DiffResist: Physics-Constrained Diffusion for Photoresist Modeling

Zixiao Wang[1], Jieya Zhou[1], Xinyun Zhang[1], Shoubo Hu[2], Farzan Farnia[1], Bei Yu[1]

[1]The Chinese University of Hong Kong [2]Huawei

Key Results

DiffResist reformulates 3D photoresist profile prediction as a physics-constrained 2D diffusion problem, achieving **state-of-the-art accuracy** (EPE-mean **2.56 nm**) with **nearly** $15\times$ **faster** inference than 3D diffusion baselines.

Motivation

Accurate 3D photoresist simulation is essential for optical lithography at advanced nodes, but existing learning-based methods face a key trade-off:

- **Direct mapping** (3D learning): demands prohibitive compute for the high-dimensional output.
- **Sequential modeling**: efficient but suffers from **error accumulation** along depth.
- **DiffResist (ours)**: a physics-constrained 2D diffusion that corrects predictions and suppresses error propagation.

Unlike open-ended video generation, 3D resist simulation benefits from **well-studied exposure physics** that act as strong priors to regularize learning.

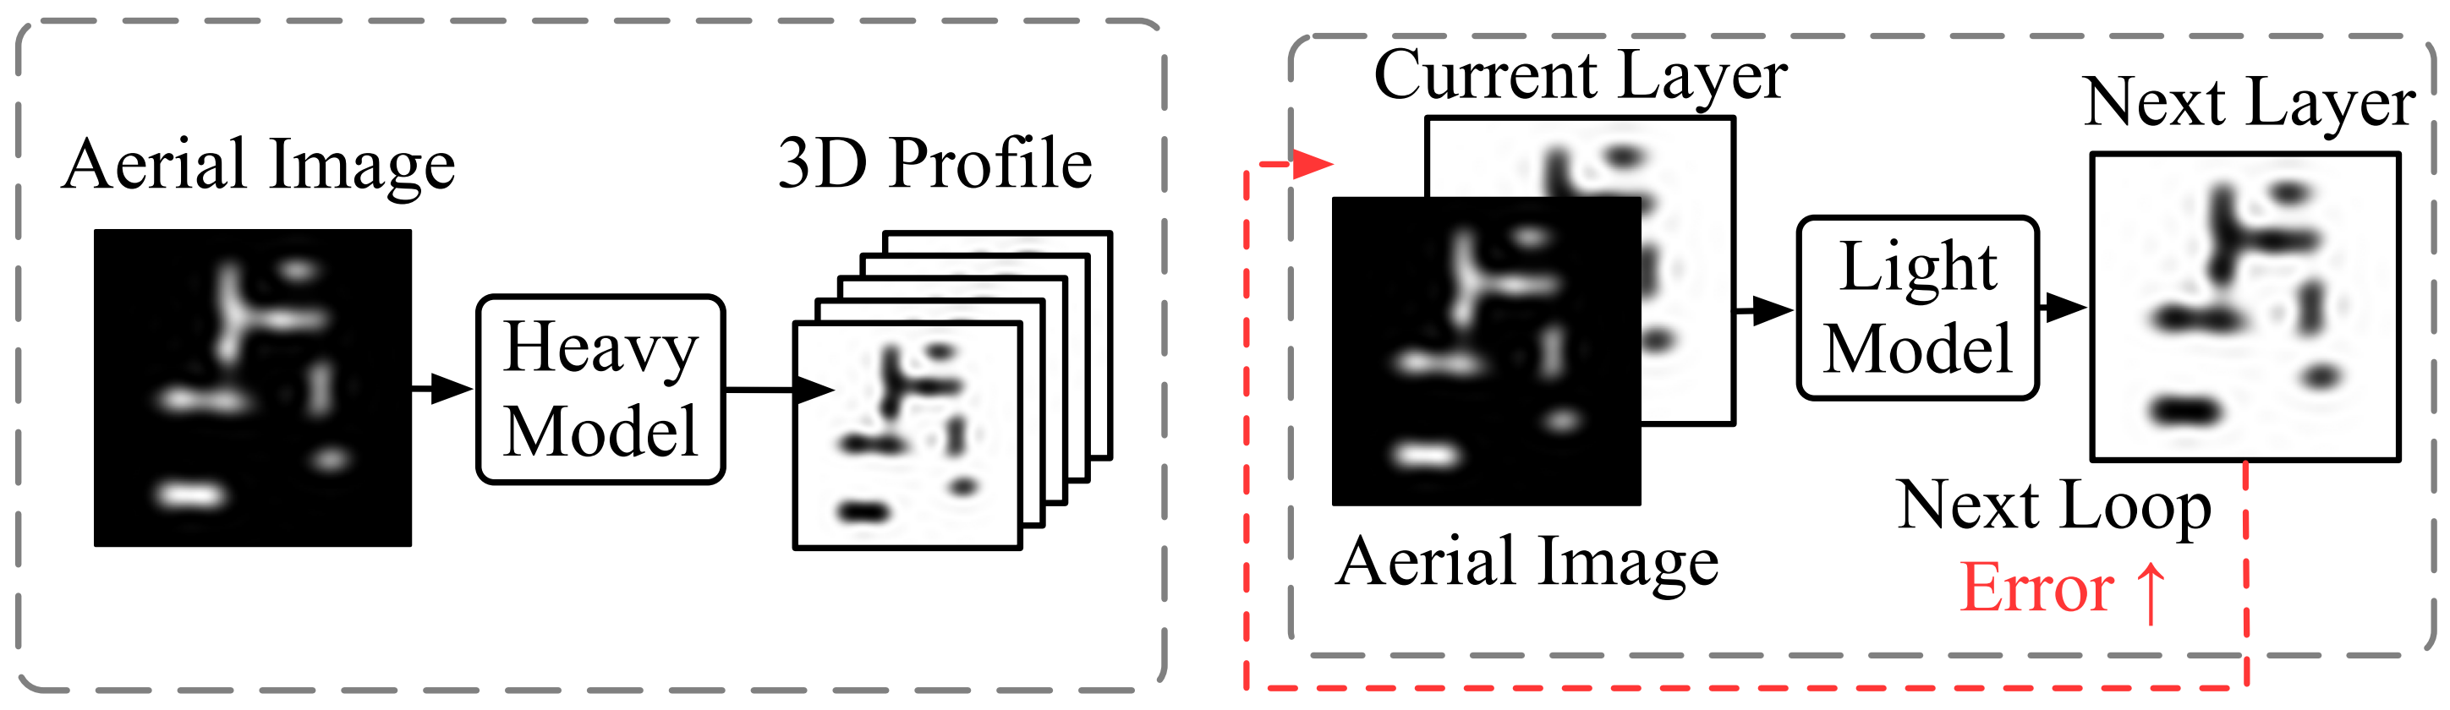

(a) Direct Mapping (b) Sequential Modeling

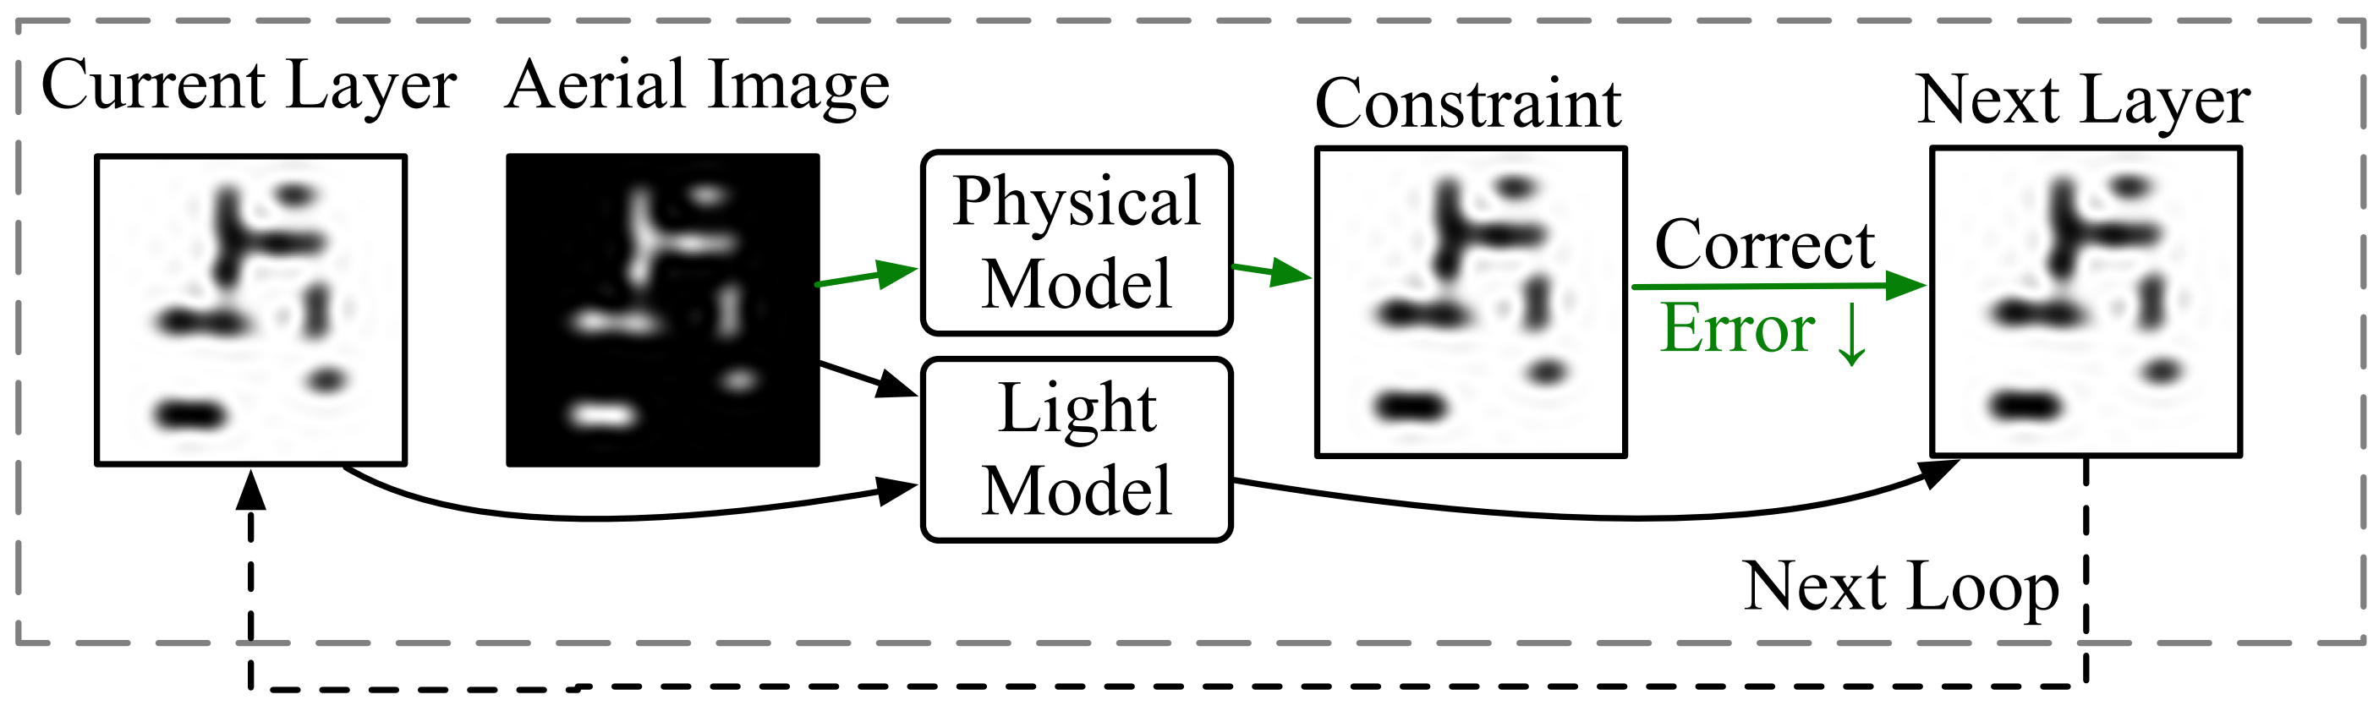

(c) DiffResist

Figure 1. Modeling paradigms. (a) Direct mapping is costly. (b) Sequential modeling accumulates errors. (c) **DiffResist** injects physical constraints to correct sequential predictions.

Resist Exposure Physics

Under the Dill model [2], intensity $\mathbf{I}(h,t)$ and inhibitor concentration $\mathbf{M}(h,t) \in [0,1]$ satisfy coupled PDEs:

$$\frac{\partial \mathbf{I}}{\partial h} = -\mathbf{I}[A\mathbf{M}+B], \qquad \frac{\partial \mathbf{M}}{\partial t} = -\mathbf{I}\mathbf{M}C, \tag{1}$$

where A, B, C are known constants. The aerial image $\mathbf{R}$ provides a **closed-form boundary condition** at the resist surface $h=0$:

