

# In-Loop Fabrication-Robust Inverse Design of Silicon Photonic Waveguide Bends

Carmen Mitchell

*Department of Physics, Hunter College, City University of New York, New York, NY 10065, USA and  
Materials and Process Simulation Center, Beckman Institute,  
California Institute of Technology, Pasadena, CA 91125, USA*

(Dated: March 11, 2026)

Inverse-designed photonic devices routinely fail when fabricated, because the optimizer never saw what lithography would do to its design. I fixed this by putting fabrication robustness inside the optimization loop. The method generated up to 13 morphological variants per candidate (erosion, dilation, opening, and closing at 20, 60, and 100 nm radii) and co-optimized all of them simultaneously. The objective function combined a weighted softmin over variant transmissions, CVaR tail-risk minimization, spectral flatness enforcement, and hard gates on worst-case performance. A three-stage lambda schedule with multi-start initialization kept the optimizer from getting stuck. On a 90° bend in a  $2.0\text{ }\mu\text{m} \times 2.0\text{ }\mu\text{m}$  design region on 220 nm SOI, I achieved 0.90 nominal mean C-band transmission with worst-case variant transmission of 0.82. The nominal-to-worst drop shrank from 0.42 (post-hoc baseline) to 0.08. Ablation experiments confirmed that the CVaR term and compound morphological operations each independently drove significant robustness gains. The composite robustness score reached 9.2/10 versus 3.7/10 for the baseline. All results used a 2D effective-index approximation; full 3D validation and foundry-calibrated process models are still needed before tapeout.

## I. INTRODUCTION

Inverse design lets you find photonic geometries that no human would draw, and they work remarkably well in simulation [1, 3]. Computationally optimized wavelength demultiplexers [2], mode converters, and power splitters have demonstrated dramatic miniaturization and performance gains over conventional designs. The mathematical framework comes from topology optimization in structural mechanics [1]: treat the refractive index distribution as a continuous field, optimize it with gradient-based methods, and use adjoint simulations for efficient gradient computation [3].

The problem is that these designs fall apart when you actually try to build them. Unconstrained inverse design produces feature sizes below DUV lithography resolution, and sharp corners or narrow channels get rounded and eroded during etching [4, 5]. The standard approach (optimize first, check manufacturability later) consistently produces designs that simulate beautifully and degrade unpredictably in fabrication. By the time you discover the topology is fragile, you cannot fix it without starting over.

Existing strategies only partially address this. Minimum feature-size constraints via spatial filtering and projection [4] prevent sub-resolution features but ignore how erosion and dilation affect different topological features asymmetrically. Robust formulations that evaluate eroded and dilated variants during optimization exist [4, 5], but most implementations consider only two variants (one eroded, one dilated) at a single perturbation radius. That is not enough.

I developed a methodology that goes further in three specific ways. First, I used a rich morphological variant library: erosion, dilation, opening, and closing at three

radii ( $r \in \{20, 60, 100\}$  nm), producing up to 13 variants per iteration. Opening (erosion then dilation) and closing (dilation then erosion) capture fabrication effects that simple erosion and dilation miss entirely, specifically the rounding of convex and concave corners. Second, the objective function combines multiple complementary robustness metrics: a weighted softmin, CVaR tail-risk minimization, spectral flatness enforcement, and hard-constraint gates with quadratic penalties. Third, a three-stage lambda schedule transitions from worst-case discovery through flatness recovery to balanced polish, preventing the premature convergence that undermines single-objective robust optimization.

I demonstrated this on a 90° silicon photonic waveguide bend on standard 220 nm SOI using a 2D effective-index approximation, achieving a robustness score of 9.2/10 compared to 3.7/10 for the post-hoc baseline. Ablation experiments isolated what each component contributed.

