DIFFUSION-GUIDED ELITE PATH INPAINTING FOR 3D ELECTROMAGNETIC COIL INITIALIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Electromagnetic coil design is a representative physical inverse-design problem: finding conductor geometries that achieve desired physical performance under geometric constraints. It underpins applications including power transfer, imaging, and induction heating. Although automatic synthesis is established for prescribed supports and parameterizations, many optimization workflows still require a manually supplied initial coil. We address this initialization gap by generating geometrically feasible candidate routes with high performance potential under a limited simulation budget. Our framework combines dual-view action diffusion with performance-preference-guided elite path reconstruction. Sequence and spatial views capture route order and three-dimensional occupancy, respectively. Quality-weighted residual learning from evaluated paths guides reconstruction while constraining deviations from a frozen geometry base. Hard depth-first search filters invalid actions and backtracks from dead ends; complete-path validation precedes evaluation and synchronous elite-archive updates. The archive maintains a pool of promising parents so that later rounds can build on evaluated candidates. We define a benchmark spanning symmetric enclosing, asymmetric enclosing, and non-enclosing scenes, three label budgets, and five repetitions. Averaged over five runs, our method leads the four learning baselines in Top-5@100 and Mean@100 across all nine settings, and in Best@100 in seven.

1 INTRODUCTION

Electromagnetic coil design is a representative physical inverse-design problem: one specifies desired physical behavior and seeks a conductor geometry that realizes it (Turner, 1986; Peeren, 2003). It underpins power transfer (Covic & Boys, 2013), magnetic resonance imaging (Turner, 1993), and induction heating (Lucía et al., 2014). Conductor geometry controls field distribution and coupling while determining whether a coil fits around the workpiece and other components. This coupling of physical response and geometric constraints also characterizes inverse design in nanophotonics (Molesky et al., 2018).

Automatic coil synthesis is established on prescribed supports (Amrein et al., 2022); shape optimization adjusts parameterized or free-form geometries (Hyun et al., 2025; Zhu et al., 2018; Takahashi, 2024), while topology optimization can construct initial coil layouts (Kaptanoglu et al., 2024). Learning-based methods also generate designs within specified representations (Curtis et al., 2023; Lalla et al., 2025). Nevertheless, expert-proposed layouts remain common in induction-coil design (Goldstein et al., 2014), and local shape optimization refines a supplied starting geometry (Takahashi, 2024). Automating refinement does not resolve how to obtain a useful starting structure for a new scene. We target this preceding step: *automated initialization for electromagnetic inverse design* (Figure 1). Specifically, we study *simulation-budgeted generative coil initialization*, constructing feasible candidate pools with high performance potential for subsequent refinement without requiring an expert to propose each structure.

This initialization task presents three challenges. First, the number of possible routes grows combinatorially, making exhaustive search impractical. Second, a coil must connect fixed terminals, avoid obstacles and self-contact, and satisfy clearance requirements. Explicit feasibility checks are therefore essential, as in constrained motion planning (Karaman & Frazzoli, 2011). Third, simulation evaluation limits the available physical labels (Giovannetti et al., 2024). The task therefore com-

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.

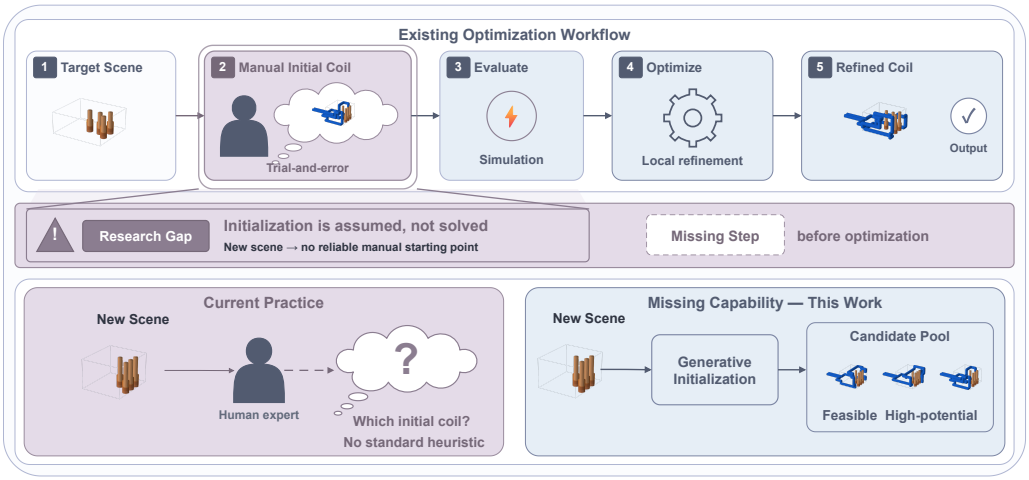

Figure 1: The initialization gap. Existing refinement workflows often rely on a manually supplied initial coil. We target this preceding step by automatically constructing geometrically feasible candidates with high performance potential for downstream optimization under a limited simulation budget.

bines constrained engineering generation (Regenwetter et al., 2022) with data-driven optimization under limited objective feedback (Trabucco et al., 2022).

We address these challenges through diffusion-guided elite path inpainting. Dual-view masked-action diffusion combines route order with spatial occupancy and anchor geometry to reconstruct local path windows. Quality-weighted residual learning uses evaluated paths to favor promising reconstructions while constraining deviations from the frozen geometry base. An elite archive supplies parents and concentrates subsequent edits on promising evaluated structures. Hard depth-first search filters invalid actions and backtracks from branches that cannot reconnect the retained path; complete-path validation precedes simulation. Thus learned performance guidance determines which legal actions to explore, while explicit checks determine which reconstructions can be accepted.

We instantiate the task in three induction-heating scenes with symmetric enclosing, asymmetric enclosing, and non-enclosing layouts. The protocol compares methods at three label budgets over five repetitions, using the best returned candidates to measure initialization potential rather than downstream refinement gains. Our contributions are:

1. **Problem formulation.** We formulate automatic three-dimensional coil initialization for electromagnetic inverse design as constrained route generation under a finite simulation budget.
2. **Method.** We combine dual-view action diffusion, performance-preference-guided elite reconstruction, and backtracking search with complete geometric validation.
3. **Evaluation design.** We define a three-scene comparison and component ablations under shared data access, budgets, and simulation objectives.

2 Problem Formulation

We study the initialization stage of electromagnetic inverse design for a fixed physical scene. A single open conductor must connect two prescribed terminals while avoiding occupied space and invalid self-contact. The desired output is a pool of geometrically feasible routes with high performance potential for later refinement. Figure 2 illustrates the route representation and geometric constraints.

To make route search discrete, we partition a closed routing domain $\Omega \subset \mathbb{R}^3$ into a regular voxel grid. Let o be its lower corner and $h > 0$ the lattice spacing, so an index $v \in \mathbb{Z}^3$ maps to $\psi(v) = o + hv$.

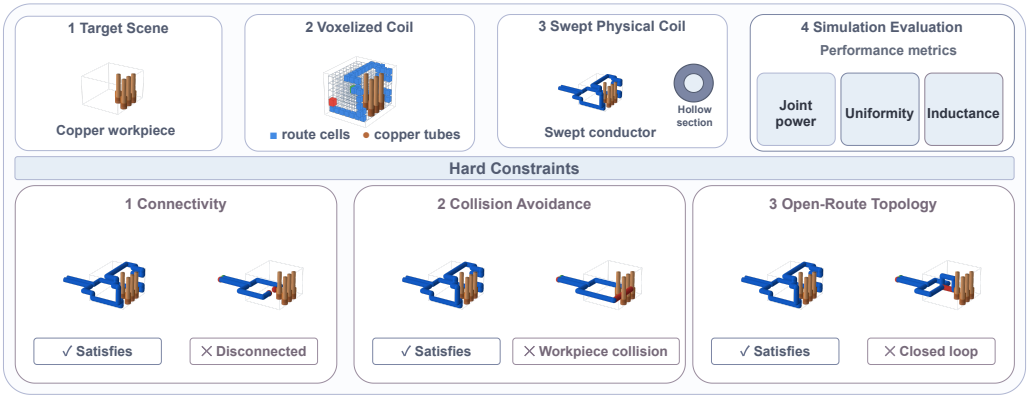

Figure 2: Coil initialization and geometric constraints (schematic). Top: a target scene, a voxel route, its swept conductor, and simulation evaluation of joint power, uniformity, and inductance. Bottom: valid routes contrasted with disconnection, workpiece collision, and a closed loop. Uniformity is quantified by nonuniformity d_U in Section 3; boundary and clearance requirements are included in Equation 3.

The nodes within the domain and their six-neighbor edges form

$$V = \{v \in \mathbb{Z}^3 : \psi(v) \in \Omega\}, \qquad E = \{(u, v) \in V^2 : \|u - v\|_1 = 1\}. \tag{1}$$

Here $\|\cdot\|_1$ is the Manhattan norm; the graph $\mathcal{G} = (V, E)$ provides axis-aligned moves of physical length h.

The terminals become fixed nodes $v_\mathrm{s}, v_\mathrm{g} \in V$. A candidate coil is an ordered centerline route

$$\gamma = (v_0, v_1, \ldots, v_L), \qquad v_i \in V, \tag{2}$$

where L counts route segments. Its physical polyline is $(\psi(v_0), \ldots, \psi(v_L))$, with fixed lead-in and lead-out segments appended before validation and evaluation. Only consecutive nodes define conductor connections; occupancy alone does not specify their traversal order.

Geometry restricts which routes are usable. Clearance-aware voxelization marks blocked nodes $\mathcal{W} \subset V$. Let $\mathrm{clr}(\gamma; \xi)$ be the minimum separation between nonconsecutive segments of the complete centerline, including fixed leads, and let $\rho > 0$ be the required separation. Other fixed conditions, including conductor dimensions, workpiece geometry, leads, materials, excitation, and simulation settings, are collected in η. Together these define $\xi = (\Omega, o, h, \mathcal{G}, v_\mathrm{s}, v_\mathrm{g}, \mathcal{W}, \rho, \eta)$ and the feasible set

$$\begin{aligned} \Gamma(\xi) = \{\gamma : v_0 = v_\mathrm{s},\ v_L = v_\mathrm{g}, \quad & v_i \in V \setminus \mathcal{W}, \\ & (v_i, v_{i+1}) \in E, \quad v_i \neq v_j\ (i < j), \\ & \mathrm{clr}(\gamma; \xi) \geq \rho, \quad \mathrm{ObsClear}(\gamma; \xi)\}, \end{aligned} \tag{3}$$

with node and edge conditions applying over their respective index ranges. Here $\mathrm{ObsClear}$ requires every segment of the complete conductor to satisfy the prescribed workpiece clearance. Thus feasibility includes domain boundaries, terminal connectivity, obstacle avoidance, a non-repeating route, and physical clearance; proximity alone does not create a branch.

Feasibility does not determine performance. Let $f(\gamma; \xi)$ be a simulation-based score, with larger values preferred, and let $\mathcal{D}$ contain available route data and any acquired labels. Given a simulation budget B and a maximum requested pool size K, a design algorithm Alg returns

$$\mathcal{C} = \mathrm{Alg}(\xi, \mathcal{D}; B, K) \subseteq \Gamma(\xi), \qquad |\mathcal{C}| \leq K. \tag{4}$$

The task is to seek high performance potential in $\mathcal{C}$ within B simulation evaluations, counting both labels supplied in $\mathcal{D}$ and new candidate evaluations. Unlabeled feasible paths may support learning without consuming physical labels. Section 3 specifies the score, budget allocation, and pool metrics for the three experimental scenes.