$$\mathbf{M}_0(\mathbf{R}) = \exp(-\mathbf{R}\,C\,T). \tag{2}$$

Pipeline Overview

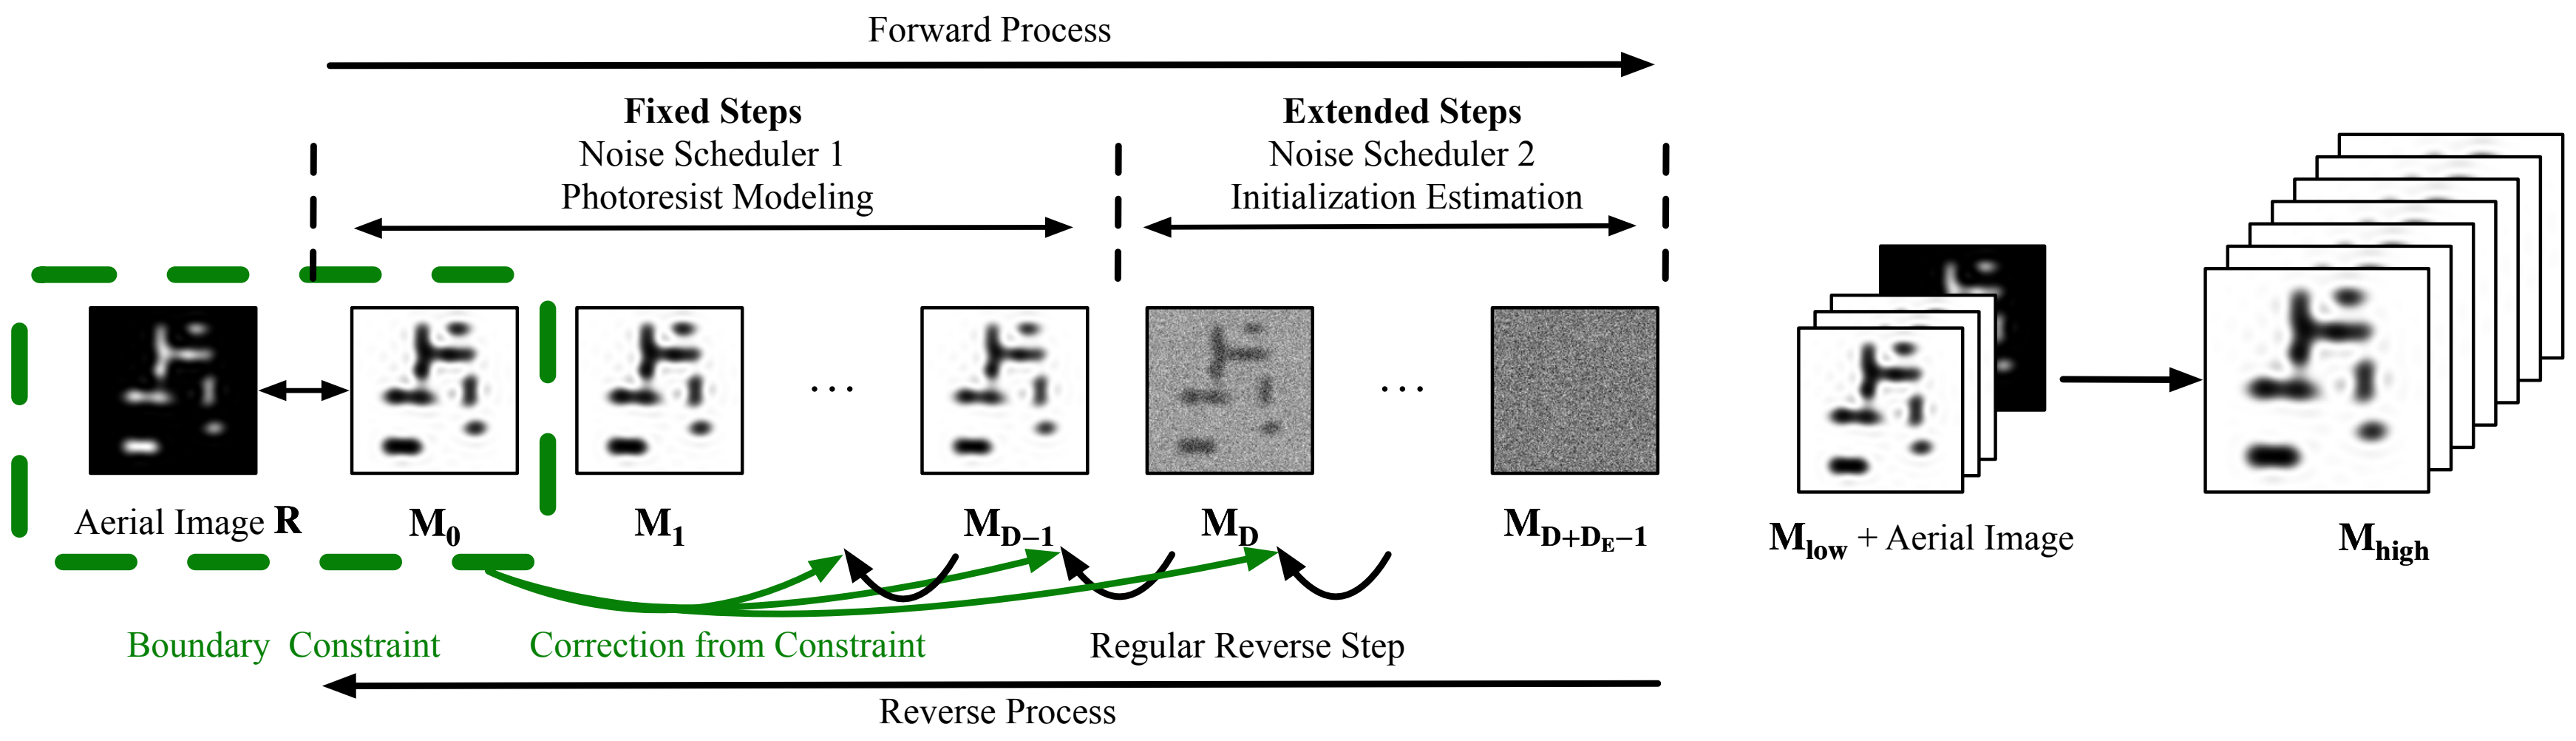

(a) 2D Diffusion Stage (b) Super-Resolution Stage

Figure 2. **DiffResist**: (a) Physics-constrained 2D diffusion predicts a low-res 3D inhibitor volume. (b) Super-resolution upsamples to target resolution.

Physics-Constrained 2D Diffusion

Forward Process (Two Stages)

We represent the 3D volume as D slices $\mathbf{M}_d \in [0,1]^{H\times W}$. **Stage 1** — Physical layers ($0 < d < D$): deterministic from data, $q(\mathbf{M}_d|\mathbf{M}_0) = \delta(\mathbf{M}_d - \mathbf{M}_d^{\text{data}})$. **Stage 2** — Warm-up ($D \le d < D+D_E$): noised to connect to $\mathcal{N}(\mathbf{0},\mathbf{I})$.

Reverse Process with Boundary Condition

Starting from $\mathbf{M}_{D+D_E-1} \sim \mathcal{N}(\mathbf{0},\mathbf{I})$, we inject the **exact boundary condition** (2):

$$\mathbf{M}_{d-1} = \underbrace{\sqrt{\bar{\alpha}_{d-1}}\,\mathbf{M}_0(\mathbf{R})}_{\text{Corrected top layer}} + \underbrace{\frac{(1-\bar{\alpha}_{d-1})\sqrt{\alpha_d}}{\sqrt{1-\bar{\alpha}_d}} \cdot \boldsymbol{\epsilon}_\theta(\mathbf{M}_d, d, \mathbf{R})}_{\text{Directional correction}} + \underbrace{\sigma_d\,\mathbf{z}}_{\text{Noise}}, \tag{3}$$

where $\boldsymbol{\epsilon}_\theta$ is a U-Net conditioned on $\mathbf{R}$.

