Generative Priors Accelerate Inverse Design for Regularized, High Performance Integrated Photonics

Taehan Kim [*1] **Junho Park** [*2] **SooHyuk Cho** [3] **Mohammad Ali** [2] **Di Liang** [2]

Abstract

Adjoint-based inverse design has revolutionized integrated photonics, yet existing workflows remain computationally inefficient due to a heavy dependence on high-quality initialization. We propose that archives of past optimization traces, typically discarded, should be repurposed as **reusable scientific memory**. We introduce a physics-compatible generative framework that learns a conditional prior over optimized device layouts to seed downstream adjoint refinement. Our model, a conditional Variational Autoencoder (cVAE), is trained on archived designs and incorporates a differentiable filter-project operator to enforce fabrication-feasible geometries. At inference, the model samples candidate templates which are ranked via a reduced-cost electromagnetic screen before final adjoint optimization. On arbitrary-ratio silicon power splitters with a matched 60-step refinement budget, our cVAE-seeded approach reduces median insertion loss by $3.2\times$ (to 0.127 dB) with no floor objective function and from 5.07 dB to 1.94 dB in the with-floor setting. Crucially, the learned prior also improves both convergence speed by 50% and the geometry-performance trade-off. Here, we demonstrate **amortization of historical optimization** into reusable priors transforms restart-heavy global search into efficient, template-guided refinement.

1. Introduction

Many scientific design problems are still solved by repeatedly calling expensive forward simulators inside iterative optimization loops. In those settings, the optimization archive itself contains valuable information: even unsuccessful or mediocre candidates encode partial structure about feasible configurations, interference patterns, and solver-preferred design motifs. Yet most pipelines retain only the final winning design and discard the rest of the search history. This wastes not only compute, but also a potential source of reusable inductive bias for future tasks.

Integrated photonics offers a concrete and demanding case study. Adjoint-based inverse design has enabled compact free-form devices for splitting, filtering, multiplexing, and routing light by optimizing over high-dimensional spatial layouts (Lalau-Keraly et al., 2013; Jensen & Sigmund, 2011; Molesky et al., 2018). However, practical workflows remain bottlenecked by two issues. First, the optimization is often highly initialization-sensitive: random density fields or weakly structured seeds can require long burn-in, frequent restarts, and expensive trial-and-error before reaching a useful interference scaffold. Second, the resulting layouts may exhibit fragmented or fabrication-fragile geometries, complicating downstream manufacturing and cleanup.

This paper focuses primarily on the *initialization bottleneck* and addresses manufacturability at the level of the learned prior while improving the final device performance. We propose to harvest historical inverse-design effort into a generative model that outputs structured, condition-aware template masks. These templates are not meant to replace Maxwell solvers. Instead, they act as a learned front end that proposes better starting points, after which conventional electromagnetic screening and adjoint optimization remain in the loop. In this sense, the method is hybrid by construction: learning provides amortized prior knowledge, while physics still determines final selection and refinement.

Our approach uses a conditional variational autoencoder (cVAE) trained on archived inverse-designed layouts paired with target specifications. For a desired condition, the model samples multiple candidate masks. A reduced-cost electromagnetic screen can then rank these candidates, and the best one is passed to a fixed-budget adjoint solver. Quantitatively, we focus on arbitrary-ratio silicon power splitters because that pipeline is the most mature and currently provides the strongest matched-budget evidence.

[*]Equal contribution [1]University of California, Berkeley, Berkeley, CA, USA [2]University of Michigan, Ann Arbor, MI, USA [3]Princeton University, Princeton, NJ, USA. Correspondence to: Di Liang <liangdi@umich.edu>.

Proceedings of the 43rd International Conference on Machine Learning Workshop on AI for Physics, Seoul, South Korea, 2026.
Copyright 2026 by the author(s).

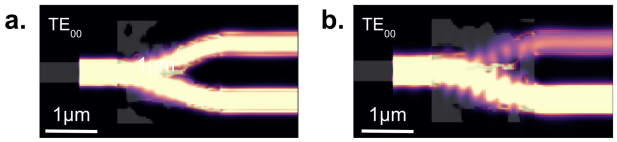

Figure 1. **Visual illustration of optical power splitter.** a) power splitter with 40:60 ratio (top to bottom) at 1310 nm wavelength at fundamental mode; b) power splitter with 20:80 ratio at 1550 nm wavelength at fundamental mode. Such devices are often use to split optical signal into multiple pathways.