3 Scenes and Evaluation Protocol

We instantiate the task in three induction-heating workpiece layouts with the same voxel spacing but different routing topology (Table 1). Rhombus has a mirror-symmetric target array and a right-side passage that permits an enclosing route. Staggered also permits enclosure, but its alternating targets break mirror symmetry. The Y layout reaches the positive-x boundary and removes the right-side passage, producing a semi-open, non-enclosing problem. The symmetry label describes the scene; no method is required to generate a symmetric coil.

Table 1: Benchmark scenes, routing spaces, and physical targets. All scenes use voxel spacing $h = 12$ mm. The blue and orange wireframes denote the search space and keep-out region; A and B mark the terminals. Targets $P_\xi^\star, \mathcal{L}^\star$ and numerical percentage errors d_P, d_U, d_L are defined below.

Scene	Routing space	Physical target
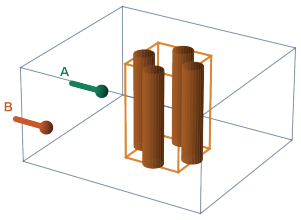 **Rhombus** Symmetric enclosing access	Domain: $120\times108\times60$ mm^3 Grid: $11\times10\times6$	$P_\xi^\star = 70$ W, $d_P \le 25$ $d_U \le 10$ $\mathcal{L}^\star = 245$ nH, $d_L \le 10$
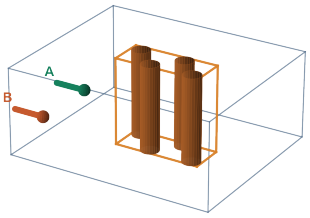 **Staggered** Asymmetric enclosing access	Domain: $144\times120\times60$ mm^3 Grid: $13\times11\times6$	$P_\xi^\star = 70$ W, $d_P \le 25$ $d_U \le 10$ $\mathcal{L}^\star = 245$ nH, $d_L \le 10$
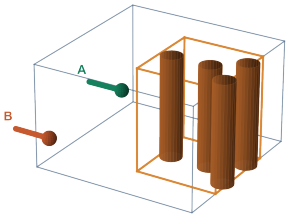 **Y** Semi-open, non-enclosing access	Domain: $96\times96\times60$ mm^3 Grid: $9\times9\times6$	$P_\xi^\star = 20$ W, $d_P \le 25$ $d_U \le 10$ $\mathcal{L}^\star = 245$ nH, $d_L \le 10$

Physical objectives. All candidates are assessed using the same simulation evaluation procedure and physical objectives. For four contact-region powers $p_1, \dots, p_4$, let $\bar{p} = \frac{1}{4}\sum_{j=1}^{4} p_j$. The scene's per-region power target is $P_\xi^\star$; simulated inductance $\mathcal{L}$ has target $\mathcal{L}^\star = 245$ nH. Define the power error, nonuniformity, and inductance error as numerical percentages:

$$d_P = 100\frac{|\bar{p} - P_\xi^\star|}{P_\xi^\star}, \quad d_U = 100\frac{\max_j p_j - \min_j p_j}{\max_j p_j}, \quad d_L = 100\frac{|\mathcal{L} - \mathcal{L}^\star|}{\mathcal{L}^\star}. \tag{5}$$

All maxima and minima range over $j = 1, \dots, 4$. Scoring requires finite $p_j \ge 0$, $\max_j p_j > 0$, and $\mathcal{L} > 0$; invalid evaluations are reported as failures, not assigned fabricated scores. The margins (in percentage points) and official score are

$$m_P = 25 - d_P, \quad m_U = 10 - d_U, \quad m_L = 10 - d_L, \qquad f(\gamma;\xi) = \min(m_P, m_U, m_L). \tag{6}$$

This max–min scalarization (Miettinen, 1999) measures the worst margin; $f \geq 0$ means all three physical thresholds are satisfied.

Data and budget. Each scene provides 10,000 feasible geometry paths, split 80/20 by path hash for geometry training and validation. Let $\mathcal{D}_n^{\text{lab}}$ denote the n physically evaluated supervision paths; nested sets $\mathcal{D}_{100}^{\text{lab}} \subset \mathcal{D}_{200}^{\text{lab}} \subset \mathcal{D}_{300}^{\text{lab}}$ use supervision source seed 0. Each label budget $n \in \{100, 200, 300\}$ includes both the 80% training and 20% validation labels. Geometry paths may overlap the labeled set, but geometry training never reads physical labels. The same labeled set is reused for five repetitions to measure training and generation variation; training seeds 0–4 are paired with search seeds 1200–1204 for our method. In the reported comparison, each run generates and evaluates $K = 100$ new candidates, giving a simulation budget $B = n + 100$ including the supplied labels. Our reported pool contains all 100 valid unique non-parent candidates, without score-based post-selection. Random-DFS uses the distinct reporting pool specified in Section 5.

Performance potential. For a reported pool $\mathcal{C}$ of 100 paths, let $f_{(1)} \geq \cdots \geq f_{(100)}$ be its scores in descending order. The suffix @100 specifies this pool size. We report

$$\text{Best@100} = f_{(1)}, \quad \text{Top-5@100} = \frac{1}{5}\sum_{i=1}^{5} f_{(i)}, \quad \text{Mean@100} = \frac{1}{100}\sum_{i=1}^{100} f_{(i)}. \tag{7}$$

Best@100 tests whether a run finds a strong initializer; Top-5@100 tests whether it offers several strong choices. Mean@100 summarizes the quality of the entire reported pool. These three metrics characterize initialization potential, not improvement after downstream refinement. Metrics are computed within each run, then summarized by mean and sample standard deviation over five seeds.

4 METHOD

4.1 OVERVIEW

To construct initial coils under this budget, our method repeatedly reconstructs local windows of evaluated paths (Figure 3). Geometry-only training first learns how feasible path windows can be reconstructed. Physics-labeled paths then train quality-weighted residual heads with a reference constraint to the frozen geometry model. During search, each round freezes the current archive, assigns offspring slots to selected unique elite parents, reconstructs masked windows with guided hard DFS, evaluates all valid unique completions, and updates the archive only after the round is complete.

4.2 FIXED-WINDOW DUAL-VIEW CONDITION

After separating the terminal connector, we compress the voxel route into M maximal axis-aligned actions $a_i = (d_i, \ell_i) \in \mathcal{A}$. Here $\mathcal{A}$ is the scene's action vocabulary, d_i is one of six signed axes, and ℓ_i is the integer run length in voxel steps. Let $\mathbf{c} = (c_0, \ldots, c_M)$ be their endpoint waypoints, distinct from the dense nodes v_i. We choose internal indices $0 < l < r < M$ with $r - l = W$; the anchors c_l, c_r enclose exactly W compressed actions, not a fixed physical length. The original actions $a_l, \ldots, a_{r-1}$ are removed and never exposed to the model at inference time.

At progress $j \in \{0, \ldots, W-1\}$, the condition contains the fixed action-waypoint prefix, the j reconstructed gap actions, the fixed suffix, and the right anchor. For grid dimensions N_x, N_y, N_z, the spatial view $S_j \in \mathbb{R}^{14 \times N_z \times N_y \times N_x}$ marks prefix/suffix occupancy and directed geometry, the current endpoint, normalized coordinates, the single right anchor, and its reachability box. The sequence view $Q_j \in \mathbb{R}^{50 \times 19}$ stores the ordered visible waypoints followed by one right-anchor row; zero-padded rows are ignored. No row or channel contains an unexecuted action from the removed parent window.

A 3D convolutional neural network (CNN) maps S_j to 64 dimensions. A single-layer unidirectional gated recurrent unit (GRU) (Cho et al., 2014) maps Q_j to 96 dimensions. Their concatenation is fused by a multilayer perceptron (MLP) into a 128-dimensional geometry state g_j. The two views are complementary: the sequence view preserves route order, whereas the spatial view exposes prefix/suffix occupancy, the current endpoint, anchor geometry, and spatial reachability that are difficult to infer from a truncated sequence alone.

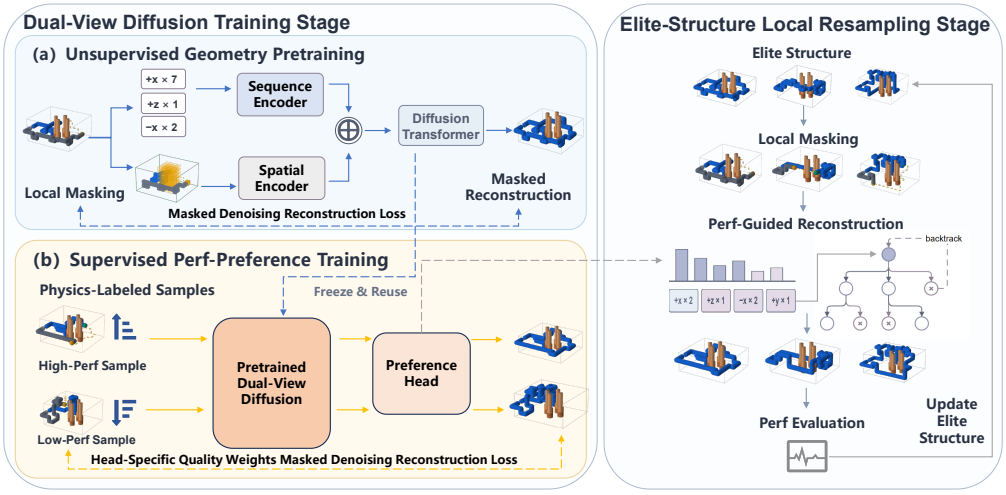

Figure 3: Framework overview. (a) Dual-view geometry pretraining learns masked path reconstruction. (b) Head-specific physical quality weights guide residual training with the geometry base frozen and a reference KL constraint. Right: elite paths undergo local masking and performance-guided reconstruction with geometric checks and backtracking. Valid unique candidates are evaluated before the elite archive is updated synchronously.

4.3 Masked-Action Diffusion and Token Mixing

The geometry base predicts $H = 4$ local action tokens; $\mathbf{b}^t = (b_0^t, \ldots, b_{H-1}^t)$ denotes their noisy values at diffusion step t. At progress j, the first $\min(H, W - j)$ positions are valid and the rest use the padding token PAD; b_q^0 is the clean target at local position q. Using learned action, position, and time embeddings of dimensions 32, 16, and 32, respectively, we concatenate

$$x_q = [g_j; e_{\text{act}}(b_q^t); e_{\text{pos}}(q); e_{\text{time}}(t)] \in \mathbb{R}^{208}. \tag{8}$$

For $q = 0, \ldots, H - 1$, stack the states as rows of $X \in \mathbb{R}^{H \times 208}$ and mix them without changing their dimension:

$$X' = X + \text{MHA}(\text{LN}(X)), \qquad \widetilde{X} = X' + \text{FFN}(\text{LN}(X')). \tag{9}$$

Here LN is layer normalization, MHA is four-head self-attention (Vaswani et al., 2017), and FFN is a $208 \to 416 \to 208$ GELU feed-forward block; PAD positions are excluded as attention keys. The base projection returns logits $z_{\text{base}} \in \mathbb{R}^{4 \times |\mathcal{A}|}$.

