

PRISM: Photonics-Informed Inverse Lithography for Manufacturable Inverse-Designed Photonic Integrated Circuits

HONGJIAN ZHOU, Arizona State University, USA

HAOYU YANG, NVIDIA Corp., USA

NICHOLAS GANGI, Rensselaer Polytechnic Institute, USA

TIANLE XU, Rensselaer Polytechnic Institute, USA

RENA HUANG, Rensselaer Polytechnic Institute, USA

JIAQI GU*, Arizona State University, USA

Photonic integrated circuits (PICs) are emerging as a key hardware platform for high-bandwidth, energy-efficient optical interconnects and for computing and sensing workloads increasingly relevant to large-scale AI systems. Recent advances in photonic inverse design have demonstrated the ability to automatically synthesize compact, high-performance photonic components that surpass conventional, hand-designed structures, offering a promising path toward scalable and functionality-rich photonic hardware. However, the practical deployment of inverse-designed PICs is bottlenecked by manufacturability: their irregular, subwavelength geometries are highly sensitive to fabrication variations, leading to large performance degradation, low yield, and a persistent gap between simulated optimality and fabricated performance. Unlike electronics, where foundries provide mature design-for-manufacturing (DFM) pipelines that correct masks and shield designers from process details, photonics lacks a systematic, flexible mask optimization flow. Fabrication deviations in photonic components cause large optical response drift and compounding error in cascaded circuits, while calibrating fabrication models remains costly and expertise-heavy, often requiring repeated fabrication cycles that are inaccessible to most designers. To bridge this gap, we introduce PRISM, a photonics-informed inverse lithography workflow that makes photonic mask optimization data-efficient, reliable, and optics-informed. PRISM (i) synthesizes compact, informative calibration patterns to minimize required fabrication data, (ii) trains a physics-grounded differentiable fabrication digital twin, enabling gradient-based optimization, and (iii) performs photonics-informed inverse mask optimization that prioritizes performance-critical features beyond geometry matching. Across multiple inverse-designed components with both electron-beam lithography and deep ultra-violet photolithography processes, PRISM significantly boosts post-fabrication performance and yield while reducing calibration area and turnaround time, enabling and democratizing manufacturable and high-yield inverse-designed photonic hardware at scale.

CCS Concepts: • **Hardware** → **Physical design (EDA)**; **Design for manufacturability**; **Emerging optical and photonic technologies**; Photonic and optical interconnect.

*Corresponding author.

Authors' Contact Information: Hongjian Zhou, hzhou144@asu.edu, Arizona State University, Tempe, Arizona, USA; Haoyu Yang, NVIDIA Corp., Austin, Texas, USA, haoyuy@nvidia.com; Nicholas Gangi, gangin2@rpi.edu, Rensselaer Polytechnic Institute, Troy, New York, USA; Tianle Xu, Xut6@rpi.edu, Rensselaer Polytechnic Institute, Troy, New York, USA; Rena Huang, HUANGZ3@rpi.edu, Rensselaer Polytechnic Institute, Troy, New York, USA; Jiaqi Gu, jiaqigu@asu.edu, Arizona State University, Tempe, Arizona, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 1557-735X/2018/8-ART111

<https://doi.org/XXXXXXX.XXXXXXX>

ACM Reference Format:

Hongjian Zhou, Haoyu Yang, Nicholas Gangi, Tianle Xu, Rena Huang, and Jiaqi Gu. 2018. PRISM: Photonics-Informed Inverse Lithography for Manufacturable Inverse-Designed Photonic Integrated Circuits. *J. ACM* 37, 4, Article 111 (August 2018), 26 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Photonic integrated circuits (PICs) are emerging as a foundational hardware platform for high-bandwidth communication [1, 2], optical interconnects [1–5], and speed-of-light computing [6–9], with growing adoption in various key use domains, e.g., artificial intelligence (AI) systems, sensing, quantum, scientific discovery, etc. Yet despite this promise, the practical design and deployment of **complex PICs at scale remain constrained by a manufacturing reality**: *what is fabricated on wafer can deviate substantially from the intended design*. For subwavelength photonic structures, e.g., gratings, inverse-designed patterns, these fine-grained structures often suffer from huge geometric deviations leading to large performance degradation and malfunction, as illustrated in Fig. 1a. This becomes especially problematic in cascaded circuits, where small component-level errors can accumulate into system-level failure and make inverse-designed layouts **difficult to manufacture reliably at scale** [10, 11].

In production-oriented deep ultraviolet (DUV) fabrication processes in photonic foundries, e.g., 193 nm technology node, inverse-designed patterns can, in general, be considered *infeasible to manufacture* due to severe distortions (e.g., rounding, blurring, bridging) that lead to **near-zero yield**. Electron-beam lithography (EBL) can provide higher patterning resolution for laboratory demonstrations, but it remains costly and slow for high-volume manufacturing and still exhibits non-ideal distortion. As illustrated in Fig. 1b, DUV fabrication induces substantially larger figure of merit (FoM) degradation and a heavier worst-case tail than EBL across representative photonic devices, underscoring the manufacturability gap for inverse-designed layouts under DUV. Figure 1c further shows that the wavelength division multiplexing transmission is jointly sensitive to lithography defocus and resist threshold, exhibiting a narrow high-performance process window. Today, designers often cope through **expensive trial-and-error**: imposing conservative design rules, sacrificing geometric degrees of freedom with smoothing, running extensive parameter sweeps to search for a feasible design, and iteratively re-taping-out designs based on sparse fabrication feedback. This workflow is slow (year-long), high-cost, expertise-heavy, and fundamentally at odds with the promise of inverse design: *it either sacrifices performance or fails to achieve robust yield*.

A natural question is *why photonics cannot simply adopt the mature design-for-manufacturing (DFM) [12] stack from electronic design automation (EDA) [13]*. In electronics, foundries routinely apply optical proximity correction (OPC) [14] and inverse lithography (ILT) [15] to optimize masks and compensate for process effects. In photonics, **two obstacles break this direct transfer**. ❶ Photonic hardware performance relies on complicated light propagation behavior; mask correction objectives based purely on geometric similarity are not sufficient, i.e., “**close geometry**” **does not imply “close function.”** Optical response arises from nonlocal electromagnetic interactions (scattering, interference, resonance, etc) and is often dominated by a subset of sensitivity-critical features, so errors in a small region can dominate performance even when global pixel-wise error is small, illustrated in Fig. 2a. Moreover, *systematic* geometric biases, such as globally biased over-/under-etch, can be especially detrimental because they coherently shift effective widths, duty cycles, and phase accumulation across the device, leading to pronounced performance degradation even with a lower L_2 fabrication error, as shown in Fig. 2b. ❷ **Accurate lithography modeling is expensive and process-specific**: foundry details are often proprietary, and purely data-driven fabrication surrogates can require large calibration datasets. When trained with limited data, black-box neural models may produce unreliable gradients, undermining gradient-based inverse




Fig. 1. Comprehensive motivation. (a) Fabrication error causes severe performance and yield degradation on inverse-designed photonic components. (b) DUV fabrication is particularly challenging for inverse-designed photonic devices, often causing substantially larger performance degradation and lower yield than EBL. Mean and std. of FoMs are marked. (c) Sensitivity map of device transmission under coupled process variations. The heatmap shows the transmission as a function of defocus coefficient and resist threshold, highlighting their joint impact on the post-fabrication response.

lithography optimization where stability depends on smoothness, generalization, and physically plausible sensitivities, as shown in Fig. 2c. These gaps motivate a new approach: a **data-efficient, physically grounded, and photonics-informed inverse lithography workflow** that closes the gap between design and manufacturing by optimizing masks for photonic function fidelity beyond pattern matching, enabling high-yield fabrication capability of inverse-designed PICs for foundries while democratizing model calibration and mask optimization through low-cost, easily adaptable customization.

In this work, we introduce PRISM, a photonics-informed neural inverse lithography framework that turns inverse-designed PIC layouts into high-yield, manufacturable hardware. PRISM workflow unified three synergistic innovations: (1) **data-efficient, identifiability-driven calibration pattern synthesis** strategy that generates compact, informative fabrication calibration patterns for training data acquisition. (2) a **physics-grounded neural fabrication model** with superior generalization and provides stable gradients for mask optimization; and (3) **photonics-informed ILT** that incorporates optical response as guidance to preserve performance-critical features beyond geometry matching. Based on the approaches in PRISM, we make the following contributions:

- **Identifiability-driven Calibration Pattern Synthesis** – We constructs a compact calibration reticle by combining (i) interpretable rectilinear/curvilinear regular motifs as controlled probes, (ii) randomized curvilinear patterns with controlled spectra for broad frequency/proximity coverage, and (iii) inverse-designed device crops for distribution matching, enabling accurate digital-twin calibration under a strict SEM/area budget.
- **Physics-grounded Differentiable Fabrication Modeling** – We train a differentiable fabrication digital twin model that generalizes under limited data and provides stable gradients for mask optimization across patterns.
- **Photonics-informed Inverse Lithography** – We formulate ILT as a function-preserving mask optimization problem and introduce EM-field-aware objectives to preserve photonic function beyond geometry fidelity.
- Across multiple inverse-designed components in EBL and 193 nm DUV photolithography, PRISM significantly boosts post-fabrication device performance and yield with substantially lower calibration cost, unlocking manufacturable inverse-designed photonic hardware at scale.




Fig. 2. Comprehensive motivation. (a) Standard mask correction is insufficient, as *close geometry* does not always lead to close optical response. (b) Low pixel-wise ℓ_2 error does not necessarily imply good photonic performance. An *unbalanced* error distribution dominated by systematic over-/under-etch introduces a global bias in geometry and can degrade the FoM significantly, despite a comparable L_2 value. (c) Neural network fabrication model trained on calibration data might have poor generalization to unseen patterns and mislead ILT optimization.

2 Preliminary

2.1 Photonic Inverse Design and Fabrication Awareness

Photonic inverse design has emerged as a powerful paradigm for synthesizing compact, high-performance photonic components by directly optimizing electromagnetic objectives, enabling structures that outperform conventional hand-designed devices. A growing body of work has extended inverse design to account for fabrication constraints, including minimum feature size control [16–19], smoothing or regularization of design variables, and robustness-aware objectives that mitigate sensitivity to process variation. These approaches have improved the practicality of inverse-designed devices by shaping the design space itself. Our work is *complementary rather than competing with fabrication-aware inverse design (FAID)* [20]. We assume a given inverse-designed layout, produced by any existing design methodology, and focus on the downstream manufacturing mask optimization. PRISM can be combined with FAID to further improve yield, or applied independently to existing inverse-designed components without redesign. To motivate mask optimization, we briefly summarize the manufacturing steps that shape the printed pattern and how they are typically modeled. Figure 3 sketches a standard pattern-transfer flow.

2.2 Fabrication Process Modeling

In modern chip fabrication, a complete process flow often consists of hundreds of steps. Patterned masks are transferred to the wafer through a sequence of processes, including lithography, development, ion implantation, etching, and polishing, all of which collectively affect the final pattern geometry on the wafer. Among these steps, lithography is one of the most critical stages because it largely determines the achievable resolution and is therefore crucial to both chip performance and manufacturing yield. In this section, we briefly introduce several key processes and discuss how they are commonly modeled and simulated in practice. Throughout this paper, we use lowercase letters for scalars, bold lowercase for vectors, $\odot$ for element-wise production, $\otimes$ for convolution, $*$ for complex conjugate, M , I , Z , W for mask, aerial, resist and final wafer image, $\mathcal{F}$ and $\mathcal{F}^{-1}$ for fast Fourier transform (FFT) and inverse FFT, respectively.




Fig. 3. Illustration of a typical pattern transfer flow in chip fabrication. Patterns are first defined by either photolithography or E-beam lithography to form a resist image on an oxide-coated silicon wafer, followed by (2) resist development, (3) etching and resist removal to transfer the pattern into the underlying film, (4) material deposition (e.g., Cu fill), and (5) chemical-mechanical polishing to planarize the surface.