2. Related Work

Adjoint Inverse Design: Adjoint-based optimization is the standard for high-dimensional photonic design (Lalau-Keraly et al., 2013; Molesky et al., 2018), but its sensitivity to initialization often necessitates many random restarts. Our work addresses this by replacing random seeds with structured generative proposals.

Machine Learning for Photonics: ML has been applied to forward surrogates and end-to-end generative design (Peurifoy et al., 2018; Tang et al., 2020). However, pure-ML approaches often struggle with the high-precision requirements of Maxwell's equations. Here, we follow a hybrid approach: using generative models to *amortize* search costs while maintaining the physical rigor of an adjoint solver in the final refinement stage.

3. Problem Setup

3.1. Physical Task: Arbitrary Power Splitting

To evaluate our generative prior, we focus on the design of **integrated photonic power splitters**. The goal is to distribute light from a single input waveguide into two output ports according to a user-specified ratio (e.g., $70:30$). Such visual is illustrated in Figure 1. In ML terms, this is a conditional generation task where the target split ratio and operating wavelength act as the condition vector c.

The design region consists of a silicon-on-insulator (SOI) slab where the spatial distribution of etched pixels determines the interference pattern of the light. The challenge is to find a geometry that achieves the target ratio while minimizing **insertion loss** (unintended power dissipation) or, in equal term, maximizing the power transmission. This is a essential "AI4Physics" problem: the relationship between the geometry and performance is governed by Maxwell's equations, making it highly non-linear and extremely sensitive to the initial "seed" geometry.

3.2. Mathematical Representation

As in Figure 11, we represent a device as a binary layout:

$$\rho \in \{0,1\}^{H\times W}, \tag{1}$$

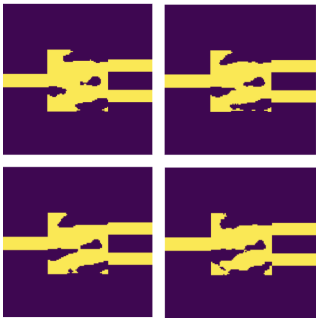

Figure 2. **Conditional samples from the trained splitter cVAE.** Four samples drawn for a fixed target condition (1430 nm with an 87:13 split target). Even before physics screening, the generated masks recover consistent global splitter structure while preserving diversity in local detail.

where $\rho_{ij} = 1$ denotes high-index material (silicon) and $\rho_{ij} = 0$ denotes the etched region (cladding, silicon dioxide). Each mask is paired with a target specification vector $c \in \mathbb{R}^d$, which contains the operating wavelength and the desired power split ratio.

The full hybrid workflow is illustrated in Figure 3, which summarizes how archived designs, conditional generation, physics screening, and adjoint refinement interact.

Given a target specification c, our goal is to produce a template mask that satisfies three requirements:

1. It should encode a physically plausible interference scaffold in 2.5D.
2. It should be biased toward fabrication-feasible geometry.
3. It should initialize downstream adjoint optimization better than an uninformed (random) seed.

4. Method

4.1. Conditional VAE as a Generative Prior

We flatten each mask into $x = \mathrm{vec}(\rho) \in \{0,1\}^D$ with $D = HW$ and train a conditional VAE

$$q_\phi(z \mid x, c) = \mathcal{N}\big(z; \mu_\phi(x,c), \mathrm{diag}(\sigma^2_\phi(x,c))\big), \tag{2}$$

$$p(z) = \mathcal{N}(0, I), \tag{3}$$

$$\ell = g_\theta(z, c) \in \mathbb{R}^D, \tag{4}$$

where ℓ denotes decoder logits.