Rather than Gaussian corruption as in continuous diffusion (Ho et al., 2020), we use $T = 8$ absorbing-mask steps (Austin et al., 2021), with clean-token retention probability $\bar{\alpha}_t = \cos^2(\pi t/(2T))$. For $t = 1, \ldots, T$, the default corruption scheme replaces valid tokens by MASK through Bernoulli draws with probability $1 - \bar{\alpha}_t$, requiring a nonempty set of masked indices M_t; PAD is excluded. Let $\omega_q(M_t)$ be nonnegative position weights summing to one over M_t. The geometry-base objective is masked cross-entropy (CE):

$$\mathcal{L}_{\text{geo}} = -\mathbb{E}\left[\sum_{q \in M_t} \omega_q(M_t) \log p_{\theta,q}(b_q^0 \mid \mathbf{b}^t, S_j, Q_j, t)\right]. \tag{10}$$

Here θ denotes the base-model parameters and $p_{\theta,q}$ is the action softmax at position q. The expectation averages training paths, uniformly sampled legal windows, all W teacher-forced prefixes, timesteps, and corruption masks. Appendix A gives the architecture and training details.

Inference initializes valid horizon positions as masked tokens and the remaining positions as PAD. Positions 1–3 may be denoised earlier and influence position 0 through the mixer, but position 0

remains masked until the final reverse step. DFS uses that position's ranking to try one action at a time. The condition is then rebuilt from the updated gap prefix, giving a receding-horizon policy rather than committing to all four predicted actions; action diffusion with receding-horizon execution is also used in robot control (Chi et al., 2023).

4.4 Quality-Weighted Physics-Preference Guidance

Physical scores are never concatenated to either view. Instead, they determine path-quality weights for four independent residual heads $r_k : \mathbb{R}^{208} \to \mathbb{R}^{|\mathcal{A}|}$ for $k \in \{\text{official}, P, U, L\}$. Each head is an MLP $208 \to 32 \to |\mathcal{A}|$ with a zero-initialized final layer, and the geometry base is frozen throughout preference training.

For path i, $w_i^{(k)} > 0$ is nondecreasing in its standardized quality target (Appendix B). Let $\ell_i^{(k)}$ be the masked negative log-likelihood (NLL) under head-guided probabilities p_k, and $d_i^{(k)}$ the Kullback–Leibler (KL) divergence $\mathrm{KL}(p_B \| p_k)$ from the detached base probabilities p_B. Both use identical conditions, noisy tokens, and timesteps, with the same masked-position weighting followed by averaging over W prefixes. The head minimizes

$$\mathcal{L}_k = \frac{1}{b} \sum_{i=1}^{b} w_i^{(k)} \ell_i^{(k)} + \frac{\lambda}{b} \sum_{i=1}^{b} d_i^{(k)}, \qquad \lambda = 1. \tag{11}$$

Here b is the path minibatch size and λ controls the reference constraint. All training paths contribute: quality weighting expresses performance preference without cross-path pairs, and the unweighted KL term constrains deviation from the geometry base. At inference, the four heads are combined uniformly:

$$z_{\text{guided}} = z_{\text{base}} + \frac{\kappa}{4} \sum_{k \in \{\text{official}, P, U, L\}} r_k(\widetilde{X}). \tag{12}$$

Each r_k is applied row-wise to $\widetilde{X}$; $\kappa \geq 0$ is the validation-selected guidance strength, distinct from the path γ. This changes the order in which DFS considers actions; it does not bypass geometric checks or update online from newly evaluated offspring.

4.5 Elite Reconstruction with Hard Backtracking

We retain the complete evaluated history, initialized with all labeled paths, and select an archive of at most ten unique elites from it. Selection combines official-score ranking with physical specialist nominations; specialists are restricted to the official top 20% of the history. Duplicate nominations are merged and remaining capacity is filled by official rank. This relates to quality-diversity archives (Mouret & Clune, 2015), using physical objectives rather than behavioral cells.

For each frozen parent, a slot considers up to the calibrated number of deterministic window attempts and stops at the first successful completion. Within an attempt, DFS must connect c_l to c_r using exactly W compressed actions. At each node, the hard predicate enumerates actions that stay in the grid, avoid obstacles and occupied prefix/suffix geometry, satisfy incremental clearance, do not enter the right anchor early, and preserve exact-arrival reachability. The diffusion policy scores and orders this hard-legal set, and DFS tries the next ranked sibling when a chosen branch reaches a dead end. Gap reconnection succeeds only when the Wth action terminates at c_r; candidate acceptance additionally requires suffix replay and full validation.

The action-waypoint suffix is replayed and a terminal connection is constructed before dense self-intersection, full-path clearance, endpoint, and hash checks. Candidates identical to the parent or any earlier candidate are rejected. Every accepted valid unique candidate is evaluated, and failed reconstructions never return a parent as an offspring. The reported pool contains 100 evaluated offspring without score-based post-selection. After all slots in the round finish, newly evaluated paths are added to the complete history and the Top-10 archive is recomputed synchronously.

5 EXPERIMENTS

5.1 BASELINES AND REPORTING

The comparison uses the same three scenes, nested label budgets, supervision source seed, five-repetition reporting protocol, and simulation evaluation procedure. **GDP** adapts graph-constrained diffusion (Shi et al., 2024) to complete voxel routes. **Pathformer-Predictor** adapts the time-series architecture of Chen et al. (2024) to score route sequences; **Voxel-Predictor** scores spatial representations. They rank pools of 10,000 and 5,000 feasible paths, respectively, with the latter a shared-pool prefix. **3D-CDM** adapts conditional 3D diffusion (Zheng et al., 2025) with supervision-derived elite prototypes and geometrically constrained path decoding. This name denotes our coil-routing adaptation, not the unmodified original model. The learning baselines use geometry-only pretraining followed by labeled adaptation. **Random-DFS**† combines n labeled paths with 100 newly evaluated random paths and reports the true-score Top-100, so its pool differs from the new-candidate pools of the other methods.

Table 2 reports within-run Best@100, Top-5@100, and Mean@100, aggregated as mean ± sample standard deviation across five repetitions. Appendix D gives per-run values for these three metrics. Appendix F visualizes the seed-0 best paths and reports their physical values.

Table 2: Baseline comparison: mean ± sample standard deviation over five repetitions; higher is better. Bold marks the highest unrounded mean among the five direct candidate methods. †Random-DFS reports true-score Top-100 from n labeled and 100 new random paths, not only new candidates; it is separated and excluded from bolding.

		Rhombus			Staggered			Y		
n	Method	Best@100	Top-5@100	Mean@100	Best@100	Top-5@100	Mean@100	Best@100	Top-5@100	Mean@100
100	Random-DFS†	−41.08 ±5.79	−45.23 ±2.76	−67.95 ±0.72	−32.08 ±5.13	−39.33 ±1.61	−71.47 ±0.82	−60.39 ±3.34	−65.93 ±1.49	−80.28 ±0.49
	GDP	−38.91 ±0.64	−39.82 ±0.94	−57.78 ±4.00	−38.87 ±3.65	−41.12 ±4.49	−76.73 ±5.06	−68.31 ±4.44	−72.16 ±2.34	−84.47 ±1.13
	Pathformer-Predictor	−37.77 ±3.09	−41.44 ±1.55	−60.19 ±1.92	−26.41 ±2.47	−30.05 ±1.44	−54.57 ±2.63	−53.24 ±4.33	−58.19 ±2.05	−78.41 ±1.43
	Voxel-Predictor	−36.91 ±3.75	−40.22 ±1.75	−60.77 ±2.67	−19.21 ±4.09	−23.87 ±3.28	−50.27 ±3.10	−49.05 ±2.90	−52.32 ±2.27	−74.10 ±1.75
	3D-CDM	**−36.53** **±2.95**	−39.57 ±2.04	−72.84 ±0.77	−26.13 ±15.09	−33.88 ±22.85	−74.03 ±7.65	−62.47 ±2.62	−66.93 ±1.43	−82.60 ±0.42
	Ours	−37.10 ±2.37	**−38.28** **±1.02**	**−52.49** **±1.86**	**−15.44** **±1.32**	**−16.69** **±2.03**	**−34.09** **±0.80**	**−47.66** **±3.86**	**−50.94** **±3.00**	**−66.43** **±1.07**
200	Random-DFS†	−38.39 ±3.77	−43.34 ±1.83	−63.25 ±0.94	−28.73 ±2.77	−34.79 ±1.68	−65.99 ±0.91	−61.63 ±2.56	−64.98 ±0.95	−78.08 ±0.32
	GDP	−45.07 ±3.65	−47.85 ±2.92	−68.91 ±1.69	−16.62 ±2.84	−19.75 ±2.16	−63.20 ±3.71	−67.79 ±5.59	−71.38 ±1.93	−83.21 ±1.53
	Pathformer-Predictor	−36.88 ±1.82	−39.88 ±1.10	−56.39 ±1.64	−19.20 ±1.51	−22.45 ±2.68	−47.83 ±1.79	−51.41 ±3.37	−56.37 ±1.60	−77.49 ±1.06
	Voxel-Predictor	−36.07 ±4.56	−39.39 ±2.94	−55.20 ±1.34	−18.86 ±3.82	−23.05 ±2.62	−45.62 ±1.14	−48.38 ±4.33	−52.50 ±2.10	−74.54 ±1.45
	3D-CDM	−33.22 ±6.88	−38.61 ±3.77	−71.71 ±1.82	−16.46 ±2.59	−20.89 ±2.96	−68.75 ±4.39	−60.22 ±4.32	−64.20 ±2.66	−82.19 ±0.70
	Ours	**−30.40** **±7.12**	**−31.40** **±6.51**	**−45.02** **±2.37**	**−13.98** **±1.74**	**−15.61** **±1.68**	**−32.10** **±1.31**	**−46.38** **±3.31**	**−48.45** **±2.22**	**−65.01** **±0.84**
300	Random-DFS†	−38.68 ±0.00	−43.80 ±1.04	−60.81 ±0.59	−15.45 ±0.00	−30.00 ±1.81	−60.81 ±0.50	−60.70 ±6.69	−63.90 ±1.46	−75.68 ±0.28
	GDP	−39.07 ±1.42	−40.87 ±1.99	−61.02 ±1.96	−17.14 ±3.36	−21.42 ±2.84	−55.70 ±3.99	−56.39 ±5.62	−62.46 ±3.51	−83.45 ±1.91
	Pathformer-Predictor	−36.47 ±1.15	−38.84 ±1.07	−53.41 ±1.61	−19.58 ±3.01	−22.55 ±2.49	−47.52 ±2.36	**−46.86** **±7.72**	−53.83 ±3.15	−76.05 ±0.95
	Voxel-Predictor	−36.97 ±2.50	−39.26 ±1.63	−53.45 ±0.79	−18.86 ±3.82	−23.02 ±2.60	−44.65 ±1.10	−49.44 ±2.40	−52.91 ±2.00	−73.65 ±2.26
	3D-CDM	−33.36 ±4.88	−36.67 ±2.65	−67.53 ±0.84	−16.91 ±2.67	−20.89 ±3.40	−69.11 ±7.48	−51.17 ±6.47	−62.96 ±2.09	−82.68 ±0.47
	Ours	**−17.08** **±8.95**	**−20.58** **±9.43**	**−44.65** **±2.00**	**−11.18** **±0.68**	**−12.04** **±0.81**	**−40.76** **±3.10**	−48.45 ±2.15	**−49.28** **±2.25**	**−57.51** **±2.41**

Results. Ours has the highest average Top-5@100 and Mean@100 in all nine scene–budget combinations among the five direct candidate methods, and the highest average Best@100 in seven. The

Best@100 exceptions are Rhombus at $n = 100$ (-37.10 versus -36.53 for 3D-CDM) and Y at $n = 300$ (-48.45 versus -46.86 for Pathformer-Predictor). The gains thus concern the quality of several candidate initializers more consistently than the single best outcome. These are descriptive comparisons, not statistical-significance or downstream-refinement claims; Random-DFS retains its post-selection interpretation.

