SIMULATOR-REFINED DIFFUSION FOR RADIO-FREQUENCY INVERSE DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion models have shown potential in inverse design of printed circuit boards (PCBs), enabling the generation of layouts conditioned on target S-parameters. Despite this promise, applying diffusion models to PCB layout generation remains challenging due to their difficulty in meeting the quantitative electromagnetic specifications. A common approach is gradient-based guidance, which biases the diffusion sampling process with the gradient of an objective used for evaluation. However, full-wave electromagnetic simulators are accurate but expensive and typically non-differentiable, whereas differentiable surrogates are informative but not always reliable. To address these limitations, this paper proposes *Simulator-Refined Diffusion* (SRD), a novel combination of a low-fidelity differentiable surrogate and a high-fidelity non-differentiable simulator within the diffusion sampling process. Unlike standard zeroth-order optimization, which requires a great number of random perturbations, our approach uses the surrogate's gradient to propose the perturbation direction while the simulator then searches based on this direction to identify an effective design update. Experimental results across different settings show that this method consistently outperforms current state-of-the-art methods, producing layouts whose simulated S-parameters match the target specifications up to *21.2% closer* for in-distribution targets and up to *19.8% closer* for out-of-distribution targets.

1 INTRODUCTION

Radio-frequency (RF) printed circuit board (PCB) layout design asks for a physical geometry whose electromagnetic response matches a given specification. These specifications are often expressed in terms of scattering parameters (S-parameters) across a frequency range, which describe how an RF signal is transmitted and reflected by the layout. Producing the desired response is a fundamental inverse problem in electromagnetics with applications spanning metallic antennas, microwave filters, and RF passives (Hassan et al., 2014; Aage & Johansen, 2017). While the mapping from geometry to response is well-defined, the inverse mapping is ill-posed and non-unique. Even for a fixed board outline, ports, and substrate, small geometric changes can alter S-parameters across frequency, while each high-fidelity evaluation requires a full-wave solution of Maxwell's equations (Aage & Johansen, 2017; Saha et al., 2026). The resulting inverse problem is both high-dimensional and expensive to verify, making it a natural candidate for machine learning methods that can learn from past designs to propose new candidates.

Conditional diffusion models offer a natural prior for this problem (Ho et al., 2020; Guo et al., 2025; Dreossi et al., 2026). These models learn the conditional distribution of layouts given a target S-parameter response and available design metadata, so at generation time, they can produce diverse candidate layouts by sampling from this distribution. Yet conditioning alone does not ensure that a generated layout reproduces the target under a full-wave simulator (Saha et al., 2026). A common approach is to guide denoising with gradients from a differentiable surrogate (Chung et al., 2023; Ye et al., 2024; Zampini et al., 2025). In RF design, using a learned response surrogate makes this guidance inexpensive, but the final design then depends on surrogate accuracy along the sequence of edits used to approach the requested response (Saha et al., 2026). Figure 1(c) illustrates this mismatch where low predicted loss can coexist with large simulated error, and our experiments show that surrogate guidance alone attains negligible improvements over unguided diffusion (Table 1).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

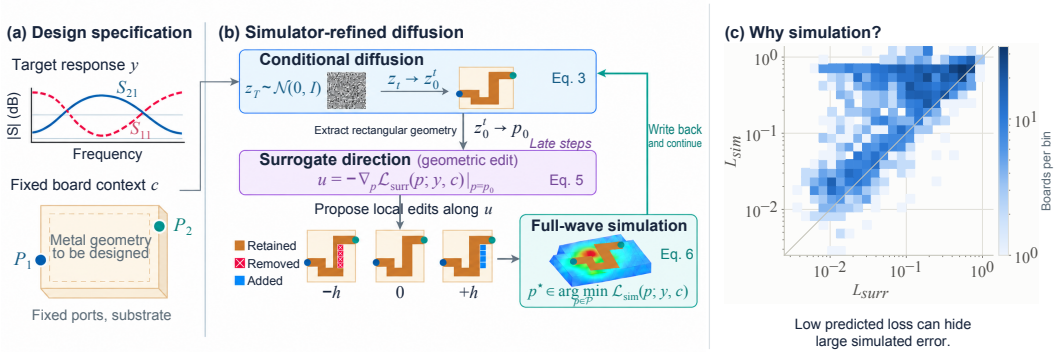

Figure 1: **Simulator-Refined Diffusion.** **(a)** Target S-parameters and board context condition generation. **(b)** At selected late denoising steps, surrogate gradients propose geometric edits (Eq. 5); the full-wave simulator selects the candidate to retain (Section 5.2). Relative to the central unedited candidate, gold-colored pixels are retained, crossed red pixels are removed, and blue-colored pixels are added. **(c)** Low correlation between the surrogate and the simulator downstream losses (RMSE) across tasks, observed especially in the low-loss regime and out of distribution regions.

This paper starts from the hypothesis that a surrogate gradient can provide an informative direction for modifying the current design even when the appropriate displacement is uncertain. A full-wave simulator can then compare a small number of signed displacements without estimating a gradient in the full layout space. This separation converts simulator use from random exploration into targeted model selection. It also suggests when to spend the budget. Early in reverse diffusion, the inferred layout remains unstable and later denoising can overwrite an expensive edit. Simulator refinement is therefore concentrated near the end of the trajectory, when the predicted geometry is sufficiently resolved and the selected modification can persist to the final design.

We instantiate this principle in *Simulator-Refined Diffusion* (SRD). At a selected denoising step, SRD extracts a structured geometric representation from the predicted clean sample and uses the surrogate to propose an edit direction. An initial full-wave simulator probe assesses whether a trial step along the proposed direction sufficiently improves the match between the simulated and target responses. The probe outcome determines whether the remaining simulation budget is allocated to evaluating additional step sizes along that direction or exploring alternative directions. SRD then writes the candidate with the lowest simulated loss back into the diffusion state before resuming denoising. The learned prior, surrogate, and simulator consequently serve complementary roles: the prior maintains plausible global structure, the surrogate focuses the search, and the simulator selects the edit that best matches the physical objective.

Contributions. This paper makes three key contributions: **(1)** it formulates RF inverse design as a budgeted inference problem in which total full-wave simulator calls are the primary resource; **(2)** it introduces SRD, a structured directional-refinement procedure that uses surrogate gradients to propose edits and simulator evaluations to select effective design updates under an explicit call budget; and **(3)** it evaluates this procedure on 250 real targets spanning 25 RF layout templates and 50 constructed out-of-distribution targets, achieving reductions in simulated S-parameter error of up to 21.2% and 19.8%, respectively, relative to the conditional diffusion baseline, *nearly twice the margin of the next best approach*.

2 Related Work

Electromagnetic and RF inverse design. Topology optimization has been used to design metallic antennas and microwave components under electromagnetic objectives (Hassan et al., 2014; Aage & Johansen, 2017). More recently, learned generative models have amortized part of this search by mapping desired responses to candidate layouts. Guo et al. (2025) synthesize EM structures from target S-parameters with a conditional diffusion model. In the RF setting, Dreossi et al. (2026) train a conditional diffusion model for the inverse design of PCB layouts and show that surrogate-based ranking can improve response matching. Our approach extends these methods by using a differentiable surrogate to propose a structured geometric direction and reserves the full-wave simulator for

selecting a signed displacement under a small call budget.