The key modeling choice is that reconstruction is computed *after* a fabrication-aware differentiable post-processing step. We first map logits to a soft density field

$$p = \mathrm{sigmoid}(\ell) \in [0,1]^D, \tag{5}$$

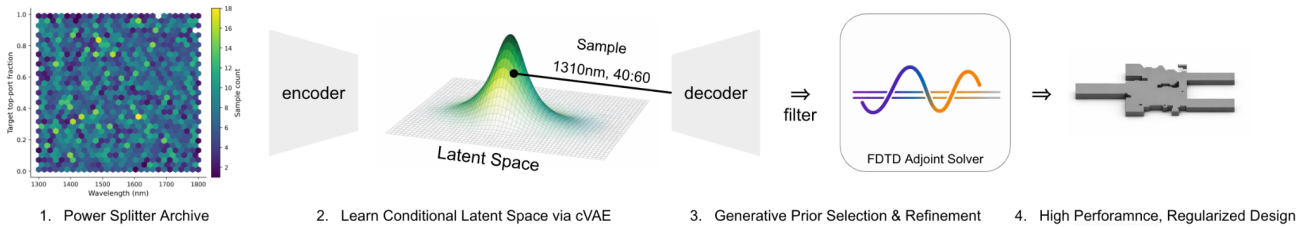

Figure 3. **Overview of the hybrid design pipeline.** Archived inverse-designed power splitters are organized by operating wavelength and target split fraction, then encoded into a conditional latent space with a cVAE. At inference, the model samples condition-aware candidate layouts, a lightweight filter or fast electromagnetic screen selects promising seeds, and a conventional FDTD adjoint solver performs final refinement. The key idea is to treat past inverse-design runs as reusable prior information that improves the starting point for downstream physics optimization.

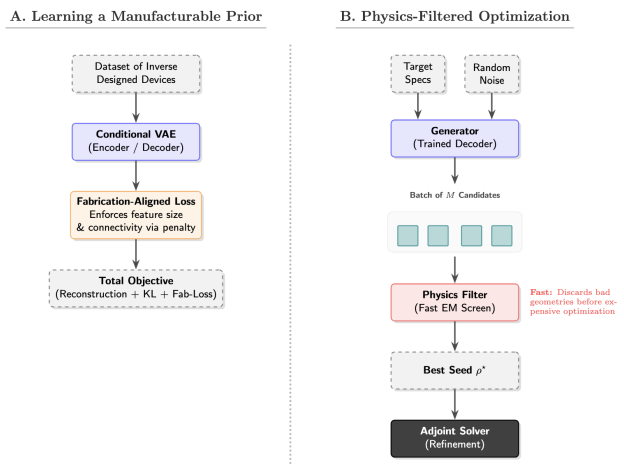

Figure 4. **Detailed pipeline of the candidate generation and screening.**

reshape to the spatial grid, apply Gaussian smoothing, and then apply a smooth projection:

$$p_f = \mathrm{Blur}_\sigma(p), \tag{6}$$

$$p_{\mathrm{proj}} = \mathrm{sigmoid}\big(k(p_f - \eta)\big). \tag{7}$$

This filter-project operator mirrors the type of minimum-feature enforcement commonly used in topology optimization. Intuitively, it discourages isolated speckles and favors clustered, connected regions before the loss is evaluated.

The training objective is

$$\mathcal{J} = \mathcal{L}_{\mathrm{rec}} + \beta\mathcal{L}_{\mathrm{KL}} + \lambda_{\mathrm{TV}}\mathrm{TV}(p_{\mathrm{proj}}) + \lambda_{\mathrm{close}}\mathcal{R}_{\mathrm{close}}(p_{\mathrm{proj}}), \tag{8}$$

where

$$\mathcal{L}_{\mathrm{rec}} = \mathrm{BCE}(p_{\mathrm{proj}}, \rho),\ \mathcal{L}_{\mathrm{KL}} = \mathrm{KL}(q_\phi(z \mid x, c) \| \mathcal{N}(0, I)).$$

The TV term encourages neighboring pixels in p_{proj} to vary smoothly as shown below.

$$\mathrm{TV}(p_{\mathrm{proj}}) = \sum_{\text{pixels } (i,j)} \left(|p_{i+1,j} - p_{i,j}| + |p_{i,j+1} - p_{i,j}|\right), \tag{9}$$

For near-binary layouts, this reduces rough edges, isolated pixels, and checkerboard patterns, acting as a soft penalty on abrupt transition of material.

As illustrated in Figure 3, we train a cVAE on a 2.5D SPINS-B, in-house-produced, archive of 6362 silicon arbitrary photonic power splitter designs, each stored as a binary device layouts paired with its condition vector (wavelength, target split ratio). The encoder and decoder use two fully connected layers of width 512 with a 64-dimensional latent space. The Appendix Figure 11 shows the total training and validation loss for the splitter cVAE over 30 epochs showing convergence.