5.2 Ablation Study

Table 3 reports Staggered at $n = 100$, with five repetitions and 100 newly evaluated candidates per run. Full is the corresponding Ours result in Table 2. *Random Reconstruction* replaces the learned reconstruction policy with a randomly initialized, frozen model, without geometry pretraining or residual guidance, while retaining performance-driven elite selection and geometry-constrained backtracking. *Sequence-only* and *Spatial-only* train and infer with one view. *w/o Elite Search* removes both initial elite selection and subsequent parent-pool updates, selecting parents uniformly from the original supervision pool. All variants retain hard DFS backtracking; Appendix E records their training and search settings.

Table 3: Ablation study on Staggered, $n = 100$: mean $\pm$ sample standard deviation over five repetitions. Each run evaluates 100 new candidates; higher is better. Bold marks the highest unrounded mean in each column.

Configuration	Best@100	Top-5@100	Mean@100
Full	$\mathbf{-15.44 \pm 1.32}$	$\mathbf{-16.69 \pm 2.03}$	$\mathbf{-34.09 \pm 0.80}$
Random Reconstruction	-21.51 ± 1.07	-22.59 ± 0.93	-41.66 ± 0.65
Sequence-only	-17.35 ± 3.60	-19.39 ± 3.11	-37.73 ± 3.16
Spatial-only	-16.07 ± 2.05	-17.21 ± 2.24	-37.34 ± 2.03
w/o Elite Search	-34.36 ± 5.72	-45.73 ± 4.16	-90.20 ± 3.38

Full has higher average Best@100, Top-5@100, and Mean@100 than the four reported variants. Jointly removing pretraining and guidance yields -21.51, -22.59, and -41.66, respectively, compared with Full's -15.44, -16.69, and -34.09. The lowest scores occur without elite search (-34.36 Best@100, -45.73 Top-5@100, and -90.20 Mean@100).

6 Conclusion

We formulated simulation-budgeted generative coil initialization as a preceding stage of electromagnetic inverse design, addressing the need for useful starting structures in workflows that commonly rely on manual initialization. Our framework constructs feasible candidate routes through performance-preference-guided elite path reconstruction. Dual-view masked-action diffusion models feasible path context, quality-weighted residuals guide local actions under a frozen-base constraint, and hard backtracking reconnects fixed anchors under exact geometric predicates. An elite archive turns selected evaluated candidates into parents for subsequent rounds without online model updates or candidate score post-selection. Across three scenes and three label budgets, the method leads the four learning baselines in average Top-5@100 and Mean@100 in all nine combinations and average Best@100 in seven. Downstream refinement is a direction for future evaluation.

AI Use Statement

AI tools assisted with manuscript writing, language polishing, and adjustments to the article's structure and layout.

References

Philipp Amrein, Feng Jia, Maxim Zaitsev, and Sebastian Littin. CoilGen: Open-source MR coil layout generator. *Magnetic Resonance in Medicine*, 88(3):1465–1479, 2022. doi: 10.1002/mrm.29294.

Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/958c530554f78bcd8e97125b70e6973d-Abstract.html.

Peng Chen, Yingying Zhang, Yunyao Cheng, Yang Shu, Yihang Wang, Qingsong Wen, Bin Yang, and Chenjuan Guo. Pathformer: Multi-scale transformers with adaptive pathways for time series forecasting. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=lJkOCMP2aW.

Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin CM Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, 2023. doi: 10.15607/RSS.2023.XIX.026.

Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pp. 1724–1734. Association for Computational Linguistics, 2014. doi: 10.3115/v1/D14-1179. URL https://aclanthology.org/D14-1179/.

Grant A. Covic and John T. Boys. Inductive power transfer. *Proceedings of the IEEE*, 101(6): 1276–1289, 2013. doi: 10.1109/JPROC.2013.2244536.

Andrew Curtis, Md Shain Shahid Chowdhury Oni, Shuntaro Inoue, Abhilash Kamineni, Regan Zane, and Nicholas Flann. Generative neural network approach to designing dynamic inductive power transfer systems. In *2023 IEEE Wireless Power Technology Conference and Expo (WPTCE)*, pp. 1–5, 2023. doi: 10.1109/WPTCE56855.2023.10216141.

Giulio Giovannetti, Nunzia Fontana, Alessandra Flori, Maria Filomena Santarelli, Mauro Tucci, Vincenzo Positano, Sami Barmada, and Francesca Frijia. Machine learning for the design and the simulation of radiofrequency magnetic resonance coils: Literature review, challenges, and perspectives. *Sensors*, 24(6):1954, 2024. doi: 10.3390/s24061954.

Rob Goldstein, William Stuehr, and Micah Black. Design and fabrication of inductors for induction heat treating. In Valery Rudnev and George E. Totten (eds.), *ASM Handbook, Volume 4C: Induction Heating and Heat Treatment*, pp. 589–606. ASM International, 2014. doi: 10.31399/asm.hb.v04c.a0005839.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Jaeyub Hyun, Darun Barazanchy, Jaspreet Pandher, Michel van Tooren, and H. Alicia Kim. Gradient-based shape optimization of induction welding coil for thermoplastic composites via level-set method. *Composites Part B: Engineering*, 306:112824, 2025. doi: 10.1016/j.compositesb.2025.112824.

Alan A. Kaptanoglu, Gabriel P. Langlois, and Matt Landreman. Topology optimization for inverse magnetostatics as sparse regression: Application to electromagnetic coils for stellarators. *Computer Methods in Applied Mechanics and Engineering*, 418:116504, 2024. doi: 10.1016/j.cma.2023.116504.

Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7):846–894, 2011. doi: 10.1177/0278364911406761.

Abderraouf Lalla, Paolo Di Barba, Sławomir Hausman, and Maria Evelina Mognaschi. Deep neural network-based design of planar coils for proximity sensing applications. *Sensors*, 25(14):4429, 2025. doi: 10.3390/s25144429.

Oscar Lucía, Pascal Maussion, Enrique J. Dede, and José M. Burdío. Induction heating technology and its applications: Past developments, current technology, and future challenges. *IEEE Transactions on Industrial Electronics*, 61(5):2509–2520, 2014. doi: 10.1109/TIE.2013.2281162.

Kaisa Miettinen. *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, Boston, 1999. doi: 10.1007/978-1-4615-5563-6.

Sean Molesky, Zin Lin, Alexander Y. Piggott, Weiliang Jin, Jelena Vucković, and Alejandro W. Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12:659–670, 2018. doi: 10.1038/s41566-018-0246-9.

Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites, 2015. URL https://arxiv.org/abs/1504.04909. Preprint.

Gerardus N. Peeren. Stream function approach for determining optimal surface currents. *Journal of Computational Physics*, 191(1):305–321, 2003. doi: 10.1016/S0021-9991(03)00320-6.

Lyle Regenwetter, Amin Heyrani Nobari, and Faez Ahmed. Deep generative models in engineering design: A review. *Journal of Mechanical Design*, 144(7):071704, 2022. doi: 10.1115/1.4053859.

Dingyuan Shi, Yongxin Tong, Zimu Zhou, Ke Xu, Zheng Wang, and Jieping Ye. Graph-constrained diffusion for end-to-end path planning. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=vuK8MhVtuu.

Toru Takahashi. Designing gradient coils with the shape derivative and the closed b-spline curves. *Magnetic Resonance Imaging*, 110:112–127, 2024. doi: 10.1016/j.mri.2024.03.042.

Brandon Trabucco, Xinyang Geng, Aviral Kumar, and Sergey Levine. Design-bench: Benchmarks for data-driven offline model-based optimization. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 21658–21676, 2022. URL https://proceedings.mlr.press/v162/trabucco22a.html.

Robert Turner. A target field approach to optimal coil design. *Journal of Physics D: Applied Physics*, 19(8):L147–L151, 1986. doi: 10.1088/0022-3727/19/8/001.

Robert Turner. Gradient coil design: A review of methods. *Magnetic Resonance Imaging*, 11(7): 903–920, 1993. doi: 10.1016/0730-725X(93)90209-V.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Xiaoyang Zheng, Junichiro Shiomi, and Takayuki Yamada. Optimizing metamaterial inverse design with 3D conditional diffusion model and data augmentation. *Advanced Materials Technologies*, 10(14):2500293, 2025. doi: 10.1002/admt.202500293.

Caoxiang Zhu, Stuart R. Hudson, Yuntao Song, and Yuanxi Wan. New method to design stellarator coils without the winding surface. *Nuclear Fusion*, 58(1):016008, 2018. doi: 10.1088/1741-4326/aa8e0a.

A Architecture and Training Details

The geometry catalog is split by path hash into 8,000 training and 2,000 validation paths. The default window-selection procedure compares $W \in \{4, 8, 12\}$ on fixed 1,000/200 subsets for five epochs before formal training. The scan first maximizes the valid, unique, non-parent completion rate. Rates within one percentage point are treated as tied; ties are broken by lower masked NLL, a lower median number of expanded nodes, and smaller W, in that order. For the selected W, validation compares window attempts $\{1, 2, 4, 8\}$ and per-attempt node limits $\{250, 500, 1000, 2000, 5000\}$. We select the setting with the smallest product of the number of attempts and the per-attempt node limit that retains at least 95% of the best success rate; ties are broken by shorter wall-clock time. The final geometry base is trained from scratch for 20 epochs with AdamW (weight decay 0.01), constant learning rate 10^{-3}, path batch size 16, and horizon four. The cosine schedule governs noise retention, $\bar{\alpha}_t = \cos^2(\pi t/16)$ for $t = 0, \ldots, 8$, not the learning rate. Training timesteps in $\{1, \ldots, 8\}$ are drawn by a stable hash of the path, gap, progress, epoch, and seed. Its checkpoint is selected by the lowest fixed-validation masked NLL; an exact tie retains the earlier epoch. Validation uses one fixed window per path with fixed timestep and mask streams.

The default mask independently corrupts valid positions and resamples empty masks; $\omega_q = 1/|M_t|$ in Equation 10. The residual CE and KL terms use the same masked-position reduction as the base objective. Every window contributes all W teacher-forced prefix conditions; a path minibatch of 16 thus contains $16W$ conditions. Unmasked clean action tokens may be visible during denoising training; this does not expose removed gap geometry in S_j or Q_j.

The spatial tensor has 14 channels and the ordered sequence has 50 rows with 19 features per row. The spatial and sequence encoders produce 64 and 96 dimensions, respectively; their fusion produces $g_j \in \mathbb{R}^{128}$. Each action state has $128 + 32 + 16 + 32 = 208$ dimensions. The action mixer uses four attention heads, zero dropout, and a $208 \to 416 \to 208$ GELU feed-forward block. GELU denotes the Gaussian error linear unit activation. The base and all residual heads consume the mixed 208-dimensional state, but residual training freezes every base parameter.

Labeled paths are split by hash into 80% training and 20% validation. Each residual head uses AdamW (weight decay 0.01) with learning rate 3×10^{-4}, path batch size 16, at most 30 epochs, patience five, and an epoch-zero fallback representing zero residual. For $\kappa \in \{0, 0.25, 0.5, 1, 2, 4\}$, the fused model is evaluated with each of the four target-specific quality weights and the same negative-log-likelihood (NLL)/KL objective. The lowest mean validation total loss selects κ; ties select the smaller value, without using offspring evaluations. Head checkpoints likewise require strict validation-loss improvement over the saved epoch-zero baseline; fallback heads remain in the four-head fusion denominator.