Diffusion guidance for inverse problems. Diffusion priors can be adapted at inference time by modifying an existing sample, conditioning on measurements, or differentiating through task-specific objectives. SDEdit injects noise into an input and denoises it to balance faithfulness and realism (Meng et al., 2022); diffusion posterior sampling incorporates gradients of a measurement-consistency objective (Chung et al., 2023); and training-free guidance generalizes such updates to off-the-shelf differentiable objectives (Ye et al., 2024). Constrained-generation methods likewise use differentiable penalties or surrogate models during sampling (Liang et al., 2025b; Zampini et al., 2025). These methods are attractive when the guidance model is reliable along the sampling trajectory. SRD addresses the setting in which the verifier is an expensive, non-differentiable simulator and the differentiable surrogate is useful mainly for proposing where to search.

Inference-time search and differentiable methods. Best-of-n sampling improves generative outputs by drawing several independent candidates and selecting among them with a scoring model (Ma et al., 2025). This form of selection is complementary to SRD as best-of-n chooses among whole trajectories, whereas SRD refines a single one, and the two can be composed by refining each sampled candidate or by refining only the best-scoring one. Tree-structured search instead reuses feedback from previous diffusion rollouts to bias further exploration (Jain et al., 2025), but requires additional sampling and reward evaluations, which can be costly when rewards come from full-wave simulations. SRD focuses on local geometric refinement within a single diffusion trajectory, coupling each simulator call to a specific edit. Another class of methods relies on zeroth-order gradient estimates of smoothed black-box objectives from function evaluations (Nesterov & Spokoiny, 2017; Zampini et al., 2025; Liang et al., 2026). These are useful when the verifier is inexpensive, but the total cost grows with the number of sampled directions, and small budgets yield noisy estimates in high-dimensional layout spaces. Differentiable perturbed optimizers instead smooth and differentiate optimization maps using random perturbations (Berthet et al., 2020). SRD does not estimate a dense simulator gradient or differentiate a perturbed optimizer. It searches a small, signed set of step sizes along one surrogate-proposed geometric direction.

3 Settings and Goals

PCB geometry and response. A printed circuit board (PCB) consists of patterned metal layers separated by dielectric substrates and connected by vias. At specified frequencies, signal transmission is determined by the metal geometry, substrate properties, and port configuration. We write $x \in \mathcal{X}$ for a metal-and-via layout and $c \in \mathcal{C}$ for the fixed design context, including port positions, substrate permittivity, and thickness. Together, (x, c) specify the board to be evaluated. We say that a layout is admissible if it is representable on the prescribed grid, lies within the board outline, and, when combined with the fixed context, yields a geometry that can be evaluated by a full-wave simulator. In practice, we represent the metal-and-via layout as a grayscale image on a grid, where each pixel encodes the local metal coverage.

For each layout-context pair (x, c), scattering parameters (S-parameters) characterize the frequency-dependent electromagnetic response of the resulting board (Pozar, 2011). For a board with P ports, $S_{ij}(f)$ is the complex ratio of the outgoing wave amplitude at port i to the incident wave amplitude at port j, with all other ports matched. The diagonal terms $S_{ii}(f)$ quantify reflection at each port, whereas the off-diagonal terms $S_{ij}(f)$ quantify transmission between distinct ports. Evaluating every port pair on a frequency grid $\{f_1, \ldots, f_M\}$ yields a response tensor in $\mathcal{Y} = \mathbb{C}^{P \times P \times M}$.

Inverse design, the goal of this work, reverses this mapping: given a target response $y \in \mathcal{Y}$ and a fixed context c, it seeks an admissible layout x whose electromagnetic response matches y.

Forward models and design losses. Computing this response requires solving Maxwell's equations on the board geometry, for example using finite elements or the method of moments (Jin, 2014; Harrington, 1993). We therefore use a full-wave simulator $\boldsymbol{F}_{\text{sim}} : \mathcal{X} \times \mathcal{C} \to \mathcal{Y}$, as the reference model for evaluating candidate designs. However, each simulation is computationally expensive and thus the community has also developed neural surrogate models $\boldsymbol{F}_{\text{surr}} : \mathcal{X} \times \mathcal{C} \to \mathcal{Y}$ that provide an inexpensive, differentiable approximation to the simulator (Dreossi et al., 2026; Saha et al., 2026).

A surrogate may be biased and may not accurately reflect the candidate layout's true S-parameters. Thus, to evaluate the quality of a layout x with respect to a target response y in context c, we define two loss functions, one based on the simulator and one based on the surrogate. For a response

distance $d : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_{\geq 0}$, define

$$\mathcal{L}_{\text{sim}}(x; y, c) = d(\boldsymbol{F}_{\text{sim}}(x, c), y)\,, \qquad \mathcal{L}_{\text{surr}}(x; y, c) = d(\boldsymbol{F}_{\text{surr}}(x, c), y)\,. \tag{1}$$

The simulator loss measures design quality; the surrogate loss approximates it.

Budgeted inverse design. Given a target response y, a fixed context c, and a budget B of full-wave simulator evaluations, the goal is to find the admissible layout whose simulated response best matches y. Let $N_{\boldsymbol{F}_{\text{sim}}}(x; y, c)$ denote the total number of calls to $\boldsymbol{F}_{\text{sim}}$ used to generate, select, and evaluate a returned layout x for the task (y, c). We formulate budgeted inverse design as

$$\begin{aligned} x_B^\star \in \arg\min_{x \in \mathcal{X}} \quad & \mathcal{L}_{\text{sim}}(x; y, c) \\ \text{subject to} \quad & N_{\boldsymbol{F}_{\text{sim}}}(x; y, c) \leq B. \end{aligned} \tag{2}$$

Here, $x_B^\star$ is the best layout attainable under budget B, and $\mathcal{L}_{\text{sim}}$ is the simulator loss defined in Eq. 1. Membership in $\mathcal{X}$ enforces the geometric and simulator-admissibility requirements defined above, whereas the constraint on $N_{\boldsymbol{F}_{\text{sim}}}$ limits the expensive full-wave evaluations. Note that x is a high-dimensional object (an image), and that Eq. 2 is a challenging high-dimensional black-box optimization problem, thus motivating the use of generative models to propose candidate layouts.

