variables. The measured cost on the production discretization is a 398.0 s pole solve followed by a 370.0 s contraction for the full 45-coordinate gradient (Supplementary Section S14).

Boundary form. For an interface Γ between isotropic media moved with normal velocity h, the numerator of (13) reduces to a Hadamard boundary integral,

$$\boldsymbol{v}^T(\partial_p\mathbf{A})\boldsymbol{u} \;\propto \int_\Gamma h\Big[(\epsilon_1-\epsilon_2)\,\mathbf{E}_{L\parallel}\cdot\mathbf{E}_{R\parallel} - (\epsilon_1^{-1}-\epsilon_2^{-1})\,D_{L\perp}D_{R\perp}\Big]\mathrm{d}A, \tag{14}$$

in which tangential $\mathbf{E}$ and normal $\mathbf{D}$ appear because those are the components continuous across a sharp dielectric boundary. Equation (14) is what makes hole radii, hole positions, and sidewalls differentiable in one framework: all are pieces of the same movable boundary. The proportionality constant and sign depend on the exact definition of $\mathbf{A}$ and on the orientation of the normal, and are pinned in this work by tests—translate one planar interface by a known amount, compare with a centred difference on identical connectivity, reverse the normal and confirm the sign flips—rather than by prose.

Mode identity. Equation (13) is a statement about one isolated pole, so every accepted step must certify that the pole it followed is the same mode. Frequency-nearest tracking is insufficient: a small shape step can legitimately move ω_r by many of the old linewidths, and what must remain continuous is the mode, not a frequency window. The optimizer ranks candidates by (i) biorthogonal field overlap with the accepted QNM on a common representation, (ii) localization in the cavity region, (iii) symmetry and polarization sector, (iv) agreement with the first-order pole prediction, and only then (v) frequency proximity as a tie-breaker. A step whose candidates clear none of the overlap and localization gates is declared unresolved and is not allowed to update the quasi-Newton history, because a silent mode switch can make a correct positive directional derivative appear to produce a negative objective change.

Finally, the raw shape gradient is mapped through a boundary metric $(I-\ell^2\partial_s^2)h=g_Q$ before it is applied, where s is arclength and ℓ a fabrication-aware smoothing length. This is a Sobolev gradient, not an added penalty term. A line search is only admissible if the applied geometry satisfies $\operatorname{dist}\big(\Gamma(\alpha),\Gamma(0)\big)\to 0$ as $\alpha\to 0$; an early revision of the multilayer pipeline violated this through polygon repair and remeshing, and is retained as a documented counterexample in Supplementary Section S18.

4 What is validated, and how

Analytical self-consistency and physical transfer are separate gates, and the distinction is load-bearing throughout this paper. Analytical Q_{ana} and FDTD Q_{FDTD} are treated as two different observables; no statement below converts one into the other.

4.1 Operator-level gates

At the analytical level we require, at every accepted geometry: a passive pole ($\operatorname{Im}\widetilde{\omega}<0$); an operator residual $\|\mathbf{A}\boldsymbol{u}\|$ at the solve tolerance; complex symmetry of $\mathbf{A}$ to the level predicted by reciprocity; convergence of $\widetilde{\omega}$ under independent refinement of the Zernike order, the vertical order, and the contour quadrature; agreement of the analytic gradient (13) with centred finite differences on identical connectivity; and clearance of the mode-identity gates of Section 3. Representative values for the production discretizations are an operator residual of 2.29×10^{-10}, reciprocity residuals between 1.1×10^{-16} and 3.8×10^{-16}, left/right null-vector equation errors below 2.35×10^{-15}, and agreement between the analytic vector–Jacobian product and a central difference of the Q derivative to 6.71×10^{-6} relative (Supplementary Section S14). These checks certify that the implementation differentiates its own truncation. They say nothing about whether that truncation's absolute Q equals a converged Maxwell solution.

Table 1: Ringdown adequacy of the four quoted endpoints, computed from the retained simulation metadata. $\mathcal{N} < 1$ means the fit observed less than one amplitude lifetime. Design23 diagnostics do not retain a source start time, so their windows are upper bounds and their $\mathcal{N}$ is optimistic. Source: `results/ringdown-adequacy.json`.

endpoint	Q_{FDTD}	$1/\gamma$ (ps)	window (ps)	$\mathcal{N}$
Quan baseline	15 449.47	24.64	22.65	0.919
Quan Phase-58	59 337.16	94.5	22.65	**0.24**
Design23 step 0	18 991.94	15.98	12	0.751
Design23 step 105	58 240.27	47.65	64	1.34

4.2 Physical transfer, and the ringdown condition

Physical claims use frozen endpoint geometries rendered into an independent FDTD solver, with the same source, monitors, domain, mesh, absorber, run time, and resonance-fitting code at both ends of a comparison. FDTD is never called by the optimizer.

