

A Unified Framework for Circuit-Aware Tapped-Inductor Inverse Design with Topology-Aware Sequence Representation

Jiadong Zhu

The Hong Kong University of Science and Technology
(Guangzhou)
Guangzhou, China
jzhu484@connect.hkust-gz.edu.cn

Zhirui Zong

The Hong Kong University of Science and Technology
(Guangzhou)
Guangzhou, China
zong@hkust-gz.edu.cn

Junjia Ru

The Hong Kong University of Science and Technology
(Guangzhou)
Guangzhou, China
jru958@connect.hkust-gz.edu.cn

Yuzhe Ma

The Hong Kong University of Science and Technology
(Guangzhou)
Guangzhou, China
yuzhema@hkust-gz.edu.cn

Abstract

Inverse design of on-chip tapped inductors is one-to-many: a target electrical response can map to many layouts, but few remain legal and effective after electromagnetic (EM) and circuit-level evaluation. Existing learning-based methods often rely on restrictive parametric codes or image-like maps that weaken layout discipline. We instead represent each inductor as a topology-aware sequence and generate it autoregressively with a decoder-only transformer conditioned on structured passive-response tokens implied by the target active circuit. A verification-aware flow then performs legality screening, minimum-edit repair, post-EM ranking, and circuit-level selection. Across three T-coil-peaked driver cases and three oscillator-tank cases, each evaluated over five seeds, our method achieves the best mean circuit-level performance in every case. It delivers the highest mean driver 3-dB bandwidths of 60.66–82.53 GHz with $2.86\times$ – $2.95\times$ bandwidth extension, and the lowest mean oscillator frequency errors of 1.84%–4.61% with the highest mean figures of merit.

ACM Reference Format:

Jiadong Zhu, Junjia Ru, Zhirui Zong, and Yuzhe Ma. 2027. A Unified Framework for Circuit-Aware Tapped-Inductor Inverse Design with Topology-Aware Sequence Representation. In *32nd Asia and South Pacific Design Automation Conference (ASPAC '27)*, January 25–28, 2027, Tokyo, Japan. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3840404.3848188>

1 Introduction

On-chip radio-frequency (RF) passives must meet electrical targets and strict layout constraints simultaneously [1, 12, 15, 24, 25]. For integrated inductors, inductance, quality factor Q , coupling, and self-resonant frequency (SRF) are tightly coupled through geometry, conductor loss, and substrate parasitics [3, 12]. Designers therefore iterate among sizing, layout realization, electromagnetic (EM) simulation, and circuit-level verification [6, 11, 19]. A tapped inductor further adds an intermediate terminal, expanding the design space while introducing tap-placement, pin-access, and legality constraints [20].

In many circuit-driven RF passive-design tasks, the starting point is a target response imposed by the surrounding active circuit, such as bandwidth extension, impedance shaping, or resonance control, rather than a preferred geometry [5, 8, 17]. The resulting inverse-design problem is response-to-layout: given a target electrical response, generate a legal tapped-inductor layout whose EM-extracted response remains useful in the surrounding circuit. This inverse mapping is inherently one-to-many: a target response may admit multiple geometries, but only a subset remains legal and competitive after physical verification.

Learning-based inverse design has been explored through direct inverse neural models for microwave filters, antennas, and on-chip EM structures [13, 21, 26], invertible neural networks for microwave component synthesis that explicitly address one-to-many inverse mappings [23], and generative models for EM structures and impedance-matching circuits [7, 14, 27]. Recent work has extended inverse design to higher-frequency RF passives and circuits [8, 9], while surrogate-assisted optimization has demonstrated circuit-aware passive search over parametric T-coil families in fixed wireline-driver contexts [22]. Nevertheless, existing formulations do not fully resolve the representation-and-realization gap in response-to-layout synthesis of tapped inductors. Low-dimensional parametric families restrict the diversity of feasible solutions, whereas image-like or pattern-centric generators do not naturally preserve ordered geometry, geometric regularity, or explicit tap semantics needed for robust layout realization. Thus, the core challenge is to build an expressive yet layout-disciplined response-to-layout framework that combines topology-aware representation and response-conditioned generation with layout realization, EM verification, and circuit-level selection.