4.2. Physics-Filtered Best-of-M Seeding

During inference, for a requested condition c, we sample M candidates $\{\rho^{(m)}\}_{m=1}^{M}$ from the trained generator and fast 2.5D FDTD modeling ranks each design,

$$s^{(m)} = \mathrm{FoM}_{\mathrm{fast}}(\rho^{(m)}, c). \tag{10}$$

And select the best design:

$$\rho^\star = \arg\max_{m \in \{1,\ldots,M\}} s^{(m)}. \tag{11}$$

Cheap physics screen allows discarding the clearly bad candidates before expensive refinement.

4.3. Downstream Adjoint Refinement

The chosen template $\rho^\star$ initializes a conventional adjoint solver. We deliberately keep this stage solver-agnostic. In our current splitter benchmark, the seed is refined by a fixed 60-step adjoint optimization in Tidy3D. **The broader point is that the generative prior changes *where the optimizer starts*, not how Maxwell's equations are solved.**

Algorithm 1 Physics-Filtered cVAE Seeding

Input: target condition c, trained generator g_θ, sample count M, fast screen $\mathrm{FoM}_{\mathrm{fast}}$, adjoint solver $\mathcal{A}$
for $m = 1$ to M **do**
  Sample $z^{(m)} \sim \mathcal{N}(0, I)$
  Decode logits $\ell^{(m)} \leftarrow g_\theta(z^{(m)}, c)$
  Build binary mask $\rho^{(m)} \leftarrow \mathrm{Binarize}(\ell^{(m)})$
  Screen candidate $s^{(m)} \leftarrow \mathrm{FoM}_{\mathrm{fast}}(\rho^{(m)}, c)$
end for
$\rho^\star \leftarrow \arg\max_m s^{(m)}$
$\hat{\rho} \leftarrow \mathcal{A}(\rho^\star, c)$
Return: refined design $\hat{\rho}$

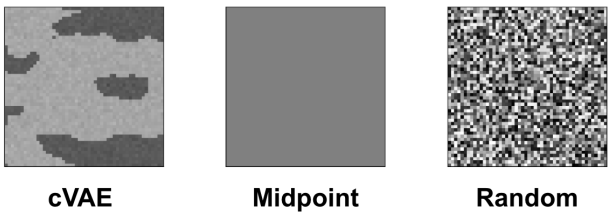

Figure 5. **Visualization of different initialization conditions used.** 3 different initialization conditions are illustarted with cVAE seed having $\alpha = 0.5$ and random being a uniform noise. Here, each pixel has some permittivity value between silicon and silicon dioxide (device material and the cladding).

5. Experimental Setup

5.1. Splitter Design Under Varying Conditions

Using the new suggested inverse-design pipeline, we design power splitters under **10 varying conditions: 2 wavelengths and 5 splitting ratio** to comprehensively evaluate the performance. For each requested condition, the model samples $M = 50$ candidate masks, ranks them with a reduced-cost 2.5D FDTD screen, and passes the top seed to a fixed 60-step adjoint refinement stage as noted before.

We compare three initializations under identical optimization budgets, as illustrated in Figure 5:

- **Random:** uninformed binary initialization,
- **Midpoint:** a uniform 0.5-density field,
- **cVAE top-1 (ours):** the best physics-screened generative seed.

Specifically, evaluation is performed on ten held-out operating conditions spanning 1310 to 1550 nm and target split fractions from 0.10 to 0.50 since these are the two main wavelengths channel for optical communications. **We also report results for two objective functions used for the adjoint refinement**: *with floor*, which includes a minimum-throughput floor term in the optimization objective, and *no floor*, which omits this term. This setup ensures a comprehensive assessment that evaluates performance across device-specific conditions while also testing the method's ability to generalize across different objective functions.

5.2. Metrics

We evaluate design performance using standard photonic metrics. The primary Figure of Merit (FoM) is **Insertion Loss (IL)**, which quantifies the total power transmitted to the output ports relative to the input power. Mathematically, it is expressed in decibels (dB) as:

$$\mathrm{IL} = -10 \log_{10} \left(\frac{P_{out,1} + P_{out,2}}{P_{in}} \right) \tag{12}$$

where P_{in} is the source power and $P_{out,i}$ is the power measured at the i-th output port. A lower IL indicates higher device efficiency.