Two-Stage Noise Schedule

$\sigma_d^2 = 0$ for physical layers (no error accumulation); standard DDPM [3] variance for warm-up steps. Training uses simplified noise-prediction loss; for $0<d<D$ the target noise is **deterministic** (known from data).

Lightweight Super-Resolution

Down-sampling $(75,147,147) \rightarrow (15,49,49)$ by $(\lambda_H,\lambda_W,\lambda_D) = (3,3,5)$. A residual network (**0.48M** params) recovers full resolution: $\mathbf{M}_{\text{high}} = g(\mathbf{M}_{\text{low}},\mathbf{R}) + \text{UP}(\mathbf{M}_{\text{low}})$, where g is a small CNN and UP is bilinear interpolation. Final printed pattern is obtained via the bulk development model.

Quantitative Results

Evaluated on LithoBench [5] (16,472 tiles) with TorchResist [4] 2000-step ground truth (single Nvidia 3090).

Baselines: 3D-DDPM, 3D-DDIM (direct 3D diffusion); SVG [1], SVG-RNN (sequential).

Table 1. Results on LithoBench. PD: Pixel Difference (%). EPE: nm.

Method	Params	FLOPs	Train (hrs)	MSE	PD	EPE-max	EPE-mean
3D-DDPM	25.61M	15.53T	146.2	4.14×10^{-2}	19.10	40.47	10.17
3D-DDIM	25.61M	15.53T	146.2	3.79×10^{-2}	18.54	39.52	9.86
SVG	8.73M	99.91G	7.3	1.41×10^{-3}	2.18	19.14	7.73
SVG-RNN	7.15M	99.82G	7.5	7.39×10^{-4}	1.02	17.06	4.62
DiffResist	6.71M	69.57G	7.0	$\mathbf{4.17\times10^{-4}}$	**0.85**	**13.64**	**2.56**

- Direct mapping (3D-DDPM/DDIM): **20× training cost**, yet fails to converge.
- **DiffResist**: best on **all metrics**, **smallest model** (6.71M), **lowest FLOPs**.
- Inference: **0.44 s** vs. 6.45 s per tile ⇒ **nearly** $15\times$ **faster**.