Photolithography Process. The lithography simulation models are designed to mimic the printing effects without performing actual lithography. In practice, the Hopkins diffraction model [21] of the partial coherence imaging system is widely adopted to approximate the printing behavior of the lithography. Its purpose is a separation of the influence of the mask and the imaging system, including the pupil function and the illumination. For numerical convenience, the Hopkins method is often stated in terms of the spatial spectrum of the aerial intensity I :

$$\mathcal{F}(I)(f, g) = \iint_{-\infty}^{\infty} \mathcal{T}((f' + f, g' + g), (f', g')) \mathcal{F}(M)(f' + f, g' + g) \mathcal{F}(M)^*(f', g') df' dg' \quad (1)$$

$$\mathcal{T}((f', g'), (f'', g'')) := \iint_{-\infty}^{\infty} \mathcal{F}(J)(f, g) \mathcal{F}(H)(f + f', g + g') \mathcal{F}(H)^*(f + f'', g + g'') df dg \quad (2)$$

where $\mathcal{T}$ is the Transmission Cross-Coefficient (TCC), the weighting factor J depends only on the effective source, and H denotes the projector transfer function. From Eqs. (1) and (2), it can be seen that $\mathcal{T}$ is independent of the mask transmission function $\mathcal{F}(M)$.

To reduce the computational overhead, a singular value decomposition (SVD) approximation is typically adopted for the Hopkins model, called the sum of coherent sources (SOCS) [22]. The basic idea is to take the SVD of the coefficient matrix. The TCC spectrum matrix can be written as $\mathcal{F}(\mathcal{T}) = \sum_i \alpha_i h_i h_i^*$, where h_i are the eigenvectors and α_i are the eigenvalues of the TCC spectrum $\mathcal{F}(\mathcal{T})$. Applying the inverse Fourier transform then yields aerial intensity by the SOCS formulation:

$$I = \sum_i \alpha_i |\mathcal{F}^{-1}(\mathcal{F}(h_i) \odot \mathcal{F}(M))|^2. \quad (3)$$

E-beam Lithography Process. Unlike photolithography, electron-beam lithography (EBL) is a maskless direct-write process whose pattern fidelity is mainly governed by electron-matter scattering in the resist/substrate stack (including forward- and back-scattering), which leads to the well-known proximity effect [23]. Therefore, EBL process modeling typically starts from a physical scattering model, where Monte Carlo simulation is widely used to track electron trajectories in the resist and substrate and to estimate the spatial distribution of deposited energy. For computational efficiency, the Monte Carlo results are often summarized by a point spread function (PSF), which is fitted using parametric models such as a double-Gaussian [24] to capture both short-range forward scattering and long-range backscattering contributions:

$$\text{PSF}(r) = \frac{\eta}{\pi\alpha^2} \exp\left(-\frac{r^2}{\alpha^2}\right) + \frac{1-\eta}{\pi\beta^2} \exp\left(-\frac{r^2}{\beta^2}\right), \quad (4)$$

where r denotes the radial distance from the beam incidence point, α and β are the characteristic scattering ranges (Gaussian widths) for the forward- and back-scattering components, respectively, and $\eta \in [0, 1]$ is the weight of the forward-scattering term (thus $1 - \eta$ corresponds to the backscattering weight). The normalization factors ensure that each Gaussian term integrates to its corresponding weight over the 2D plane, so that the total PSF integrates to unity. Given an incident dose map derived from the layout, the energy deposition density can then be computed by convolving the dose map with the PSF.

Development Process. After exposure, development models predict the final resist profile (e.g., critical dimension, sidewall angle, resist loss) from the aerial image. A widely used approach is to convert the aerial image into a latent image (effective exposure in resists) [25], then use a rate-based dissolution model [26] to describe how fast each resist location dissolves in developer. In practice, this is often implemented as: (1) a fixed or variable threshold/contour model [27] for fast approximation; or (2) a more accurate dissolution-rate model where the local dissolution speed depends on exposure level, capturing effects such as resist contrast, diffusion/blur, and post-exposure bake smoothing.

Etching Process. Given the developed resist pattern, etch modeling describes how it transfers into films, predicting etch depth, CD bias, and profile distortions. In most process simulators, etching is formulated as a moving-boundary/surface-evolution problem [28], where the material surface advances at a local etch rate determined by simplified plasma physics and feature geometry.

Deposition Process. After pattern transfer, deposition models [29] predict conformality and topography growth on 3D features (e.g., step coverage and gap fill). Most models combine a global deposition rate with geometry-dependent effects (orientation, shadowing, and transport/reaction limits in narrow trenches), and iteratively update the surface to estimate post-deposition thickness nonuniformity.

Polishing Process. To recover planarity for subsequent layers, chemical-mechanical polishing (CMP) models [30] predict post-polish topography (e.g., dishing and erosion), which is critical for multilayer patterning and lithography focus control. Typical models combine a global removal component (pressure, speed, slurry) with layout-dependent effects driven by local pattern density; simulators often compute an effective local contact/pressure from a windowed density map and iteratively update the surface to obtain post-CMP thickness and planarity metrics.

2.3 Inverse Lithography in EDA and Uniqueness in Photonic ILT

The lithography process is a critical step in chip manufacturing, as it transfers the mask pattern onto the wafer and largely determines the critical dimension, performance, and yield. In EDA, DFM techniques such as optical proximity correction and ILT are mature and widely deployed, forming a closed-loop interface between design and manufacturing. Recent work has leveraged data-driven and neural surrogates [31–34] to accelerate lithography modeling and enable gradient-based ILT [35–40] under increasingly complex process conditions [41]. These methods have been highly effective for rectilinear, logic-style layouts and geometry-centric objectives.

However, photonic ILT introduces several unique challenges and motivations. (i) Inverse-designed photonic layouts are often *not manufacturing-feasible* under production DUV nodes: severe diffraction-driven distortions (rounding, blurring, bridging) can cause dramatic FoM degradation and a heavy worst-case tail, leading to near-zero yield for some devices. (ii) Unlike logic OPC, where geometric fidelity is a strong proxy, photonic performance is governed by nonlocal electromagnetic interactions; thus, “**close geometry**” **does not imply** “**close function**.” Small but structured errors or systematic biases (e.g., global over-/under-etch) can dominate optical response even when pixel-wise error is low. (iii) Accurate lithography modeling remains expensive and highly process-specific. Detailed foundry information (e.g., calibrated optical process parameters and OPC recipes) is often




Fig. 4. Overview of proposed PRISM flow to accelerate photonic inverse design convergence.

proprietary and not readily accessible to designers, *especially when most photonic foundries do not provide OPC/ILT services*. This creates a major barrier to designing high-performance, compact, inverse-designed photonic devices, which are typically highly sensitive to fabrication imperfections. The lack of a reliable process model also hinders fabrication-aware inverse design, making model acquisition a central challenge. A faithful model benefits both sides: it helps foundries diagnose and improve processes, enhance yield, and develop advanced PDKs, while enabling designers to produce robust photonic devices under realistic manufacturing non-ideality.

To address the aforementioned problem, PRISM brings the power of ILT to integrated photonic circuits with photonics-specific adaptation by incorporating physics/optics in the data synthesis, model design, and mask optimization workflow, bridging the simulation-fabrication gap in photonic inverse design through a photonics-aware design-technology co-optimization loop.

3 Proposed PRISM Framework

Inverse-designed photonics breaks two assumptions that standard EDA ILT relies on: (i) mask geometric fidelity is the sole objective, and (ii) process modeling can be treated as a solved, foundry-provided service. In photonics, neither holds. Optical function is governed by a small set of sensitivity-critical structures, so minimizing pixel-wise geometry error can still yield a large performance drop; meanwhile, many photonic designers do not have access to calibrated OPC recipes or high-volume process characterization. PRISM is designed around a single goal: **turn photonic curvilinear mask optimization into a data-efficient, stable, and optical function-preserving design-technology co-optimization loop that both photonic designers and foundries can benefit from low-cost yield augmentation**.

Figure 4 encapsulates the key paradigm of PRISM to achieve this goal: close the loop between a locally calibratable fabrication digital twin and a physics-guided inverse mask optimization engine, so that photonic ILT can become a first-class, accessible new design stage introduced into general yield-oriented PIC design closure.

The remainder of this section details the three pillars of Fig. 4. In Sec. 3.1, we synthesize a calibration set for identifiability under budget constraints: a compact pattern dataset that prioritizes curvilinear motifs and proximity interactions that are most informative for inverse-designed photonics, maximizing process learnability per unit cost in fabrication, metrology, and labeling. In Sec. 3.2, we learn a physically-grounded, parametric fabrication digital twin that captures complicated stochastic fabrication processes and, crucially, provides stable gradients to guide inverse mask optimization. Finally, Sec. 3.3 closes the loop by integrating the differentiable fabrication digital twin into a physics-guided photonic ILT procedure that operates in latent topology space, enforcing manufacturability while explicitly protecting performance-critical features under process variations, beyond deterministic pixel-wise geometry matching. The ILT-corrected layouts are then used for downstream simulation, verification, and yield assessment, and can be iteratively refined, together with the fabrication digital twin model, enabling accelerated design convergence.




Fig. 5. (a) Illustration of regular, random, and photonic device patterns for fabrication model learning. (b) Design images and post-fab wafer images form data/label pairs for fabrication model training.

3.1 Curvilinear Calibration Pattern Synthesis for Photonic ILT

A practical bottleneck for photonic ILT is **calibration**: *learning an accurate fabrication digital twin from limited post-fab measurements*, since chip fabrication area, metrology (e.g., SEM imaging), and label annotation are all expensive. PRISM formulates calibration as an identifiability-driven pattern synthesis problem, we aim to select a compact set of layouts whose fabricated outcomes maximally reflect the dominant fabrication distortions encountered in inverse-designed photonic masks.

Complementary Pattern Synthesis Strategies in PRISM. PRISM provides three calibration pattern categories, each targeting a distinct aspect of the process response:

- ① **Regular patterns as controlled probes:** We include both rectilinear and curvilinear photonic motifs (lines/slots, bends, tapers, junctions, corners) and their compositions. These patterns provide *interpretable* responses and isolate specific mechanisms (e.g., linewidth bias, corner rounding, curvature-dependent bias), making them efficient anchors for model learning.
- ② **Randomized freeform patterns as spectral coverage probes:** Inverse-designed photonics contains rich multi-scale structure; capturing this requires coverage across spatial frequencies and feature sizes. We therefore generate randomized curvilinear patterns with controlled power spectra to span a broad frequency distribution, from smooth low-frequency contours to high-frequency edge/corner detail, while varying feature widths and spacings. These patterns efficiently expose non-local proximity interactions and process nonlinearities that are difficult to excite using uncontrolled random patterns.
- ③ **Inverse-designed device crops as distribution-matching probes:** Finally, we include patterns sampled directly from inverse-designed devices (or cropped regions of interest). This aligns calibration with the true data distribution seen at deployment, reducing domain shift and improving on-task accuracy when the goal is yield augmentation for inverse-designed devices.

Pattern Layout Arrangement. All patterns are tiled into a compact calibration reticle with controlled spacing. As shown in Fig. 5a, each pattern occupies a $3 \mu\text{m} \times 3 \mu\text{m}$ window, and adjacent patterns are separated by a $2 \mu\text{m}$ spacing to mitigate mutual interference (e.g., proximity and loading effects). The pattern size is chosen based on the SEM field-of-view (typically a few microns), so that each SEM image can fully capture one pattern reticle, thereby reducing the need for precise alignment during SEM image stitching.

"Virtual Foundry" as Controlled Study Vehicle. To systematically evaluate PRISM across pattern choices and model designs, we create a *virtual foundry* as a controlled proxy to represent an

arbitrary fabrication flow. The virtual foundry serves as a synthetic fabrication data generator/oracle, producing paired mask–wafer examples with adjustable process settings and variations, allowing us to (i) study generalization under limited calibration, (ii) isolate how calibration distributions affect downstream ILT, and (iii) conduct algorithmic ablation studies that would be prohibitively expensive in real fabrication. Importantly, the virtual foundry is not a claim about any specific foundry; it is a *blackbox hidden testbed* for stress-testing the proposed calibration methodology and quantifying design-process co-optimization behavior.