The action vocabulary contains every signed-axis run length supported by the scene grid (Table 4). Its size therefore changes with the scene, while the encoder and 208-dimensional action-state contracts remain fixed.

Table 4: Scene-specific compressed-action support. Each range applies to both signs of its axis.

Scene	x steps	y steps	z steps	$\lvert\mathcal{A}\rvert$
Rhombus	1–10	1–9	1–5	48
Staggered	1–12	1–10	1–5	54
Y	1–8	1–8	1–5	42

A.1 Run Records

Each run records its configuration, input and checkpoint hashes, random seeds, evaluation outputs, and archive trajectory. The five repetition indices 0–4 correspond to training seeds 0–4 and search seeds 1200–1204; they share the same labeled pool.

A.2 Diagram and Architecture

Figure 3 summarizes training and elite reconstruction. Geometry pretraining uses clean actions as self-supervised targets without physical labels. Ordered waypoint features and spatial occupancy are encoded by the GRU and CNN, then concatenated and fused by an MLP to condition the action-token attention/FFN mixer. The preference module comprises four independent residual heads trained with quality-weighted denoising and reference KL. The search tree represents legally

filtered action proposals and sibling backtracking; complete-path validation precedes physical evaluation.

B Quality-Weighted Residual Training

For target k, let $y_i^{(k)}$ be the official score or the corresponding physical margin. Thus $y_i^{(\text{official})} = f(\gamma_i;\xi)$ and $y_i^{(k)} = m_k(\gamma_i;\xi)$ for $k \in \{P,U,L\}$, where γ_i is labeled path i. Only training labels determine the mean μ_k, population standard deviation σ_k, and weight normalizer:

$$\begin{aligned} u_i^{(k)} &= \frac{y_i^{(k)} - \mu_k}{\max(\sigma_k, 10^{-6})}, \quad \widetilde{w}_i^{(k)} = \operatorname{clip}(\exp(u_i^{(k)}/\tau), 0.5, 2), \\ w_i^{(k)} &= \frac{\widetilde{w}_i^{(k)}}{\operatorname{mean}_{\text{train}} \widetilde{w}^{(k)}}, \qquad \tau = 1. \end{aligned} \tag{13}$$

The implementation clips in log space before exponentiation. Here $u_i^{(k)}$ is standardized quality, τ is the quality-weight temperature, $\widetilde{w}_i^{(k)}$ is the clipped unnormalized weight, and $\operatorname{mean}_{\text{train}}$ averages it over training paths. The operator $\operatorname{clip}(x,a,b) = \min(b, \max(a,x))$ bounds a scalar. The bounds apply before mean normalization, not to the final weights. Validation reuses the training statistics and normalizer; all training paths, including the bottom quartile, receive positive weight. No preference pairs or additional elite-only denoising term are used.

For a fixed prefix j, noisy horizon $\mathbf{b}^t$, timestep t, and masked position q, define $p_B = \operatorname{softmax}(z_{\text{base},q})$ and $p_k = \operatorname{softmax}(z_{\text{base},q} + r_k(\widetilde{X}_q))$. Here subscript q selects a row; $p_B(a)$ and $p_k(a)$ are the corresponding categorical probabilities for action a, with the shared inputs suppressed for readability. The reference term is

$$\mathrm{KL}(p_B \| p_k) = \sum_{a \in \mathcal{A}} p_B(a)\big[\log p_B(a) - \log p_k(a)\big]. \tag{14}$$

The base is detached, and probabilities are normalized over the full action vocabulary. Cross-entropy (CE) and KL first use the masked-position reduction from Equation 10, then average the W prefix conditions within each path. Equation 11 quality-weights only the NLL term; KL is not multiplied by $w_i^{(k)}$.

Four heads are trained independently, with only the current head receiving gradients. At each epoch, a shared deterministic draw $u \in [0,1)$ selects index $\lfloor uN_i \rfloor$ in each path's list of N_i legal window starts. Heads reuse the same window, timestep, and mask streams; validation fixes these streams across epochs. Frozen geometry states may be cached by path hash and window endpoints. No newly evaluated search candidate is used to update model parameters or recalibrate the quality weights.

C Archive and Exact-Window Search

The archive selects at most ten unique parents from the complete evaluated history, including all n labeled paths. With q_o official nominations and q_s nominations per physical specialist, $q_o + 3q_s = 10$. A specialist must belong to the first $\lceil 0.2|\mathcal{H}| \rceil$ paths in the official ranking of history $\mathcal{H}$. Official ranking uses descending score then ascending path hash; specialist ranking uses descending margin, descending official score, then ascending hash. Normalized waypoint hashes remove duplicates, and official rank fills unused capacity. Each round freezes the archive and allocates at most one slot per selected unique parent. Failed reconstructions never return the parent as an offspring; the comparison reports 100 successfully generated and evaluated candidates per run. All successful candidates in a round are evaluated before synchronously rebuilding the archive from the updated history.

Algorithm 1 gives the local search semantics. The action-order query performs eight denoising steps, applies Equation 12, and ranks legal actions by $z^{t=1}_{\text{guided},0}(a) + T_{\text{order}} g_a$. The first term is the final guided logit of action a at horizon position zero; g_a are independent standard Gumbel draws from a deterministic request seed. The nonnegative order temperature T_{order} scales ranking noise and is distinct from the quality-weight temperature τ. The hard predicate includes fixed-suffix occupancy, obstacle and incremental clearance checks, and exact-arrival reachability; it forms the

legal set before policy ranking. REPLAYANDCONNECT replays the fixed waypoint suffix and constructs a terminal connection by Manhattan axis permutations; FULLVALIDATE checks the resulting complete path. The fixed context and duplicate hash refer to normalized action waypoints, not an unchanged terminal connector. The completion-acceptance controller records attempt termination when a rejected completion ends the current attempt; this is denoted by ATTEMPTSTOPPED in the pseudocode. In Algorithm 1, Π is the frozen base-plus-head policy with fixed κ, P is the current waypoint prefix, $\mathcal{S}$ is the DFS stack, and $\perp$ denotes failure. The suffix $\mathbf{c}_{r:M} = (c_r, \ldots, c_M)$ is fixed; $P \,\|\, a$ appends the endpoint reached by action a, and $\mathrm{last}(P)$ returns its last waypoint. NOVELHASH checks against the evaluated history and candidates already accepted during the current round; the search budget bounds expanded nodes in this attempt.

Algorithm 1 Preference-guided exact-W gap completion

Require: parent waypoints $\mathbf{c}$; window indices l, r with $W = r - l$; policy Π; search budget
Ensure: a valid novel reconstruction or $\perp$

$P \leftarrow (c_0, \ldots, c_l)$
$\mathcal{S} \leftarrow [(P, 0)]$ ▷ DFS stack stores reconstructed prefix and depth
while $\mathcal{S} \neq \emptyset$ and budget remains **do**
  $(P, j) \leftarrow \textsc{Pop}(\mathcal{S})$
  if $j = W$ **then**
    if $\mathrm{last}(P) = c_r$ **then**
      $\gamma' \leftarrow \textsc{ReplayAndConnect}(\mathbf{c}, P, [l, r])$
      if $\textsc{FullValidate}(\gamma')$ and $\textsc{NovelHash}(\gamma')$ **then**
        return γ'
      end if
      if ATTEMPTSTOPPED **then**
        return $\perp$
      end if
    end if
    continue
  end if
  $(S_j, Q_j) \leftarrow \textsc{LeakageSafeCondition}(P, \mathbf{c}_{r:M})$
  $\mathcal{A}_{\text{legal}} \leftarrow \textsc{LegalActions}(P, \mathbf{c}_{r:M}, c_r, W - j - 1)$
  $\mathcal{A}_{\text{ranked}} \leftarrow \textsc{GuidedActionOrder}(\Pi, S_j, Q_j, j, \mathcal{A}_{\text{legal}})$
  for all $a \in \mathcal{A}_{\text{ranked}}$ in reverse order **do**
    $\textsc{Push}(\mathcal{S}, (P \,\|\, a, j + 1))$
  end for
end while
return $\perp$ ▷ all ranked siblings failed: backtrack/attempt failure

D EXTENDED RESULTS

The raw tables expose all 270 scene–budget–method–repetition cells. Each seed cell has two lines: Best@100/Top-5@100; Mean@100. Repetitions 0–4 use the training/search seed mapping in Appendix A. All scores are computed before rounding.

Table 5: Per-run results for Rhombus. Cell lines: Best@100/Top-5@100; Mean@100. †Random-DFS uses union Top-100.

n	Method	seed 0	seed 1	seed 2	seed 3	seed 4
100	Random-DFS†	-46.38/-48.38 -68.11	-46.17/-47.54 -68.83	-32.26/-41.69 -67.16	-39.46/-43.61 -68.39	-41.11/-44.95 -67.26
	GDP	-38.69/-39.86 -56.73	-39.61/-40.28 -57.19	-39.55/-40.60 -62.49	-38.57/-40.14 -60.47	-38.15/-38.21 -52.03
	Pathformer-Predictor	-42.38/-43.57 -57.00	-37.01/-42.08 -60.06	-38.49/-41.67 -61.37	-37.12/-40.10 -61.96	-33.84/-39.77 -60.54
	Voxel-Predictor	-34.73/-39.50 -59.03	-38.52/-40.90 -63.48	-31.85/-37.72 -60.91	-37.82/-40.51 -57.25	-41.64/-42.44 -63.18
	3D-CDM	-32.67/-36.65 -72.44	-39.86/-41.75 -73.49	-34.41/-38.51 -71.79	-37.28/-41.07 -72.79	-38.45/-39.88 -73.66
	Ours	-33.11/-36.77 -51.34	-38.98/-39.06 -50.08	-36.82/-37.70 -52.39	-38.47/-39.14 -54.24	-38.14/-38.72 -54.38

Table 5: Per-run results for Rhombus (continued).

n	Method	seed 0	seed 1	seed 2	seed 3	seed 4
200	Random-DFS†	-31.75/-40.31 -61.96	-39.23/-43.20 -63.07	-40.00/-43.75 -63.19	-39.92/-45.07 -63.44	-41.05/-44.37 -64.58
	GDP	-44.25/-48.47 -67.81	-45.00/-47.76 -70.58	-46.03/-48.26 -66.59	-50.12/-51.46 -70.28	-39.95/-43.32 -69.31
	Pathformer-Predictor	-38.66/-41.71 -56.67	-37.01/-39.78 -54.48	-37.77/-39.69 -54.94	-37.12/-39.48 -57.65	-33.84/-38.73 -58.21
	Voxel-Predictor	-34.73/-39.95 -54.57	-39.93/-41.57 -57.01	-28.71/-34.26 -56.22	-37.82/-40.19 -53.98	-39.13/-40.99 -54.22
	3D-CDM	-39.16/-42.28 -74.22	-37.92/-39.87 -69.07	-24.51/-36.76 -71.60	-37.45/-41.18 -71.94	-27.05/-32.97 -71.70
	Ours	-25.27/-28.51 -48.45	-38.45/-38.67 -46.06	-24.51/-24.94 -42.65	-37.93/-38.12 -44.89	-25.87/-26.77 -43.02
300	Random-DFS†	-38.68/-44.10 -60.19	-38.68/-44.54 -61.48	-38.68/-43.71 -60.71	-38.68/-44.58 -61.35	-38.68/-42.06 -60.30
	GDP	-38.69/-39.43 -60.81	-41.59/-44.24 -62.16	-38.29/-39.35 -58.47	-38.32/-40.84 -63.59	-38.44/-40.49 -60.09
	Pathformer-Predictor	-36.69/-40.16 -55.60	-37.01/-39.50 -54.13	-37.09/-38.71 -51.30	-37.12/-38.45 -53.37	-34.44/-37.35 -52.65
	Voxel-Predictor	-34.73/-39.89 -53.07	-38.09/-38.86 -54.06	-34.35/-36.81 -52.63	-37.29/-39.48 -52.99	-40.39/-41.28 -54.50
	3D-CDM	-37.23/-38.30 -66.43	-30.77/-36.26 -68.08	-26.05/-32.24 -68.33	-35.31/-37.76 -67.99	-37.43/-38.78 -66.83
	Ours	-31.09/-35.49 -46.37	-8.16/-13.36 -46.20	-11.36/-13.60 -42.07	-15.03/-16.19 -42.94	-19.78/-24.26 -45.68