We address this problem through three coupled elements. First, a tapped inductor is encoded as an ordered outer-contour sequence of direction, length, width, and tap tokens, enabling segment-wise geometric variation within the studied tapped-octagonal family without reducing the layout to a small predefined parameter vector. Second, rather than conditioning directly on dense complex scattering-parameter (S -parameter) tensors, we use cross-attention [18] to condition an autoregressive transformer decoder on structured, frequency-resolved passive tokens defined over a task-specific grid. Each per-frequency token encodes inductance L , coupling k , and an ordinal Q preference, while a separate global token encodes the SRF preference.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

ASPAC '27, Tokyo, Japan

© 2027 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2923-2/2027/01

<https://doi.org/10.1145/3840404.3848188>

Autoregressive decoding is well matched to the ordered geometry because early segment decisions constrain later ones. Third, generated candidates undergo a verification-aware flow comprising legality screening, fixed-topology minimum-edit repair, post-EM passive ranking, and final circuit-level selection. Together, these elements yield a unified end-to-end inverse-design flow.

In summary, the contributions are as follows:

- We introduce a topology-aware outer-contour sequence representation that preserves traversal order, variable geometry, and explicit tap semantics without restricting layouts to a small predefined parametric family.
- We propose structured frequency-resolved conditioning and autoregressive sequence generation that separate frequency-local inductance, coupling, and quality-factor preferences from a frequency-global SRF preference.
- We establish a unified inverse-design workflow that integrates candidate generation, legality screening, minimum-edit repair, post-EM passive ranking, and final circuit-level selection.
- We validate the proposed method on T-coil-peaked driver and oscillator-tank cases, demonstrating superior realizability and the best circuit-level performance across all evaluated cases.

2 Preliminaries

2.1 Tapped Inductors in RF Networks

Integrated spiral inductors are widely used in RF front ends for matching, tuned loading, resonance, and broadband passive shaping, with their inductance, loss, and self-resonant behavior strongly dependent on geometry, substrate coupling, and conductor parasitics [1, 12, 15, 24, 25]. By introducing an intermediate terminal, tapped inductors provide additional electrical degrees of freedom for bandwidth extension, pole-zero shaping, resonant tanks, and multi-port impedance control [5, 17, 20]. Their topology, geometry, and tap location must therefore be jointly determined according to the surrounding circuit requirements, making their synthesis inherently response-driven.

2.2 Passive Characterization and Circuit-Level Validation

Integrated inductors and tapped coils are frequency-dependent networks typically characterized by S -parameters and derived passive metrics including effective inductance, quality factor Q , coupling, and SRF [3, 12, 15]. These metrics summarize energy storage, loss, and inter-port interaction over frequency. Because geometry jointly affects conductor resistance, parasitic capacitance, coupling, and substrate loss, these metrics exhibit inherent trade-offs.

After grounding the tap, let $Z(f) \in \mathbb{C}^{2 \times 2}$ be the resulting two-port impedance matrix. The outer and inner terminals are denoted P and S , respectively. For $u \in \{S, P\}$, the effective inductance and port quality factor are

$$L_u(f) = \frac{\Im\{Z_{uu}(f)\}}{2\pi f}, \quad Q_u(f) = \frac{\Im\{Z_{uu}(f)\}}{\Re\{Z_{uu}(f)\}}. \quad (1)$$

The coupling metric $k(f)$ is extracted as

$$k(f) = \sqrt{\frac{\Im\{Z_{SP}(f)\}\Im\{Z_{PS}(f)\}}{\Im\{Z_{SS}(f)\}\Im\{Z_{PP}(f)\}}}, \quad (2)$$




Figure 1: Sequence representation of a tapped spiral inductor.




Figure 2: Centerline-based versus outline-based width-aware representations.

whenever the denominator is positive. The SRF is operationally the first frequency at which either $L_S(f)$ or $L_P(f)$ becomes nonpositive [3, 10, 16, 24, 25].

However, passive metrics alone do not fully determine circuit utility. Idealized target specifications omit loss and parasitics and thus serve as electrical references rather than directly realizable passives [1, 2]. Moreover, candidates with similar passive error can behave differently in a surrounding network. A useful inverse design must therefore yield a legal layout, preserve the requested response after inductor-level EM extraction, and satisfy the circuit-level objective [6, 11].

3 Proposed Conditional Generative Framework

3.1 Topology-Aware Sequence Representation

Instead of predicting layout pixels or continuous polygon coordinates directly, we represent each inductor by a discrete token sequence over ordered outer-contour segments:

$$\mathbf{y} = [\text{SOP}], (\text{DIR}, L, W), \dots, [\text{TAP_IDX}], i_{\text{tap}}, [\text{EOP}], \quad (3)$$

where [SOP] and [EOP] denote the start and end tokens, respectively, [TAP_IDX] introduces the tap index, and DIR, L , and W encode segment direction, contour length, and inward metal width, respectively. The variable i_{tap} identifies the tapped segment. The mapping from the geometry to the sequence is illustrated in Figure 1. The gold polyline traces the outer contour, the light-blue arrows indicate the eight-direction traversal order, and the blue square markers identify the P terminal, the S terminal, and the tap location. Traversal starts

at the outer P terminal, follows the ordered contour, and ends at the inner S terminal, with the tap location encoded relative to this traversal order.