Qualitative Results

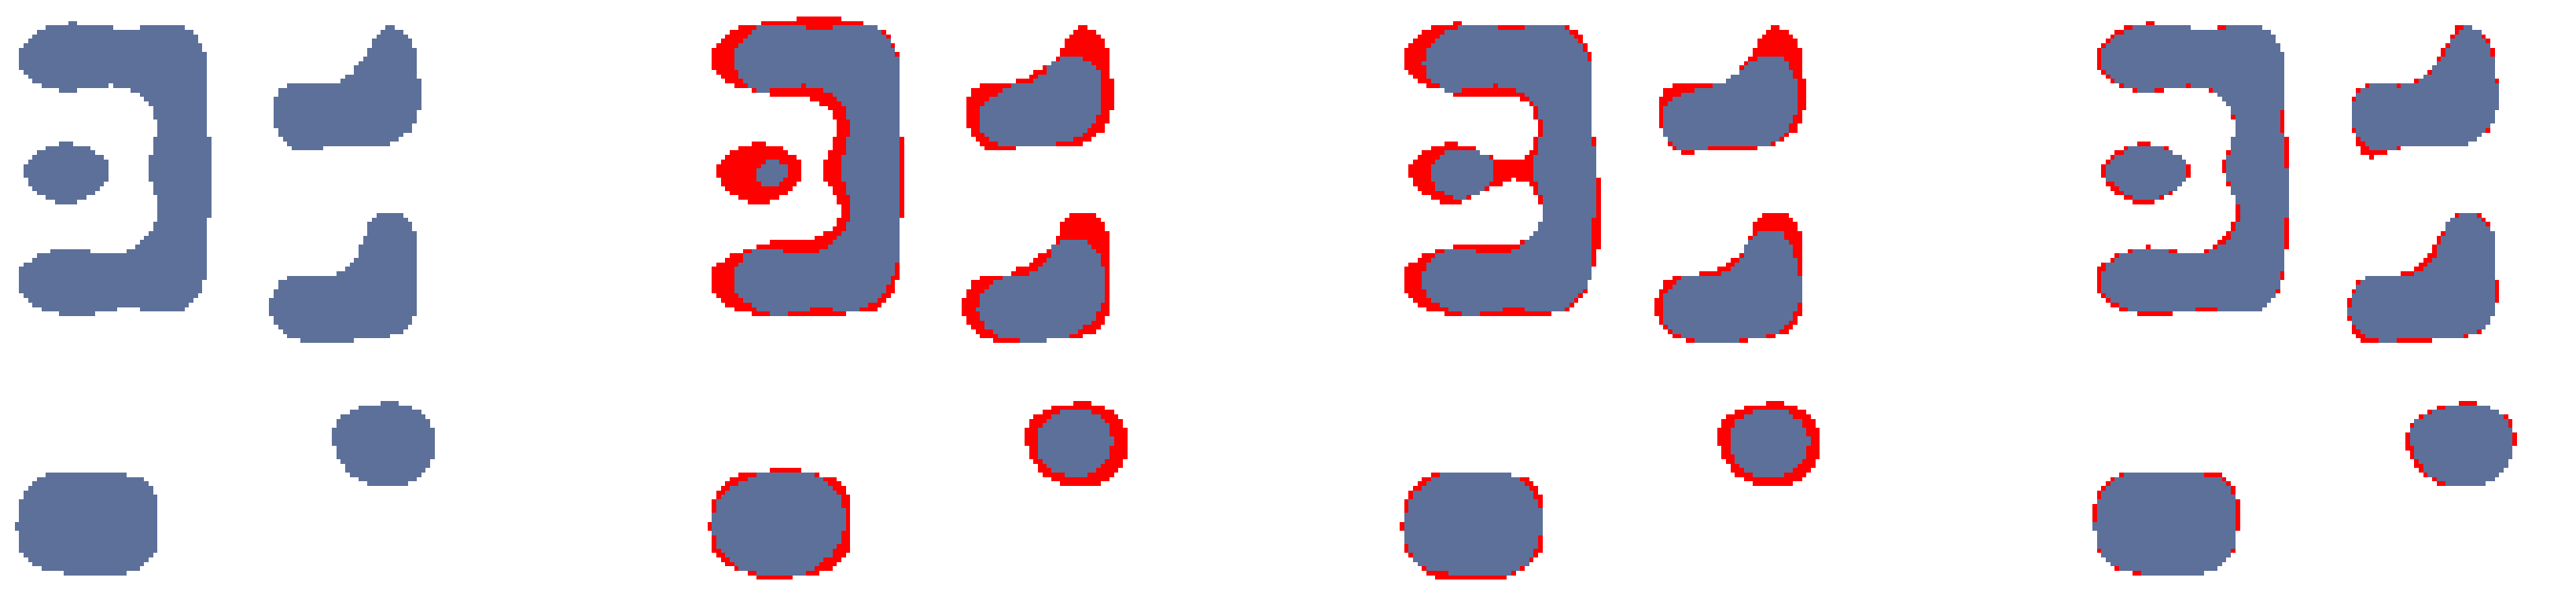

(a) Ground Truth (b) SVG (c) SVG-RNN (d) **DiffResist**

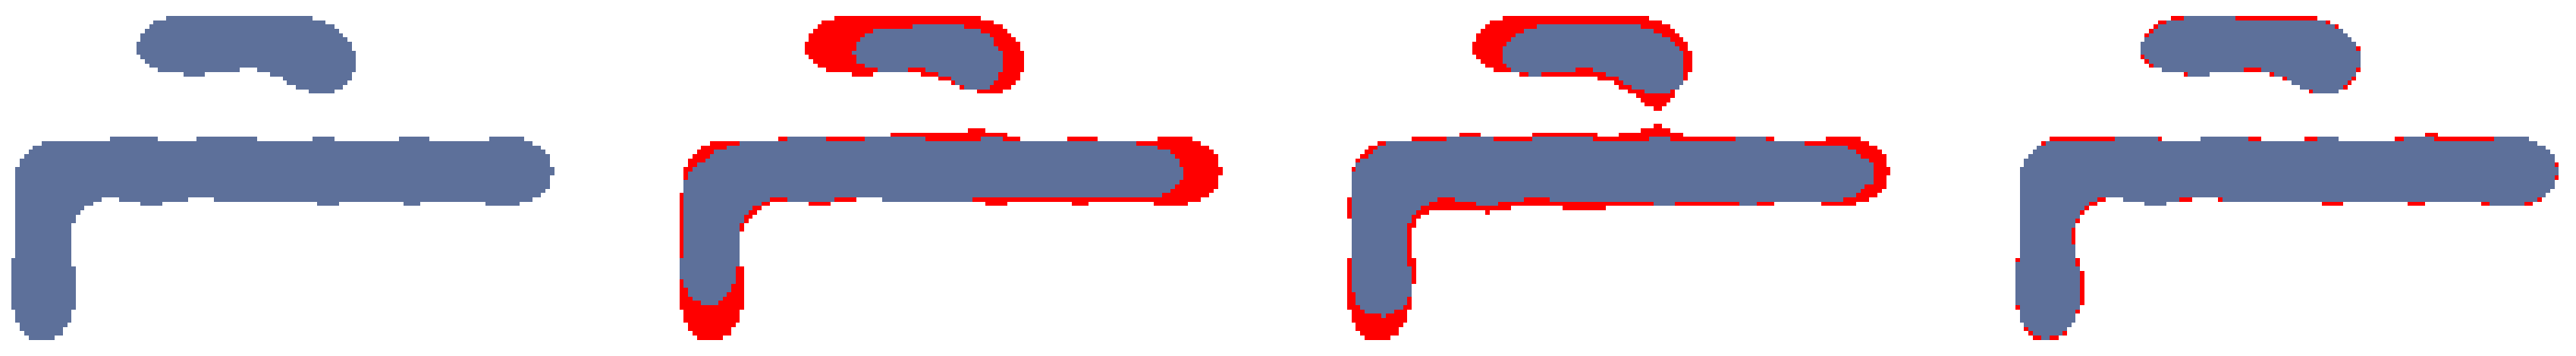

(e) Ground Truth (f) SVG (g) SVG-RNN (h) **DiffResist**

Figure 3. Predicted printed images (difference from ground truth highlighted). **DiffResist** produces the closest match.