Table 6: Per-run results for Staggered. Cell lines: Best@100/Top-5@100; Mean@100. †Random-DFS uses union Top-100.

n	Method	seed 0	seed 1	seed 2	seed 3	seed 4
100	Random-DFS†	-26.36/-38.49 -72.71	-27.78/-38.51 -70.52	-36.20/-39.76 -71.08	-31.90/-37.92 -71.32	-38.16/-41.95 -71.71
	GDP	-45.33/-47.96 -84.10	-36.63/-37.37 -79.53	-36.98/-37.84 -72.95	-37.41/-39.05 -71.85	-37.99/-43.35 -75.20
	Pathformer-Predictor	-27.63/-30.29 -53.16	-26.40/-31.38 -53.32	-22.20/-27.64 -55.54	-28.47/-30.04 -52.10	-27.36/-30.89 -58.72
	Voxel-Predictor	-17.60/-22.77 -50.95	-22.40/-26.53 -49.37	-14.01/-21.67 -49.28	-24.23/-28.04 -55.09	-17.80/-20.33 -46.67
	3D-CDM	-16.25/-22.95 -68.27	-23.07/-24.59 -68.43	-19.13/-21.78 -75.15	-52.77/-74.68 -86.76	-19.41/-25.40 -71.56
	Ours	-17.70/-20.09 -33.74	-14.57/-14.90 -33.24	-15.45/-16.93 -35.31	-14.46/-15.63 -33.74	-15.01/-15.92 -34.41
200	Random-DFS†	-23.76/-33.13 -66.45	-29.99/-36.01 -66.66	-29.99/-36.01 -64.79	-29.91/-32.77 -65.26	-29.99/-36.01 -66.81
	GDP	-17.94/-19.04 -63.03	-20.27/-23.49 -64.88	-15.62/-19.27 -67.03	-16.68/-17.92 -57.13	-12.59/-19.02 -63.93
	Pathformer-Predictor	-17.60/-19.46 -46.22	-19.28/-26.45 -48.20	-21.03/-22.28 -48.64	-20.31/-23.33 -50.20	-17.80/-20.72 -45.88
	Voxel-Predictor	-17.60/-22.18 -46.08	-20.68/-25.10 -44.83	-14.01/-21.19 -45.24	-24.23/-26.45 -47.39	-17.80/-20.30 -44.57
	3D-CDM	-13.33/-17.70 -69.29	-17.97/-19.80 -61.37	-17.25/-20.24 -69.00	-14.26/-25.72 -71.71	-19.50/-20.97 -72.40
	Ours	-14.19/-14.49 -32.17	-12.38/-13.72 -33.25	-16.70/-17.59 -31.14	-12.56/-17.14 -33.48	-14.06/-15.12 -30.46
300	Random-DFS†	-15.45/-27.69 -61.14	-15.45/-28.39 -59.97	-15.45/-31.41 -60.75	-15.45/-31.41 -61.19	-15.45/-31.12 -60.98
	GDP	-12.52/-17.44 -51.34	-17.41/-24.33 -58.76	-15.34/-19.70 -51.31	-19.35/-23.54 -58.50	-21.10/-22.12 -58.58
	Pathformer-Predictor	-17.60/-19.46 -47.08	-17.24/-22.89 -44.95	-21.03/-23.28 -50.15	-24.23/-26.06 -49.77	-17.80/-21.04 -45.66
	Voxel-Predictor	-17.60/-22.18 -46.08	-20.68/-24.96 -44.01	-14.01/-21.19 -44.87	-24.23/-26.45 -45.10	-17.80/-20.30 -43.20
	3D-CDM	-15.41/-20.57 -76.19	-15.34/-18.41 -61.07	-16.07/-20.36 -65.75	-16.09/-18.44 -64.63	-21.64/-26.69 -77.92
	Ours	-11.46/-13.06 -41.93	-10.26/-11.15 -35.84	-12.10/-12.65 -44.14	-11.17/-11.93 -40.15	-10.92/-11.42 -41.75

Table 7: Per-run results for Y. Cell lines: Best@100/Top-5@100; Mean@100. †Random-DFS uses union Top-100.

n	Method	seed 0	seed 1	seed 2	seed 3	seed 4
100	Random-DFS†	-65.28/-68.22 -79.60	-56.12/-64.32 -80.20	-59.04/-65.52 -80.81	-61.14/-65.17 -80.09	-60.37/-66.41 -80.71
	GDP	-72.71/-75.11 -84.40	-64.12/-68.78 -82.89	-71.52/-72.68 -85.18	-63.00/-71.23 -84.02	-70.19/-73.00 -85.86
	Pathformer-Predictor	-47.60/-55.69 -77.12	-51.59/-58.29 -80.05	-54.61/-56.69 -79.74	-52.94/-59.59 -78.15	-59.44/-60.70 -77.00
	Voxel-Predictor	-49.73/-54.17 -76.90	-52.25/-54.91 -73.73	-45.46/-49.87 -74.05	-46.67/-50.22 -72.09	-51.15/-52.44 -73.72
	3D-CDM	-65.26/-68.97 -82.47	-59.95/-64.96 -82.45	-62.99/-67.15 -82.70	-64.61/-66.76 -83.25	-59.55/-66.80 -82.11
	Ours	-45.43/-49.06 -66.84	-43.03/-46.88 -65.66	-53.03/-54.56 -66.34	-49.64/-52.32 -65.29	-47.16/-51.89 -68.02
200	Random-DFS†	-64.13/-66.22 -78.01	-60.51/-65.15 -78.15	-64.13/-65.48 -78.15	-58.12/-64.01 -77.61	-61.25/-64.05 -78.50
	GDP	-67.82/-70.89 -85.02	-58.09/-68.44 -81.54	-70.73/-71.41 -84.48	-71.30/-73.41 -83.09	-70.98/-72.74 -81.93
	Pathformer-Predictor	-56.05/-58.63 -77.10	-51.69/-56.41 -78.83	-47.36/-54.43 -76.27	-49.09/-55.40 -76.90	-52.86/-56.99 -78.35
	Voxel-Predictor	-49.73/-52.26 -75.68	-52.25/-55.23 -76.45	-47.36/-52.22 -73.94	-41.37/-49.44 -73.47	-51.19/-53.33 -73.16
	3D-CDM	-62.94/-64.49 -81.37	-59.05/-63.81 -81.67	-64.10/-66.16 -82.10	-61.75/-66.65 -82.88	-53.25/-59.92 -82.94
	Ours	-45.14/-46.66 -65.37	-45.47/-48.98 -64.52	-42.07/-45.70 -66.04	-48.61/-50.01 -65.25	-50.61/-50.93 -63.85
300	Random-DFS†	-63.69/-64.62 -75.49	-63.69/-64.62 -75.95	-63.69/-64.33 -75.30	-63.69/-64.62 -75.92	-48.72/-61.30 -75.72
	GDP	-56.46/-65.68 -84.50	-55.22/-60.15 -81.63	-58.63/-63.00 -85.67	-63.55/-65.76 -84.18	-48.10/-57.72 -81.27
	Pathformer-Predictor	-49.97/-56.28 -77.11	-36.47/-51.22 -75.26	-55.32/-57.47 -76.72	-41.37/-50.15 -74.90	-51.19/-54.00 -76.26
	Voxel-Predictor	-49.73/-52.26 -77.00	-52.25/-53.87 -73.34	-47.36/-51.64 -70.64	-46.67/-50.89 -73.76	-51.19/-55.90 -73.51
	3D-CDM	-47.72/-64.18 -83.25	-61.47/-65.39 -82.61	-50.99/-63.47 -81.98	-44.15/-61.58 -82.95	-51.49/-60.15 -82.60
	Ours	-51.86/-52.75 -61.52	-46.52/-47.46 -55.72	-46.96/-47.48 -55.80	-47.76/-48.43 -56.54	-49.14/-50.27 -57.95

Random-DFS reporting pool. The historical Random-DFS comparison first takes the n supervision paths and the first 100 new random paths that do not duplicate supervision or each other. It ranks their union by true score and reports the first 100; ties use source and sample ID. Shared supervision can yield identical best scores across repetitions, explaining zero Best@100 standard deviations in some cells. This outcome-based post-selection is an offline reference, not the unfiltered distribution of 100 random new candidates. It is retained in the main comparison with an explicit marker, rather than treated as a like-for-like generated pool.

The frozen artifacts retain sample/evaluation correspondence, source membership, generation failures, and per-round archive states. The five repetitions share supervision rather than representing five independent labeled datasets. Neither equal evaluation counts nor the score tables establish equal computational cost. No statistical significance or downstream optimization improvement is inferred.

E Ablation Settings and Results

The Full row uses the Staggered $n = 100$ runs already reported in the main comparison, with training seeds 0–4 and search seeds 1200–1204. The four reported variants use the same scene and supervision source seed 0, with the same pairing of repetition indices. Their 100 labeled paths are split 80/20 by path hash; geometry training, where enabled, uses the 8,000/2,000 split of 10,000 feasible paths. Each reported pool contains 100 valid unique non-parent candidates, without mixing in supervision paths or true-score post-selection.

Variant definitions.

- **Full**: the dual-view, quality-weighted residual method with elite reconstruction and hard DFS, using the main-comparison checkpoints and settings.
- **Random Reconstruction**: use a randomly initialized, frozen base without geometry training, bypass residual heads, and set $\kappa = 0$. Randomness refers to model initialization. Dual

views, the action mixer, performance-based elite selection, and dynamic archive updates are retained.

- **Sequence-only**: disable the spatial branch in training and inference, replacing its output with zeros; train the single-view base and residual heads from scratch.
- **Spatial-only**: analogously disable the sequence branch throughout training and inference. Both single-view variants preserve the 128-dimensional geometry state and 208-dimensional action state.
- **w/o Elite Search**: select parents from the original 100 labeled paths using stable seed/round/path-hash sampling, without score-based initial selection or subsequent offspring entering the parent pool. An archive may be recorded but does not determine future parents.

Random Reconstruction uses a frozen random base for each seed. The single-view variants train their respective models from scratch; w/o Elite Search uses a separately trained dual-view base and residual heads. All reported variants retain backtracking.

Fixed settings for the four variants. All four use $W = 4$, at most two window attempts per slot, and 250 expanded nodes per attempt, without window or DFS-budget scans. Windows are uniform; token corruption is Bernoulli with PAD excluded. Training uses uniform masked cross-entropy without position-zero or timestep weighting, an action horizon of four, and eight diffusion steps. The action-order temperature is 1 and auxiliary decoding is stochastic. Where active, the residual heads are uniformly fused with $\kappa = 1$; Random Reconstruction uses $\kappa = 0$. Dynamic archives use four official and two nominations per physical specialist, with specialist eligibility restricted to the official top 20%. Each round selects at most ten unique parents, one slot each, and updates only after its candidates are evaluated. w/o Elite Search substitutes its frozen uniform parent pool for this selection.