For the DUV fabrication setting, we emulate the photolithography step using the sum of coherence system (SOCS) representation of the Hopkins imaging model, a standard abstraction in computational lithography that is both physically grounded and computationally efficient. Specifically, we construct the DUV lithography model using *Torchlitho* [42], an open-source framework for computing the transmission cross-coefficient (TCC) kernel from the illumination source J and the projector system transfer function H . The model parameters, including the wavelength, numerical aperture (NA), and related process parameters, are calibrated based on preliminary SEM data collected from our previous AIM Photonics MPW SiPho chip tape-out fabricated with 193-nm DUV lithography. The calibrated Hopkins-based model fits the annotated SEM labels reasonably well, with acceptable residual errors, suggesting that it captures the dominant lithography-induced distortions observed in real foundry data to some extent. Using the decomposed SOCS kernels, we then evaluate the aerial image intensity via Eq. (3). We obtain the developed resist pattern by thresholding the aerial image, followed by a parametrized pattern-transfer surrogate that emulates common non-ideal fabrication effects, including shot-to-shot dose variation, stochastic resist/etch defects, resolution-limited line bridging, and curvature-dependent etching threshold variation. Concretely, we apply a set of morphological and geometric operators that form a minimal “distortion basis.” We remove small isolated resist islands, discarding foreground components whose area is below a preset limit. It models stochastic resist residues and defect-like speckles that would not reliably transfer through etching. We then apply the morphological *closing* operation to seal narrow gaps and form unintended connections, mimicking *bridging* phenomena caused by insufficient resolution, resist scumming, or etch microloading. Next, we apply the morphological *opening* operation to eliminate tiny protrusions and thin connections, which approximates the breakage of fragile features and the removal of narrow resist necks during pattern transfer. Finally, we model systematic curvature-dependent bias (e.g., corner rounding and over-/under-etch) by computing a signed distance field and offsetting the effective etch front by a curvature-dependent shift, so highly curved boundaries experience stronger local bias than near-straight segments.

For the EBL fabrication setting, the exposure step is modeled by an energy deposition map produced by electron forward scattering in the resist and long-range backscattering from the substrate. The Monte-Carlo scattering results can be summarized by a parametric point spread function (PSF). Specifically, we use the *three-Gaussian plus exponential (3G+exp)* model [43]:

$$P(r) = \frac{1}{\pi(1 + \eta + \eta' + \eta'')} \left[\frac{1}{\alpha^2} \exp\left(-\frac{r^2}{\alpha^2}\right) + \frac{\eta}{\beta^2} \exp\left(-\frac{r^2}{\beta^2}\right) + \frac{\eta'}{\gamma^2} \exp\left(-\frac{r^2}{\gamma^2}\right) + \frac{\eta''}{2\gamma_2^2} \exp\left(-\frac{r}{\gamma_2}\right) \right], \quad (5)$$

where r denotes the distance from an exposure point; α is the forward-scattering range; β is the backscattering range; γ captures mid-range scattering; and γ_2 controls the exponential tail. The dimensionless ratios η , η' , and η'' weight the backscattering, mid-range, and exponential components relative to the forward-scattering term, respectively. Given a (pixelized) dose map $d(\mathbf{r})$, the deposited energy map $E(\mathbf{r})$ is computed by convolving d with the PSF:

$$E(\mathbf{r}) = (P \otimes d)(\mathbf{r}). \quad (6)$$

We calibrated these parameters by fitting the outputs of the PreFab model [44], which was trained on real EBL fabrication data. Specifically, we queried PreFab with a diverse set of input patterns and used the resulting input-output pairs to tune our EBL virtual foundry. We then reuse the same virtual-fabrication pipeline as in the DUV setting (development + pattern-transfer surrogate), with the resist threshold set to 0.5 to obtain the developed resist pattern.

Across both DUV and EBL, these parameterized operators provide a controllable source of nonlocal blur and nonlinear stochastic pattern-transfer effects, enabling thorough studies and yield assessment without requiring repeated physical fabrication. It should be noted that the current virtual foundry is intended as a feasibility investigation focusing on the dominant and practically important sources of fabrication distortion, particularly lithography and etching/development effects. As more SEM images become available and more comprehensive models of additional fabrication processes, such as development, etching, polishing, and layout-dependent process effects, are incorporated, the virtual foundry can be further refined to improve its predictive accuracy.

3.2 Learning a Differentiable Fabrication Digital Twin

The photonic inverse lithography procedure in PRISM flow embeds the fabrication digital twin inside a gradient-based inverse mask optimization loop. Therefore, the digital twin must satisfy a stronger requirement than forward-only predictors: beyond accurate wafer image prediction on nominal (uncorrected) masks, it must provide **physically consistent and stable gradients with respect to mask perturbations, and remain reliable on the out-of-distribution, ILT-corrected masks encountered during optimization.**

Digital Twin Architecture Design. Because lithography largely determines pattern fidelity and introduces the dominant non-locality in printing, we adopt a *factorized* differentiable fabrication model consisting of (i) a *lithography module* that maps the mask M to a continuous exposure field (i.e., aerial image in DUV lithography or the energy deposition map in EBL), and (ii) a *thresholding module* that summarizes downstream development/etch transfer by projecting this field into a binary (or soft-binary) wafer pattern. This interpretable decomposition injects physically meaningful prior knowledge, reduces learning difficulty under limited calibration data, and yields reliable gradients for the downstream ILT loop.

Formally, given an intended design mask $M \in \mathbb{R}^{S_i \times S_i}$, the predicted wafer image $W \in \mathbb{R}^{S_o \times S_o}$ is obtained as follows,

$$W = \Phi_{\theta=(\theta_I, \theta_T)}(M) = \psi(I_{\theta_I}(M) - T_{\theta_T}(M)), \quad \psi(x) = \frac{\tanh(2\beta x) + 1}{2}, \quad (7)$$

where $I_{\theta_I}(M) \in \mathbb{R}^{S_o \times S_o}$ denotes the predicted continuous exposure field produced by the lithography module with learnable parameters θ_I , and $T_{\theta_T}(M) \in \mathbb{R}^{S_o \times S_o}$ denotes a threshold map produced by the threshold module with learnable parameters θ_T , and $\psi(\cdot)$ is an element-wise differentiable projection function with binarization sharpness controlled by β .

Motivated by the need for both accurate forward prediction and reliable gradients in ILT, we study two complementary parameterizations of the lithography module I_{θ_I} : (1) a **data-driven neural twin** that uses a U-Net backbone to learn the multi-scale nonlocal mask-to-exposure mapping; and (2) a **physics-parameterized twin** that instantiates I_{θ_I} with a physics-grounded DUV/EBL model with learnable optical parameters, e.g., Zernike coefficients in DUV SOCS kernels or variables in a double-Gaussian EBL PSF function, to capture system imperfections.

We **hypothesize** that, while a data-driven neural twin may achieve competitive forward prediction accuracy with enough calibration data, a physics-grounded twin will offer better generalization and reliable gradient in the data-limited regime, particularly for challenging DUV lithography




Fig. 6. Our proposed photonic inverse lithography flow performs variation-aware, photonics-informed mask optimization.

and for the out-of-distribution (OOD) masks produced during iterative ILT. In the later evaluation section, we evaluate both digital twin designs, comparing not only wafer prediction accuracy but also downstream ILT correction quality and yield.

Physics-Grounded Data-Efficient Learning. To fully leverage the collected and annotated fabrication data, we partition each labeled mask–wafer pair into many training samples via *patch-based supervision*: we extract densely overlapping patches with a sliding window and apply fabrication-invariant augmentations (translations, rotations, and flips), which preserve lithography physics while improving data efficiency.

Lithography processes are intrinsically nonlocal; each pixel is coupled to its neighborhood. Naïvely training on isolated patches introduces artificial boundary conditions: pixels near the crop boundary lack the surrounding context required to predict the correct exposure/printing outcome, leading to biased learning and unreliable artifacts. To respect this physics, each training sample uses an *expanded input* of size $S_i = \text{patch_size} + \text{context_size}$, while supervision is applied only on the center cropped region in the output of size $S_o = \text{patch_size}$, as illustrated in Fig. 5b. Equivalently, the model predicts on the full field, but we treat only the valid interior region as trustworthy and discard boundary pixels. The context_size is chosen to exceed the effective receptive field of the model, ensuring that every supervised output pixel is computed from sufficient neighborhood information.

We train the fabrication digital twin by minimizing the mean square error (MSE) loss between the prediction W and the labeled wafer image $W^* \in \mathbb{R}^{S_o \times S_o}$:

$$\mathcal{L}_{\text{MSE}}(W, W^*) = \frac{1}{N S_o^2} \sum_{n=1}^N \left\| \Phi_{\theta}(M^{(n)}) - W^{*(n)} \right\|_2^2, \quad (8)$$

where N is the number of training samples.

For computational efficiency in both training and downstream ILT, we optionally downsample the patch resolution from the native $1000 \text{ px}/\mu\text{m}$ in the GDS mask, e.g., to $250 \text{ px}/\mu\text{m}$, trading spatial detail for speed with minimal impact on predicted pattern fidelity in our study.

3.3 Photonics-Informed Inverse Lithography for Mask Correction

Photonic ILT *differs from traditional pattern-matching in one key respect*: the objective is not merely to reproduce the nominal geometry pixel-by-pixel, but to recover the subset of structures that most strongly govern optical function under variations, while maintaining mask printability under a given process. Algorithm 1 describes this differentiable, photonics-informed optimization loop by combining (i) a well-conditioned levelset-based topology parameterization, (ii) explicit manufacturability priors, and (iii) performance-aware weighting derived from adjoint sensitivity.

Algorithm 1 Photonics-informed, variation-aware ILT in latent topology space

```

1: Input: target mask  $M_0 \in [0, 1]^{H_0 \times W_0}$ , target wafer label  $W^*$  (or desired wafer target), fabrication
   twin  $\Phi_\theta(\cdot; \beta, v)$ , weight map  $\Gamma \in [0, 1]^{H_0 \times W_0}$ 
2: Hyperparameters: latent resolution  $r_\ell$ , original resolution  $r_0$ , sharpness  $\alpha$ , blur kernel  $G_{\text{blur}}$ ,
   inference sharpness  $\beta$ , process variation scale  $v = (v_{\min}, v_{\text{nom}}, v_{\max})$ , step size  $\eta$ , iterations  $K$ 
3: Output: corrected mask  $M$ 
4: Initialize latent field at full resolution.
5:  $\rho_0 \in \mathbb{R}^{H_0 \times W_0} \leftarrow 0.1 \cdot \mathbb{1}[M_0 > 0.5] - 0.1 \cdot \mathbb{1}[M_0 \leq 0.5]$ 
6: Downsample to latent topology variable:
7:  $\rho \leftarrow \mathcal{D}(\rho_0; s), \quad s = r_\ell / r_0$  ▷ Eq. (9)
8: for  $k = 0$  to  $K - 1$  do
9:   (1) Mask printability constraints.
10:   $\tilde{\rho} \leftarrow \rho \otimes G_{\text{blur}}$  ▷ Gaussian blur in latent space
11:  (2) Project latent variable to a differentiable mask.
12:   $\hat{M} \leftarrow \mathcal{D}(\tilde{\rho}; s), \quad s = r_0 / r_\ell$  ▷ upsample, Eq. (10)
13:   $M \leftarrow (1 + \exp(-\alpha \hat{M}))^{-1}$  ▷ sigmoid re-parameterization, Eq. (10)
14:  (3) Patch-based forward prediction (context-aware).
15:  Extract  $B$  overlapping input patches  $M_B \in \mathbb{R}^{B \times S_i \times S_i}$  from  $M$  ▷ Fig. 6
16:  Predict wafer patches with variation injection  $v$ :  $W_B \leftarrow \frac{1}{3} \sum_{v_i \in v} \Phi_\theta(M_B; \beta, v_i)$  ▷ Eq. (11)
17:  Stitch  $W_B$  into full wafer  $W \in [0, 1]^{H_0 \times W_0}$  by averaging overlaps
18:  (4) ILT objective and update.
19:   $\mathcal{L} \leftarrow \|\Gamma \odot (W - W^*)\|_4^4$  ▷ weighted  $\ell_4$ , Eq. (12)
20:   $\rho \leftarrow \rho - \eta \nabla_\rho \mathcal{L}$  ▷ backprop through stitching and  $\Phi_\theta$  and apply update
21: return  $M$ 