A quality factor extracted from a time-domain ringdown is only constrained if the recorded window is long compared with the amplitude lifetime of the mode being fitted. With $\gamma = \omega/(2Q)$ the amplitude decay rate, the relevant figure of merit is the number of lifetimes observed, $\mathcal{N} = (t_{\mathrm{run}} - t_{\mathrm{start}})\gamma$. We audited all 35 retained FDTD observations in this way (Table 1); 24 of them, including three of the four endpoints that carry the quoted gains, were fitted over less than one lifetime.

This single number replaces the vaguer hedging that a project record accumulates. It says exactly what is uncertain and in which direction: a sub-lifetime exponential fit is poorly conditioned and, empirically in this data set, biased high. The clearest internal evidence is that a 12 ps diagnostic partway along the second device's trajectory reports $Q_{\mathrm{FDTD}} = 156\,520$ at step 61, whereas the 64 ps analyses that follow it report values near 58 240.27 for geometries with a higher analytical Q (Fig. 5). We therefore treat both device gains as provisional and quote them only together with Table 1.

4.3 Predeclared promotion panel

Promotion from provisional to validated requires a matched convergence panel that was declared before any of it was run and is recorded in `results/publication-revalidation-plan.json` (Table 2). Each endpoint is evaluated over the full factorial of four binary axes, giving 16 variants per endpoint and 64 tasks in total. The gain is promoted only if the same mode is identified in every variant, every fit passes its diagnostics, and the optimized Q exceeds the initial Q in *every* executed variant; the central value is then chosen by the predeclared refinement rule. Unreliable fits are retained and marked, never silently dropped. The panel is presently predeclared not submitted, and no build command in the source repository is able to submit it.

Table 2: Predeclared matched convergence panel. Full factorial over four binary axes per endpoint, 64 tasks total, with nine required artifacts per task. Source: `results/publication-revalidation-plan.json`.

axis	Quan–Lončar	Design23
steps per wavelength	28 / 35	18 / 20
run time (ps)	60 / 90	64 / 96
PML layers	12 / 16	12 / 16
padding	1.25 / 1.5625	1.0 / 1.25

The audit of Table 1 also revises the panel itself. At the Phase-58 amplitude lifetime of 94.5 ps, the predeclared 60 ps and 90 ps run times correspond to only $\mathcal{N} = 0.63$ and 0.95 lifetimes. The declared axis is thus not by itself sufficient to certify a Q of that size, and the promotion run must extend the run-time axis until $\mathcal{N} \gtrsim 3$ at the optimized endpoint. We record this as a required amendment rather than quietly reinterpreting the original declaration.

5 Device I: air-clad Quan–Lončar nanobeam

The first device is the canonical deterministic-design nanobeam cavity [11]: a mirror-symmetric silicon beam, 700 nm×220 nm in cross-section, refractive index 3.46, suspended in air, with 15 elliptical through-holes per side and a resonance near 1 502.32 nm. The free variables are the 15 positive-side gaps and the two semi-axes of each hole, i.e. 45 coordinates. Because the cladding is homogeneous, the reference Green tensor is the one of Eq. (7) with a single air cladding, and the production discretization is a $P2$ hole basis with one vertical order against a $y3$–$z1$ waveguide cross-section basis.

What the optimizer did. The accepted design is the third step of a constrained ascent in which each step was a scaled L^2 displacement of at most 1 nm, with a maximum single-coordinate move of about 0.34 nm. Reconstructing both endpoint geometries and reducing them to per-hole descriptors (Fig. 3, upper row) shows that this is *local refinement* and nothing more: the mean in-plane aspect ratio moves from 1 to 1.0086, with a maximum of 1.016 near the seventh hole; the largest transverse semi-axis change over all holes is 0.819 nm; the transverse air fraction $2r_y/W$ moves from a maximum of 0.3465 to 0.3482; and the mean hole area changes by a factor of 0.9937. The only large motion is longitudinal: hole centres move by up to 24.6 nm, outward at the mirror end and inward near the centre, i.e. a mild stretch of the taper envelope. The device family is unchanged.