We use the outer contour because centerline-based representations become awkward when adjacent segments have different widths. As shown in Figure 2, outward expansion can create gaps or corner spurs, highlighted by the blue and red triangles, respectively, whereas an explicit outer boundary enables deterministic inward expansion and corner stitching. This better preserves the polygonal geometry and avoids excess-metal trimming and unnecessary perturbations to current crowding and local parasitic coupling.

In this work, DIR is instantiated as an eight-direction vocabulary $\{N, NE, E, SE, S, SW, W, NW\}$, corresponding to octagonal outer contours. This matches common RFIC practice, where octagonal spirals offer a practical compromise between circular and square layouts and can reduce corner-related current crowding and series resistance relative to square spirals while remaining convenient for regular polygonal layout generation [4, 10, 25]. The sequence formalism is not restricted to eight directions; replacing the DIR vocabulary with a different set of allowable directions adapts it to other types of polygonal inductors.

This sequence representation is a deliberate middle ground between low-dimensional parameterizations and rasterized geometry maps. Compared with parameter sets, it preserves a substantially larger inverse solution space and can describe diverse yet still legalizable layouts. Relative to image-like encodings, it retains explicit ordering, local direction semantics, and tap-location semantics.

3.2 Problem Formulation

Let $\mathcal{R}$ denote the target electrical response induced by a given target circuit case, and let $\mathbf{c}$ denote the condition tokens derived from it. The problem is to use $\mathbf{c}$ as the conditioning input and generate a tapped-inductor sequence $\mathbf{y}$ such that the realized geometry $\mathcal{G}(\mathbf{y})$ remains valid under the downstream realization flow and its EM response $\mathcal{S}(\mathcal{G}(\mathbf{y}))$ matches $\mathcal{R}$ as closely as possible. In circuit terms, the goal is to generate a passive layout to meet the requirements implied by the circuit-level objective.

3.3 Framework Overview

As illustrated in Figure 3, the proposed framework is organized into five stages. First, a target active circuit defines the desired electrical behavior of the tapped inductor. Second, the target S -parameter response is converted into a structured passive specification and then encoded into condition tokens. Third, a conditioned autoregressive transformer generates a candidate pool of topology-aware inductor sequences. Fourth, these raw candidates pass through legality filtering and repair so that only realizable layouts continue downstream. Finally, the post-EM candidates are ranked in the passive space, inserted into the surrounding circuit, and compared by circuit-level simulation to select the winner. Thus, the framework targets circuit-level utility rather than layout realizability or passive-response fidelity alone.

3.4 Target Construction and Condition Encoding

We ground the tap terminal in the target three-port response, recover the corresponding two-port impedance matrix, and calculate Equation (1) and Equation (2) on a task-specific frequency grid

$\mathcal{F} = \{f_1, f_2, \dots, f_M\}$. The grid spans 5–100 GHz for broadband T-coil peaking, whereas oscillator tasks use target-centered local windows. The frequency-resolved rows are

$$\tilde{\mathbf{c}}_f = [L_S(f), L_P(f), k(f), Q^{\text{lv}}(f)], \quad f \in \mathcal{F}, \quad (4)$$

and the global condition row is

$$\tilde{\mathbf{c}}_{\text{global}} = [0, 0, 0, \text{SRF}^{\text{lv}}]. \quad (5)$$

Here Q^{lv} and SRF^{lv} denote five-level ordinal preferences. At each frequency, Q^{lv} is calibrated from training-split empirical quantiles of

$$q_{\text{bal}}(f) = \max(\min(Q_S(f), Q_P(f)), 0), \quad (6)$$

whereas SRF^{lv} is calibrated from the empirical SRF quantiles over the dataset. These levels encode relative preferences for lower loss and higher SRF, avoiding the need to assign nonphysical exact Q or SRF values to an ideal target.

As shown in Figure 4, a shared per-frequency token encoder maps each continuous $[L_S, L_P, k]$ tuple through a multilayer perceptron (MLP) into $\mathbb{R}^{d_{\text{model}}}$ and additively combines the result with learned Q^{lv} and frequency-position embeddings. The shared encoder provides a common local electrical encoding, while frequency-position embeddings distinguish frequency roles. A separate global encoder is used because SRF characterizes the response globally. The resulting stack of condition tokens is

$$\mathbf{c} = [\mathbf{c}_{f_1}, \mathbf{c}_{f_2}, \dots, \mathbf{c}_{f_M}, \mathbf{c}_{\text{global}}] \in \mathbb{R}^{(M+1) \times d_{\text{model}}}, \quad (7)$$

which preserves task-relevant frequency variation without the dimensionality and phase-coordinate sensitivity of raw S -parameters.