```

Manufacturability-aware Levelset Topology Parameterization. Directly optimizing the full-resolution binary mask is poorly conditioned: it encourages high-frequency pixel artifacts and yields unmanufacturable masks. Given the target design mask $M_0 \in [0, 1]^{H_0 \times W_0}$ at the original resolution, we initialize a latent levelset field $\rho_0 \in \mathbb{R}^{H_0 \times W_0}$ around zero according to the nominal inverse designed mask M_0 , so that foreground and background are represented by small positive (0.1) and negative values (-0.1), respectively. To improve optimization efficiency, instead of optimizing ρ_0 at the original resolution, we downsample the latent levelset knots to a lower resolution and optimize $\rho \in \mathbb{R}^{H \times W}$ as the latent topology variable:

$$\rho = \mathcal{D}(\rho_0; s) = \text{Interp}_{\text{bilinear}}(\rho_0; s), \quad s = r_\ell / r_0, \quad (9)$$

where $\mathcal{D}(\cdot)$ denotes the bilinear interpolation operation, r_0 and r_ℓ are the original and latent resolutions (in px/ μm), and s is the corresponding scale factor.

After obtaining the latent levelset knots ρ , at each optimization iteration, we first apply a Gaussian blur to enforce spatial smoothness of the levelset function and suppress high-frequency artifacts (e.g., pixel-level jaggies and unrealistically sharp corners), which also encourages manufacturable feature sizes. We then interpolate this levelset function back to the original mask resolution using bilinear interpolation and convert it to a differentiable mask via a sigmoid-based soft binarization:

$$\hat{M} = \mathcal{D}(\rho \otimes G; s), \quad s = r_0 / r_\ell, \quad M = 1 / (1 + \exp(-\alpha \hat{M})), \quad (10)$$

where G is a Gaussian smoothing kernel, α is the sharpness factor (set to 200 in this work). By construction, $M(x, y) \in [0, 1]$. Since our fabrication prediction model takes inputs of size

$S_i \times S_i$, we further slice M into partially overlapping patches before prediction, as shown in Fig. 6. Specifically, we extract B patches and group them into mini-batches for parallel inference, denoted by $M_B \in \mathbb{R}^{B \times S_i \times S_i}$. The model then produces the corresponding batch of wafer patches :

$$W_B = \frac{1}{3} \sum_{v_i \in v} \Phi_\theta(M_B; \beta, v_i) \quad (11)$$

where β is the inference-time sharpness in Eq. (7) (fixed to $s = 20$ in our experiments), and v is the process-variation vector (nominal + corner cases $(v_{min}, v_{nom}, v_{max}) = (0.85, 1.0, 1.15)$), and we scale the T_{θ_T} output, thereby providing a simple way to emulate dose/resist/etch variations in the fabrication process. We then stitch all predicted patches back to the full predicted wafer image $W \in [0, 1]^{H_0 \times W_0}$ by averaging in the overlapping regions and calculate the ILT loss.

Photonics-Informed ILT Objective Beyond Pixel-wise Matching. Finally, we optimize the latent variable by minimizing an ℓ_4 loss guided by a user-specified photonic-informed weight map $\Gamma \in \mathbb{R}^{H_0 \times W_0}$:

$$\mathcal{L}_4(W, W^*) = \|\Gamma \odot (W - W^*)\|_4^4. \quad (12)$$

In this study, we guide ILT using the normalized *adjoint sensitivity* map of the device figure of merit, i.e., $\Gamma = |\partial \text{FoM} / \partial M_0|$ (normalized to $[0.5, 1.5]$). Intuitively, the adjoint sensitivity identifies mask regions where small geometric perturbations induce the largest FoM change, and thus where faithful recovery matters most. This **embeds photonic insight (design-intent) directly into ILT**: the optimizer concentrates correction effort on performance-critical features, while allowing FoM-insensitive regions to relax strict pixel-wise fidelity and instead deform to accommodate neighboring corrections, improving overall FoM-driven parametric yield. Additionally, the adjoint sensitivity is precomputed only once on the intended target design mask and applied as a weighting map to the loss calculation. The additional computational cost mainly comes from an element-wise weighting operation, whose overhead is negligible.

Finally, the corrected mask is legalized through morphology operations (opening and closing) to help satisfy the mask-rule check (MRC) and improve mask printability. However, this post-processing step does not guarantee the removal of all rule violations, especially when the inverse-designed device itself contains intrinsically small or closely spaced features that violate the target PDK design-rule check (DRC). Incorporating DRC/MRC fixing directly into the inverse-design or ILT loop is an important future research topic and is beyond the main scope of this work.

4 Evaluation Results

This section evaluates PRISM along two coupled axes: **(i) fabrication digital twin prediction accuracy** and **(ii) downstream effectiveness for photonic ILT and yield recovery**. While wafer-image prediction accuracy is a necessary prerequisite, ILT places a stricter requirement on the gradient reliability of the digital twin. We therefore evaluate both forward prediction quality and robustness on ILT-generated masks, and then provide insights on how different model designs and device properties impact end-to-end ILT effectiveness. Our experiments are conducted on a Linux server with a 64-core AMD EPYC 9554 CPU (base frequency ~ 3.1 GHz) and four NVIDIA RTX PRO 6000 Blackwell GPUs, each with 98GB VRAM.

Digital Twin Training Dataset. In practice, SEM imaging and reliable labeling are the dominant bottlenecks for calibration. Hence, we focus on learning high-performing digital twin predictors under a constrained data volume budget. Following Sec. 3.1, we synthesize three calibration datasets that reflect complementary coverage of the mask space: *regular patterns* (interpretable regular motifs), *frequency-diverse random patterns*, and *inverse-designed device samples*. These choices allow us to study the cost-accuracy trade-off of calibration, and to identify which pattern families best support robust digital twin learning. Specifically, for both the regular pattern and random




Fig. 7. Various Si/SiN inverse-designed, subwavelength photonic device corpus used as benchmarks. Design region size is marked on top. All 2.5D devices are optimized using the open-source tool MAPS [45] using the effective index (Si: 2.78, SiN: 1.7) with SiO₂ as cladding. 3D grating coupler is inverse designed using Tidy3d.

pattern datasets, we select 10 patterns of size $3\mu\text{m} \times 3\mu\text{m}$ as the training set. We place patterns with a $2\mu\text{m}$ spacing to mitigate pattern-to-pattern proximity interference, resulting in an *effective patterned area* of $10 \times (3 \times 3) = 90\mu\text{m}^2$, while the *total occupied area* of the calibration block is approximately $(5 \times 5) \times 10 = 250\mu\text{m}^2$. Similarly, we select representative photonic-device samples with a comparable effective patterned area ($\approx 90\mu\text{m}^2$) to match the calibration budget, including a bend, a Bragg grating, a 1×3 splitter, an MDM, and a TDM.

Digital Twin Test Dataset. To assess both wafer image prediction accuracy and generalization on ILT masks, our test set comprises two subsets: (**Test 1**) a photonic-device dataset containing inverse-designed devices and several regular, process-variation-sensitive photonic structures, and (**Test 2**) an ILT-corrected photonic device dataset. Notably, the inverse-designed devices in the test set are generated using different random seeds from those used in the photonic-device training set, ensuring non-overlapping design instances. In addition, Test 1 further includes extra device families beyond the training distribution, such as the 2.5D silicon devices shown in Fig. 7. For Test 2, we directly apply established physics-based ILT pipelines to generate corrected masks: the ICCAD-2013 SOCS kernels [21] for DUV and a PSF-based model for EBL [24]. Note that these pipelines are different from (and not calibrated against) our virtual foundry settings, as *our virtual fab is strictly hidden*. Evaluating these ILT-corrected masks allows us to probe model robustness on mask distribution drift induced by ILT.

Evaluation Metrics for Wafer Image Prediction. Given the predicted wafer image W and the ground-truth wafer image W^* , let N denote the total number of pixels. We define MSE, pixel accuracy (PA), and intersection-over-union (IoU) as

$$\text{MSE} = \frac{1}{N} \|W^* - W\|^2; \quad \text{PA} = \frac{W^* \cap W}{W^*}; \quad \text{IoU} = \frac{W^* \cap W}{W^* \cup W}, \quad (13)$$

where $\cap$ and $\cup$ denote pixel-wise intersection and union.

Fabrication Digital Twin Model. CNN. We adopt a CNN [46] architecture similar to [47], but remove the fully connected layers and keep only convolutional layers. It is a lightweight fully-convolutional model with 5 blocks of (3×3 Conv + BN + ReLU) and channel widths $\{32, 64, 128, 128, 128\}$ followed by a 1×1 convolution and a sigmoid output, totaling $\approx 0.4\text{M}$ parameters. **UNet** [48]. A lightweight UNet is implemented (depth=3, base channel=8) with bilinear upsampling and skip connections, using double 3×3 conv blocks per stage and a $3 \times 3 + 1 \times 1 + \text{Sigmoid}$ head, totaling $\approx 0.2\text{M}$ parameters. **FNO** [49]. The model combines a reduced Fourier neural operator branch (average pooling by 16, up to 32×32 Fourier modes, 64-channel output) with a shallow spatial convolutional

Table 1. Comparison of four digital twin models trained on three calibration datasets and tested on inverse-designed photonic devices fabricated with **EBL**, under both **uncorrected** and **ILT-corrected** conditions.

Training dataset		CNN (#params:~ 0.4M)			FNO (#params:~ 0.4M)			UNet (#params:~ 0.2M)			PSF (#params:5)		
		MSE	PA	IoU	MSE	PA	IoU	MSE	PA	IoU	MSE	PA	IoU
Normal patterns	Test1	1.83E-03	0.998	0.992	5.73E-03	0.994	0.976	1.14E-03	0.999	0.995	1.65E-03	0.998	0.993
	Test2	2.45E-03	0.998	0.990	6.56E-03	0.993	0.972	1.34E-03	0.999	0.994	2.02E-03	0.998	0.991
Random patterns	Test1	1.51E-03	0.998	0.994	2.55E-03	0.997	0.989	9.84E-04	0.999	0.996	9.64E-04	0.999	0.996
	Test2	2.22E-03	0.998	0.991	3.15E-03	0.997	0.987	1.27E-03	0.999	0.995	1.37E-03	0.999	0.994
Photonic devices	Test1	2.23E-03	0.998	0.991	1.90E-03	0.998	0.992	9.99E-04	0.999	0.996	9.82E-04	0.999	0.996
	Test2	2.81E-03	0.997	0.988	2.55E-03	0.997	0.989	1.33E-03	0.999	0.994	1.39E-03	0.999	0.994

Table 2. Comparison of four digital twin models trained on three calibration datasets and tested on inverse-designed photonic devices fabricated with **193 nm DUV**, under both **uncorrected** and **ILT-corrected** conditions.