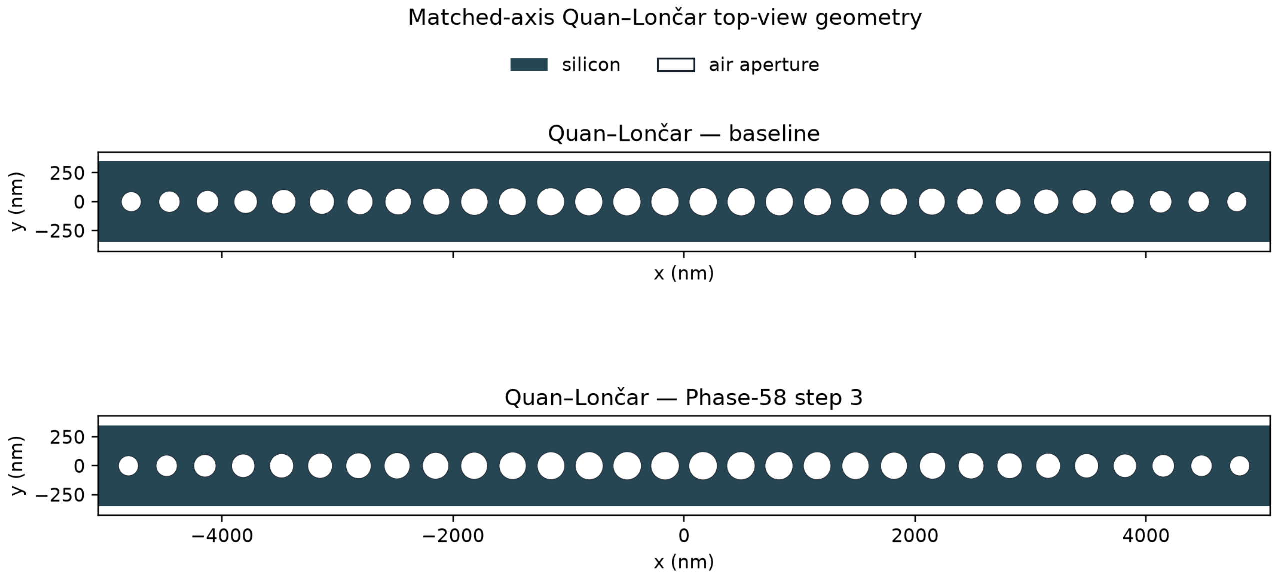

Figure 2: Matched-axis top views of the frozen Quan–Lončar endpoints. Silicon is blue, apertures white. Both panels are reconstructed from the retained seed files (`quan_baseline.json`, `quan_optimized_phase58.json`); no field array enters the drawing. The two geometries are visually almost identical, which is the result: see Fig. 3 for the per-hole differences.

Result. Under matched 24 ps independent FDTD analyses of the two frozen geometries, the fitted quality factor rises from 15 449.47 at 1 502.32 nm to 59 337.16 at 1 499.97 nm, a factor of 3.8407. The corresponding analytical values are 7 455.16 and 18 037.3, a factor of 2.42: the surrogate under-predicts the absolute FDTD quality factor at both ends, by factors of 0.483 and 0.304 respectively, and under-predicts the gain as well.

Condition on this number. Both FDTD analyses used identical sources, monitors, domain, mesh, absorber and run time, so the *ratio* is a matched comparison. The absolute values are not converged: at $Q_{\mathrm{FDTD}} = 59\,337.16$ the amplitude lifetime is 94.5 ps while the usable ringdown window was only 22.65 ps, i.e. $\mathcal{N} = 0.24$ lifetimes (Table 1). A fit over a quarter of a lifetime constrains the decay rate weakly and, in this data set, biases it low—hence Q high. This gain is therefore reported as a provisional candidate pending the panel of Table 2 with the run-time amendment noted there.

A separate eight-step continuation of the same ascent reached an analytical Q of 30 307.26 but was never evaluated in FDTD. It is retained in the record under a distinct name and is not part of any claim here.

6 Device II: multilayer Design23 cavity

The second device tests whether the construction survives a substantially harder reference problem. The stack is, from the bottom: an SiO_2 substrate half-space ($n = 1.45$); a 300 nm layer containing the patterned Si_3N_4 strip ($n = 2.0$) with air elsewhere; 200 nm anthracene ($n = 1.8$); 200 nm PVA ($n = 1.5$); and an air half space. The strip is 556.264 nm wide and carries 42 elliptical through-holes per side, giving 126 free coordinates. The resonance sits near 792.347 nm.

Three ingredients are needed here that the first device does not require. The reference symbol is stratified rather than homogeneous, built from the Fresnel recursion and vertical kernels of Section 2; because the stack breaks z reflection, the symmetry projector retains a fixed (x, y) sector while carrying both z parities. The finite-width strip is then added by a compact Dyson step, and the E_y-dominant guided feed mode appears as a simple pole β_f of that step. Its Laurent terms,

$$\mathbf{G}_{\mathrm{wg}}(\beta) = \frac{R_+}{\beta - \beta_f} + \frac{R_-}{\beta + \beta_f} + \text{regular}, \qquad R_- = -R_+^T, \tag{15}$$

are subtracted before the longitudinal quadrature and restored exactly in real space as ordered right- and left-going feedthrough propagation, with a factor one half at coincident nodes. Attempting the contour integral without this subtraction is not merely inaccurate: the finite-hole self term alone changes the guided block by about 8 % (Supplementary Section S10).