For the two single-view variants and w/o Elite Search, base training uses AdamW for 20 epochs, learning rate 10^{-3}, and path batch 16, selecting the lowest fixed-validation masked NLL. Their residual training freezes the base and uses AdamW at 3×10^{-4}, path batch 16, at most 30 epochs, patience five, and epoch-zero fallback. Quality-weight temperature and reference KL weight are both 1; the loss is quality-weighted denoising plus $\mathrm{KL}(p_B \| p_k)$, not a pairwise preference objective. Quality statistics are fitted on training labels and reused for validation. Checkpoints are selected on validation loss. No search offspring updates model parameters.

Reporting. Table 8 reports Best@100, Top-5@100, and Mean@100 for the 25 displayed runs. Each run uses a reported pool of 100 candidates; the main table aggregates each metric as mean and sample standard deviation across five repetitions.

Table 8: Per-run ablation results for Staggered, $n = 100$. Cell lines: Best@100/Top-5@100; Mean@100. Higher is better.

Configuration	seed 0	seed 1	seed 2	seed 3	seed 4
Full	-17.70/-20.09 -33.74	-14.57/-14.90 -33.24	-15.45/-16.93 -35.31	-14.46/-15.63 -33.74	-15.01/-15.92 -34.41
Random Reconstruction	-21.91/-23.39 -42.06	-21.88/-22.81 -42.49	-20.97/-22.18 -41.34	-22.81/-23.38 -40.80	-19.99/-21.18 -41.59
Sequence-only	-15.68/-17.76 -35.93	-16.10/-20.32 -39.67	-15.30/-16.96 -38.51	-15.88/-17.46 -33.28	-23.77/-24.43 -41.27
Spatial-only	-16.50/-18.27 -38.03	-15.86/-17.55 -35.73	-19.24/-20.01 -40.16	-13.76/-14.08 -35.06	-15.01/-16.15 -37.74
w/o Elite Search	-35.62/-43.51 -90.06	-32.79/-51.34 -92.27	-37.39/-47.67 -84.40	-40.54/-45.80 -92.47	-25.46/-40.34 -91.80

The source archives include run-level metrics, configurations, checkpoint identities, budget ledgers, and independent audits of the original variant batch and the additional 500 candidates from the joint ablation. They are compact results archives, not self-contained backups of all candidate files and model weights.

F SEED-0 BEST STRUCTURES AND PHYSICAL VALUES

Each panel shows the highest-scoring path in the seed-0 reported pool, not a five-run average. Power is $\bar{p}$, the mean of the four region powers; nonuniformity is d_U (%); inductance is $\mathcal{L}$ (nH). All values in a panel come from the same path. Blue paths include fixed leads; copper objects denote workpieces. Rendering thickness is schematic; path centerlines are unchanged. Views and scales are fixed within each scene. Higher Best@100 and lower nonuniformity are preferred; power and inductance are compared with their targets. †Random-DFS reports true-score Top-100 from supervision plus 100 new random candidates; each group identifies its best path's source.

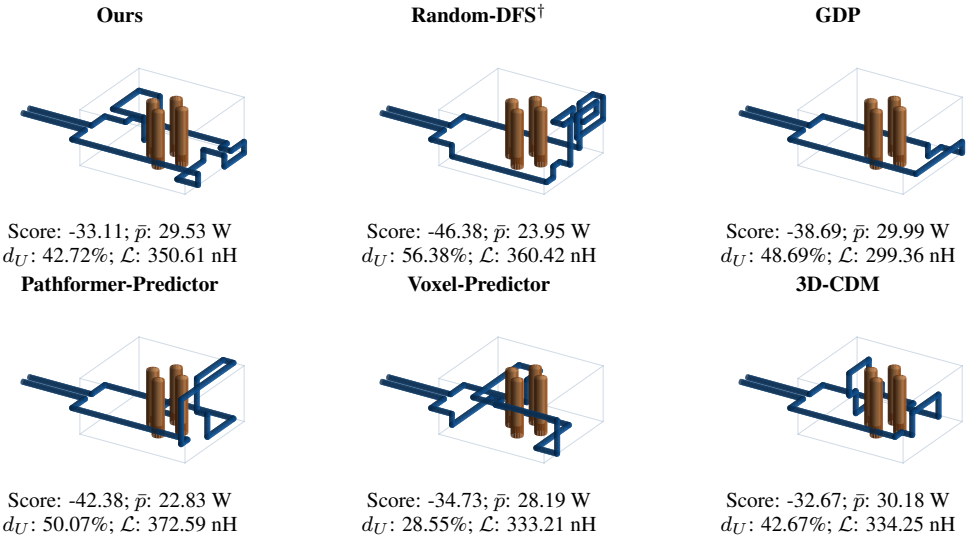

Figure 4: Seed-0 best structures: Rhombus, $n = 100$.

Table 9: Seed-0 best-path scores and physical values for Rhombus, $n = 100$ (Figure 4).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-33.11	29.53	42.72	350.61
Random-DFS†	-46.38	23.95	56.38	360.42
GDP	-38.69	29.99	48.69	299.36
Pathformer-Predictor	-42.38	22.83	50.07	372.59
Voxel-Predictor	-34.73	28.19	28.55	333.21
3D-CDM	-32.67	30.18	42.67	334.25

†Random-DFS best-path source: supervision pool.

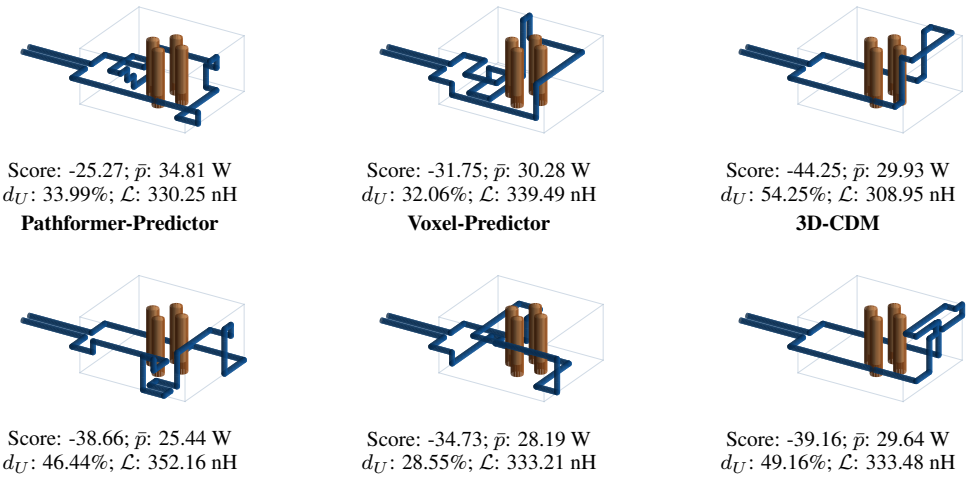

Figure 5: Seed-0 best structures: Rhombus, $n = 200$.

Table 10: Seed-0 best-path scores and physical values for Rhombus, $n = 200$ (Figure 5).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-25.27	34.81	33.99	330.25
Random-DFS†	-31.75	30.28	32.06	339.49
GDP	-44.25	29.93	54.25	308.95
Pathformer-Predictor	-38.66	25.44	46.44	352.16
Voxel-Predictor	-34.73	28.19	28.55	333.21
3D-CDM	-39.16	29.64	49.16	333.48

†Random-DFS best-path source: new random candidates.

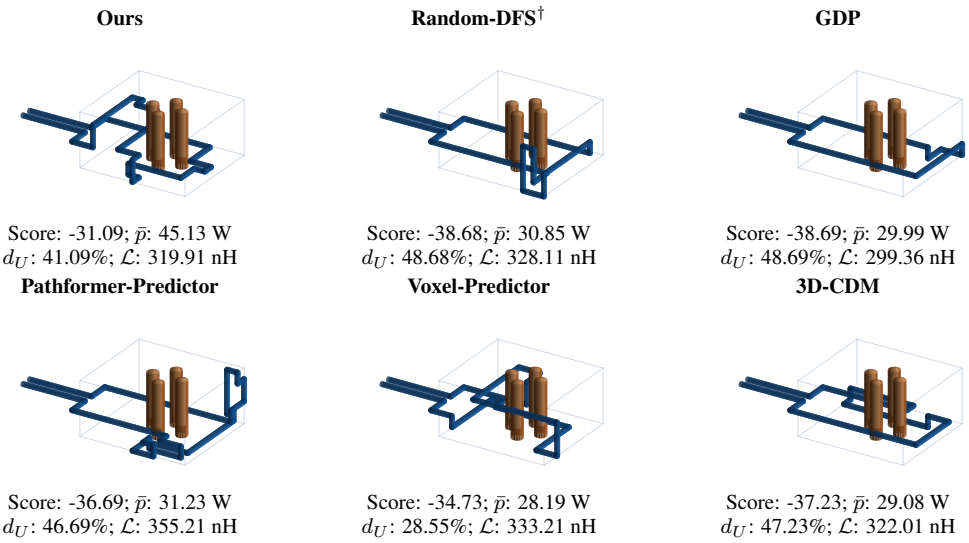

Figure 6: Seed-0 best structures: Rhombus, $n = 300$.

Table 11: Seed-0 best-path scores and physical values for Rhombus, $n = 300$ (Figure 6).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-31.09	45.13	41.09	319.91
Random-DFS†	-38.68	30.85	48.68	328.11
GDP	-38.69	29.99	48.69	299.36
Pathformer-Predictor	-36.69	31.23	46.69	355.21
Voxel-Predictor	-34.73	28.19	28.55	333.21
3D-CDM	-37.23	29.08	47.23	322.01

†Random-DFS best-path source: supervision pool.

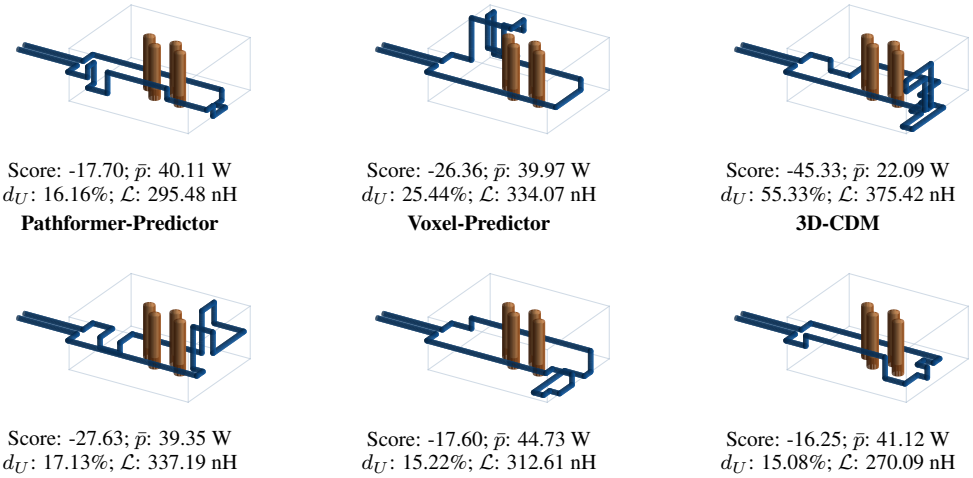

Figure 7: Seed-0 best structures: Staggered, $n = 100$.