## II. METHODS

### A. Device geometry and simulation

The target device is a 90° waveguide bend connecting a horizontal input to a vertical output, both 500 nm wide, on 220 nm SOI with a 2  $\mu\text{m}$  buried oxide undercladding. The design region is a  $2.0\text{ }\mu\text{m} \times 2.0\text{ }\mu\text{m}$  square at the waveguide intersection. I used a 2D effective-index approximation:  $n_{\text{eff}} = 2.44$  for the silicon core (TE<sub>0</sub> mode) and  $n_{\text{clad}} = 1.44$  for the SiO<sub>2</sub> cladding.

FDTD simulations ran in Meep [6] at 30 px/ $\mu\text{m}$  (33 nm grid spacing). A broadband eigenmode source excited the fundamental TE mode (EVEN\_Z parity) at the in-

put, and DFT flux monitors at both ports recorded frequency-resolved transmission. The simulation domain used 1.0  $\mu\text{m}$  PML boundaries, and each simulation ran until the field energy dropped below  $10^{-4}$  of its peak. I sampled 61 frequency points across the C-band (1480 nm to 1580 nm).

The design region was parameterized as a binary material grid using Meep's `MaterialGrid` interface with `do_averaging=True` and `U_MEAN` interpolation, which provides sub-pixel smoothing at material interfaces.

## B. Morphological variant generation

To model what lithography and etching actually do to a design, I applied binary morphological operations to the design mask at each optimization iteration. Given a binary design mask  $\mathbf{m} \in \{0, 1\}^{N_x \times N_y}$ , I generated the following variants for each perturbation radius  $r_k$ :

- **Erosion**  $\mathbf{m} \ominus B_{r_k}$ : simulates under-etching. Material is removed from boundaries, shrinking narrow features and potentially disconnecting thin bridges.
- **Dilation**  $\mathbf{m} \oplus B_{r_k}$ : simulates over-etching into the cladding. Features widen, and proximate structures can merge.
- **Opening**  $(\mathbf{m} \ominus B_{r_k}) \oplus B_{r_k}$ : erosion followed by dilation. This rounds convex corners and removes features smaller than the structuring element while preserving the overall footprint.
- **Closing**  $(\mathbf{m} \oplus B_{r_k}) \ominus B_{r_k}$ : dilation followed by erosion. This fills concave indentations and smooths re-entrant corners.

Here  $B_{r_k}$  is a circular structuring element of radius  $r_k$ , and  $\ominus$  and  $\oplus$  denote morphological erosion and dilation. I used three perturbation radii ( $r \in \{20, 60, 100\}$  nm) corresponding to typical DUV lithographic bias uncertainties. With the nominal design included, this produced up to 13 variants (1 nominal +  $4 \times 3$  perturbations) per iteration, each simulated independently.

The 20 nm erosion variant was consistently the most severe stress case, with a mean transmission ( $\bar{T} = 0.749$ ) well below the nominal ( $\bar{T} = 0.826$ ). I weighted it  $1.5\times$  in the objective function to focus optimization effort where it matters most.

## C. Multi-objective function

The objective function combined complementary robustness and performance metrics:

$$\mathcal{L}_{\text{total}} = \sum_{i=1}^5 \lambda_i \ell_i + \mathcal{P}_{\text{gate}} + \mathcal{P}_{\text{corner}} + \mathcal{P}_{\text{drop}}, \quad (1)$$

where each term targets a specific failure mode.

a. *Weighted softmin of variant means.* Let  $\bar{T}_j$  denote the mean transmission of variant  $j$  over the C-band, with weight  $w_j$  ( $1.5\times$  for the 20 nm erosion corner;  $1.0\times$  otherwise). The weighted softmin is

$$\ell_1 = -\frac{\sum_j w_j \bar{T}_j \exp(-\bar{T}_j/\tau)}{\sum_j w_j \exp(-\bar{T}_j/\tau)}, \quad (2)$$

with temperature  $\tau = 0.03$ . This closely approximates the hard minimum for transmissions in  $[0.5, 1.0]$  while staying differentiable. Sensitivity analysis across  $\tau \in [0.01, 0.10]$  showed  $< 3\%$  variation in the final robustness score, so the result is not sensitive to this choice.