3.5 Transformer-Based Candidate Generation

The decoder-only transformer in Figure 5 generates the sequence in Equation (3) from [SOP] to [EOP]. Before entering the decoder stack, tokens in the partially generated prefix are mapped to learned token embeddings and combined with positional embeddings. Each decoder block applies masked self-attention, cross-attention to $\mathbf{c}$, and a feed-forward MLP. Self-attention captures geometric dependencies because early segment decisions constrain later lengths, widths, and tap placement, whereas cross-attention injects the target electrical specification at each step, allowing different geometry tokens to attend to frequency-specific conditions or the global SRF preference. After the decoder stack, layer normalization and a linear vocabulary projection produce next-token logits. At each decoding step, nucleus sampling draws a token from the top- p -filtered distribution to extend the prefix. Decoding continues until [EOP] to complete one candidate, and repeated sampling yields a diverse candidate pool for downstream legality filtering, repair, and verification.

3.6 Legality Filtering and Minimum-Edit Repair

All candidates undergo geometry-aware legality screening and, when necessary, minimum-edit repair before EM simulation. Geometry checks reject malformed sequences and self-intersections not confined to the terminal segment rather than aggressively rewriting them, because repairing such cases would amount to topology redesign and obscure the quality of the generated sample itself. For near-legal samples, a topology-preserving integer linear program (ILP) maintains the segment count, direction order, topology, and tap index while adjusting only integer-valued segment lengths and widths. Let N denote the number of ordered contour segments; hat-denoted quantities are the raw generated values, whereas ℓ_i and w_i are




Figure 3: Overview of the response-to-layout generation and verification flow.




Figure 4: Frequency-resolved condition construction and encoding on a task-specific grid.

the repaired integer variables for $i = 0, \dots, N-1$. The repair solves

$$\begin{aligned} \min_{\ell, w} \quad & \sum_{i=0}^{N-1} (|\ell_i - \hat{\ell}_i| + |w_i - \hat{w}_i|) \\ \text{s.t.} \quad & \ell_i \geq \ell_{\min}, \quad w_{\min} \leq w_i \leq w_{\max}, \quad w_0 = w_1, \\ & \ell_i, w_i \in \mathbb{Z}, \end{aligned} \quad (8)$$

where $\ell_{\min}$ and $[w_{\min}, w_{\max}]$ specify the length and width bounds, respectively. In the implementation, each absolute-value term is converted to an equivalent linear form using a standard nonnegative auxiliary variable. Under our layout convention, the first two segments jointly form the feed-corner polygon used for terminal construction; imposing $w_0 = w_1$ maintains a uniform junction width and avoids an artificial width discontinuity that could exacerbate current crowding and terminal parasitics.




Figure 5: Decoder-only autoregressive transformer with cross-attention to task-specific condition tokens.

Because the direction sequence is fixed during repair, the relative orientations of critical segment pairs are fixed, allowing the following three interaction sets to be precomputed:

$$\begin{aligned} E_{\text{sp}} &= \{(i, i+8) \mid i = 1, \dots, N-9\}, \\ E_{\text{in}} &= \{(i, i+4) \mid i = N-8, \dots, N-5\}, \\ E_{\text{tail}} &= \{(N-8, N-1), (N-7, N-1), (N-6, N-1)\}. \end{aligned} \quad (9)$$

The sets E_{sp} , E_{in} , and E_{tail} cover adjacent-turn spacing, inner-core approach, and tail interactions, respectively. For each pair (i, j) , $\Delta_{ij}(\ell) = \sum_{m=0}^{N-1} a_{ijm} \ell_m$ denotes the oriented perpendicular separation of segment j from the supporting line of segment i . The fixed coefficient $a_{ijm} \in \{0, \pm 1, \pm 1/\sqrt{2}\}$ is the signed projection of segment

m onto the inward normal of segment i , extracted from the generated contour under the fixed direction sequence. With minimum spacing $s_{\min}$, legality requires

$$\begin{aligned} w_i &\leq \Delta_{ij}(\ell) - s_{\min}, & (i, j) \in E_{\text{sp}}, \\ w_i + w_j &\leq \Delta_{ij}(\ell) - s_{\min}, & (i, j) \in E_{\text{in}}, \\ w_i + \gamma_{ij} w_j &\leq \Delta_{ij}(\ell) - s_{\min}, & (i, j) \in E_{\text{tail}}. \end{aligned} \quad (10)$$

The tail width-projection coefficient γ_{ij} is 0 for a single-segment constraint and $1/\sqrt{2}$ for the corner-coupled case. These precomputed coefficients keep all spacing constraints affine and compact. Adapting the formulation to another polygon family requires replacing the projection constants and pair rules.