Training dataset		CNN (#params:~ 0.4M)			FNO (#params:~ 0.4M)			UNet (#params:~ 0.2M)			SOCS (#params:12)		
		MSE	PA	IoU	MSE	PA	IoU	MSE	PA	IoU	MSE	PA	IoU
Normal patterns	Test1	1.45E-02	0.985	0.947	1.55E-02	0.984	0.944	3.66E-03	0.996	0.987	3.78E-03	0.996	0.986
	Test2	2.79E-02	0.972	0.905	2.80E-02	0.972	0.905	6.35E-03	0.994	0.978	7.03E-03	0.993	0.976
Random patterns	Test1	2.12E-02	0.979	0.923	1.63E-02	0.984	0.943	3.99E-03	0.996	0.986	3.27E-03	0.997	0.988
	Test2	3.60E-02	0.964	0.877	4.24E-02	0.958	0.868	7.99E-03	0.992	0.972	7.03E-03	0.993	0.976
Photonic devices	Test1	1.29E-02	0.987	0.953	1.09E-02	0.989	0.961	3.70E-03	0.996	0.987	4.25E-03	0.996	0.985
	Test2	2.87E-02	0.971	0.905	2.71E-02	0.973	0.910	7.85E-03	0.992	0.973	1.69E-02	0.983	0.941

branch (two 3×3 convolutions), then upsamples and concatenates the features and predicts the output with a $3 \times 3 + 1 \times 1$ convolution plus sigmoid, for five convolution layers plus one Fourier layer and roughly $\approx 0.4M$ parameters. **Learnable-SOCS for DUV** [22]. We assume the imaging conditions are known (typical DUV $\lambda = 193$ nm and $NA = 1.35$) as well as the illumination source. The pupil function is modeled with a fixed amplitude defined by the NA cutoff and a *learnable* phase parameterized by a small set of Zernike coefficients. The only learnable parameters are the Zernike coefficient vector $\theta = \mathbf{c} \in \mathbb{R}^{N_Z}$ (with N_Z basis functions, $N_Z = 12$), which defines the pupil phase aberration and is optimized from data, while the source map and NA-limited amplitude remain fixed. Given the source support, we construct the SOCS/TCC operator by stacking source-weighted, shifted pupil functions into a matrix and performing a direct SVD decomposition to obtain a low-rank set of SOCS kernels (top singular vectors) and their weights (squared singular values). These (data-calibrated) SOCS kernels are then applied in the Fourier domain to efficiently simulate aerial images for ILT. **Double-Gaussian PSF model for EBL** [24]. For EBL, we use a physics-motivated double-Gaussian PSF model and approximate the aerial pattern by convolving the mask with a learnable PSF. The learnable parameters include the mixture weight parameter η (mapped to a valid nonnegative mixing ratio) and four inverse standard deviations $\sigma_{ax}^{-1}, \sigma_{ay}^{-1}, \sigma_{bx}^{-1}, \sigma_{by}^{-1}$. Physically, the first Gaussian captures the narrow/core blur of the beam, while the second Gaussian models a broader tail due to long-range scattering/proximity effects; η controls their relative contribution, and the σ parameters determine the blur strength and anisotropy along the two axes.

4.1 Evaluation of Fabrication Digital Twins on Wafer Image Prediction

Table 1 and Table 2 compare wafer-image prediction accuracy for the EBL and 193nm DUV processes. We evaluate both (i) uncorrected device masks (**Test1**) and (ii) ILT-corrected masks produced during optimization (**Test2**), since the latter reflects the out-of-distribution inputs that a digital twin must handle when embedded in the ILT loop.

EBL vs. DUV. Across all methods, **EBL is substantially easier to model** than 193nm DUV: on both Test1 and Test2, all candidates achieve near-saturated segmentation metrics (PA/IoU typically

Table 3. Performance and yield improvement from photonic inverse lithography using different fabrication digital twin models under EBL process.

Device	Device performance (FoM)				w/o ILT				PRISM-based ILT (UNet)				PRISM-based ILT (PSF)			
	before fab.	after fab.			L2 error	yield _{80%}	yield _{95%}		L2 error	after fab.	yield _{80%}	yield _{95%}	L2 error	after fab.	yield _{80%}	yield _{95%}
Bending	$ S_{21} ^2: 0.964$	$ S_{21} ^2: 0.964 \pm 0.001$			1.573E+2	100%	100%		1.091E+2	$ S_{21} ^2: 0.963 \pm 0.001$	100%	100%	1.183E+2	$ S_{21} ^2: 0.964 \pm 0.010$	100%	100%
Crossing	$ S_{11} ^2: 0.948$	$ S_{11} ^2: 0.947 \pm 0.001$			1.700E+2	100%	100%		1.128E+2	$ S_{11} ^2: 0.948 \pm 0.001$	100%	100%	1.081E+2	$ S_{11} ^2: 0.948 \pm 0.001$	100%	100%
Optical diode	$ S_{21} ^2: 0.948$	$ S_{21} ^2: 0.939 \pm 0.002$			2.852E+2	100%	100%		1.715E+2	$ S_{21} ^2: 0.947 \pm 0.002$	100%	100%	1.667E+2	$ S_{21} ^2: 0.947 \pm 0.001$	100%	100%
	$ S_{12} ^2: 0.000$	$ S_{12} ^2: 0.000 \pm 0.000$								$ S_{12} ^2: 0.000 \pm 0.000$				$ S_{12} ^2: 0.000 \pm 0.000$		
1x3 Splitter	$ S_{21} ^2: 0.332$	$ S_{21} ^2: 0.325 \pm 0.001$			3.953E+2	100%	100%		2.297E+2	$ S_{21} ^2: 0.333 \pm 0.001$	100%	100%	2.417E+2	$ S_{21} ^2: 0.329 \pm 0.001$	100%	100%
	$ S_{11} ^2: 0.326$	$ S_{11} ^2: 0.328 \pm 0.001$								$ S_{11} ^2: 0.324 \pm 0.001$				$ S_{11} ^2: 0.325 \pm 0.002$		
	$ S_{41} ^2: 0.323$	$ S_{41} ^2: 0.324 \pm 0.001$								$ S_{41} ^2: 0.324 \pm 0.001$				$ S_{41} ^2: 0.326 \pm 0.001$		
1x2 MDM mux	$ S_{21} ^2: 0.933$	$ S_{21} ^2: 0.932 \pm 0.001$			2.805E+2	100%	100%		1.644E+2	$ S_{21} ^2: 0.933 \pm 0.001$	100%	100%	1.683E+2	$ S_{21} ^2: 0.932 \pm 0.001$	100%	100%
	$ S_{11} ^2: 0.940$	$ S_{11} ^2: 0.929 \pm 0.002$								$ S_{11} ^2: 0.939 \pm 0.001$				$ S_{11} ^2: 0.938 \pm 0.001$		
1x2 WDM mux	$ S_{21} ^2: 0.973$	$ S_{21} ^2: 0.940 \pm 0.016$			5.876E+2	86%	18%		3.400E+2	$ S_{21} ^2: 0.964 \pm 0.012$	100%	97%	3.663E+2	$ S_{21} ^2: 0.965 \pm 0.010$	100%	96%
	$ S_{11} ^2: 0.973$	$ S_{11} ^2: 0.861 \pm 0.071$								$ S_{11} ^2: 0.957 \pm 0.020$				$ S_{11} ^2: 0.957 \pm 0.024$		
1x2 TTS	$ S_{21} ^2: 0.824$	$ S_{21} ^2: 0.619 \pm 0.149$			4.806E+2	49%	13%		2.659E+2	$ S_{21} ^2: 0.801 \pm 0.048$	98%	82%	2.645E+2	$ S_{21} ^2: 0.803 \pm 0.044$	98%	80%
	$ S_{11} ^2: 0.840$	$ S_{11} ^2: 0.827 \pm 0.011$								$ S_{11} ^2: 0.827 \pm 0.017$				$ S_{11} ^2: 0.824 \pm 0.020$		
Bragg grating	$\lambda_c: 1.512$	$\lambda_c: 1.516 \pm 1.536e-07$			3.176E+2	100%	0%		2.380E+2	$\lambda_c: 1.513 \pm 3.158e-06$	100%	100%	2.374E+2	$\lambda_c: 1.515 \pm 1.080e-06$	100%	70%
	$BW: 0.074$	$BW: 0.073 \pm 2.439e-08$								$BW: 0.074 \pm 4.5658e-08$				$BW: 0.074 \pm 1.769e-08$		
	$T_{\min}: 0.089$	$T_{\min}: 0.010 \pm 1.108e-06$								$T_{\min}: 0.087 \pm 2.315e-06$				$T_{\min}: 0.094 \pm 1.573e-06$		
	$\Pi_{\text{ph}}: 0.487$	$\Pi_{\text{ph}}: 0.474 \pm 2.462e-06$								$\Pi_{\text{ph}}: 0.494 \pm 1.033e-05$				$\Pi_{\text{ph}}: 0.486 \pm 2.433e-06$		

≈ 0.998 - 0.999) with low regression error ($\text{MSE} \sim 10^{-3}$). In contrast, **DUV prediction is more challenging**: for CNN and FNO, MSE increases by roughly an order of magnitude relative to EBL (from $\sim 10^{-3}$ to $\sim 10^{-2}$), and IoU drops noticeably (down to ≈ 0.87 - 0.96). This gap is consistent with the stronger diffraction-limited resolution and longer-range interactions in DUV, which induce more nonlinear and nonlocal behavior.

Model Robustness and OOD Generalization. Within each process, *UNet and the physics-grounded digital twins are the most robust*, while CNN and FNO are less stable, especially on DUV and on ILT-corrected inputs. Concretely, on DUV Test1, UNet/SOCS achieve $\text{MSE} \sim (3\text{-}4) \times 10^{-3}$ with $\text{IoU} \approx 0.985$ - 0.988 , whereas CNN/FNO exhibit markedly larger errors. On the more challenging DUV Test2, CNN/FNO degrade further ($\text{IoU} \approx 0.87$ - 0.91 , MSE up to $\sim 4 \times 10^{-2}$), whereas UNet remains comparatively stable. Overall, these trends suggest that (i) strong multi-scale spatial modeling (UNet) improves prediction stability, and (ii) embedding physical structure (PSF/SOCS) provides additional robustness and generalizability when the input mask distribution shifts due to ILT.

Impacts of Calibration Dataset on Model Generalization. Finally, we evaluate how calibration pattern choice impacts generalization to various device structures and ILT-corrected masks. Overall, all three datasets synthesized within the PRISM framework (regular, frequency-diverse random patterns, and device samples) provide sufficiently informative coverage to learn the core mask-to-wafer mapping and achieve strong performance on both Test1 and Test2. Differences across datasets are modest and scenario-dependent: frequency-diverse random patterns can improve robustness to distribution shift in some cases, while regular primitives offer a particularly favorable cost-performance trade-off due to their ease of fabrication, metrology, and annotation. Collectively, these results validate the PRISM calibration strategy: a compact, budget-aware pattern set can support accurate digital twin learning and generalization to the mask distributions encountered in downstream ILT.

4.2 Evaluation of PRISM on Photonic Inverse Lithography

In the ILT experiments, we evaluate post-fabrication photonic performance under both EBL and DUV processes, comparing *uncorrected designs* against *ILT-corrected masks*, where each mask is optimized for 100 iterations. For each device, we measure the transmission at each output port, e.g., $|S_{21}|^2$ for the transmission from port 1 to port 2, and characterize the distribution of post-fabrication performance under process variations.

Process Variation Sampling in Virtual Foundry. We model fabrication variations by sampling the virtual-fab parameters from Gaussian distributions and generating an ensemble of post-fabrication wafer images. Let ξ denote a random variation instance, $W(\xi)$ the corresponding

Table 4. Performance and yield improvement from photonic inverse lithography using different fabrication digital twin models under 193 nm DUV process.