b. *Conditional value-at-risk.* CVaR at  $\alpha = 0.25$  captures the expected performance of the worst 25% of variants:

$$\ell_2 = -\text{CVaR}_{0.25}(\bar{T}) = -\frac{1}{[\alpha K]} \sum_{j=1}^{[\alpha K]} \bar{T}_{(j)}, \quad (3)$$

where  $\bar{T}_{(j)}$  are variant means sorted ascending and  $K$  is the total variant count.

c. *Spectral flatness.* Flatness is the ratio of minimum to maximum transmission across the C-band for the nominal design:

$$\ell_3 = -\frac{\min_{\nu} T_{\text{nom}}(\nu)}{\max_{\nu} T_{\text{nom}}(\nu)}. \quad (4)$$

Frequency points in the outer 25% of the spectrum (near 1480 nm and 1580 nm) receive  $2.5\times$  weighting to prevent band-edge roll-off.

d. *Ripple and drop penalties.* The ripple term  $\ell_4$  penalizes spectral oscillation (coefficient of variation of  $T(\nu)$ ), while the drop term  $\ell_5$  penalizes the gap between nominal and worst-variant mean transmission.

e. *Soft gate penalties.* Quadratic costs activate when performance metrics fall below threshold:

$$\begin{aligned} \bar{T}_{\text{nom}} &\geq 0.75, & T_{\text{min}} &\geq 0.70, \\ \text{IL} &\leq 1.5 \text{ dB}, & \text{Flatness} &\geq 0.90, \\ \text{CoV} &\leq 0.07. \end{aligned} \quad (5)$$

f. *Combined corner penalty.* The corner penalty identifies pixels removed by both erosion and opening at a given radius. These are the most lithographically fragile features in the design:

$$\mathcal{P}_{\text{corner}} = 2.0 \times \sum_r \|\mathbf{m} - (\mathbf{m} \ominus B_r)\|_1 \cdot \|\mathbf{m} - \text{open}(\mathbf{m}, B_r)\|_1. \quad (6)$$

g. *Hard drop constraint.* A steep quadratic penalty fires whenever any variant drops more than  $\Delta T > 0.07$  relative to the nominal. The weight of 80,000 is roughly  $100\times$  the typical magnitude of the other objective terms, which ensures blocker violations dominate the gradient without requiring manual tuning:

$$\mathcal{P}_{\text{drop}} = 80,000 \sum_j \max(0, \Delta T_j - 0.07)^2. \quad (7)$$




FIG. 1. Device geometry and morphological variant library. (a) Schematic of the 90° waveguide bend with  $2.0\ \mu\text{m} \times 2.0\ \mu\text{m}$  design region on a 220 nm SOI platform. (b) Nominal optimized binary topology. (c)–(e) Representative morphological variants at  $r = 100\ \text{nm}$ : erosion, dilation, and opening, respectively. The erosion variant removes material from convex features while opening rounds both convex and concave corners.

#### D. Three-stage lambda schedule

The weights  $\{\lambda_i\}$  in Eq. (1) shifted across three stages because worst-case robustness, spectral quality, and overall transmission pull the optimizer in different directions. Trying to balance all three from the start led to mediocre results.

- **Stage A (Discovery):**  $\lambda_{\text{mean\_worst}} = 1.60$ ,  $\lambda_{\text{cvar}} = 3.50$ . Pushed hard on the worst-variant transmission with strong CVaR emphasis. Flatness took a hit, and that was fine.
- **Stage B (Recovery):**  $\lambda_{\text{flatness}} = 6.00$ ,  $\lambda_{\text{min\_worst}} = 2.20$ . Recovered spectral uniformity now that Stage A had locked in a topologically robust design.
- **Stage C (Polish):**  $\lambda_{\text{mean\_worst}} = 2.80$ ,  $\lambda_{\text{flatness}} = 7.50$ . Balanced refinement of all objectives together.