4 Guided Diffusion for Inverse Design

Conditional diffusion. Diffusion models learn a layout prior by reversing a gradual corruption process (Ho et al., 2020; Song et al., 2021). Let z_0 be a continuous representation of a clean layout. Under the standard variance schedule $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$, a noisy state is $z_t = \sqrt{\bar{\alpha}_t}\, z_0 + \sqrt{1 - \bar{\alpha}_t}\, \epsilon$, with $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ (Liang et al., 2025a). A noise predictor $\epsilon_\theta(z_t, t; y, c)$, which equivalently parameterizes the conditional score, guides the reverse process from Gaussian noise toward a layout conditioned on the target response and board context. At each step, it yields a clean-layout prediction

$$z_0^t = \frac{z_t - \sqrt{1 - \bar{\alpha}_t}\, \epsilon_\theta(z_t, t; y, c)}{\sqrt{\bar{\alpha}_t}}. \tag{3}$$

This prediction provides the representation on which design objectives and edits can be evaluated during sampling. It remains continuous; conversion to an admissible PCB is required before full-wave evaluation. Conditioning encourages response matching but does not directly capture the simulator objective in Eq. 2.

Loss guidance. Guided diffusion supplements the learned denoising process with a target-dependent loss gradient (Chung et al., 2023; Ye et al., 2024; Dreossi et al., 2026). Because the full-wave simulator does not expose this gradient, a standard approach uses the differentiable surrogate. Let $\widetilde{\mathcal{L}}_{\text{surr}}$ denote its loss extended to the continuous layout representation, with a differentiable decoder when needed. The guidance direction is

$$g_t = -\lambda_t \nabla_{z_t} \widetilde{\mathcal{L}}_{\text{surr}}(z_0^t; y, c), \qquad \lambda_t \geq 0, \tag{4}$$

where the gradient is taken through the clean prediction and λ_t controls the guidance strength. This makes response-directed updates inexpensive, but requires a differentiable path to the surrogate loss; hard geometry extraction and binary rasterization do not provide that path.

Limitations. While differentiability makes guidance computable, it does not ensure that it improves the simulated response. Low surrogate loss can coexist with substantially larger simulator loss (Figure 1(c)). Differentiability alone establishes neither accurate local directions nor reliable step sizes. Optimizing $\mathcal{L}_{\text{surr}}$ can therefore favor layouts that do not improve $\mathcal{L}_{\text{sim}}$, particularly when edits leave the surrogate's training distribution. This motivates using simulator feedback to assess proposed changes while limiting its cost under the budget B.

5 Simulator-Refined Diffusion

This section introduces *Simulator-Refined Diffusion* (SRD) which leverages a full-wave simulator within the diffusion sampling process and spends its calls on selecting among a small set of candidate

edits rather than relying solely on random perturbations. At selected denoising steps, SRD first extracts a geometric representation of the predicted layout and uses the surrogate to propose an edit direction. It then uses simulator feedback from a trial displacement along the surrogate descent direction to determine the subsequent search strategy. If following this direction can sufficiently reduce the simulator loss $\mathcal{L}_{\text{sim}}$, SRD tests additional step sizes along the same direction. Otherwise, it explores other directions. The candidate with the lowest simulator loss is then written back into the diffusion state. The method consists of two key components:

1. **Edit space:** Unlike guidance applied to continuous diffusion states, SRD operates on a low-dimensional geometric representation composed of rectangles. This representation is also required by the simulator and fabrication workflows.
2. **Surrogate+simulator edits:** SRD uses surrogate gradients to propose a search direction. The simulator evaluates candidate edits and selects the one with the lowest simulator loss.

A key feature of SRD is that it requires no additional training, since the diffusion prior and the surrogate are used as given and guidance happens exclusively at sampling time. This also allows the model to generalize to unseen targets and contexts, since the surrogate is used only to propose edits and the simulator is used to select among them.

Concretely, during generation, at each reverse step t the model predicts a clean sample z_0^t from the current (noisy) state z_t. Then, at pre-defined refinement steps, SRD maps this prediction to a structured geometry $p_0 = G(z_0^t)$, which extracts a layout description required by the PCB simulator, queries the surrogate for a direction u, and evaluates a small set of candidates induced by u with the simulator. From this set, SRD selects the best candidate $p^\star$, writes it back into the clean prediction, and re-noises before continuing to z_{t-1}:

$$z_t \longrightarrow z_0^t \xrightarrow{G} p_0 \xrightarrow{\boldsymbol{F}_{\text{surr}}} u \xrightarrow{\boldsymbol{F}_{\text{sim}}} p^\star \longrightarrow z_0^t \longrightarrow z_{t-1}.$$

Figure 1, in the Introduction, illustrates schematically this procedure. The following subsections describe the components of SRD in detail.

5.1 Structured edit space

At each refinement step t, the diffusion model predicts a clean layout $z_0^t \in \mathbb{R}^{H\times W}$ from the noisy state z_t using Eq. 3, where H and W denote the image height and width. The predicted clean sample is a continuous-valued representation of the layout, but the simulator requires an explicit geometric description. To bridge this gap, we define a geometry map G that converts the prediction into a structured representation $p_0 = G(z_0^t)$. We implement G using a CenterNet model (Zhou et al., 2019) trained on rasterized layouts to predict rectangle centers and dimensions.

Edits could in principle be applied to z_0^t and mapped through G for evaluation, but continuous state edits result in inefficient allocation of downstream simulator calls. Pixel-space perturbations affect the simulator input through the learned map G, so their magnitude does not directly control the resulting changes in rectangle positions or dimensions. Instead, SRD operates on the structured description that G produces. For a layout with n extracted rectangles, p_0 contains the two center coordinates, width, and height of each rectangle. An edit is thus a displacement in $\mathbb{R}^{4n}$ rather than in $\mathbb{R}^{H\times W}$. This representation allows SRD to control rectangle positions and dimensions directly, producing explicit candidate geometries for simulator evaluation.

5.2 Surrogate-proposed, simulator-selected edits

Operating in rectangle space significantly reduces the dimensionality and provides direct control over rectangle positions. However, choosing which edits to evaluate remains challenging. Random perturbations can be used to estimate an update direction, but the perturbations themselves are sampled without reference to the design objective. This can squander limited simulator calls on candidates that offer little or no improvement in response matching. To address this issue, SRD uses *the surrogate* to propose an edit direction. In practice, $\boldsymbol{F}_{\text{surr}}$ is a state-of-the-art forward model pretrained on the same corpus as the diffusion prior (Section 6); this surrogate provides a descent direction on the design objective,

$$u = -\nabla_p \mathcal{L}_{\text{surr}}(p; y, c)\big|_{p=p_0}. \tag{5}$$

Computing this direction through the surrogate remains inexpensive in execution since it is natively differentiable, so the cost is negligible in comparison to the full-wave simulator.

Although the surrogate provides a *suggested* direction for improvement, a decrease in surrogate loss may not translate into a decrease in simulator loss (Figure 1 (c)). Errors in the predicted response can distort the loss gradient, particularly when edits move the geometry away from the training distribution. To mitigate this issue, SRD first evaluates a trial displacement $p_0 + hu$ with the simulator, where h is a fixed step size. The probe provides direct evidence about this proposed step and determines how subsequent candidates are generated. We consider three ways of using this evidence, which differ in how the remaining candidates are constructed once the probe result is known.