Trajectory. A pure BFGS ascent on $\log Q$ accepted 105 of 106 attempted steps, raising the analytical quality factor from 8 958.74 to 717 003, a factor of 80.03. Independent FDTD diagnostics were taken every three steps (Fig. 5); they do not follow that scale. The exact retained endpoints rise from 18 991.94 at step 0 to 58 240.27 at step 105, a factor of 3.0666. The initial point is the step-0 endpoint of *this* trajectory; a Q of 20 942.06 belonging to a

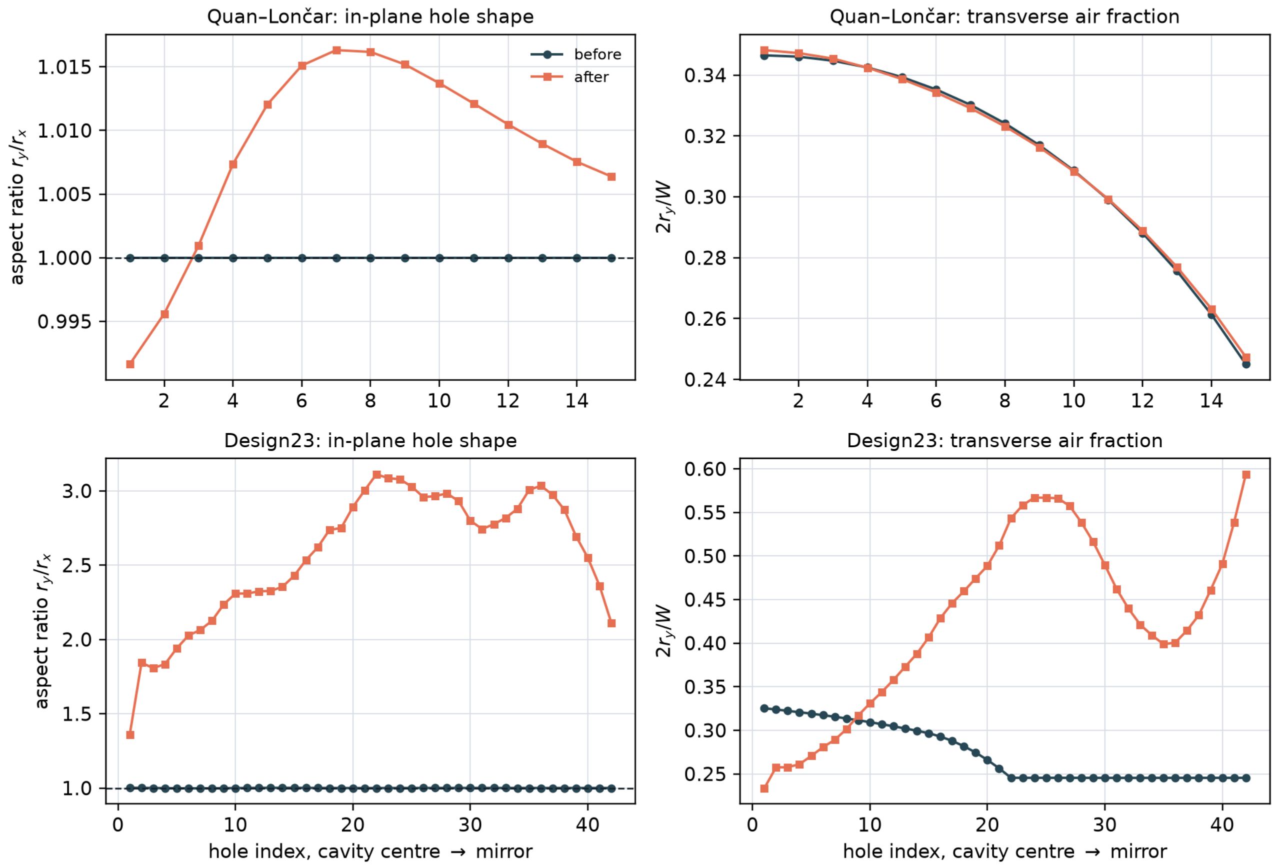

Figure 3: Per-hole shape descriptors before and after optimization, plotted from the cavity centre outward. Note the vertical scales: the Quan–Lončar aspect ratio moves by about one percent (upper left) whereas the Design23 aspect ratio moves by up to a factor of 3.112 (lower left), and the Design23 transverse air fraction reaches 0.5938 of the beam width (lower right). The two campaigns therefore exercise different optimization regimes: local refinement and device restructuring. Source: `results/geometry-metrics.json`.

separately qualified seed lineage exists in the record and is deliberately not substituted for it.