Table 12: Seed-0 best-path scores and physical values for Staggered, $n = 100$ (Figure 7).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-17.70	40.11	16.16	295.48
Random-DFS†	-26.36	39.97	25.44	334.07
GDP	-45.33	22.09	55.33	375.42
Pathformer-Predictor	-27.63	39.35	17.13	337.19
Voxel-Predictor	-17.60	44.73	15.22	312.61
3D-CDM	-16.25	41.12	15.08	270.09

†Random-DFS best-path source: new random candidates.

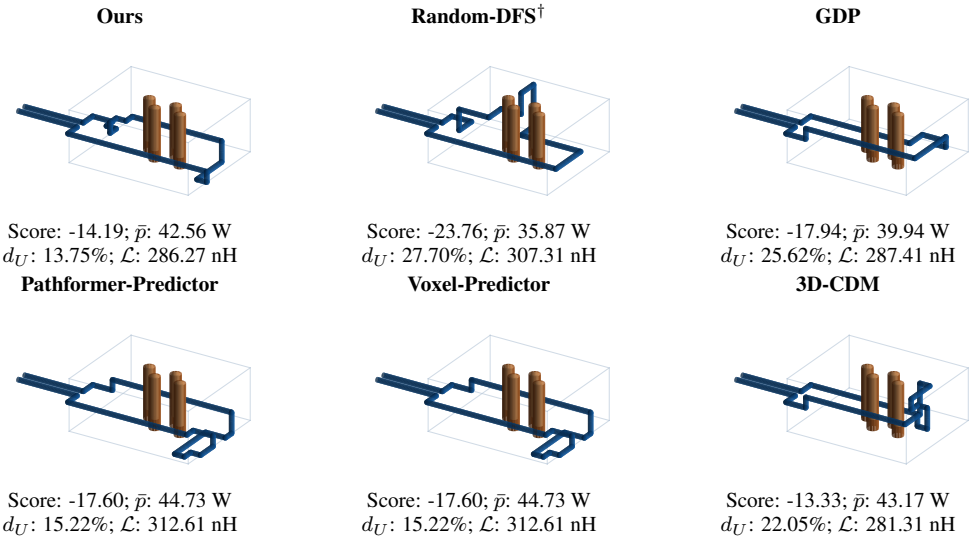

Figure 8: Seed-0 best structures: Staggered, $n = 200$.

Table 13: Seed-0 best-path scores and physical values for Staggered, $n = 200$ (Figure 8).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-14.19	42.56	13.75	286.27
Random-DFS†	-23.76	35.87	27.70	307.31
GDP	-17.94	39.94	25.62	287.41
Pathformer-Predictor	-17.60	44.73	15.22	312.61
Voxel-Predictor	-17.60	44.73	15.22	312.61
3D-CDM	-13.33	43.17	22.05	281.31

†Random-DFS best-path source: new random candidates.

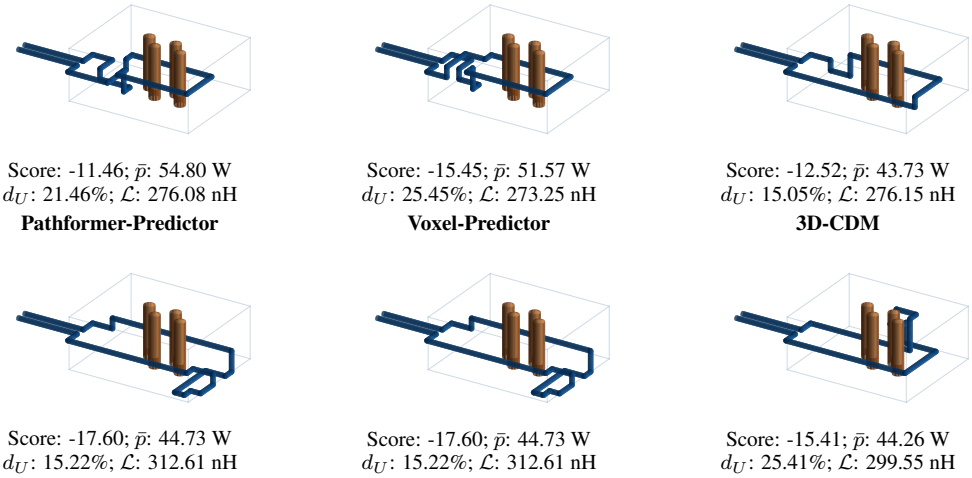

Figure 9: Seed-0 best structures: Staggered, $n = 300$.

Table 14: Seed-0 best-path scores and physical values for Staggered, $n = 300$ (Figure 9).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-11.46	54.80	21.46	276.08
Random-DFS†	-15.45	51.57	25.45	273.25
GDP	-12.52	43.73	15.05	276.15
Pathformer-Predictor	-17.60	44.73	15.22	312.61
Voxel-Predictor	-17.60	44.73	15.22	312.61
3D-CDM	-15.41	44.26	25.41	299.55

†Random-DFS best-path source: supervision pool.

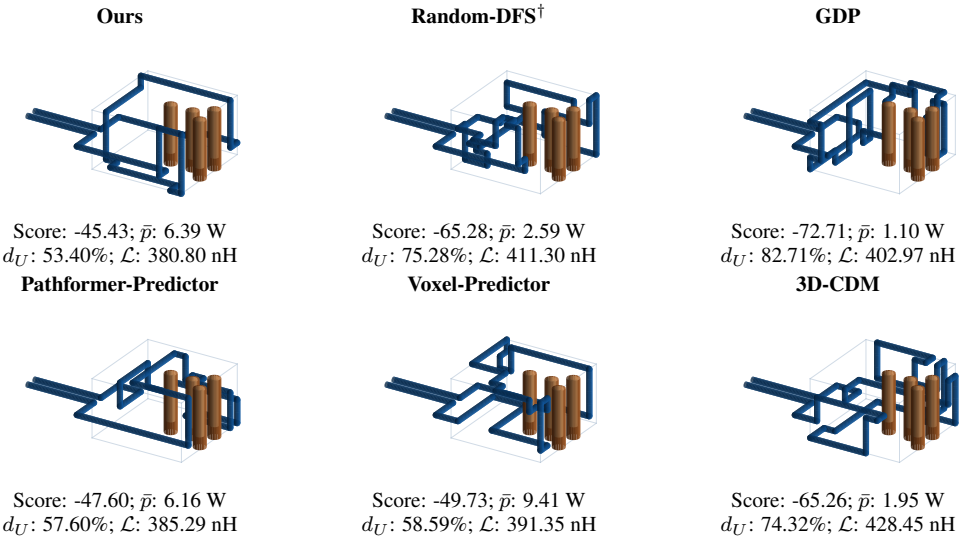

Figure 10: Seed-0 best structures: Y, $n = 100$.

Table 15: Seed-0 best-path scores and physical values for Y, $n = 100$ (Figure 10).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-45.43	6.39	53.40	380.80
Random-DFS†	-65.28	2.59	75.28	411.30
GDP	-72.71	1.10	82.71	402.97
Pathformer-Predictor	-47.60	6.16	57.60	385.29
Voxel-Predictor	-49.73	9.41	58.59	391.35
3D-CDM	-65.26	1.95	74.32	428.45

†Random-DFS best-path source: supervision pool.

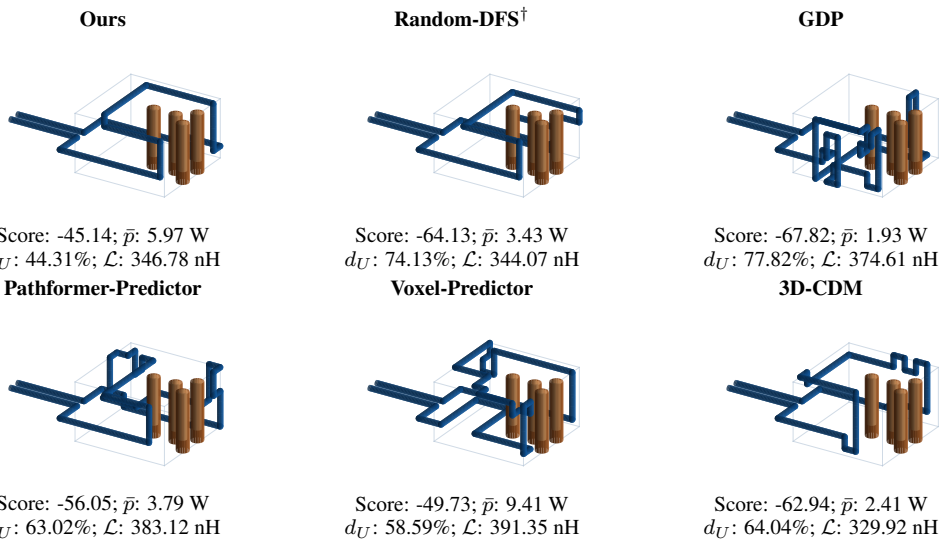

Figure 11: Seed-0 best structures: Y, $n = 200$.

Table 16: Seed-0 best-path scores and physical values for Y, $n = 200$ (Figure 11).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-45.14	5.97	44.31	346.78
Random-DFS†	-64.13	3.43	74.13	344.07
GDP	-67.82	1.93	77.82	374.61
Pathformer-Predictor	-56.05	3.79	63.02	383.12
Voxel-Predictor	-49.73	9.41	58.59	391.35
3D-CDM	-62.94	2.41	64.04	329.92

†Random-DFS best-path source: supervision pool.

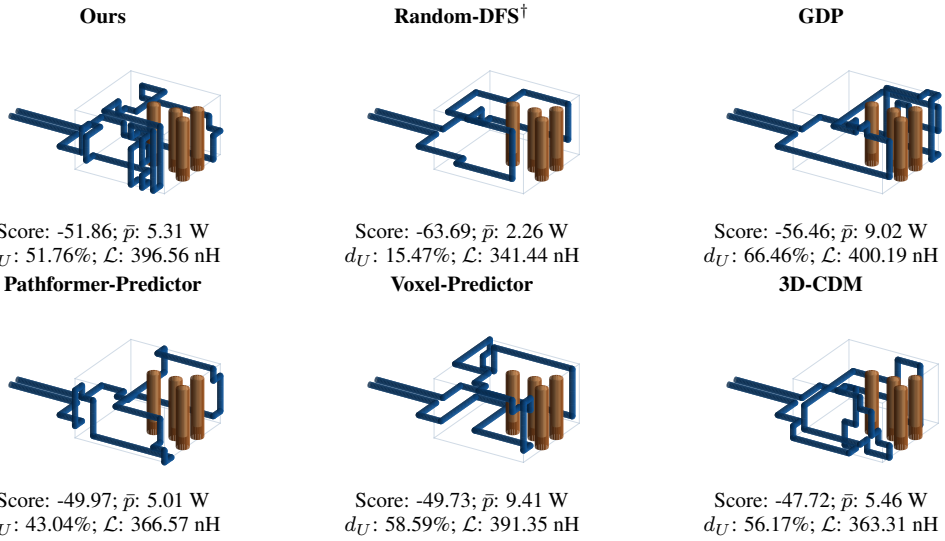

Figure 12: Seed-0 best structures: Y, $n = 300$.

Table 17: Seed-0 best-path scores and physical values for Y, $n = 300$ (Figure 12).

Method	Best@100	Power (W)	Nonuniformity (%)	Inductance (nH)
Ours	-51.86	5.31	51.76	396.56
Random-DFS†	-63.69	2.26	15.47	341.44
GDP	-56.46	9.02	66.46	400.19
Pathformer-Predictor	-49.97	5.01	43.04	366.57
Voxel-Predictor	-49.73	9.41	58.59	391.35
3D-CDM	-47.72	5.46	56.17	363.31

†Random-DFS best-path source: supervision pool.