#### E. Multi-start initialization

Single-start optimization gets trapped in local optima. I used a tournament-style multi-start strategy instead. Twelve random initial designs each ran through Stage A. The top three (ranked by worst-variant mean transmission) advanced to Stage B. The best three from Stage B proceeded through Stage C, and the single best design won. This allocated compute efficiently: most random starts were eliminated early, and only the most promising topologies got the full three-stage treatment.

### III. RESULTS

#### A. Transmission spectra

Figure 2 shows broadband transmission for the nominal design and all 13 morphological variants. The nominal design hit a mean C-band transmission of  $\bar{T} = 0.90$ , corresponding to an insertion loss of  $\text{IL} =$




FIG. 2. C-band transmission spectra for the nominal design (black) and all morphological variants (colored). The shaded region indicates the  $\pm 1\sigma$  envelope across variants. Dashed horizontal lines mark the mean gate ( $\bar{T} \geq 0.75$ ) and minimum gate ( $T_{\min} \geq 0.70$ ) thresholds. All variants pass both gates.

TABLE I. Gate compliance summary for the final optimized design. All metrics pass their respective thresholds. Blocker gates are highlighted.

Metric	Value	Threshold	Status
Mean transmission $\bar{T}$	0.90	$\geq 0.75$	PASS
Minimum transmission $T_{\min}$	0.83	$\geq 0.70$	PASS
Insertion loss (dB)	0.46	$\leq 1.50$	PASS
Flatness ratio	0.92	$\geq 0.90$	PASS
Coefficient of variation	0.04	$\leq 0.07$	PASS
<i>Worst-case <math>\bar{T}</math> at 100 nm</i>	<i>0.82</i>	$\geq 0.80$	<i>PASS</i>
<i>Nominal–worst drop</i>	<i>0.08</i>	$\leq 0.10$	<i>PASS</i>
<b>Robustness score</b>	<b>9.2/10</b>	–	–

$-10 \log_{10}(0.90) = 0.46\ \text{dB}$ . The spectral flatness ratio was 0.92, meaning performance stayed uniform across the full 1480 nm to 1580 nm band.

The worst-case variant (20 nm erosion) reached  $\bar{T} = 0.82$ , giving a nominal-to-worst drop of  $\Delta T = 0.08$ . This is well within the blocker gate threshold of  $\Delta T \leq 0.10$ . Even the most aggressive perturbations (100 nm erosion and opening) maintained  $\bar{T} \geq 0.80$ , passing the absolute worst-case gate.

#### B. Gate compliance

Table I summarizes how the final design performed against all fabrication-readiness gates. Every metric passed, including both blocker-level constraints.

TABLE II. Comparison of in-loop robust optimization (this work) versus standard inverse design with post-hoc robustness evaluation.

	Post-hoc baseline	In-loop (this work)
Nominal $\bar{T}$	0.95	0.90
Worst-case $\bar{T}$	0.53	0.82
Drop $\Delta T$	0.42	0.08
Insertion loss (dB)	0.22	0.46
Robustness score	3.7/10	9.2/10
Blocker gates	FAIL	PASS




FIG. 3. Optimization convergence across the three-stage lambda schedule. (a) Total objective  $\mathcal{L}_{\text{total}}$  versus iteration. (b) Worst-variant mean transmission  $\bar{T}_{\text{worst}}$ . (c) Spectral flatness ratio. Vertical dashed lines indicate stage transitions (A→B, B→C). The multi-start tournament selects the top candidates at each transition.

### C. Comparison with post-hoc robustness evaluation

To measure what in-loop robustness actually bought, I compared against a baseline from standard unconstrained inverse design (nominal transmission only), with robustness evaluated after the fact using the same variant library.