One diagnostic, at step 70, reports $Q_{\mathrm{FDTD}} = 436.3$ with a resonance-fit error of 0.0725, some two orders of magnitude above the surrounding scatter, together with a wavelength excursion of nearly 30 nm. It is excluded from every aggregate statistic by a fit-quality rule declared in advance—resonance error above 0.01—and is retained and marked in Fig. 5 rather than deleted. It is the only exclusion; the remaining 32 of 33 observations are used as recorded.

What the optimizer did. Unlike the first device, this is a restructuring of the cavity rather than a refinement of it (Figs. 3 and 4, lower row). The mean longitudinal semi-axis contracts from 76.157 nm to 45.96 nm while the mean transverse semi-axis expands from 76.166 nm to 118.16 nm, reaching 165.1 nm. The mean in-plane aspect ratio therefore moves from 1.0001—circular to four decimal places—to 2.5613, with a maximum of 3.112. Near-circular holes become transverse slots spanning up to 0.5938 of the beam width, leaving dielectric rails of 113 nm on each side, down from 187.6 nm. The mean hole *area* changes by only a factor of 0.9501: the optimizer reshaped the air inclusions at nearly constant air filling rather than removing or adding material. The mirror section also contracts, the outermost hole centre moving inward by 268 nm through the accumulation of small pitch reductions.

Spectral and modal signature. The restructuring is accompanied by a blueshift of 21.8 nm (792.347 nm $\to$ 770.497 nm) and by a halving of the effective mode volume, from 5.047 to 2.652 $(\lambda/n)^3$, a factor of 0.525. A concentration of field energy into the air slots at fixed period would produce exactly this pair of signatures, and the geometry change is what one would design if one wanted an air-type rather than a dielectric-type mode. We report this as a *characterized geometric and spectral restructuring consistent with* a change of mode character, and explicitly not as a proven band-edge swap: the present evidence is a correlation of hole shape, wavelength and mode volume, and the diagnostic that would settle it—a band structure of the terminal mirror cell, or a decomposition of the retained

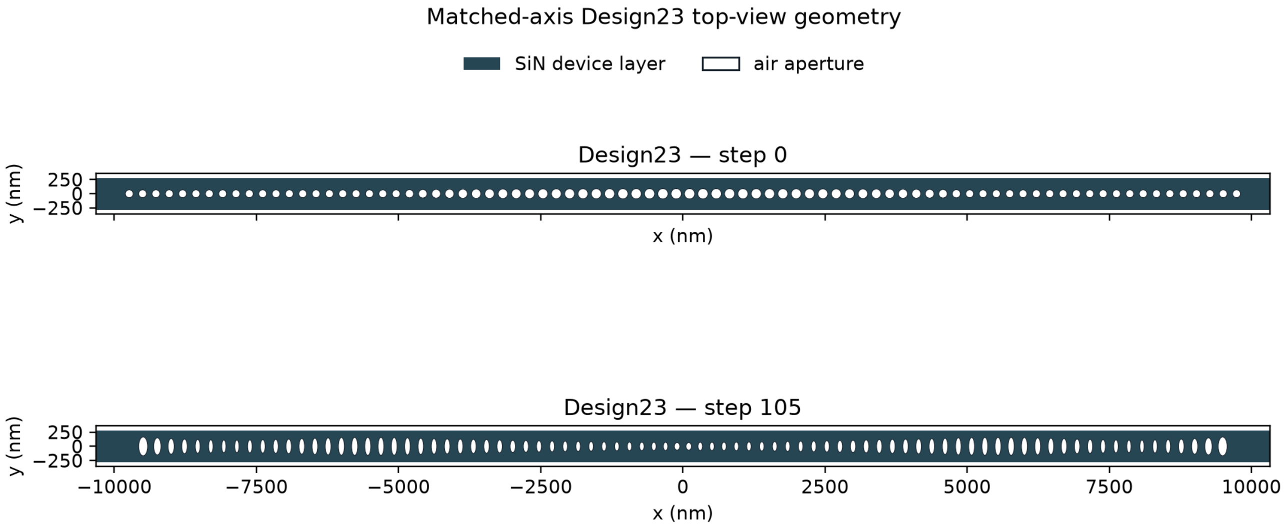

Figure 4: Matched-axis Design23 geometry before and after 105 accepted analytical steps. Si_3N_4 is blue, apertures white. The source is the pair of retained 128-entry geometry vectors (126 free coordinates plus the two fixed cross-section dimensions).

field energy between the dielectric and the air regions—was not part of this campaign. A useful counterpoint is that $\lambda/2p$ evaluated on the mean pitch is 1.689 at step 0 and 1.686 at step 105: the ratio of wavelength to period is essentially unchanged, so arithmetic on the period alone cannot establish a band swap either way.