Device	Device performance (FoM)				w/o ILT				PRISM-based ILT (UNet)				PRISM-based ILT (SOCS)			
	before fab.	after fab.			L2 error	yield _{80%}	yield _{90%}		L2 error	after fab.	yield _{80%}	yield _{90%}	L2 error	after fab.	yield _{80%}	yield _{90%}
Bending	$ S_{21} ^2: 0.964$	$ S_{21} ^2: 0.755\pm0.052$			7.381E+2	51%	0%		3.032E+2	$ S_{21} ^2: 0.963\pm0.008$	100%	100%	2.766E+2	$ S_{21} ^2: 0.963\pm0.004$	100%	100%
Crossing	$ S_{11} ^2: 0.948$	$ S_{11} ^2: 0.718\pm0.051$			7.275E+2	11%	0%		3.142E+2	$ S_{11} ^2: 0.944\pm0.006$	100%	100%	2.735E+2	$ S_{11} ^2: 0.944\pm0.005$	100%	100%
Optical diode	$ S_{21} ^2: 0.948$	$ S_{21} ^2: 0.384\pm0.011$			1.093E+3	0%	0%		4.384E+2	$ S_{21} ^2: 0.866\pm0.087$	88%	78%	3.963E+2	$ S_{21} ^2: 0.886\pm0.069$	90%	85%
	$ S_{12} ^2: 0.000$	$ S_{12} ^2: 0.122\pm0.006$								$ S_{12} ^2: 0.005\pm0.009$				$ S_{12} ^2: 0.004\pm0.008$		
1x3 Splitter	$ S_{21} ^2: 0.332$	$ S_{21} ^2: 0.089\pm0.026$			1.544E+3	0%	0%		6.410E+2	$ S_{21} ^2: 0.302\pm0.009$	98%	50%	5.724E+2	$ S_{21} ^2: 0.312\pm0.008$	100%	90%
	$ S_{11} ^2: 0.326$	$ S_{11} ^2: 0.072\pm0.016$								$ S_{11} ^2: 0.323\pm0.018$				$ S_{11} ^2: 0.323\pm0.020$		
	$ S_{41} ^2: 0.323$	$ S_{41} ^2: 0.115\pm0.012$								$ S_{41} ^2: 0.322\pm0.021$				$ S_{41} ^2: 0.318\pm0.014$		
1x2 MDM mux	$ S_{21} ^2: 0.933$	$ S_{21} ^2: 0.696\pm0.031$			1.004E+3	0%	0%		4.602E+2	$ S_{21} ^2: 0.924\pm0.008$	100%	100%	3.709E+2	$ S_{21} ^2: 0.928\pm0.008$	98%	98%
	$ S_{11} ^2: 0.940$	$ S_{11} ^2: 0.430\pm0.071$								$ S_{11} ^2: 0.898\pm0.040$				$ S_{11} ^2: 0.901\pm0.054$		
1x2 WDM mux	$ S_{21} ^2: 0.973$	$ S_{21} ^2: 0.215\pm0.045$			2.093E+3	0%	0%		9.409E+2	$ S_{21} ^2: 0.736\pm0.117$	28%	0%	7.792E+2	$ S_{21} ^2: 0.838\pm0.131$	54%	40%
	$ S_{11} ^2: 0.973$	$ S_{11} ^2: 0.131\pm0.031$								$ S_{11} ^2: 0.620\pm0.308$				$ S_{11} ^2: 0.676\pm0.342$		
1x2 TTS	$ S_{21} ^2: 0.824$	$ S_{21} ^2: 0.376\pm0.014$			1.752E+3	0%	0%		8.412E+2	$ S_{21} ^2: 0.424\pm0.179$	0%	0%	6.510E+2	$ S_{21} ^2: 0.563\pm0.241$	24%	16%
	$ S_{11} ^2: 0.840$	$ S_{11} ^2: 0.441\pm0.022$								$ S_{11} ^2: 0.651\pm0.086$				$ S_{11} ^2: 0.548\pm0.267$		
Bragg grating	$\lambda_c: 1.512$ $BW: 0.074$ $T_{min}: 0.089$ $\Pi_{pb}: 0.487$	$\lambda_c: 1.545\pm9.126e-06$ $BW: 0.0524\pm5.641e-06$ $T_{min}: 0.363\pm1.972e-03$ $\Pi_{pb}: 0.456\pm7.426e-04$			8.872E+2	0%	0%		5.602E+2	$\lambda_c: 1.513\pm1.453e-05$ $BW: 0.068\pm3.341e-06$ $T_{min}: 0.120\pm4.652e-04$ $\Pi_{pb}: 0.485\pm8.320e-04$	28%	0%	3.646E+2	$\lambda_c: 1.513\pm1.414e-05$ $BW: 0.073\pm3.434e-06$ $T_{min}: 0.094\pm2.438e-04$ $\Pi_{pb}: 0.487\pm3.988e-04$	90%	79%

fabricated wafer pattern, and $\mathcal{F}(\cdot)$ a photonic figure of merit such as $|S_{ij}|^2$. We draw $\xi \sim \mathcal{N}(0, \Sigma)$ and evaluate $\mathcal{F}(W(\xi))$ over multiple samples. Specifically, the sampled parameters are

$$\begin{aligned}
 \text{Dose: } & d(\xi) \sim \mathcal{N}(d_0, (0.02 d_0)^2), \\
 \text{EBL lithography: } & \alpha(\xi) \sim \mathcal{N}(\alpha_0, (0.05 \alpha_0)^2), \\
 \text{DUV lithography: } & h(\xi) \sim \mathcal{N}(0, \sigma_h^2), \quad \sigma_h = 25, \\
 & K_{\text{SOCS}}(\xi) = K_{\text{SOCS}}(h(\xi)), \\
 \text{Post-lithography: } & t(\xi) \sim \mathcal{N}(t_0, (0.02 t_0)^2), \\
 & k_{\text{open}}(\xi) \sim \mathcal{P}_{\text{open}}, \quad k_{\text{close}}(\xi) \sim \mathcal{P}_{\text{close}}, \quad \tau_{\kappa}(\xi) \sim \mathcal{P}_{\kappa}, \\
 \text{Ensemble: } & \{W(\xi_i)\}_{i=1}^{N_s}, \quad \{\mathcal{F}(W(\xi_i))\}_{i=1}^{N_s}.
 \end{aligned} \tag{14}$$

where d_0 is the nominal exposure dose, α_0 is the nominal forward-scattering parameter in the EBL PSF model (Eq. (5)), h denotes the defocus value used to generate SOCS kernels $K_{\text{SOCS}}(\xi)$ in DUV, and t_0 is the nominal resist threshold in the threshold stage. Moreover, k_{open} and k_{close} are the structuring-element sizes used in morphological open/close operations, and τ_{κ} denotes the curvature-based etching threshold; $\mathcal{P}_{\text{open}}$, $\mathcal{P}_{\text{close}}$, and $\mathcal{P}_{\kappa}$ represent user-specified sampling distributions (e.g., uniform sampling).

For each sample ξ_i , we (i) draw dose $d(\xi_i)$ and process-specific lithography parameters, (ii) generate the aerial image/energy deposition map, and (iii) draw post-lithography parameters (e.g., $t(\xi_i)$, $k_{\text{open}}(\xi_i)$, $k_{\text{close}}(\xi_i)$, $\tau_{\kappa}(\xi_i)$) to obtain the post-fabrication wafer image $W(\xi_i)$ via the virtual-fab pipeline. We then evaluate the device performance $\mathcal{F}(W(\xi_i))$ over $\{W(\xi_i)\}_{i=1}^{N_s}$ to estimate the post-fabrication performance distribution and yield.

Yield Definition. Given a target performance value f^* (e.g., the original design's $|S_{ij}|^2$) and sampled post-fab performance $f_i = \mathcal{F}(W(\xi_i))$, we define the *parametric yield* at level $p \in (0, 1)$ as the fraction of samples achieving at least a p -relative performance:

$$\text{yield}_p = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathbf{1}[f_i \geq p f^*]. \tag{15}$$

For example, $\text{yield}_{80\%}$ counts the percentage of samples whose post-fabrication performance is no worse than 80% of the target.

PRISM-based ILT Result Summary. Table 3 and Table 4 report ILT results on both inverse-designed photonic devices and regular yet process-sensitive devices under EBL and 193nm-DUV processes. Based on the digital-twin study in Sec. 3.2, we report results using the best-performing twins for




Fig. 8. Nominal-case comparison of PRISM-based ILT for an inverse-designed 1×2 Si WDM mux under EBL and 193nm-DUV processes. We show the mask, the wafer image generated by our virtual fab, and the error map, together with the resulting performance metrics ($|S_{21}|^2$ and $|S_{31}|^2$). The inverse-designed device exhibits a substantially larger performance degradation under the DUV process than under EBL. For DUV, the physics-based SOCS lithography model yields more faithful wafer patterns and better recovered device performance than the learned U-Net surrogate. In contrast, for EBL, the learned U-Net and the analytical PSF model achieve comparable pattern fidelity and FoM recovery.

each setting: **PRISM-based ILT (UNet)** and **PRISM-based ILT (PSF)** for EBL, and **PRISM-based ILT (SOCS)** for DUV.

Across devices, **EBL** exhibits mild process-induced degradation for the evaluated structures, and both PRISM-based ILT (UNet) and PRISM-based ILT (PSF) often achieve comparable yield improvements. This suggests that, for these device classes, the EBL fabrication process is easier to learn, making both data-driven and physics-parametrized digital twins effective engines in ILT.

In contrast, under **DUV** the uncorrected designs can suffer severe post-fabrication degradation, particularly for highly resonant grating-like geometries (e.g., WDM mux / Bragg grating), where $\text{yield}_{80\%} / \text{yield}_{90\%}$ can collapse to near zero in many cases. In this regime, PRISM-based ILT substantially improves yield, with **PRISM-based ILT (SOCS)** consistently providing the most reliable recovery in the most process-sensitive cases. Critically, the DUV results support our hypothesis that, under limited calibration data, a learnable physics-grounded digital twin generalizes more reliably and provides more stable gradients for the ILT loop, particularly on the out-of-distribution




Fig. 9. Qualitative and spectra comparison of the uncorrected design and PRISM-based ILT results on DUV-fabricated Si Bragg grating. We show the masks, corresponding post-fabrication wafer images, and edge error maps for the nominal design, as well as PRISM-based ILT (U-Net) and PRISM-based ILT (SOCS), together with the post-fabrication transmission spectra under sampled process variations (blue) overlaid with the nominal design response (red), illustrating the impact of ILT on pattern fidelity and spectral robustness.

masks generated during iterative correction, whereas *over-parameterized neural twins can achieve good forward prediction yet fail to provide a reliable first-order oracle for effective ILT optimization*. Figure 11 visualizes PRISM-based ILT (SOCS) outcomes under 193 nm DUV, showing the original design masks M_0 , the optimized masks M , and the corresponding change maps $M - M_0$, which highlight the structured, physically meaningful mask adjustments introduced during correction.

Case Study 1: Inverse-designed WDM mux. Figure 8 contrasts the WDM device under EBL and DUV in the nominal setting, showing the input mask, virtual-fabricated wafer image, and resulting device performance. Under EBL, ILT requires only modest boundary refinements. Both UNet and PSF variants achieve strong recovery. Under DUV, however, the correction demands substantial mask modifications to counter diffraction-limited imaging and long-range proximity effects, and PRISM-based ILT (SOCS) provides more reliable recovery than UNet thanks to its superior generalization and reliable gradient flow due to explicit encoding of diffraction physics.

Case Study 2: Spectrum-sensitive Bragg grating. Figure 9 highlights a complementary regime: the Bragg grating is structurally regular, yet its spectral response is highly sensitive to small fabrication perturbations due to its long dimension that accumulates error and resonant working principle. Under 193 nm DUV, diffraction-induced blur and long-range proximity effects distort the effective index and duty cycle, directly perturbing the coupling coefficient and accumulated phase, which produces large passband/stopband shifts. The mask/wafer/spectrum triplet in Fig. 9 illustrates why a physics-grounded digital twin is critical for gradient-based ILT: PRISM-based ILT (UNet) exhibits large post-fabrication deviation from the intended topology, suggesting optimization drift into an OOD mask regime, whereas PRISM-based ILT (SOCS) produces more structured corrections with interpretable assist-like features and yields higher-fidelity post-fab wafer patterns with improved duty-cycle and teeth height recovery. Consistent with this, SOCS-based ILT yields spectra that more closely track the target response under sampled process variations.