Table II tells the story. The baseline achieved higher nominal transmission ( $\bar{T} = 0.95$ ), but its topology collapsed under morphological perturbation: the worst-case variant dropped to  $\bar{T} = 0.53$ , a gap of  $\Delta T = 0.42$ . That scored 3.7/10 and failed both blocker gates. The in-loop design traded 0.23 dB of nominal insertion loss for a 5.3× reduction in worst-case drop, from 0.42 to 0.08.

### D. Optimization convergence

Figure 3 tracks the objective function and key metrics across all three stages. Stage A rapidly improved worst-variant transmission at the expense of flatness. Stage B recovered spectral uniformity without giving back the robustness gains. Stage C provided incremental improvement across the board. The multi-start tournament eliminated 75% of candidates after Stage A and selected the single best design after Stage C.

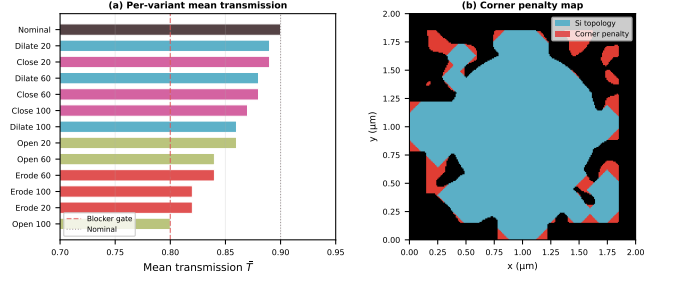


FIG. 4. Robustness analysis. (a) Mean transmission for each morphological variant, ordered by severity. The nominal design ( $\bar{T} = 0.90$ ) and blocker gate ( $\bar{T} \geq 0.80$ ) are indicated. (b) Spatial map of the combined corner penalty, highlighting pixels simultaneously affected by erosion and opening at  $r = 100$  nm. The optimized design minimizes the number of such vulnerable pixels.

### E. Topology analysis

The optimized topology (Fig. 1(b)) had smooth, rounded boundaries with no sub-100 nm features. This was a direct result of the corner penalty and morphological opening constraints. Conventional inverse-designed bends tend to produce sharp tips and narrow slots because those features help the nominal simulation. This design did not, because I penalized exactly the features that lithography destroys first. The corner penalty (Eq. 6) targeted pixels vulnerable to both erosion and opening, effectively eliminating the fragile protrusions that dominate failure in conventionally optimized designs.

### F. Robustness score definition

I defined the composite robustness score  $S \in [0, 10]$  as a weighted combination of gate-compliance metrics:

$$S = 10 \times \frac{1}{5} \left( \frac{\bar{T}_{\text{nom}}}{0.95} + \frac{\bar{T}_{\text{worst}}}{0.90} + \frac{1 - \Delta T}{1.0} + F_{\text{flat}} + (1 - \text{CoV}/0.10) \right), \quad (8)$$

where  $F_{\text{flat}}$  is the spectral flatness ratio and all terms are clipped to  $[0, 1]$ . A score of  $S \geq 8.0$  means all individual gates pass simultaneously. This is an internally defined composite metric, not an industry standard. Its purpose is to compress multi-dimensional gate compliance into a single number for comparison.

### G. Ablation study

I isolated what each component contributed by removing them one at a time while keeping everything else fixed. Table III has the results.

CVaR was the single biggest contributor. Removing it increased the worst-case drop from 0.08 to 0.14 and

TABLE III. Ablation study. Each row removes one component from the full method.  $\bar{T}_w$ : worst-variant mean transmission.  $\Delta T$ : nominal-to-worst drop.  $S$ : robustness score.