The **directional** variant evaluates additional step sizes along u if the probe sufficiently reduces the simulator loss, and along $-u$ otherwise. The intuition is that the surrogate may identify a useful search axis even when its preferred sign is inaccurate. Reversing the direction allows the simulator to test this possibility while keeping the search focused. However, this strategy cannot explore edits outside the surrogate-defined axis, and a failed probe does not guarantee that the opposite direction will improve the loss.

The **adaptive** variant continues along u only if the probe sufficiently reduces the simulator loss. Otherwise, it falls back to random perturbations of the rectangle parameters. This reduces reliance on the surrogate when its proposed step is ineffective and explores directions outside the initial search axis, at the cost of less focused exploration.

Finally, the **hybrid** variant evaluates both probe-selected directional candidates and random perturbations at each refinement step. Combining these proposals retains the information supplied by the surrogate while considering more diverse geometric edits. This broader search significantly increases the required simulator call budget, but can provide the strongest results in many cases.

For each variant, SRD selects the candidate with the lowest simulator loss from its candidate set $\mathcal{P}$:

$$p^\star \in \arg\min_{p \in \mathcal{P}} \mathcal{L}_{\text{sim}}(p; y, c). \tag{6}$$

Each candidate set includes the unedited geometry p_0, ensuring that the selected geometry has a simulator loss no greater than that of p_0. Thus, the refinement quality monotonically improves (or remains the same). Algorithm 1 (A.2) summarizes the full sampling procedure for all three variants.

Remark. While the search strategies specify how to propose candidate edits, we must decide when to apply refinement and how many candidates to evaluate at each step. These choices are highly coupled under a limited simulator budget. Earlier refinement can offer greater flexibility, as more denoising steps remain to accommodate layout changes. However, the clean prediction may still be unstable, making rectangle extraction less reliable. Subsequent denoising may also overwrite the selected edit, leading to reduced downstream improvement under a matched budget. Later refinement operates on a more stable geometry but can reduce opportunities to recover from a poor design. The candidate set size also introduces a trade-off: larger sets allow more step sizes or directions to be compared but require more simulator calls, whereas smaller sets reduce budget cost but may miss useful edits. In our experiments, SRD consistently outperforms all baselines across the tested refinement settings and candidate budgets (Section 6.2).

6 Experiments

Benchmark. All methods are evaluated on a test set of 250 samples from 25 layout templates, with 10 samples per template, as shown in Figure 2. An out-of-distribution (OOD) set of 50 constructed S-parameters is also introduced to expand the range of S-parameters specifications beyond the test set. Each sample specifies a board's complex S-parameters at 51 frequency points over 1–20 GHz, feed positions, and substrate properties. Both sets include one- and two-port configurations. Ground-truth S-parameters are computed with Palace (AWS Center for Quantum Computing, 2023), a full-wave 3D finite-element electromagnetic solver built on the MFEM library (Anderson et al., 2021). Simulations are performed at 51 frequency points over 1–20 GHz, with a 900 s time limit per solve. We evaluate generated layouts

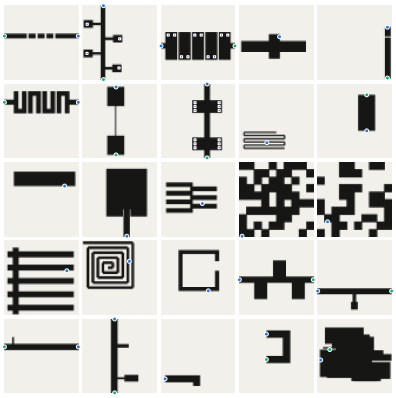

Figure 2: Representative boards.

Method	Test set					OOD set				
	Mean loss↓	Median loss↓	Reduction (%)↑	Sim. success (%)↑	Palace calls	Mean loss↓	Median loss↓	Reduction (%)↑	Sim. success (%)↑	Palace calls
CD (Dreossi et al., 2026)	0.3222	0.2820	0.00	<u>99.6</u>	1	0.5418	0.6120	0.00	**98.0**	1
GGD (Ye et al., 2024)	0.3174	0.2687	1.49	98.4	1	0.5249	0.6121	3.12	<u>96.0</u>	1
DPO (Zampini et al., 2025)	0.2819	0.2441	12.51	**100.0**	7	0.4821	0.5403	11.02	<u>96.0</u>	17
SRD (adaptive)	<u>0.2608</u>	**0.1977**	<u>19.06</u>	**100.0**	7	0.4444	<u>0.4338</u>	17.98	**98.0**	17
SRD (directional)	0.2640	<u>0.1980</u>	18.06	**100.0**	7	**0.4346**	**0.4082**	**19.79**	**98.0**	17
SRD (hybrid)	**0.2539**	**0.1977**	**21.20**	**100.0**	11	<u>0.4361</u>	0.4368	<u>19.51</u>	<u>96.0</u>	29

Table 1: Comparison on the test and OOD sets. DPO and all SRD variants apply guidance during the last 2 denoising steps on the test set and the last 4 on the OOD set. Palace calls count the selection budget at each guided step plus one final Palace evaluation; CD and GGD use only the final evaluation. Reductions are relative to the mean loss of CD. Simulation success is the percentage of generated boards with a valid Palace result. For each performance metric, bold and underlined values indicate the best and second-best results, respectively.

using the RMSE $\mathcal{L}_{\text{sim}}(x; y, c) = \|\boldsymbol{F}_{\text{sim}}(x, c) - y\|_{\text{F}}/(P\sqrt{2M})$, where $M = 51$ is the number of frequency points and P is the number of ports. The Frobenius norm is taken over all port pairs and frequencies. Lower values indicate closer agreement with the target.

Pretrained models and data. All methods share the same two frozen networks from Dreossi et al. (2026). The diffusion prior is a conditional UNet-based denoiser over 64×64 grayscale metal masks of an 8×8 mm board, conditioned on the complex target response (all port pairs at $M{=}51$ frequency points over 1–20 GHz), the port positions, and the substrate properties (permittivity, loss tangent, and thickness). The differentiable surrogate $\boldsymbol{F}_{\text{surr}}$ is a ViT-based forward model that maps a layout, ports, and substrate description to the complex response on the same frequency grid. Both the surrogate and diffusion models are trained on the same corpus of $\sim$1.26M procedurally generated PCBs spanning the same family of parameterized templates, each labeled with full-wave S-parameters obtained from Palace simulations. Neither network is fine-tuned for any of the following experiments and all guidance happens at sampling time. Architectures, training hyperparameters, and compute are detailed in Dreossi et al. (2026).

Competing methods. We evaluate our method, SRD, alongside three baselines:

(1) **Conditional diffusion (CD)** (Dreossi et al., 2026): CD generates PCB layouts conditioned on target S-parameters.
(2) **Gradient-guided diffusion (GGD)** (Ye et al., 2024): GGD adapts diffusion posterior sampling to guide denoising with gradients of the S-parameter loss calculated by surrogate.
(3) **Differentiable Perturbed Optimizer (DPO)** (Zampini et al., 2025): DPO uses random perturbations to estimate update directions from simulator evaluations.