3.7 Passive Ranking and Circuit-Level Selection

After legality screening and repair, surviving candidates are realized as layouts and evaluated by EM simulation. The following correlation analysis both supports the frequency-resolved condition representation and motivates the post-EM ranking proxy. In normalized frequency-resolved $[L_S, L_P, k]$ space over $\mathcal{F}$, we define the passive matching error as

$$e_{\text{pass}} = \sqrt{\frac{1}{D} \sum_{r=1}^D \left(z(x_r^{\text{cand}}) - z(x_r^{\text{tgt}}) \right)^2}, \quad (11)$$

where x_r^{cand} and x_r^{tgt} are the r -th components of the candidate and target passive vectors, $D = 3|\mathcal{F}|$ is the passive-vector dimensionality, and $z(\cdot)$ uses the training-split mean and standard deviation.

A smaller e_{pass} indicates higher fidelity to the target passive response. To capture passive quality beyond $[L_S, L_P, k]$ matching, we define

$$q_{\text{score}} = \frac{1}{|\mathcal{F}|} \sum_{f \in \mathcal{F}} \log_{10}(1 + q_{\text{bal}}(f)), \quad (12)$$

which summarizes the balanced Q behavior in Equation (6) over $\mathcal{F}$, and a larger value indicates lower loss. The hybrid score is

$$s_{\text{hyb}} = e_{\text{pass}} - \alpha z(q_{\text{score}}) - \beta z(\log_{10} \text{SRF}), \quad (13)$$

where lower values are preferred, with the final term additionally rewarding higher SRF. We evaluate this proxy by inserting the EM-extracted S -parameters of 1,000 randomly collected inductor designs into a T-coil-peaked driver and correlating each passive score with circuit-level 3-dB bandwidth. With $\alpha = \beta = 1$, Table 1 shows stronger Pearson and Spearman correlations as Q and SRF are added. The negative signs indicate that smaller passive scores predict larger bandwidths. Figure 6 visualizes the complete hybrid relationship, where each point represents one EM-evaluated layout, and the fitted line summarizes the overall trend.

Accordingly, post-EM candidates are ranked using Equation (13).

The highest-ranked candidates are inserted into the external circuit with their EM-extracted S -parameters. The final winner is selected by the task-specific circuit objective: 3-dB bandwidth for the T-coil-peaked driver, and oscillation-frequency accuracy together with oscillator figure of merit (FOM) for the oscillator tank.

4 Experimental Results

4.1 Setup

All experiments run on Linux with an Intel Xeon Gold 6348 CPU, four NVIDIA A30 GPUs, and 512 GB of memory. Cadence EMX extracts inductor S -parameters, and Spectre uses the TSMC N28 model

Table 1: Correlation between circuit-level 3-dB bandwidth and passive ranking proxies over 1,000 EM-evaluated designs.

Ranking proxy	Pearson	Spearman
e_{pass}	-0.401	-0.455
$e_{\text{pass}} - z(q_{\text{score}})$	-0.532	-0.550
$e_{\text{pass}} - z(q_{\text{score}}) - z(\log_{10} \text{SRF})$	-0.595	-0.581




Figure 6: Circuit-level 3-dB bandwidth versus the hybrid score over 1,000 randomly collected EM-evaluated designs.

deck for circuit evaluation. Process rules constrain metal widths to 2–12 μm and require a minimum spacing of 1 μm . The dataset contains 10,155 tapped octagonal spiral inductors with paired geometry sequences and EM responses, using a 95%/5% training/validation split. All circuit benchmarks use independently constructed ideal target responses that are not drawn from either split.

The generator uses six decoder layers, $d_{\text{model}} = 256$, eight heads, dropout 0.15, and block size 256. AdamW training uses 50 epochs and a learning rate of 2.5×10^{-4} ; inference uses top- $p = 0.8$ and temperature 0.85. We compare against four baselines: conditional generative adversarial network (cGAN)-based inverse generation [7, 27], conditional invertible neural network (cINN)-based inverse modeling [23], covariance matrix adaptation evolution strategy (CMA-ES) search with a learned surrogate [22], and deterministic direct inverse regression [13, 21]. All methods use the same condition, output sequence, legality rules, EM extraction, passive ranking, and circuit selection. The fixed budget is 10 unique samples, 8 legality-first EM candidates, and 6 post-EM circuit finalists; Direct emits one deterministic candidate. If fewer than eight candidates survive filtering/repair, all survivors are evaluated; if none survive, the run is reported as a failure.