Case Study 3: Variation-aware ILT and Adjoint Gradient-based ILT Guidance. Figure 10




Fig. 10. Comparison of the $|S_{21}|^2$ distributions for a 1×2 WDM mux after ILT with progressively added variation-aware optimization and adjoint gradient-based guidance.

presents an ablation study validating two mechanisms for improving robustness: **(i) variation-aware ILT**, where the ILT objective is optimized over three representative process corners (min/nominal/max), and **(ii) adjoint-guided ILT**, where a normalized adjoint sensitivity map reweights the geometry loss to emphasize performance-critical regions. Compared to nominal-only ILT without weighting, each mechanism improves the post-fabrication performance distribution, and their combination yields the most consistent gains, improving average device FoM while reducing tail risk under process variations.

Case Study 4: ILT-Challenging Instances Under Variation (WDM and TTS). Despite these improvements, a small subset of devices remains challenging to fully recover the yield under significant DUV process variation. In particular, the WDM and thermally tunable switch (TTS) still exhibit limited yield even with SOCS-based ILT, e.g., $\text{yield}_{80\%} = 54\%$ for the WDM mux and $\text{yield}_{80\%} = 24\%$ for TTS. These are the hardest cases in our benchmark because they exhibit intrinsic sensitivity due to strong scattering and resonance, where fabrication errors accumulate over long optical paths and induce nontrivial drift on their sharp spectra. These cases point to an important insight: fully closing the design-fabrication gap for inverse designed photonics may require moving beyond post hoc ILT on fixed designs toward *fabrication-aware inverse design (FAID) from the outset*. ILT primarily seeks to recover the original design intent under a given process model, but it cannot fundamentally reduce a device's intrinsic sensitivity to process variations. Fabrication-aware inverse design (FAID) is an orthogonal direction that improves robustness *by construction*. In the following section, we further show that FAID with explicit minimum-feature-size control can work synergistically with PRISM, producing masks that are easier to correct and yielding improved post-fabrication performance after ILT.

4.3 "FAID-ILT Synergy": Impacts of Device Area and MFS on Yield and ILT Effectiveness

Building on the ILT study in Sec. 4.2, we now shift perspective from ILT algorithm-centric evaluation to a device-centric question: *which device geometries are intrinsically more recoverable under fabrication distortions, and how does this interact with ILT?* While ILT can compensate for systematic process bias, its effectiveness depends strongly on the geometric characteristics of the underlying design. To answer this, we systematically study a family of 1×2 WDM multiplexers with varying *design region area* and *minimum feature size (MFS)* constraints. The design region area trades off device compactness and available functional degrees of freedom, while MFS (enforced via Gaussian-smoothing during levelset-based inverse design) limits the smallest printable features, representing a fabrication-aware inverse-design (FAID) approach.




Fig. 11. 193 nm DUV mask correction results using PRISM-based ILT (SOCS): for each device, the original design mask M_0 , the PRISM-corrected mask, and the change on masks are shown.

Table 5. Performance robustness and ILT effectiveness evaluated on an inverse-designed 1×2 WDM mux with various device footprints and minimum feature sizes (MFS).

1x2 WDM design region size		5 × 5μm ²				6.5 × 6.5μm ²				8 × 8μm ²			
Controlled MFS (nm)		50	100	200	300	50	100	200	300	50	100	200	300
Ideal design mask	$ S_{21} ^2$	0.957	0.968	0.951	0.929	0.983	0.980	0.980	0.970	0.994	0.996	0.997	0.988
	$ S_{31} ^2$	0.915	0.946	0.923	0.926	0.965	0.969	0.969	0.973	0.992	0.995	0.994	0.993
Post 193nm-DUV fab	$ S_{21} ^2$	0.313	0.167	0.221	0.399	0.068	0.119	0.144	0.082	0.305	0.322	0.263	0.227
	$ S_{31} ^2$	0.036	0.031	0.110	0.094	0.193	0.141	0.278	0.081	0.007	0.005	0.013	0.022
PRISM-based ILT (SOCS)	$ S_{21} ^2$	0.743	0.856	0.785	0.801	0.864	0.772	0.823	0.902	0.708	0.745	0.904	0.907
	$ S_{31} ^2$	0.770	0.536	0.732	0.567	0.812	0.631	0.742	0.824	0.512	0.675	0.857	0.857

Table 5 reveals a three-way trade-off between **design freedom (area)**, **nominal performance (FoM)**, and **fabrication robustness / ILT recoverability**: Smaller design regions reduce available degrees of freedom and thus lower the best attainable nominal FoM under ideal fabrication. Increasing MFS further restricts the effective design space, typically degrading the peak nominal FoM by suppressing fine-scale features. However, for sufficiently large design regions (e.g., $8 \times 8 \mu\text{m}^2$), increasing MFS substantially improves post-fabrication yield after DUV ILT, even when nominal FoM decreases slightly. In particular, larger MFS values (e.g., mfs = 300) applied in FAID eliminate fragile sub-resolution features. As a result, the generated masks lie closer to a “printable space” of the DUV process, making the ILT correction target more achievable. These results suggest a practical guideline for fabrication-aware photonic inverse design under DUV: Allocate sufficient design area to preserve functional degrees of freedom, and impose a moderate-to-large minimum feature size so that the geometry stays in a manufacturable regime where ILT remains effective. In this view, FAID and ILT are complementary: FAID improves robustness *by construction*, while ILT compensates residual systematic distortions. Together, they yield designs that are both high-performing and reliably recoverable under process variations.

4.4 PRISM-based ILT on Complicated-Function SiN Inverse-Designed Photonic Devices

To validate the scalability of PRISM and to examine the process sensitivity of low-index-contrast silicon nitride (SiN) devices, we perform mask ILT on two representative SiN components: a 1×5 MDM mux ($17 \mu\text{m} \times 9 \mu\text{m}$) and a 1×4 WDM mux ($24 \mu\text{m} \times 24 \mu\text{m}$).

Table 6 summarizes the nominal-case results under both EBL and 193 nm DUV virtual fabrication. Under the EBL process, the post-fabrication device performance remains close to the ideal design for both devices, reflecting the high pattern fidelity of EBL. Under the DUV process, the MDM mux still exhibits only a moderate degradation, whereas the WDM mux suffers a severe performance

Table 6. Performance comparison of SiN 1×5 MDM mux and 1×4 WDM mux fabricated under EBL and 193nm-DUV processes (nominal case). EBL achieves high pattern fidelity and preserves the designed performance, whereas DUV introduces much larger pattern distortions, leading to a pronounced performance degradation. PRISM-based ILT can resume the post-fab device performance for both EBL and DUV.

Device		1×5 MDM mux					1×4 WDM mux			
Transmission		TE0	TE1	TE2	TE3	TE4	1530nm	1550nm	1570nm	1590nm
Ideal design mask		0.974	0.961	0.955	0.942	0.946	0.957	0.964	0.952	0.960
EBL	Design w/o ILT	0.974	0.961	0.955	0.942	0.945	0.952	0.953	0.940	0.958
	PRISM-based ILT (UNet)	0.974	0.961	0.955	0.942	0.946	0.954	0.964	0.950	0.959
DUV	Design w/o ILT	0.936	0.925	0.924	0.901	0.912	0.057	0.016	0.021	0.036
	PRISM-based ILT (SOCS)	0.972	0.952	0.953	0.940	0.942	0.915	0.843	0.867	0.916




Fig. 12. Nominal-case 193nm-DUV virtual-fabrication results for two SiN devices: (a) a 1×5 MDM ($17 \mu\text{m} \times 9 \mu\text{m}$) and (b) a 1×4 WDM ($24 \mu\text{m} \times 24 \mu\text{m}$). For each device, we visualize the input mask M , the virtual-fabricated wafer image W , and the error map $W - M$, along with the corresponding S-parameters. Owing to the lower index contrast (SiN/SiO₂) and high-order TE mode profiles, the SiN MDM is relatively insensitive to nominal DUV process distortions and maintains high performance even without correction. In contrast, the grating-based WDM is highly process-sensitive and fails without correction, whereas PRISM-based ILT (SOCS) significantly reduces pattern errors and recovers the desired spectral response.

drop after fabrication. This discrepancy is mainly attributed to the grating-based (pseudo-periodic) filtering mechanism of the WDM, whose spectral response is highly sensitive to duty-cycle errors, linewidth bias, and edge-placement distortions, as visualized in Fig. 12. Encouragingly, applying PRISM-based ILT (SOCS) substantially reduces the printed-pattern error and largely restores the desired WDM transmission response. These results suggest that low-index-contrast devices with broader modal profiles (e.g., MDMs) tend to be intrinsically more robust to DUV process distortions, while grating-like, phase-sensitive structures (e.g., WDMs) require stronger lithography-aware correction; in such cases, a physics-grounded predictor enables more reliable recovery under DUV.




Fig. 13. 193nm DUV virtual fabrication results for an inverse-designed 3D Si vertical grating coupler optimized with Tidy3D. We compare the uncorrected design and the PRISM-based ILT (SOCS) mask, together with their corresponding virtual DUV-fabricated wafer images W , error maps $W - M_0$, and the 3D FDTD-simulated field and coupling transmission.

4.5 PRISM-based ILT on 3D Inverse-Designed Devices

To further demonstrate PRISM's applicability beyond 2.5D benchmarks, we evaluate an inverse-designed 3D Si vertical grating coupler that couples an out-of-plane Gaussian beam into the in-plane fundamental mode, effectively acting as a *photonic via* for future 3D very large-scale photonic integration (VLPI). This device is very challenging to design in such a compact footprint, making inverse design an essential enabler technique. Figure 13 reports the nominal-case 193nm DUV virtual fabrication results using 3D FDTD simulation, comparing the uncorrected mask and PRISM-based ILT (SOCS). In the ideal simulation, the coupler achieves a peak coupling efficiency of 0.394. Under DUV without correction, the coupling efficiency collapses to 0.10 (74.6% relative drop). With PRISM-based ILT (SOCS), we successfully recover the coupling efficiency to 0.337. These results indicate that PRISM has the capability of boosting the manufacturability of intrinsically process-sensitive 3D PIC building blocks.

5 Conclusion

We presented PRISM, a photonics-informed inverse lithography framework for curvilinear, inverse-designed photonic components that treats *device performance*, not mere geometric fidelity, as the primary objective, enabling a dramatic yield boost (from near-zero to near-unity in representative cases) under realistic fabrication conditions across both EBL and 193 nm DUV processes. Our study delivers three advances that shift ILT from *rectilinear physics-agnostic geometry repair* to *curvilinear functionality-preserving yield optimization* in inverse-designed photonics. (i) We elevate *gradient reliability* to a first-class requirement and show that physics-grounded digital twins can provide stable first-order guidance for ILT on the out-of-distribution masks generated during iterative correction. (ii) We demonstrate that this capability remains practical under realistic, scarce-data conditions by synthesizing compact, physics-guided calibration motifs that efficiently capture non-ideal fabrication behaviors at low fabrication, metrology, and labeling cost, lowering the barrier for both designers and foundries. (iii) We further inject physics priors via adjoint sensitivity weighted, variation-aware ILT objectives, explicitly optimizing yield and worst-case performance rather than relying on physics-agnostic, geometry-only correction. Overall, PRISM advances photonic design-technology co-optimization by elevating photonics ILT to a first-class EPDA stage and bridging the simulation-to-fabrication gap for inverse-designed, process-sensitive photonic components. By enabling scalable, high-yield photonic manufacturing under both EBL and 193 nm DUV, PRISM

accelerates the lab-to-fab translation of inverse-designed photonic devices and expedites the path toward large-scale photonic integration.

Acknowledgments

This work was in partial support by the Empire State Development's Division of Science, Technology and Innovation (NYSTAR) Focus Center at RPI, C210117. The authors would like to thank the support of the SMART fellowship program.