Additionally, we track **Splitting Error**, defined as the absolute deviation from the target split ratio:

$$\epsilon_{split} = \left| \frac{P_{out,1}}{P_{out,1} + P_{out,2}} - T \right| \tag{13}$$

where T is the target transmission fraction (e.g., 0.3 for a $30 : 70$ split). To **quantify optimization efficiency**, we measure the number of adjoint steps required to reach the top 20% of the best-found FoM.

Manufacturability Index. To quantify the geometric complexity and fabrication feasibility of the generated masks, we use a dimensionless perimeter-to-area ratio:

$$\mathcal{M} = \frac{\sqrt{A}}{P} \tag{14}$$

where P is the total perimeter of the silicon-cladding interface and A is the total area of the silicon regions. A lower $\mathcal{M}$ indicates a more "compact" and contiguous geometry with fewer high-frequency features or isolated "islands," which are typically difficult to resolve in standard lithographic processes. By tracking $\mathcal{M}$ alongside performance, we can evaluate the ability of the generative prior to provide seeds that reside in a more favorable region of the performance-manufacturability Pareto front.

6. Results

6.1. The Learned Prior Produces Structured Splitter Templates

Before any physics screening or adjoint refinement, the cVAE already captures recognizable splitter motifs. Figure 6 shows samples drawn for a fixed target condition. The masks are not identical, which indicates meaningful diversity, but they consistently recover a connected central mixing region

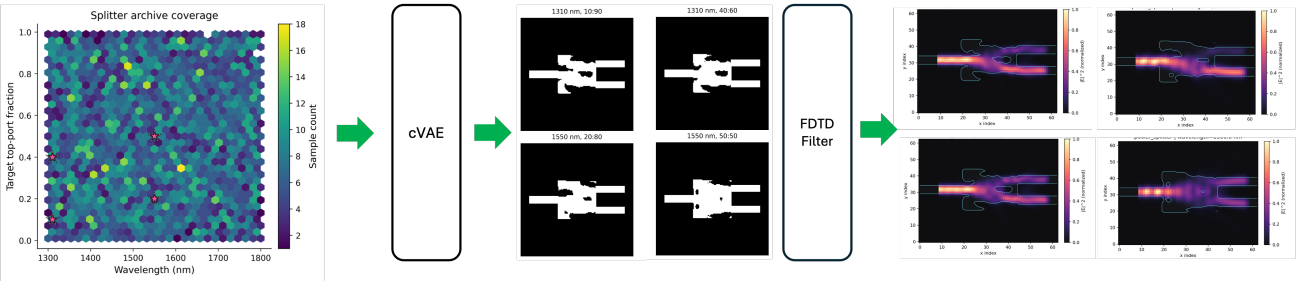

Figure 6. **cVAE sampling of design seeds.** 4 different designs are randomly sampled from the latent space and simulated through forward 2.5D FDTD. The FDTD results confirm that these designs are physics-informed and functional at exact condition that we sampled these designs at.

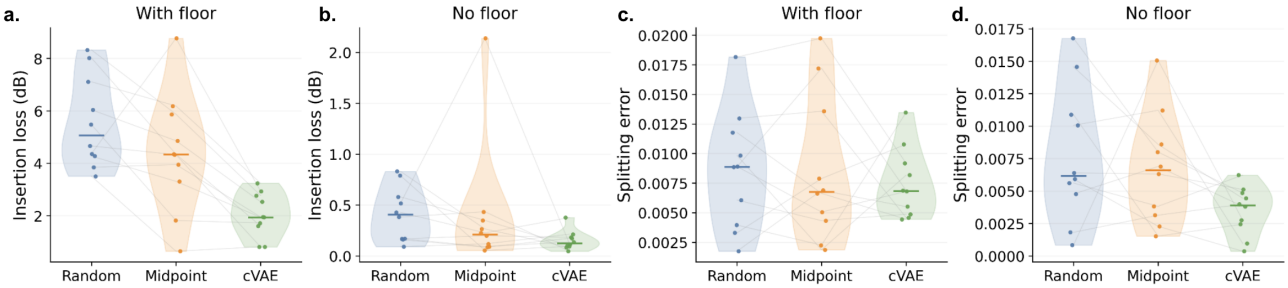

Figure 7. **Endpoint splitter performance under matched refinement budgets.** Insertion loss and splitting error are reported for random initialization, midpoint initialization, and cVAE-seeded refinement, under both objective families. The cVAE prior gives the clearest gains in insertion loss while keeping split error competitive and tightly distributed.