Our method, **Simulator-Refined Diffusion (SRD)**, uses surrogate gradients to propose search directions and simulator evaluations to estimate update directions. The *adaptive* strategy uses perturbations along the surrogate descent direction if an initial probe sufficiently reduces $\mathcal{L}_{\text{sim}}$, and random perturbations otherwise. The *directional* strategy follows or reverses the surrogate descent direction based on the probe. The *hybrid* strategy uses these directional perturbations together with random perturbations within each guidance step.

6.1 Comparison across Methods

Table 1 shows that all SRD variants outperform the baselines in both mean and median S-parameter loss $\mathcal{L}_{\text{sim}}$ on both sets. GGD provides only modest loss reductions over CD (1.49% and 3.12%), whereas DPO achieves larger improvements (12.51% and 11.02%) through simulator feedback. With the same Palace call budgets as DPO, SRD increases these reductions to 19.06% and 17.98% on the test and OOD sets, respectively. These results demonstrate that surrogate gradients help identify more promising candidate edits than random perturbations, while simulator evaluations determine which edits most effectively reduce the S-parameter loss.

On the test set, SRD (hybrid) considers both perturbations along the surrogate direction and random perturbations within each guidance step. This larger candidate set helps identify better edits, yielding the lowest mean S-parameter loss. On the OOD set, the three variants perform similarly, with SRD (directional) achieving the lowest mean loss. The diffusion model cannot produce favorable samples with low S-parameter loss before guidance is applied. Therefore, larger perturbations may offer greater opportunities for improvement. SRD spends part of its Palace call budget assessing whether to follow the surrogate direction, leaving fewer calls to evaluate larger perturbations. This allocation induces slightly higher loss than the directional and hybrid variants.

Next, Figure 3 compares the mean S-parameter loss across 25 templates grouped by difficulty. SRD's gains over CD are most pronounced on medium and hard templates, where CD produces larger response errors. SRD also improves on easy templates, although the loss levels reached by CD are already low, leaving relatively less room for improvement.

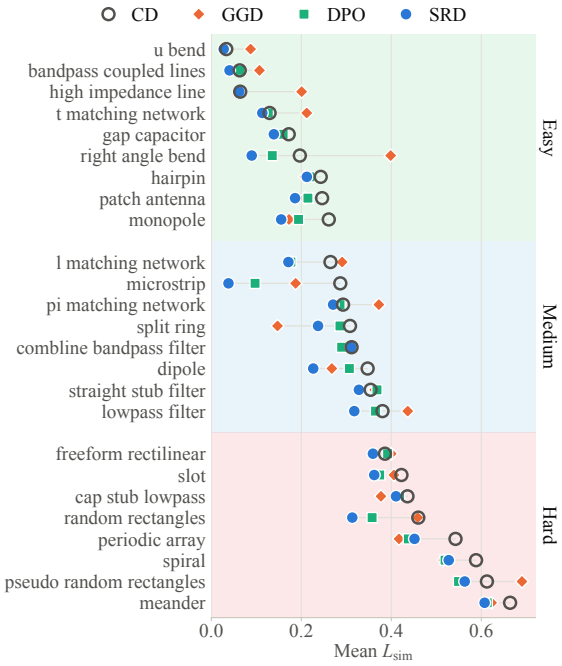

Figure 3: Mean S-parameter loss by template on the test set. Shaded regions indicate the easy, medium, and hard target groups. Lower loss is better.

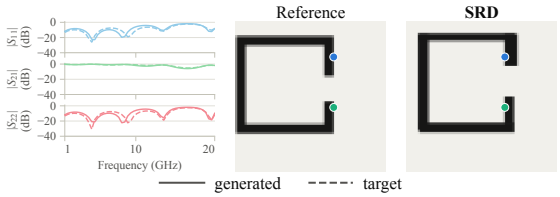

Figure 4: Reference (center) and SRD (right) layouts with their target S-parameters (left).

Finally, we report an example layout generated by SRD alongside the reference layout and their S-parameter responses target (Figure 4). Notice how the SRD proposed layout produces responses that are closer to the target values (Figure 4(left)). Interestingly, when comparing the reference and the generated layouts, one can note that while the port locations remain fixed, the metal geometry generated shifts relative to them. In the reference layout, the ports lie on the right edges of the vertical metal segments. In the SRD layout, they lie on the left edges. Throughout the generated set, we observe similar results. This artifact illustrates the "creative" ability of SRD in searching for good layouts that achieve the target S-parameter response. In other words, SRD is capable of constructing novel layouts for a given design objective.

Appendix A.3 illustrates how SRD perturbs the layout during the denoising steps. Appendix A.4 further examines how SRD improves response accuracy while retaining proximity to the learning distribution. Appendix A.5 provides further qualitative comparisons across methods and SRD.

6.2 Quality–Cost Trade-off

Figure 5 shows that all SRD variants consistently outperform CD and DPO in three different number of guidance steps settings. SRD (hybrid) achieves the lowest mean S-parameter loss in each setting with a larger Palace call budget, which allows evaluating both directional and random perturbations. The right panel illustrates the efficiency of

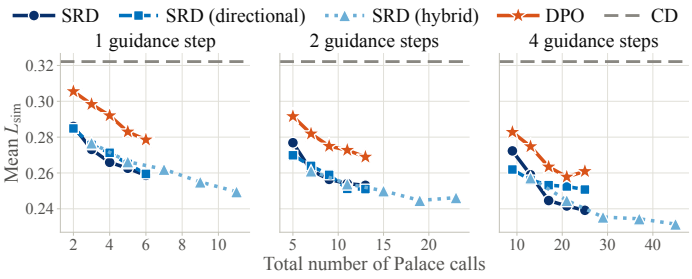

Figure 5: Quality–cost trade-off across methods.

Method	Perturbations		Candidate selection		Performance		
	Surrogate direction	Random	Palace	Surrogate	Mean loss ↓	Median loss ↓	Reduction (%) ↑
Ablations	×	✓✓	×	×	0.4064	0.4058	−26.14
	×	✓✓	×	✓✓	0.3263	0.2986	−1.28
	✓✓	×	×	✓✓	0.2873	0.2284	10.83
CD	×	×	×	×	0.3222	0.2820	0.00
DPO	×	✓✓	✓✓	×	0.2819	0.2441	12.51
SRD	✓✓	✓	✓✓	×	**0.2608**	**0.1977**	**19.06**

Table 2: Ablations on the test set. Double checkmarks indicate components that are always used. A single checkmark indicates conditional use, and crosses indicate inactive components. Candidates use surrogate descent directions or random perturbations; selection uses Palace evaluations or surrogate predictions. Mean loss reductions are relative to CD.

adaptive strategy selection. At 21 Palace calls, SRD achieves lower loss than both variants by adaptively selecting the more effective search strategy. However, it performs slightly worse than either at 13 calls because each guidance step uses two initial calls to evaluate the current design and a probe, leaving fewer calls for further exploration. Furthermore, the middle and right panels show how budget allocation affects performance. At nine Palace calls, SRD achieves lower loss with two guidance steps than with four. Under this tight budget, SRD can benefit more from evaluating additional candidates per step than from increasing the number of guidance steps.