We evaluate three cases each for the differential T-coil-peaked driver and oscillator tank in Figure 7. Table 2 lists the case-specific parasitic capacitance C_p and no-inductor reference 3-dB bandwidth $\text{BW}_{3\text{dB}}$ for the driver cases, as well as tank capacitance C_t and target oscillation frequency f_{osc} for the oscillator cases. Conditioning follows the circuit objective. For the T-coil-peaked driver, conditioning uses a grid from 5 to 100 GHz in 5 GHz increments and the highest ordinal Q and SRF preference levels, with $\alpha = \beta = 1$ in hybrid passive ranking. For the oscillator tank, conditioning uses target-centered frequencies spaced 1 GHz apart within ± 2 GHz of each case's target f_{osc} , the highest ordinal Q preference, and an SRF level determined relative to the target f_{osc} of each case.

For the driver in Figure 7(a), finalists are evaluated using Spectre alternating-current (AC) analysis, from which the differential voltage




Figure 7: Circuit contexts: (a) T-coil-peaked driver and (b) oscillator tank.

Table 2: Benchmark setups for the T-coil-peaked driver and oscillator tank.

T-coil-peaked driver			Oscillator tank		
Case	C_p (fF)	Ref. BW _{3dB} (GHz)	Case	C_t (fF)	Target f_{osc} (GHz)
1	58	28.45	1	83.10	28
2	88	25.51	2	65.48	39
3	158	20.56	3	36.68	60

gain is $G_{diff}(f) = 20 \log_{10} |(V_{OUT,P} - V_{OUT,N}) / (V_{IN,P} - V_{IN,N})|$. The 3-dB bandwidth is the first frequency at which $G_{diff}(f)$ falls 3 dB below its low-frequency value; the corresponding threshold is approximately -5.29 dB. Bandwidth extension is the achieved bandwidth normalized by the corresponding no-inductor reference value in Table 2. For the oscillator tank in Figure 7(b), finalists are evaluated using Spectre periodic steady-state (PSS) and periodic noise (PNoise) analyses. The 10-MHz-offset figure of merit is reported as $FOM_{10M} = -\mathcal{L}_{10M} + 20 \log_{10} [f_{osc} / (10 \text{ MHz})] - 10 \log_{10} [P_{DC} / (1 \text{ mW})]$, where $\mathcal{L}_{10M}$ is the 10-MHz-offset phase noise and P_{DC} is the DC power drawn from the 0.5 V supply.

4.2 T-Coil-Peaked Driver Results

Table 3 reports mean $\pm$ standard deviation over five independent sampling seeds for every stochastic method, and boldface marks the best stochastic mean. The first three metrics characterize layout realizability and post-EM passive fidelity. *Legal* counts raw candidates that need no repair; these are retained first in the EM shortlist of at most eight candidates, with remaining slots filled by ILP-repaired candidates counted as *Repair*. *Cond. RMSE* is the minimum post-EM normalized condition-space root-mean-square error (RMSE) among the EM-evaluated shortlist, which covers the frequency-resolved $[L_S, L_P, k, Q^{lv}]$ and the global SRF^{lv} , with the ordinal preference terms penalized only when the realized levels fall below their requested preferences. This is the most direct task-aligned error because it compares the realized passive response against the same target-side features that define the condition. The final two columns measure circuit-level bandwidth.

Our method produces the most directly legal layouts in every case, and when repair is needed, the mean edit $\sum_i (|\Delta \ell_i| + |\Delta w_i|)$ is only 1.00–1.50 μm for our method, versus up to 11.25 μm for the baselines. The corresponding repair step requires less than 0.1 s, which is negligible. More directly legal candidates indicate better geometric alignment, allowing more samples to preserve the generator’s original intent. Our method also achieves the lowest mean

Table 3: Results for the T-coil-peaked driver over five seeds.