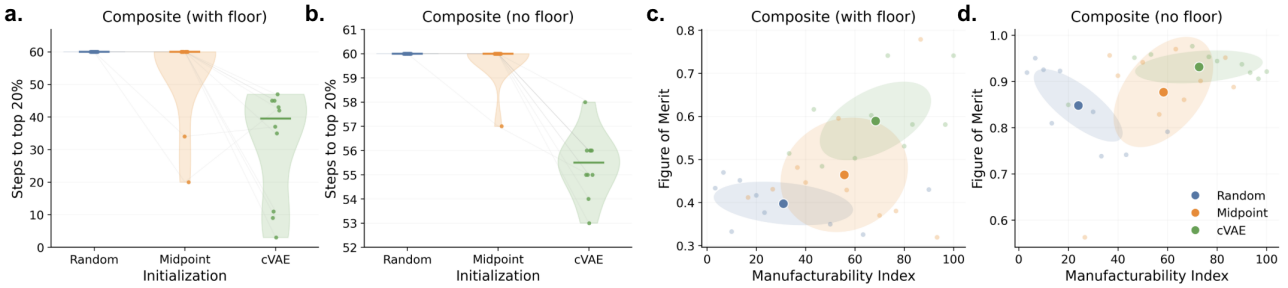

Figure 8. **Convergence and geometry-performance analysis.** (a,b) FoM versus manufacturability index for the with-floor and no-floor objective families, respectively; (c,d) steps required to reach the top 20% of the condition-wise optimum.

and continuous output branches. This is the qualitative signature we want from a generative prior: it should capture reusable interference scaffolds rather than memorizing one exact layout while still having a physics-informed design that we could verify in 2.5D forward simulation.

6.2. cVAE Seeding Improves Endpoint Performance Under Matched Budgets

Figure 7 summarizes the main quantitative splitter results. The strongest effect is on insertion loss. In the no-floor setting, cVAE-selected seeds reduce median insertion loss by $3.2\times$ to 0.127 dB relative to random initialization. In the with-floor setting, the median insertion loss decreases from 5.07 dB for random initialization to 1.94 dB for cVAE seeding, with midpoint initialization remaining intermediate.

The effect on splitting error is smaller than the effect on throughput, but the cVAE results remain competitive and noticeably tighter. This pattern is useful from an optimization perspective: the learned prior primarily helps the solver discover a good transmission-preserving interference scaffold, and then the fixed adjoint budget fine-tunes the remaining ratio mismatch. In other words, the model does not merely produce prettier masks; it systematically changes the starting basin of the physics optimization.

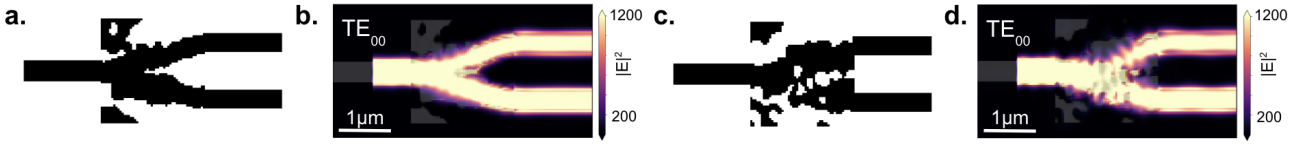

Figure 9. **Optimized Layout** Optimized layouts for a 1310 nm, 40:60 splitter under the no floor objective. The cVAE-seeded design exhibits a more contiguous and regular topology than random initialization, with a manufacturability score of 70.0 versus 23.3 and a lower insertion loss of 0.046 dB versus 0.382 dB while having a cleaner mode at the output ports.

6.3. The Prior Also Improves Convergence and the Geometry-Performance Trade-off

Figure 8 shows that cVAE seeds do more than improve final endpoints. Under both objective families, they reach the top 20% of the condition-wise optimum faster than random initialization by 50%. This means the benefit appears early in the trajectory rather than only at the end. The prior is therefore changing the direction of the optimization, not simply regularizing the final solution.