6.3 Ablations

To evaluate the contributions of surrogate directions and simulator feedback, we vary the perturbation strategy and candidate selector in Table 2. With random perturbations, surrogate-based selection yields a mean loss of 0.3263, slightly above CD (0.3222), while Palace-based selection reduces it to 0.2819. This comparison supports the use of simulator feedback for selecting effective edits, since surrogate predictions may not reliably track the simulated loss.

Surrogate gradients also improve candidate generation. Under surrogate-based selection, surrogate descent directions yield a mean loss of 0.2873, compared with 0.3263 for random perturbations. SRD (adaptive) combines surrogate-proposed directions with Palace-based selection and a conditional fallback to random perturbations. It achieves the lowest mean and median losses among the configurations in this ablation: 0.2608 and 0.1977, respectively.

7 Limitations and Conclusion

This paper introduced Simulator-Refined Diffusion (SRD) for radio-frequency printed circuit board (PCB) inverse design. SRD is a diffusion model that incorporates surrogate-based search directions and full-wave simulator feedback. At selected denoising steps, SRD extracts a structured geometric representation of the predicted clean layout and uses a surrogate gradient to propose an edit direction. Then, a full-wave simulator probe assesses this direction and guides the generation of candidate edits. Across 250 test targets spanning 25 layout templates and 50 constructed out-of-distribution targets, SRD variants reduce mean simulated S-parameter loss by up to 21.2% and 19.8%, respectively, relative to a state-of-the-art conditional diffusion model for PCB design. Our proposed approach also outperforms other baselines at matched simulator-call budgets on both sets and an extensive ablation study supported the complementary roles of surrogate directions and simulator-based selection. These results are significant and suggest that the dual use of the surrogate and simulator feedback can become an important tool for the PCB inverse design process.

Building on these promising results, future work could extend SRD's rectangle-based representation to capture finer geometric features and more complex topologies. Additionally, exploring search directions beyond those proposed by the surrogate could further expand the range of PCB designs the method can produce.

AI USE STATEMENT

In this paper, we used AI tools to assist in creating assets for diagrams and figures, for accelerating code development pipelines, and in a limited role for refining the written exposition. We have not used generative AI tools in any way to support development of the methodology, in any role contributing to mathematical correctness of the work, nor as an author of the paper. All AI assisted work has been carefully reviewed, and we take full responsibility for the content of the work.

ETHICS STATEMENT

This work studies computational inverse design of RF printed-circuit-board layouts and does not introduce a dataset of people or make decisions about individuals. Its intended benefit is to improve electromagnetic-response matching under a constrained full-wave simulation budget. Potential risks include dual-use applications of improved RF design tools and the computational and energy cost of model training and repeated finite-element simulation. The method also remains subject to simulation error and should not be treated as sufficient validation for fabrication or safety-critical deployment. Appropriate downstream use requires domain review, manufacturing checks, regulatory compliance where applicable, and physical validation.

REPRODUCIBILITY STATEMENT

The problem definition and budget formulation are given in Section 3; the sampling procedure is specified in Section 5 and Algorithm 1; and the benchmark, Palace metric, baselines, and guidance settings are described in Section 6 and Table 1. Appendix A.1 records the available simulator settings, and Appendix A.2 gives the complete sampling pseudocode.

REFERENCES

Niels Aage and Villads Egede Johansen. Topology optimization of microwave waveguide filters. *International Journal for Numerical Methods in Engineering*, 112(3):283–300, 2017. doi: 10.1002/nme.5551. URL https://doi.org/10.1002/nme.5551.

Robert Anderson, Julian Andrej, Andrew Barker, Jamie Bramwell, Jean-Sylvain Camier, Jakub Cerveny, Veselin Dobrev, Yohann Dudouit, Aaron Fisher, Tzanio Kolev, et al. Mfem: A modular finite element methods library. *Computers & Mathematics with Applications*, 81:42–74, 2021.

AWS Center for Quantum Computing. Palace: PArallel LArge-scale Computational Electromagnetics, 2023. URL https://github.com/awslabs/palace.

Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach. Learning with differentiable perturbed optimizers. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 9508–9519. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/6bb56208f672af0dd65451f869fedfd9-Paper.pdf.

Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=OnD9zGAGT0k.

Tommaso Dreossi, Christopher M Bryant, Hao Liu, Nathan Mirman, Noah Kessler, Michael Frei, and Harish Krishnaswamy. Inverse design of multi-layer sub-pixel-resolution rf passives through grayscale diffusion with flexible s-parameter conditioning. *arXiv preprint arXiv:2605.08233*, 2026.

Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering*, 79(11):1309–1331, 2009.

Yingqing Guo, Emir Ali Karahan, Zihao Li, Zijian Shao, Zaixi Zhang, Mengdi Wang, and Kaushik Sengupta. Dall-em: Generative ai with diffusion models for new design space discovery and target-to-electromagnetic structure synthesis. In *2025 IEEE/MTT-S International Microwave Symposium - IMS 2025*, pp. 926–929, 2025. doi: 10.1109/IMS40360.2025.11103838.

Roger F. Harrington. *Field Computation by Moment Methods*. Wiley-IEEE Press, 1993. ISBN 978-0-7803-1014-8. Reissue of the 1968 edition.

Emadeldeen Hassan, Eddie Wadbro, and Martin Berggren. Topology optimization of metallic antennas. *IEEE Transactions on Antennas and Propagation*, 62(5):2488–2500, 2014. doi: 10.1109/TAP.2014.2309112. URL https://ieeexplore.ieee.org/document/6750741.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Vineet Jain, Kusha Sareen, Mohammad Pedramfar, and Siamak Ravanbakhsh. Diffusion tree sampling: Scalable inference-time alignment of diffusion models. *Advances in Neural Information Processing Systems*, 38:161584–161623, 2025. doi: 10.52202/085713-4871. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/d6484394c4cb5e1f4ecad8d90b912025-Abstract-Conference.html.

Jian-Ming Jin. *The Finite Element Method in Electromagnetics*. Wiley-IEEE Press, 3 edition, 2014. ISBN 978-1-118-57136-1.

Jinhao Liang, Jacob K Christopher, Sven Koenig, and Ferdinando Fioretto. Simultaneous multi-robot motion planning with projected diffusion models. *arXiv preprint arXiv:2502.03607*, 2025a.

Jinhao Liang, Yixuan Sun, Anirban Samaddar, Sandeep Madireddy, and Ferdinando Fioretto. Chance-constrained flow matching for high-fidelity constraint-aware generation. *arXiv preprint arXiv:2509.25157*, 2025b.

Jinhao Liang, Sven Koenig, and Ferdinando Fioretto. Simulation-informed diffusion for decentralized multi-robot motion planning. *arXiv preprint arXiv:2605.27697*, 2026.

Nanye Ma, Shangyuan Tong, Haolin Jia, Hexiang Hu, Yu-Chuan Su, Mingda Zhang, Xuan Yang, Yandong Li, Tommi Jaakkola, Xuhui Jia, and Saining Xie. Scaling inference time compute for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2523–2534, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Ma_Scaling_Inference_Time_Compute_for_Diffusion_Models_CVPR_2025_paper.html.

Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=aBsCjcPu_tE.

Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017. doi: 10.1007/s10208-015-9296-2. URL https://doi.org/10.1007/s10208-015-9296-2.

David M. Pozar. *Microwave Engineering*. John Wiley & Sons, 4 edition, 2011. ISBN 978-0-470-63155-3.

Aristides G Requicha. Representations for rigid solids: Theory, methods, and systems. *ACM Computing Surveys (CSUR)*, 12(4):437–464, 1980.

Piyush Saha, Evan Newell, Hanna O'Leary, Arun Natarajan, and Alireza Aghasi. K-TRAIL: Simulator-guided generative design of EM/RF circuits, 2026. URL https://arxiv.org/abs/2609.23183.

Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.

Haotian Ye, Haowei Lin, Jiaqi Han, Minkai Xu, Sheng Liu, Yitao Liang, Jianzhu Ma, James Zou, and Stefano Ermon. Tfg: Unified training-free guidance for diffusion models. *Advances in Neural Information Processing Systems*, 37:22370–22417, 2024.

Stefano Zampini, Jacob K Christopher, Luca Oneto, Davide Anguita, and Ferdinando Fioretto. Training-free constrained generation with stable diffusion models. *Advances in Neural Information Processing Systems*, 38:30993–31024, 2025. doi: 10.52202/085713-0918. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/2774a3b52d436b5930da660dd2b32b3a-Abstract-Conference.html.

Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019.

A Appendix

A.1 Simulator configuration and evaluation

Candidate and final layouts are converted into solid models (Requicha, 1980), meshed with `gmsh` (Geuzaine & Remacle, 2009), and evaluated with Palace (AWS Center for Quantum Computing, 2023). Each Palace invocation returns the complete S-matrix for all ports at 51 uniformly spaced frequencies from 1 to 20 GHz; internal frequency-sweep solves are included in that invocation.

Geometry and materials. Boards measure 8×8 mm and use a 64×64 raster. CenterNet (Zhou et al., 2019) extracts the rectangular signal geometry. The signal layer and full-board ground plane are each 35 μm thick and modeled as conductive copper volumes with $\sigma = 5.8 \times 10^7$ S/m. Substrate relative permittivity ε_r, dielectric loss tangent $\tan\delta$, and thickness are listed in Table 3. Vias have a 0.3 mm diameter and a 0.6 mm pad diameter.

Table 3: Substrate properties used in simulation.

Material	ε_r	$\tan\delta$	Thickness (μm)
Air (suspended)	1.00	0	203.2
Rogers RO4003	3.55	0.0027	203.2
FR-4 (S1000H)	4.60	0.011	200.0

Ports and boundaries. Each feed uses a 50 Ω lumped port spanning the substrate between the ground and signal layers, with width $\min(0.45\,\mathrm{mm}, \text{trace width})$ and excitation along $+z$. The air box extends 2 mm laterally and above the board, with no air below it. The five side and top faces use second-order absorbing conditions; the bottom face has no explicit boundary tag and uses the solver's natural perfect-magnetic-conductor condition.

Mesh and solver. We use `gmsh` 4.15.2 and Palace v0.16.0. Meshes use first-order tetrahedral geometry with characteristic lengths of 50 μm–1 mm and no adaptive mesh refinement. Failed meshing attempts are retried with a 30 μm minimum size and then uniform sizing. Palace uses order-2 Nédélec elements, the SuperLU_DIST direct solver, and an adaptive reduced-order frequency sweep with tolerance 10^{-2}. Each solve uses one MPI rank on an NVIDIA A10G GPU with a 900 s timeout.

A.2 COMPLETE SAMPLING ALGORITHM

Algorithm 1 SRD Sampling

Input: y, c; K reverse sampling steps $t_K > \cdots > t_1$; refinement at the final N steps; step size h absolute loss reduction threshold $\varepsilon_{\mathrm{imp}}$; budget B; variant v; candidate counts $m_{\mathrm{d}}, m_{\mathrm{r}}$
Output: PCB layout $\widehat{x}$

$z \sim \mathcal{N}(0, \mathbf{I})$
for $k = K, \ldots, 1$ **do**
  $\epsilon \leftarrow \epsilon_\theta(z, t_k; y, c)$
  $z_0^{t_k} \leftarrow \textsc{PredictClean}(z, \epsilon, t_k)$
  if $k \leq N$ **and** budget permits refinement **then**
    $p_0 \leftarrow G(z_0^{t_k})$
    $u \leftarrow -\nabla_p \mathcal{L}_{\mathrm{surr}}(p; y, c)|_{p=p_0}$
    $p_1 \leftarrow p_0 + hu$
    $(\ell_0, \ell_1) \leftarrow \textsc{Simulate}(p_0, p_1)$
    $s \leftarrow +1$ **if** $\ell_0 - \ell_1 > \varepsilon_{\mathrm{imp}}$ **else** -1
    if $v =$ adaptive **and** $s = -1$ **then**
      $\mathcal{Q} \leftarrow \textsc{Random}(p_0)$
    else
      $\mathcal{Q} \leftarrow \textsc{Directional}(p_0, su)$
      if $v =$ hybrid **then**
        $\mathcal{Q} \leftarrow \mathcal{Q} \cup \textsc{Random}(p_0)$
      end if
    end if
    $p^\star \leftarrow \textsc{SimulateAndSelect}(\{p_0, p_1\} \cup \mathcal{Q})$
    $z_0^{t_k} \leftarrow \textsc{WriteBack}(z_0^{t_k}, p^\star)$
    $z \leftarrow \textsc{ReNoise}(z_0^{t_k}, \epsilon, t_k)$
  end if
  $z \leftarrow \textsc{SamplingStep}(z, \epsilon, t_k)$
end for
$\widehat{x} \leftarrow \textsc{DecodeLayout}(z)$
Evaluate $\mathcal{L}_{\mathrm{sim}}(\widehat{x}; y, c)$
return $\widehat{x}$

DIRECTIONAL generates m_{d} candidates along the supplied direction using step size h, while RANDOM generates m_{r} candidates using the configured random perturbations. SIMULATEANDSELECT returns the candidate with the lowest simulator loss among those evaluated, including the baseline and probe.

A.3 GEOMETRIC PERTURBATIONS

SRD generates candidate geometries by perturbing the centers, widths, and heights of the extracted rectangles. Palace evaluates these candidates. SRD selects the candidate with the lowest simulator loss and writes its geometry back into the clean prediction. Figure 6 shows denoising steps 17–19 of a 20-step sampling process. SRD and CD share the same layout at step 17. At steps 18 and 19, SRD adds metal along the existing boundaries and widens the conducting strip. The port locations remain fixed. In contrast, CD produces no visible change over these steps.

A.4 EFFECT ON THE DIFFUSION PRIOR

For each method, we compute the L_2 distance from each generated layout to its nearest training sample in DINOv3 embedding space (Siméoni et al., 2025). Figure 7 compares these distances with those of CD. SRD retains distances close to CD for many samples while achieving the lowest mean and median S-parameter losses among the four compared methods (Table 1). These results suggest that SRD improves response accuracy with limited changes to the learned distribution.