Case	Method	Legal	Repair	Cond. RMSE	BW _{3dB} (GHz)	BW ext.
1	Ours	6.60$\pm$1.34	1.40 $\pm$ 1.34	1.050$\pm$0.070	82.53$\pm$1.09	2.90$\pm$0.04$\times$
	cGAN	0.80 $\pm$ 1.79	7.20 $\pm$ 1.79	1.219 $\pm$ 0.032	53.44 $\pm$ 4.11	1.88 $\pm$ 0.15 $\times$
	cINN	0.80 $\pm$ 0.84	7.20 $\pm$ 0.84	1.110 $\pm$ 0.037	55.37 $\pm$ 11.76	1.95 $\pm$ 0.41 $\times$
	Search	0.80 $\pm$ 1.10	7.20 $\pm$ 1.10	1.213 $\pm$ 0.105	49.46 $\pm$ 17.85	1.74 $\pm$ 0.63 $\times$
	Direct	1	0	1.526	38.24	1.34 $\times$
2	Ours	3.60$\pm$0.55	4.40 $\pm$ 0.55	1.051$\pm$0.021	73.03$\pm$5.35	2.86$\pm$0.21$\times$
	cGAN	2.40 $\pm$ 3.36	5.60 $\pm$ 3.36	1.309 $\pm$ 0.185	60.73 $\pm$ 20.98	2.38 $\pm$ 0.82 $\times$
	cINN	0.80 $\pm$ 0.45	7.20 $\pm$ 0.45	1.097 $\pm$ 0.072	55.38 $\pm$ 4.18	2.17 $\pm$ 0.16 $\times$
	Search	0.00 $\pm$ 0.00	8.00 $\pm$ 0.00	1.274 $\pm$ 0.109	48.95 $\pm$ 16.06	1.92 $\pm$ 0.63 $\times$
	Direct	0	1	1.143	33.34	1.31 $\times$
3	Ours	6.40$\pm$0.89	1.60 $\pm$ 0.89	0.993$\pm$0.033	60.66$\pm$3.46	2.95$\pm$0.17$\times$
	cGAN	0.00 $\pm$ 0.00	8.00 $\pm$ 0.00	1.177 $\pm$ 0.010	45.21 $\pm$ 0.26	2.20 $\pm$ 0.01 $\times$
	cINN	2.00 $\pm$ 0.00	6.00 $\pm$ 0.00	1.145 $\pm$ 0.046	47.14 $\pm$ 1.79	2.29 $\pm$ 0.09 $\times$
	Search	0.20 $\pm$ 0.45	7.60 $\pm$ 0.55	1.100 $\pm$ 0.023	53.42 $\pm$ 4.44	2.60 $\pm$ 0.22 $\times$
	Direct	1	0	1.513	29.01	1.41 $\times$

* Direct is deterministic and emits one candidate; *Repair* is not bolded because it is reported for interpretation only and has no universal better-or-worse direction.




Figure 8: Gain curves for the T-coil-peaked driver.

condition-space RMSE, indicating the highest post-EM fidelity to the requested electrical conditions.

Our method achieves the highest mean bandwidth in all three cases, reaching 82.53, 73.03, and 60.66 GHz for Cases 1–3; the corresponding improvements over the strongest baseline in each case are 49.1%, 20.3%, and 13.6%. The same advantage persists after normalization to the no-inductor reference, with bandwidth extension ratios of 2.90 $\times$, 2.86 $\times$, and 2.95 $\times$.

The gain curves in Figure 8 provide a complementary view of the same trend. For each method and case, the plotted curve corresponds to the seed with median BW_{3dB} across repeated experiments, and the gray dashed line denotes the 3-dB-down threshold. Our method consistently pushes the gain roll-off furthest to the right, supporting wider bandwidth for the T-coil-peaked drivers. Some baselines remain competitive over limited frequency ranges, but their earlier high-frequency roll-offs result in lower bandwidth.

4.3 Oscillator Tank Results

With our method’s layout realizability and passive fidelity established in Section 4.2, this section reports only the key circuit-level metrics. Table 4 reports mean $\pm$ standard deviation over five independent sampling seeds for each stochastic method. The f_{osc} column reports the obtained oscillation frequency, which should be as close as possible to the corresponding targets in Table 2, while the relative frequency error is computed per seed as $e_f = |f_{osc} - f_{target}| / f_{target} \times 100\%$. Our method achieves the lowest mean frequency errors of 4.61%, 2.69%, and 1.84%, while also attaining the highest mean FOM values of 194.25, 197.22, and 195.53 in Cases 1–3, respectively.

Table 4: Circuit-level oscillator tank results over five seeds.