Ablation Study & Error Propagation

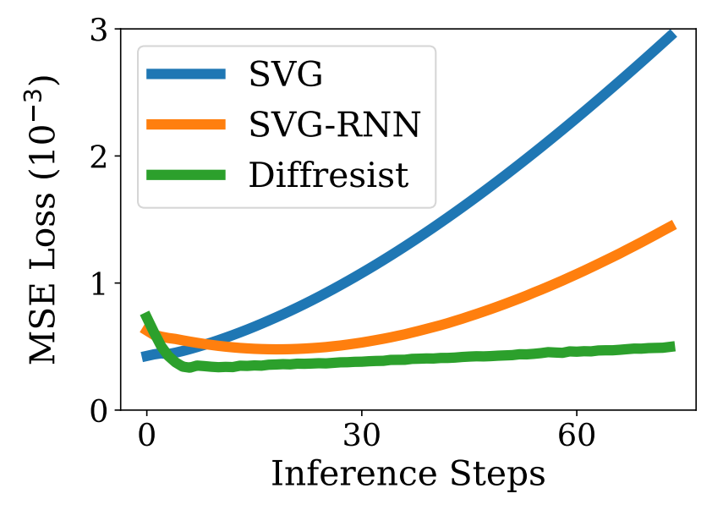

(a) Layer-wise MSE

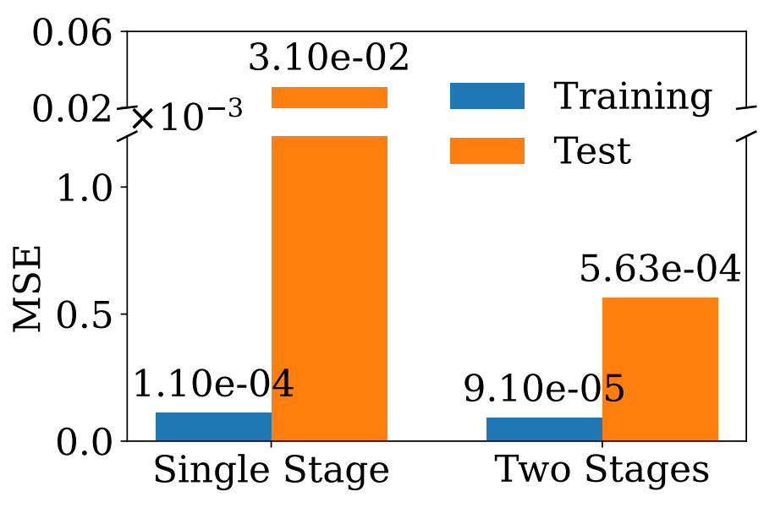

(b) Noise schedule

Figure 4. Ablation: (a) Layer-wise MSE — baseline errors grow rapidly; **DiffResist** stays stable via boundary condition. (b) Two-stage vs. linear noise schedule — similar training loss, but two-stage yields lower test MSE.