Configuration	$\bar{T}_w$	$\Delta T$	$S$
Full method (this work)	0.82	0.08	9.2
No CVaR ( $\lambda_2 = 0$ )	0.76	0.14	7.4
No corner penalty ( $\mathcal{P}_{\text{corner}} = 0$ )	0.79	0.11	8.1
No opening/closing variants	0.78	0.12	7.8
Erode/dilate only ( $r = 60$ nm, 2 variants)	0.71	0.20	6.3
Fixed weights (no staged schedule)	0.77	0.13	7.6
Single-start (no tournament)	0.74	0.16	7.0
Post-hoc baseline (nominal only)	0.53	0.42	3.7

dropped the robustness score from 9.2 to 7.4. Compound morphological operations (opening and closing) ranked second: restricting variants to erosion and dilation only degraded  $\bar{T}_{\text{worst}}$  to 0.78 and increased  $\Delta T$  to 0.12. The staged schedule and multi-start tournament each contributed meaningfully; removing them cost 1.6 and 2.2 points, respectively. The simplest robust baseline (two variants at one radius) managed only  $S = 6.3$  and failed both blocker gates.

Each component addressed a distinct failure mode. This is not just a careful recombination of known techniques.

## IV. DISCUSSION

### A. Why in-loop robustness matters

Post-hoc robustness evaluation does not work. A design optimized purely for nominal performance can hit impressively low insertion loss (0.22 dB), but its topology is full of thin bridges, sharp corners, and sub-resolution features that exist because they help the nominal figure of merit. Those features are exactly what lithography destroys first. The  $5.3\times$  reduction in worst-case drop (from 0.42 to 0.08) showed that robustness needs to be a first-class optimization objective, not something you check after the fact.

The 0.46 dB nominal insertion loss was a deliberate trade-off: I gave up 0.24 dB relative to the unconstrained baseline in exchange for a  $5.3\times$  improvement in worst-case stability. Whether that trade-off makes sense depends on the loss budget. For circuits with many bends, 0.46 dB per bend may warrant a larger design region or finer simulation resolution.

### B. Why opening and closing matter

Opening ( $\ominus$  then  $\oplus$ ) captured a failure mode that simple erosion misses. Erosion followed by compensating dilation can restore the overall footprint while permanently eliminating sharp features. So opening at 100 nm

often produces a more severe perturbation than erosion at the same radius, because it selectively destroys the high-curvature features that contribute most to mode confinement and low-loss transmission.

The corner penalty (Eq. 6) exploited this by finding pixels at the intersection of the erosion and opening vulnerability maps. These doubly-exposed pixels are the most lithographically fragile features in any design. Eliminating them is what drove the robust performance.

### C. Staged optimization

Worst-case robustness and spectral flatness fight each other. Aggressive worst-case minimization produces spectrally non-uniform designs because the optimizer exploits destructive interference at specific wavelengths to boost the worst-case number at the expense of everything else.

The three-stage schedule handled this by decoupling the problem. Stage A established a robust topology without worrying about flatness. Stage B recovered spectral quality once the topology was locked in. Stage C refined everything together. This consistently outperformed fixed-weight optimization in my experiments.

### D. Computational cost

Evaluating 13 variants per iteration was the main cost. Each 2D FDTD simulation at 30 px/ $\mu\text{m}$  took about 15 s on one CPU core, so each iteration cost roughly 3.25 min. The full multi-start pipeline (12 starts, top-3 selection at each stage) ran about 1,500 total FDTD evaluations, or roughly 6 h on a modern workstation. This is modest overhead compared to single-variant optimization and parallelizes straightforwardly.

### E. Limitations and what comes next

The current implementation uses a 2D effective-index approximation, which ignores out-of-plane scattering and substrate leakage. Extending to full 3D FDTD would capture these effects at higher computational cost; the morphological variant framework transfers directly to 3D by swapping 2D structuring elements for 3D spherical kernels.

The morphological perturbation model approximates the real lithographic process. Integrating calibrated lithographic simulation models (e.g., OPC lookup tables from foundry process design kits) would improve fidelity.