The same figure also highlights the geometry-performance relationship. Here, FoM is defined as $\sqrt{T \cdot (1 - SE)}$, where T denotes transmission, and SE denotes Splitting Error; such FoM convention is to jointly reward high transmisison/throughput and low splitting error. cVAE solutions occupy a tighter region of the FoM-versus-manufacturability plane than random seeds and avoid much of the spread associated with restart-heavy search. **Figure 8 c and d essentially show that the devices produced from the cVAE seeds were not only performing better, but also has high manufacutrabiliy the same time, breaking the trade off between the performance and manufacturability.**

In our qualitative inspections, these solutions are more contiguous and less fragmented than the designs produced from uninformed initial conditions. Figure 9 shows an example of high manufacturability - high performance layout produced from cVAE-seed and a low manufacturability - mid performance layout produced from random seed. Intuitively, we can see that the layout produced from random seed is harder to manufacture. Also, looking at the mode profile for both, we see that the mode propagation at the output ports has more oscillation, illustrating more unstable behavior.

7. Discussion

The main conceptual point is that inverse-design archives should be treated as datasets, not as by-products. A historical archive encodes where the physics solver repeatedly finds viable interference scaffolds and what geometry families are compatible with target specifications. A conditional generative model can distill that information into a reusable prior, and a cheap physics screen can connect the latent model back to electromagnetic reality before high-fidelity refinement.

This framing matters beyond photonics. Many AI4Physics problems share the same structure: expensive solver-in-the-loop optimization, repeated tasks with related conditions, and large archives of prior search results. In such settings, learning a prior over *good starting points* may be more practical and more reliable than attempting to replace the full simulator end-to-end.

There are several limitations in the current draft. First, the paper focuses on a single device family, so the broader transferability of the prior-learning recipe still needs to be tested more extensively. Second, the reduced-cost physics screen can only help if it correlates well with high-fidelity performance, so stronger calibration experiments would improve the paper. Third, the present workshop version focuses on simulation-time efficiency and geometry regularity rather than fabrication results. These are natural directions for the next iteration.

8. Conclusion

We presented a hybrid machine-learning and physics pipeline that amortizes past photonic inverse-design effort into a reusable generative prior. A cVAE trained on archived layouts proposes structured, condition-aware template masks; an optional fast electromagnetic screen selects a strong seed; and a standard adjoint solver performs the final refinement. On matched-budget arbitrary-ratio splitter benchmarks, this substantially improves insertion loss and convergence relative to uninformed initialization, while the splitter training curves indicate a stable learning regime for the prior itself. More broadly, the work suggests a systematic loop where historical optimization can be reused as learned prior knowledge, not thrown away after a single run.

Impact Statement

This work aims to reduce the computational cost of physics-based design automation and to improve the reliability of inverse-designed photonic components. Positive impacts include faster exploration of photonic design spaces, better use of existing simulation archives, and potentially more energy-efficient optical systems. The main risk is over-trusting generative outputs without adequate physical verification. We therefore emphasize a hybrid workflow in

which learned proposals are always screened and refined with explicit electromagnetic simulation before deployment.

References

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

Christopher M. Lalau-Keraly, Samarth Bhargava, Owen D. Miller, and Eli Yablonovitch. Adjoint shape optimization applied to electromagnetic design. *Optics Express*, 21(18):21693–21701, 2013.

Sean Molesky, Zin Lin, Alexander Y. Piggott, Weiliang Jin, Jelena Vucković, and Alejandro W. Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12(11):659–670, 2018.

John Peurifoy, Yichen Shen, Li Jing, Yi Yang, Fidel Cano-Renteria, Brendan G. DeLacy, John D. Joannopoulos, Max Tegmark, and Marin Soljačić. Nanophotonic particle simulation and inverse design using artificial neural networks. *Science Advances*, 4(6):eaar4206, 2018.

Alexander Y. Piggott, Jesse Lu, Konstantinos G. Lagoudakis, Jan Petykiewicz, Thomas M. Babinec, and Jelena Vučković. Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer. *Nature Photonics*, 9(6):374–377, 2015.

Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, volume 28, 2015.

Yingheng Tang, Keisuke Kojima, Toshiaki Koike-Akino, Ye Wang, Pengxiang Wu, Youye Xie, Mohammad H. Tahersima, Devesh K. Jha, Kieran Parsons, and Minghao Qi. Generative deep learning model for inverse design of integrated nanophotonic devices. *Laser & Photonics Reviews*, 14(12):2000287, 2020.