Condition on this number. The two endpoints were *not* analyzed under matched settings: step 0 used a 12 ps ringdown and step 105 used 64 ps, i.e. $\mathcal{N} = 0.751$ and 1.34 lifetimes respectively. The mismatch is in the direction that would *inflate* the earlier value and therefore *deflate* the quoted gain, but it is a mismatch and the gain is provisional for that reason in addition to the panel requirement.

7 Directions transfer; magnitudes do not

The two campaigns provide an unusually direct test of what a truncated open-system model is good for, because the surrogate errs in *opposite directions* on the two devices and, more strikingly, in opposite directions at the two ends of a single trajectory.

Magnitudes do not transfer. On the first device the surrogate under-predicts the FDTD quality factor at both endpoints, by factors of 0.483 and 0.304. On the second it begins by under-predicting, $Q_{\mathrm{ana}}/Q_{\mathrm{FDTD}} = 0.472$ at step 0, and ends by over-predicting by more than an order of magnitude, 12.31 at step 105. Across the 32 qualified observations the ratio spans 0.472 to 12.98, a range of 27.5× (Fig. 6, right). No single calibration constant, and no monotone reparameterization of one observable into the other, is available. Any procedure that used the analytical Q as an estimate of the physical Q—for reporting, for stopping, or for accepting a step—would have been badly misled, in both directions, by this data.

Table 3: Agreement between the analytical surrogate and independent FDTD on the Design23 trajectory. Spearman ρ tests monotone ordering without assuming a shared scale; the sign test asks whether consecutive diagnostics move the same way, with a two-sided binomial p against a fair coin. Source: `results/surrogate-fidelity.json`.

	all qualified	12 ps	64 ps
points n	32	23	9
Spearman ρ	0.774	0.911	0.417
sign agreement	19/31	14/22	5/8
sign test p	0.28	0.29	0.73
$Q_{\mathrm{ana}}/Q_{\mathrm{FDTD}}$ span	27.5×	6.36×	1.17×

Part of that drift is a change in the FDTD observable rather than in the surrogate: the protocol changed from 12 ps to 64 ps between steps 67 and 80, and the Q values that a short ringdown assigns to a long-lived mode are biased high (Section 4). Splitting the trajectory at that boundary is therefore mandatory, and we report every aggregate for the two segments separately as well as together (Table 3). The bias drift survives the split: it runs from 0.472 to 3 within the 12 ps segment alone, before the protocol changes at all.

Ordering does transfer, at trajectory scale. Within the 12 ps segment—23 observations under one fixed FDTD protocol—the rank correlation between the two observables is $\rho = 0.911$. The surrogate therefore orders designs along the trajectory in close agreement with the independent solver even while mis-stating every indi-

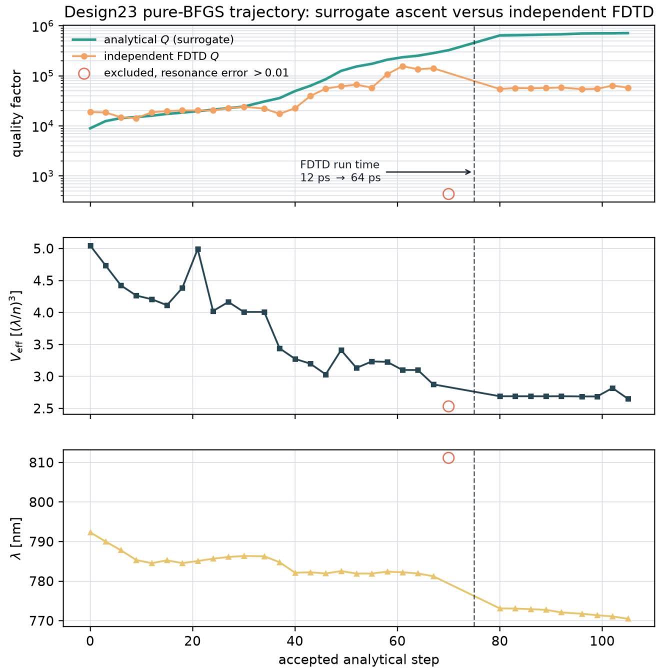

Figure 5: The 105-step trajectory. Top: analytical and independent FDTD quality factor on a logarithmic axis. Middle and bottom: effective mode volume and resonance wavelength. The dashed line marks the change of FDTD protocol from 12 ps to 64 ps; the open red marker is the single observation excluded by the predeclared fit-quality rule. Analytical and FDTD Q are distinct observables plotted on a shared axis for comparison of *trend*, not of value.

vidual value by a drifting factor. This, and not any absolute agreement, is what makes the model useful as an optimization objective.