References

- [1] Yuyang Wang, Songli Wang, Robert Parsons, Asher Novick, Vignesh Gopal, Kaylx Jang, Anthony Rizzo, Chia-Pin Chiu, Kaveh Hosseini, Tim Tri Hoang, Sergey Shumarayev, and Keren Bergman. Silicon photonics chip i/o for ultra high-bandwidth and energy-efficient die-to-die connectivity. In *Proc. CICC*, pages 1–8, 2024.
- [2] Yvain Thonnart, Stéphane Bernabé, Jean Charbonnier, et al. Popstar: a robust modular optical noc architecture for chiplet-based 3d integrated systems. In *Proc. DATE*, 2020.
- [3] Yaoyao Ye, Jiang Xu, Baihan Huang, Xiaowen Wu, Wei Zhang, Xuan Wang, Mahdi Nikdast, Zhehui Wang, Weichen Liu, and Zhe Wang. 3-D Mesh-Based Optical Network-on-Chip for Multiprocessor System-on-Chip. *IEEE TCAD*, 32(4):584–596, 2013.
- [4] Ebadollah Taheri, Sudeep Pasricha, and Mahdi Nikdast. Resipi: A reconfigurable silicon-photonics 2.5d chiplet network with pcms for energy-efficient interposer communication. In *Proc. ICCAD*, 2022.
- [5] Yan-Lin Chen, Wei-Che Tseng, Wei-Yao Kao, and Yao-Wen Chang. A general wavelength-routed optical networks-on-chip model with applications to provably good customized and fault-tolerant topology designs. In *Proc. ICCAD*, pages 1–7, 2023.
- [6] Hailong Zhou, Jianji Dong, Junwei Cheng, Wenchan Dong, Chaoran Huang, Yichen Shen, Qiming Zhang, Min Gu, Chao Qian, Hongsheng Chen, et al. Photonic matrix multiplication lights up photonic accelerator and beyond. *Light: Science & Applications*, 11(1):30, 2022.
- [7] Yichen Shen, Nicholas C. Harris, Scott Skirlo, et al. Deep learning with coherent nanophotonic circuits. *Nature Photonics*, 2017.
- [8] Shanshan Yu Shiyue Hua, Erwan Divita et al. An integrated large-scale photonic accelerator with ultralow latency. *Nature*, 2025.
- [9] Mikhail Bernadskiy Sufi R. Ahmed, Reza Baghdadi et al. Universal photonic artificial intelligence acceleration. *Nature*, 2025.
- [10] Pingchuan Ma, Zhengqi Gao, Amir Begovic, Meng Zhang, Haoyu Yang, Haoxing Ren, Rena Huang, Duane S Boning, and Jiaqi Gu. Boson⁻¹: Understanding and enabling physically-robust photonic inverse design with adaptive variation-aware subspace optimization. In *Proc. DATE*, pages 1–7, 2025.
- [11] Hongjian Zhou, Haoyu Yang, Haoxing Ren, Joaquin Matres, and Jiaqi Gu. End-to-End Physical Design Automation Flow for Yield-Optimized Inverse-Designed Large-Scale Electronic-Photonic Integrated Circuits. In *Proc. DAC*, 2026.
- [12] Srinivas Raghvendra and Philippe Hurat. Dfm: Linking design and manufacturing. In *18th International Conference on VLSI Design held jointly with 4th International Conference on Embedded Systems Design*, pages 705–708. IEEE, 2005.
- [13] Guyue Huang, Jingbo Hu, Yifan He, Jialong Liu, Mingyuan Ma, Zhaoyang Shen, Juejian Wu, Yuanfan Xu, Hengrui Zhang, Kai Zhong, et al. Machine learning for electronic design automation: A survey. *ACM TODAES*, 26(5):1–46, 2021.
- [14] Chris Spence. Full-chip lithography simulation and design analysis: how opc is changing ic design. In *Emerging Lithographic Technologies IX*, volume 5751, pages 1–14. SPIE, 2005.
- [15] Linyong Pang. Inverse lithography technology: 30 years from concept to practical, full-chip reality. *Journal of Micro/Nanopatterning, Materials, and Metrology*, 20(3):030901–030901, 2021.
- [16] Mingkun Chen, Jiaqi Jiang, and Jonathan A Fan. Design space reparameterization enforces hard geometric constraints in inverse-designed nanophotonic devices. *ACS Photonics*, 7(11):3141–3151, 2020.
- [17] Erez Gershnel, Mingkun Chen, Chenkai Mao, Evan W Wang, Philippe Lalanne, and Jonathan A Fan. Reparameterization approach to gradient-based inverse design of three-dimensional nanophotonic devices. *ACS Photonics*, 10(4):815–823, 2022.
- [18] Alec M Hammond, Ardavan Oskooi, Mo Chen, Zin Lin, Steven G Johnson, and Stephen E Ralph. High-performance hybrid time/frequency-domain topology optimization for large-scale photonics inverse design. *Optics Express*, 30(3):4467–4491, 2022.
- [19] Erfan Khoram, Xiaoping Qian, Ming Yuan, and Zongfu Yu. Controlling the minimal feature sizes in adjoint optimization of nanophotonic devices using b-spline surfaces. *Optics Express*, 28(5):7060–7069, 2020.

- [20] Alexander Y Piggott, Jan Petykiewicz, Logan Su, and Jelena Vučković. Fabrication-constrained nanophotonic inverse design. *Scientific reports*, 7(1):1786, 2017.
- [21] Shayak Banerjee, Zhuo Li, and Sani R Nassif. Iccad-2013 cad contest in mask optimization and benchmark suite. In *Proc. ICCAD*, pages 271–274. IEEE, 2013.
- [22] Nick Cobb. Sum of coherent systems decomposition by svd. 1995.
- [23] Jianyun Zhou and XiaoMin Yang. Monte carlo simulation of process parameters in electron beam lithography for thick resist patterning. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena*, 24(3):1202–1209, 2006.
- [24] M Gentili, L Grella, A Lucchesini, L Luciani, L Mastrogiacomio, and P Musumeci. Energy density function determination in very-high-resolution electron-beam lithography. *Journal of Vacuum Science & Technology B: Microelectronics Processing and Phenomena*, 8(6):1867–1871, 1990.
- [25] Frederick H Dill, William P Hornberger, Peter S Hauge, and Jane M Shaw. Characterization of positive photoresist. *IEEE Transactions on electron devices*, 22(7):445–452, 1975.
- [26] James W Thackeray, Roger A Nassar, Robert Brainard, Dario Goldfarb, Thomas Wallow, Yayi Wei, Jeff Mackey, Patrick Naulleau, Bill Pierson, and Harun H Solak. Chemically amplified resists resolving 25 nm 1: 1 line: space features with euv lithography. In *Emerging Lithographic Technologies XI*, volume 6517, pages 394–404. SPIE, 2007.
- [27] John Randall, Kurt G Ronse, Thomas Marschner, Anne-Marie Goethals, and Monique Ercken. Variable-threshold resist models for lithography simulation. In *Optical Microlithography XII*, volume 3679, pages 176–182. SPIE, 1999.
- [28] Demetre J Economou. Modeling and simulation of plasma etching reactors for microelectronics. *Thin Solid Films*, 365(2):348–367, 2000.
- [29] Andreas Lundbäck and Lars-Erik Lindgren. Modelling of metal deposition. *Finite Elements in Analysis and Design*, 47(10):1169–1177, 2011.
- [30] Hyunseop Lee, Dasol Lee, and Haedo Jeong. Mechanical aspects of the chemical mechanical polishing process: A review. *International journal of precision engineering and manufacturing*, 17(4):525–536, 2016.
- [31] Guojin Chen, Wanli Chen, Yuzhe Ma, Haoyu Yang, and Bei Yu. Damo: Deep agile mask optimization for full chip scale. In *Proc. ICCAD*, pages 1–9, 2020.
- [32] Wei Ye, Mohamed Baker Alawieh, Yibo Lin, and David Z Pan. Lithogan: End-to-end lithography modeling with generative adversarial networks. In *Proc. DAC*, pages 1–6, 2019.
- [33] Mingjie Liu, Haoyu Yang, Brucek Khailany, and Haoxing Ren. An adversarial active sampling-based data augmentation framework for ai-assisted lithography modeling. In *Proc. ICCAD*, pages 1–9. IEEE, 2023.
- [34] Yibo Lin, Yuki Watanabe, Taiki Kimura, Tetsuaki Matsunawa, Shigeki Nojima, Meng Li, and David Z Pan. Data efficient lithography modeling with residual neural networks and transfer learning. In *Proc. ISPD*, pages 82–89, 2018.
- [35] Haoyu Yang, Shuhe Li, Yuzhe Ma, Bei Yu, and Evangeline FY Young. Gan-opc: Mask optimization with lithography-guided generative adversarial nets. In *Proc. DAC*, pages 1–6, 2018.
- [36] Bientian Jiang, Lixin Liu, Yuzhe Ma, Hang Zhang, Bei Yu, and Evangeline FY Young. Neural-ilt: Migrating ilt to neural networks for mask printability and complexity co-optimization. In *Proc. ICCAD*, pages 1–9, 2020.
- [37] Binwu Zhu, Su Zheng, Ziyang Yu, Guojin Chen, Yuzhe Ma, Fan Yang, Bei Yu, and Martin DF Wong. L2o-ilt: Learning to optimize inverse lithography techniques. *IEEE TCAD*, 43(3):944–955, 2023.
- [38] Guojin Chen, Ziyang Yu, Hongduo Liu, Yuzhe Ma, and Bei Yu. Develset: Deep neural level set for instant mask optimization. *IEEE TCAD*, 42(12):5020–5033, 2023.
- [39] Haoyu Yang, Zongyi Li, Kumara Sastry, Saumyadip Mukhopadhyay, Mark Kilgard, Anima Anandkumar, Brucek Khailany, Vivek Singh, and Haoxing Ren. Generic lithography modeling with dual-band optics-inspired neural networks. In *Proc. DAC*, pages 973–978, 2022.
- [40] Guojin Chen, Haoyu Yang, Haoxing Mark Ren, Bei Yu, and David Z Pan. Differentiable edge-based opc. In *Proc. ICCAD*, pages 1–9, 2024.
- [41] Shuyuan Sun, Fan Yang, Bei Yu, Li Shang, Dian Zhou, and Xuan Zeng. Efficient ilt via multigrid-schwartz method. In *Proc. DAC*, pages 1–6, 2024.
- [42] Guojin Chen, Hao Geng, Bei Yu, and David Z Pan. Open-source differentiable lithography imaging framework. In *DTCO and Computational Patterning III*, volume 12954, pages 133–142. SPIE, 2024.
- [43] Wei Liu, Wenze Yao, Chengyang Hou, Hongcheng Xu, Haojie Zhao, Yiqin Chen, Huigao Duan, and Jie Liu. Hnu-eb1: A software toolkit for electron beam lithography simulation and optimization. In *2021 International Workshop on Advanced Patterning Solutions (IWAPS)*, pages 1–4. IEEE, 2021.
- [44] PreFab Photonics. PreFab: Artificial nanofabrication of integrated photonic circuits. <https://github.com/PreFab-Photonics/PreFab>, 2026.
- [45] Pingchuan Ma, Zhengqi Gao, Haoyu Yang Meng Zhang, Mark Ren, Rena Huang, Duane S. Boning, and Jiaqi Gu. MAPS: Multi-Fidelity AI-Augmented Photonic Simulation and Inverse Design Infrastructure. In *Proc. DATE*, 2025.

- [46] Purwono Purwono, Alfian Ma'arif, Wahyu Rahmaniar, Haris Imam Karim Fathurrahman, Aufaclav Zatu Kusuma Frisky, and Qazi Mazhar ul Haq. Understanding of convolutional neural network (cnn): A review. *International Journal of Robotics and Control Systems*, 2(4):739–748, 2023.
- [47] Dusan Gostimirovic, Dan-Xia Xu, Odile Liboiron-Ladouceur, and Yuri Grinberg. Deep learning-based prediction of fabrication-process-induced structural variations in nanophotonic devices. *ACS Photonics*, 9(8):2623–2633, 2022.
- [48] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proc. MICCAI*, pages 234–241, 2015.
- [49] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *Proc. ICLR*, 2021.