A. Appendix: Details on Model Architecture and Pipeline

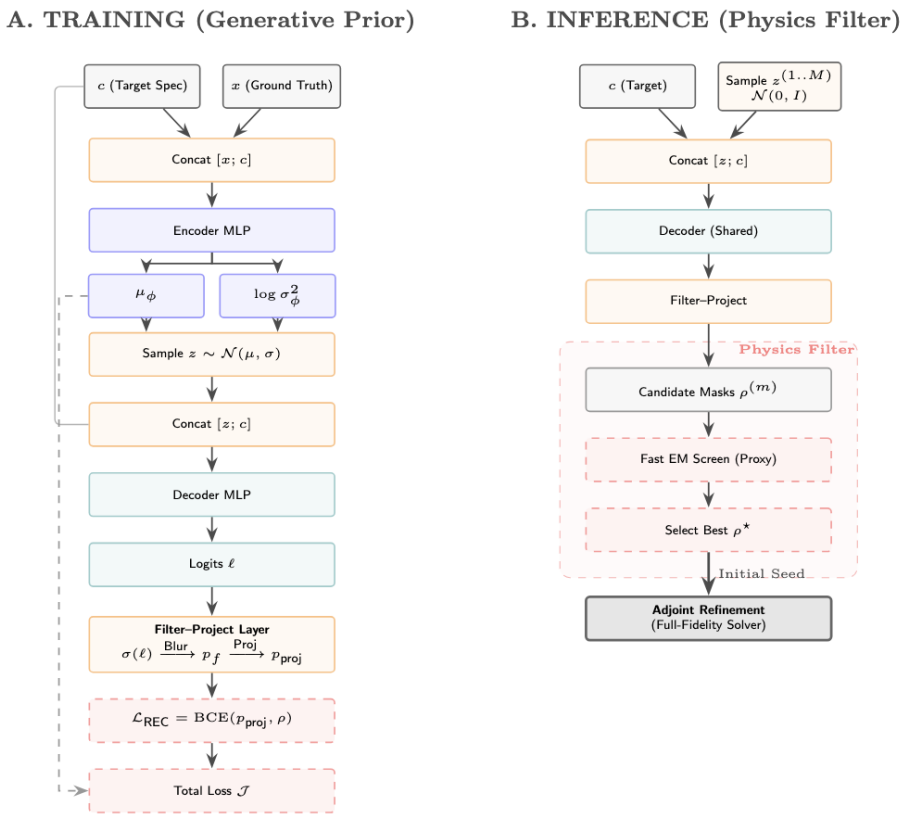

Figure 10. More Details on Model Architecture and Pipeline

B. Appendix: Training Log

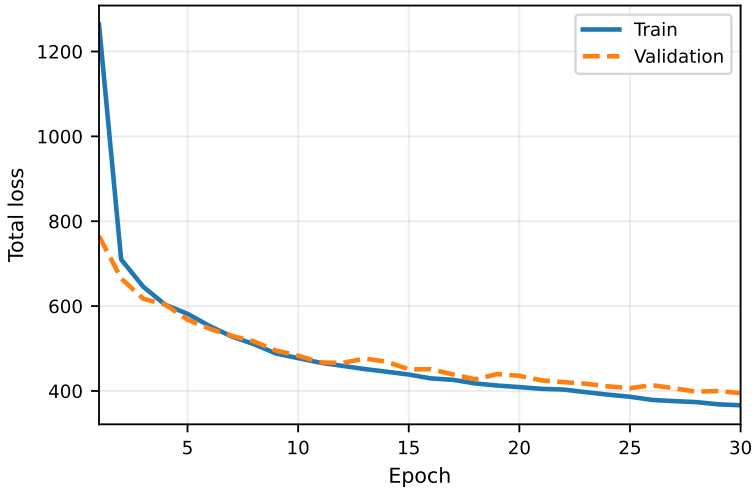

Figure 11. **Power splitter cVAE training history.** Total training and validation loss over 30 epochs using the archived splitter dataset. Both curves decrease smoothly and remain closely matched, indicating stable optimization.

C. Appendix: Additional Implementation Notes

The current splitter system uses a 64-dimensional latent space and two fully connected layers of width 512 in both encoder and decoder. At inference time, the cVAE samples 50 candidates for each requested condition, scores them with a reduced-cost electromagnetic screen, and refines the best one for 60 adjoint steps. The workshop draft focuses on the verified splitter benchmark and the stability of the learned splitter prior.