But not step by step. We state the limit of this claim as precisely as the data allows. Consecutive diagnostics are three optimizer steps apart, and asking whether the two observables move in the same direction between them gives 19 agreements out of 31 comparisons—61 %, with a two-sided binomial p of 0.28 against a fair coin. Restricting to the larger moves does not help. The per-step sign agreement in this data set is therefore *not* distinguishable from chance, and we do not claim it. The direction-transfer evidence is the rank correlation over the trajectory together with the two matched endpoint comparisons, not a step-level correspondence: small analytical improvements are not resolvable against FDTD fit scatter, and a user of this method should expect to verify progress over tens of steps rather than one.

A sharper historical instance. The clearest single demonstration that the correct treatment of the continuum is what carries direction transfer comes from a repair to the second device's solver. A saved geometry direction was evaluated three ways: independent FDTD reported an endpoint ratio of 1.0007628—flat; the pre-repair an-

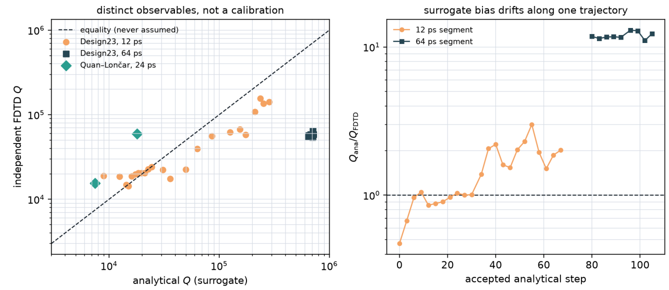

Figure 6: Left: analytical against independent FDTD quality factor on matched logarithmic axes. The diagonal is a reference line, not a model—the two axes carry different observables. Right: the ratio $Q_{\mathrm{ana}}/Q_{\mathrm{FDTD}}$ against optimizer step, showing that the surrogate bias is not a constant calibration factor but drifts by 27.5× along a single trajectory. Source: `results/surrogate-fidelity.json`.

alytical operator reported 3.6045899, an apparent 3.6× improvement; and the repaired pole-subtracted operator reported 1.0000101, correctly flat. The surrogate's value of a direction can be wrong by a factor of three and a half when the feed pole is mishandled, and it is the operator construction of Section 2—not a fitted correction—that makes it right.

8 Discussion

The practical conclusion of this work is narrower than numerical agreement between solvers, and more useful. A reduced outgoing model can supply productive ascent directions in a high-dimensional geometry space while its absolute linewidth is wrong by an order of magnitude—and while the sign and size of that error drift along the very trajectory it is driving. The two roles should therefore be separated by construction: analytical convergence certifies the local reduced problem and supplies the gradient, and a sparse independent solver certifies transfer and determines every reported physical number. Our results support that division and would not support any weaker one.

This has a concrete consequence for how such an optimization is run. Because the surrogate bias drifts, a stopping criterion or a step-acceptance test phrased in absolute analytical Q is unsafe; because per-step sign agreement with an independent solver is not resolvable (Section 7), a periodic independent check placed every few steps cannot certify individual steps either. What the evidence does support is independent verification at trajectory scale—at the endpoints, under matched settings, with a ringdown long compared with the fitted lifetime.

The two devices probe genuinely different failure modes, which is why running both matters more than optimizing either one further. In the air-clad cavity, basis order and the weak radiation channels dominate the error budget. In the multilayer cavity, the stratified continuum quadrature, the feed-pole subtraction, mode tracking through a large geometry change, and geometry transfer to an independent mesh are all additional and independent sources of error; the historical instance in Section 7, in which a mishandled feed pole turned a flat direction into an apparent 3.6× gain, shows how consequential that machinery is. Success on both is evidence for operator-level portability rather than for a second optimization of the same device.

The two campaigns also exercise different optimization regimes, and the contrast is informative in itself. The first device was refined: aspect ratios moved by one percent and the device family was preserved. The second was restructured: circular holes became transverse slots at nearly constant air area, the mode volume halved, and the resonance blueshifted by 21.8 nm. That a gradient built from an analytically continued continuum can carry a design that far from its seed—while an independent solver still confirms a gain—is the more surprising of the two outcomes, and it is also the one whose mode-character interpretation we have deliberately left open pending a diagnostic that can settle it.

What this work does not establish should be equally clear. It does not establish fabrication robustness, experimental quality factors, or agreement of the gradients themselves between the analytical and FDTD descriptions; each of those is a separate study. It does not establish converged absolute quality factors for either device: the panel of Table 2, amended for run time as described, is required for that, and until it is executed the two gains remain PROVISIONAL in the accompanying claims ledger. And it does not establish that the second device's mode changed character, only that its geometry, wavelength, and mode volume all moved together in the way such a change would produce.