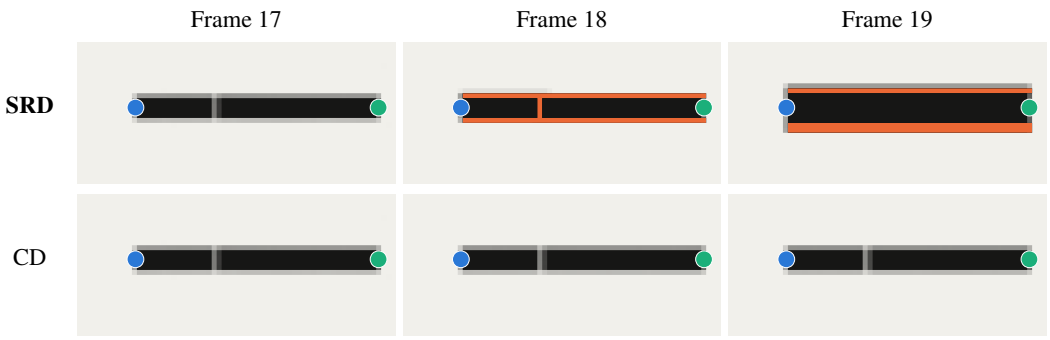

Figure 6: Intermediate layouts at denoising steps 17–19 of a 20-step sampling process for SRD (top) and CD (bottom). Both methods start from the same Frame-17 layout. Black denotes metal. Orange marks the geometric perturbations applied by SRD. Blue and aqua dots mark ports 1 and 2.

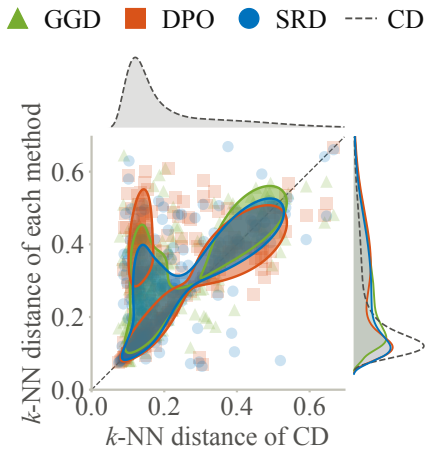

Figure 7: Nearest-neighbor distances ($k = 1$) to the training set in DINOv3 embedding space. Each point represents one sample. In the scatter plot, the horizontal axis reports the CD distance. The vertical axis reports the distance for the corresponding GGD, DPO, or SRD layout on the same target and seed. Colored contours show two-dimensional kernel density estimates (KDEs). Marginal KDE curves summarize the distance distributions along each axis. Points on the dashed diagonal indicate unchanged proximity to the training set; points above and below indicate layouts that are farther from and closer to it. This comparison shows how much each method moves layouts relative to the prior, and should be read alongside the losses in Table 1.

A.5 Qualitative Design Comparison

Figure 8 compares an example target PCB with layouts generated by CD, GGD, DPO, and SRD and their S-parameter responses. CD exhibits a pronounced transmission dip near 15 GHz that is absent from the target response. GGD and DPO remove this dip but retain substantial deviations in the reflection magnitudes. SRD more closely matches both the reflection and transmission responses across the frequency range, illustrating how geometric refinement improves electromagnetic agreement with the target.

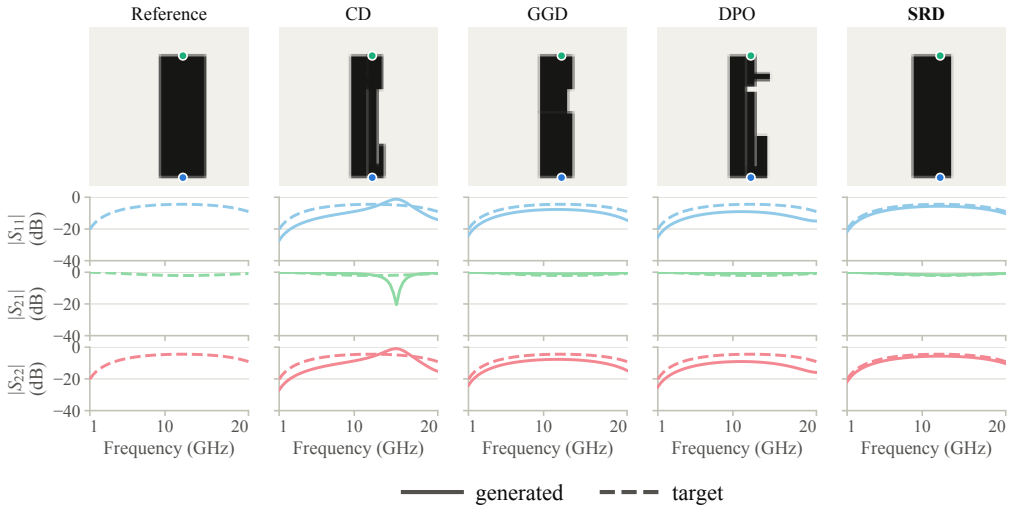

Figure 8: PCB layouts and their S-parameter magnitude responses. Solid and dashed curves indicate generated and target responses, respectively.

Figure 9 shows example PCB layouts generated by the adaptive, directional, and hybrid SRD variants. In these examples, the adaptive and directional variants often produce similar overall structures, with differences in conductor widths and local branch geometry. The hybrid variant shows more pronounced geometric differences in several cases.

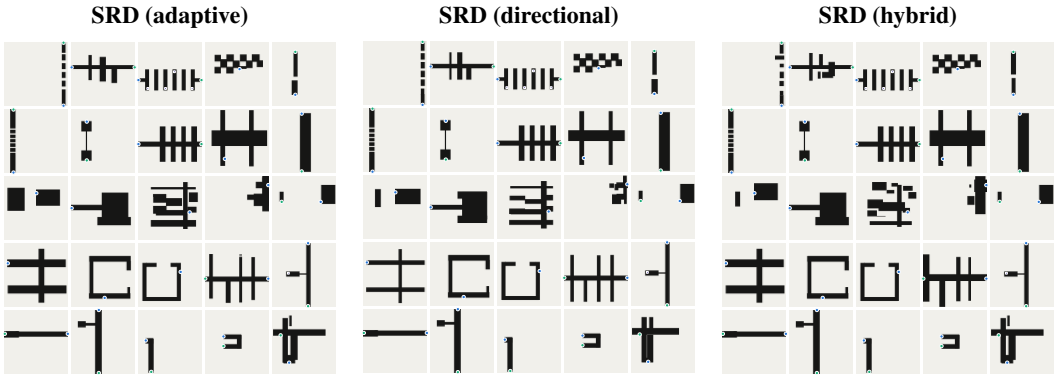

Figure 9: Example PCB layouts generated by the three SRD variants. Each panel contains 25 layouts.