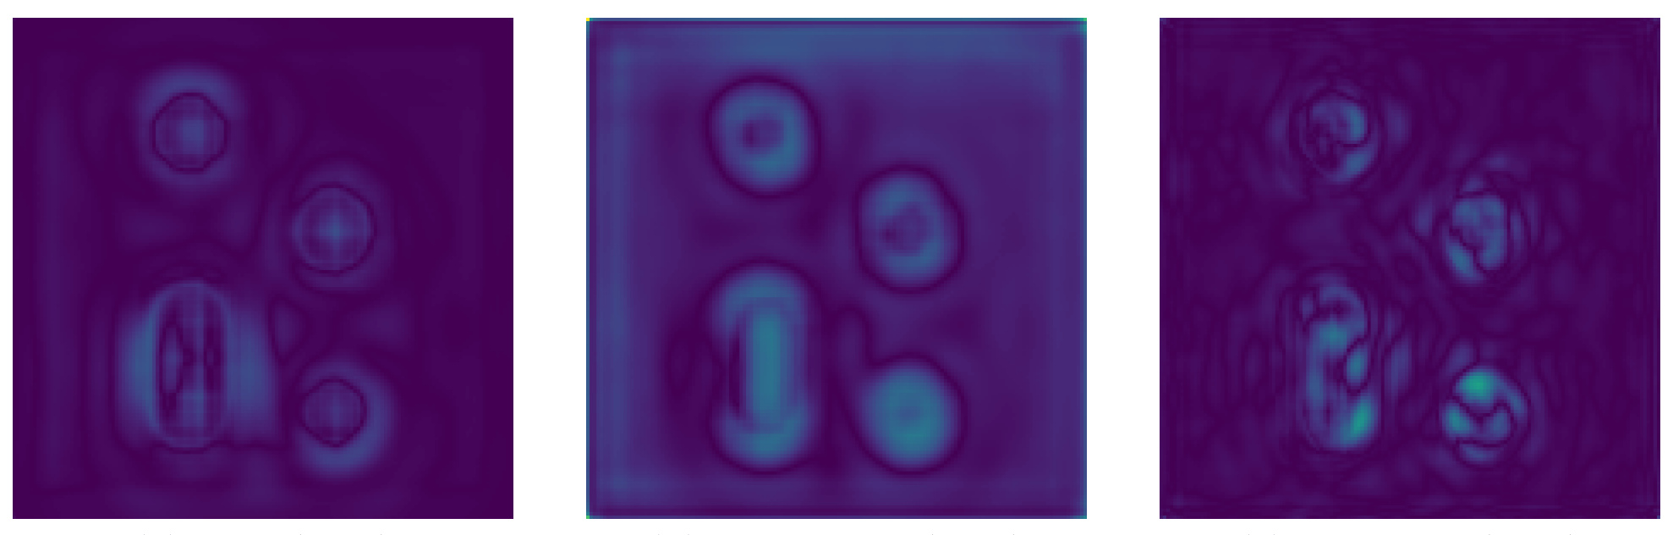

(a) SVG (first) (b) SVG-RNN (first) (c) **DiffResist** (first)

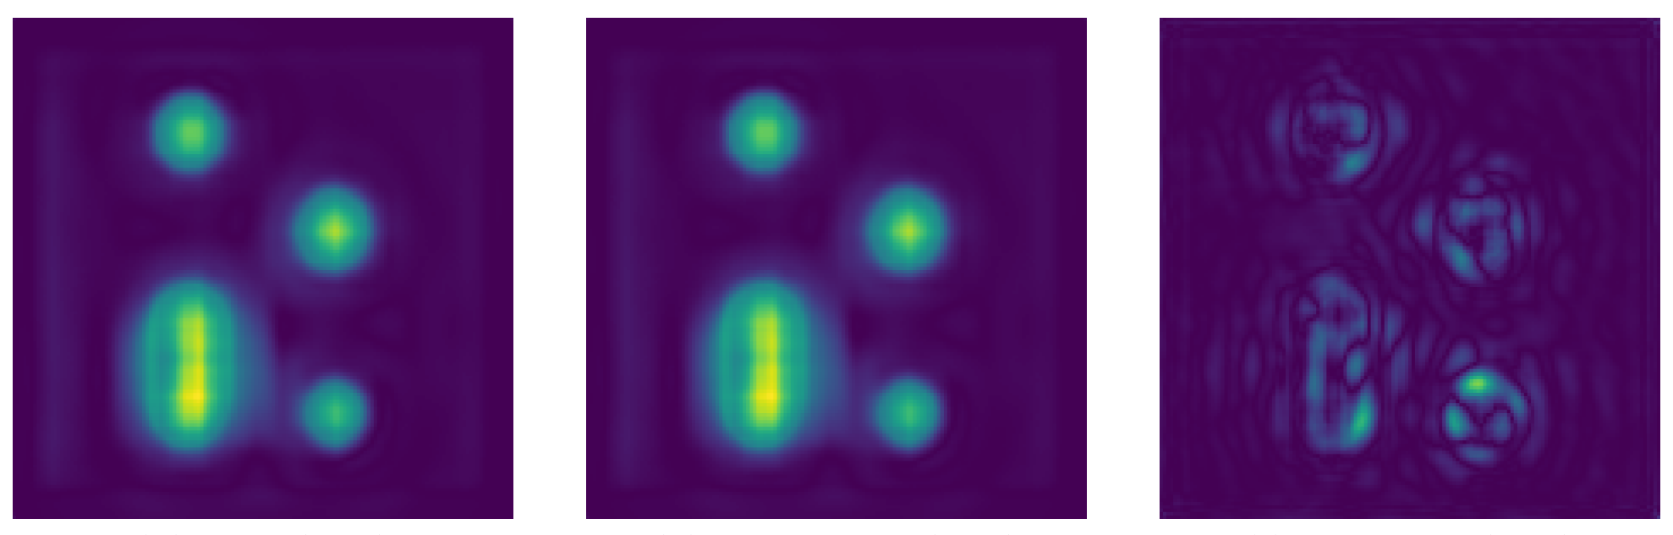

(d) SVG (last) (e) SVG-RNN (last) (f) **DiffResist** (last)

Figure 5. Error heatmaps at the **first** (top) / **last** (bottom) inference step. Baseline errors amplify across depth; **DiffResist** stays bounded by the physical boundary condition.

Conclusion

DiffResist reformulates 3D resist prediction as a physics-constrained 2D diffusion problem. By aligning diffusion steps with physical depth, injecting the resist–air boundary condition, and confining stochasticity to a warm-up segment, it mitigates error accumulation. With lightweight super-resolution, it achieves **state-of-the-art accuracy** and **nearly** $15\times$ **faster inference** on LithoBench.

References

[1] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. In *International Conference on Machine Learning*, pages 1174–1183, 2018.

[2] 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.

[3] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.

[4] Zixiao Wang, Jieya Zhou, Su Zheng, Shuo Yin, Kaichao Liang, Shoubo Hu, Xiao Chen, and Bei Yu. Torchresist: Open-source differentiable resist simulator. In *SPIE Advanced Lithography + Patterning*, 2025.

[5] Su Zheng, Haoyu Yang, Binwu Zhu, Bei Yu, and Martin Wong. LithoBench: Benchmarking AI Computational Lithography for Semiconductor Manufacturing. *Advances in Neural Information Processing Systems*, 36, 2024.

Design, Automation and Test in Europe Conference (DATE) 2026, Verona, Italy