9 Methods summary

Analytical poles are located by continuation from a known symmetry sector along the frozen contour homotopy of Eq. (8) and accepted only when the operator residual, the passive sign convention, complex symmetry, and the branch-overlap test all pass. Gradients are validated against both coordinate and random-direction central differences at the production discretization; the representative agreement is 6.71×10^{-6} relative on the Q derivative. Geometry optimization uses bound-constrained quasi-Newton updates on $\log Q$ with a trust region, a Sobolev boundary metric, and a line search that rejects any candidate whose branch overlap falls below 0.999. Exact thresholds, modal orders, quadrature rules, and optimizer settings are tabulated in Supplementary Sections S13 and S14.

FDTD is never called by the optimizer and never supplies a gradient. Frozen endpoint geometries are rendered into independent simulations and analyzed with the same resonance-fitting code at both ends of any comparison. The publication protocol, including the predeclared factorial panel and its promotion gate, is enumerated in Supplementary Section S15. No historical fit is promoted because it agrees with the analytical trend, and no analytical value is substituted for a physical one.

Data and code availability

The manuscript source, evidence tables, figure generators, complete derivation archive, and implementation snapshot are maintained in a provenance-preserving repository. Every number printed in this paper and its Supplementary Information is generated from a machine-readable record under `results/` and injected into the LaTeX source by a single exporter, so that a replaced measurement changes the document without any prose being retyped. Raw field arrays are retained outside version control and indexed by absolute path, size, and SHA-256 checksum. Reviewer and archival access, together with the released subset, will be specified at submission.

Acknowledgments

Author list, affiliations, funding, and computational-resource statements are in preparation. The authors will disclose the role of automated coding and language tools in accordance with the target venue's policy.

References

[1] Jakob S. Jensen and Ole Sigmund. Topology optimization for nano-photonics. *Laser & Photonics Reviews*, 5(2):308–321, 2011.

[2] Sean Molesky, Zin Lin, Alexander Y. Piggott, Weiliang Jin, Jelena Vuckovic, and Alejandro W. Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12:659–670, 2018.

[3] Tyler W. Hughes, Momchil Minkov, Ian A. D. Williamson, and Shanhui Fan. Adjoint method and inverse design for nonlinear nanophotonic devices. *ACS Photonics*, 5(12):4781–4787, 2018.

[4] Momchil Minkov, Ian A. D. Williamson, Lucio C. Andreani, Dario Gerace, Beicheng Lou, Alex Y. Song, Tyler W. Hughes, and Shanhui Fan. Inverse design of photonic crystals through automatic differentiation. *ACS Photonics*, 7(7):1729–1741, 2020.

[5] Momchil Minkov, Vincenzo Savona, and Dario Gerace. Photonic crystal slab cavity simultaneously optimized for ultra-high Q/V and vertical radiation coupling. *Applied Physics Letters*, 111(13):131104, 2017.

[6] Lucio Claudio Andreani and Dario Gerace. Photonic-crystal slabs with a triangular lattice of triangular holes investigated using a guided-mode expansion method. *Physical Review B*, 73:235114, 2006.

[7] Philippe Lalanne, Wei Yan, Kevin Vynck, Christophe Sauvan, and Jean-Paul Hugonin. Light interaction with photonic and plasmonic resonances. *Laser & Photonics Reviews*, 12(5):1700113, 2018.

[8] Christophe Sauvan, Jean-Paul Hugonin, Ivan S. Maksymov, and Philippe Lalanne. Theory of the spontaneous optical emission of nanosize photonic and plasmon resonators. *Physical Review Letters*, 110:237401, 2013.

[9] Philip Trost Kristensen, Kathrin Herrmann, Francesco Intravaia, and Kurt Busch. Modeling electromagnetic resonators using quasinormal modes. *Advances in Optics and Photonics*, 12(3):612–708, 2020.

[10] Philippe Lalanne, Wei Yan, Alexandre Gras, Christophe Sauvan, Jean-Paul Hugonin, Mondher Besbes, Guillaume Demesy, M. D. Truong, Boris Gralak, Frederic Zolla, Andre Nicolet, Felix Binkowski, Lin Zschiedrich, Sven Burger, Jorn Zimmerling, Rob Remis, Paul Urbach, H. T. Liu, and Thomas Weiss. Quasinormal mode solvers for resonators with dispersive materials. *Journal of the Optical Society of America A*, 36(4):686–704, 2019.

[11] Qimin Quan and Marko Lončar. Deterministic design of wavelength scale, ultra-high Q photonic crystal nanobeam cavities. *Optics Express*, 19(19):18529–18542, 2011.