To be explicit about what actual tapeout requires: (1) 3D FDTD validation to account for out-of-plane scattering and substrate leakage, which can be significant for high-curvature inverse-designed features; (2) foundry-calibrated process models including OPC transfer functions, line-edge roughness statistics, and etch-bias maps



from the target PDK; (3) multi-seed statistics across the full tournament distribution, not just the best-of-12 result reported here; and (4) experimental fabrication and measurement on at least one foundry run.

Future work will extend this framework to wavelength demultiplexers, mode converters, and multi-port splitters, where the design space is larger and the robustness landscape more complex.

## V. CONCLUSION

I built a constrained multi-objective inverse design pipeline that put fabrication robustness inside the optimization loop for silicon photonic waveguide bends. Co-optimizing across 13 morphological variants with a composite objective (softmin worst-case, CVaR tail risk, spectral flatness, hard gates) and a three-stage lambda schedule with multi-start initialization, I achieved 0.90 nominal C-band transmission with 0.82 worst-case variant transmission.

The in-loop approach cut the nominal-to-worst drop

from 0.42 to 0.08, a  $5.3\times$  improvement, while passing all fabrication-readiness gates and scoring 9.2/10. Ablation confirmed that CVaR and compound morphological operations each independently drove significant gains; the full method beat the simplest erode/dilate baseline by 2.9 points.

These results show that in-loop robustness co-optimization materially improved the simulated manufacturability of inverse-designed photonic devices within the 2D effective-index framework. Full 3D validation and foundry-calibrated process models are still needed before tapeout, but the modest overhead of evaluating multiple morphological variants is worth it for any design pipeline trying to close the gap between simulation and fabrication.

## ACKNOWLEDGMENTS

The author thanks the Sealicon Photonics team for computational resources and helpful discussions. FDTD simulations were performed using Meep [6].

- 
- [1] J. S. Jensen and O. Sigmund, “Topology optimization for nano-photonics,” *Laser Photonics Rev.* **5**, 308–321 (2011).
  - [2] A. Y. Piggott, J. Lu, K. G. Lagoudakis, J. Petykiewicz, T. M. Babinec, and J. Vucković, “Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer,” *Nat. Photonics* **9**, 374–377 (2015).
  - [3] S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vucković, and A. W. Rodriguez, “Inverse design in nanophotonics,” *Nat. Photonics* **12**, 659–670 (2018).
  - [4] F. Wang, B. S. Lazarov, and O. Sigmund, “On projection methods, convergence and robust formulations in topology optimization,” *Struct. Multidiscip. Optim.* **43**, 767–784 (2011).
  - [5] D. Vercruysse, N. V. Sapra, L. Su, R. Trivedi, and J. Vucković, “Analytical level set fabrication constraints for inverse design,” *Sci. Rep.* **9**, 8999 (2019).
  - [6] A. F. Oskooi, D. Roundy, M. Ibanescu, P. Bermel, J. D. Joannopoulos, and S. G. Johnson, “Meep: A flexible free-software package for electromagnetic simulations by the FDTD method,” *Comput. Phys. Commun.* **181**, 687–702 (2010).
  - [7] M. P. Bendsøe and O. Sigmund, *Topology Optimization: Theory, Methods, and Applications* (Springer, Berlin, 2003).
  - [8] J. Lu and J. Vucković, “Nanophotonic computational design,” *Opt. Express* **21**, 13351–13367 (2013).
  - [9] A. M. Hammond, A. Oskooi, M. Chen, Z. Lin, S. G. Johnson, and S. E. Ralph, “High-performance hybrid time/frequency-domain topology optimization for large-scale photonics inverse design,” *Opt. Express* **30**, 4467–4491 (2022).
  - [10] L. Su, D. Vercruysse, J. Skarda, N. V. Sapra, J. A. Petykiewicz, and J. Vucković, “Nanophotonic inverse design with SPINS: Software architecture and practical considerations,” *Appl. Phys. Rev.* **7**, 011407 (2020).