Case	Method	f_{osc} (GHz)	e_f (%)	FOM _{10M}
1	Ours	29.16±1.38	4.61±4.42	194.25±1.28
	cGAN	30.94±1.26	10.49±4.49	193.97±1.11
	cINN	31.21±1.52	11.48±5.44	193.74±0.96
	Search	32.74±2.14	16.92±7.65	193.53±0.12
	Direct	33.08	18.16	192.51
2	Ours	37.95±0.56	2.69±1.43	197.22±3.42
	cGAN	37.61±0.39	3.57±1.00	193.65±0.26
	cINN	40.12±1.26	3.49±2.33	196.26±1.76
	Search	40.69±2.90	6.70±4.78	195.81±1.71
	Direct	40.29	3.32	194.16
3	Ours	58.89±0.78	1.84±1.29	195.53±0.37
	cGAN	49.97±0.45	16.72±0.75	193.07±0.25
	cINN	58.26±0.38	2.90±0.63	194.36±1.04
	Search	63.48±3.58	6.18±5.47	194.18±1.03
	Direct	49.72	17.13	192.60

* Direct is deterministic and emits one candidate.

Across all experimental cases, the end-to-end runtimes of all methods are approximately 5–5.5 min, with EM and circuit-level simulations accounting for about 98% of the total. Thus, our method improves candidate quality at essentially fixed cost. Overall, the proposed method is effective not only for broadband impedance shaping in T-coil-peaked drivers but also for prescribed-frequency resonance placement in oscillator tanks, achieving the best mean performance in every evaluated case.

5 Conclusion

This paper presents a unified response-to-layout framework integrating topology-aware sequences, frequency-resolved conditioning, autoregressive generation, and simulation-based verification. Adapting the framework to a new process requires process-specific EM data and design rules, while the sequence and conditioning interfaces and verification flow remain unchanged. Untapped hexagonal or rectangular inductors would further simplify training and inference because their representations use a smaller direction vocabulary, eliminate tap semantics, and require simpler legality constraints. Beyond spiral inductors, the formulation can support a wider range of RF passive and circuit-aware inverse-design tasks.

References

- [1] J. N. Burghartz and B. Rejaei. 2003. On the Design of RF Spiral Inductors on Silicon. *IEEE Transactions on Electron Devices (TED)* 50, 3 (2003), 718–729. doi:10.1109/TED.2003.810474
- [2] Cadence System Analysis. 2021. Inductor Parasitic Capacitance Limits Upper Operating Frequencies. Accessed 2026-04-06. <https://resources.system-analysis.cadence.com/blog/inductor-parasitic-capacitance-limits-upper-operating-frequencies>
- [3] Yu Cao, Robert A. Groves, Xuejue Huang, Noah D. Zamdmer, Jean-Olivier Plouchart, Richard A. Wachnik, Tsu-Jae King, and Chenming Hu. 2003. Frequency-Independent Equivalent-Circuit Model for On-Chip Spiral Inductors. *IEEE Journal of Solid-State Circuits* 38, 3 (2003), 419–426. doi:10.1109/JSSC.2002.808285
- [4] S. Chaki, S. Aono, N. Andoh, Y. Sasaki, N. Tanino, and O. Ishihara. 1995. Experimental Study on Spiral Inductors. In *IEEE/MTT-S International Microwave Symposium (IMS)*. 753–756.
- [5] Sherif Galal and Behzad Razavi. 2003. 10-Gb/s Limiting Amplifier and Laser/Modulator Driver in 0.18- μ m CMOS Technology. *IEEE Journal of Solid-State Circuits* 38, 12 (2003), 2138–2146. doi:10.1109/JSSC.2003.818567
- [6] Reinier Gonzalez-Echevarria, Elisenda Roca, Rafael Castro-López, Francisco V. Fernández, Javier J. Sieiro, José María López-Villegas, and Neus Vidal. 2017. An Automated Design Methodology of RF Circuits by Using Pareto-Optimal Fronts of EM-Simulated Inductors. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* 36, 1 (2017), 15–26. doi:10.1109/TCAD.2016.2564362
- [7] Zheming Gu, Da Li, Yunlong Wu, Yudi Fan, Chengting Yu, Hongsheng Chen, and Er-Ping Li. 2023. A Solution to the Dilemma for FSS Inverse Design Using Generative Models. *IEEE Transactions on Antennas and Propagation (TAP)* 71, 6 (2023), 5100–5109. doi:10.1109/TAP.2023.3266053
- [8] Emir Ali Karahan, Zheng Liu, Aggraj Gupta, Zijian Shao, Jonathan Zhou, Uday Khankhoje, and Kaushik Sengupta. 2024. Deep-learning enabled generalized inverse design of multi-port radio-frequency and sub-terahertz passives and integrated circuits. *Nature Communications* 15, 1 (Dec. 2024), 10734. doi:10.1038/s41467-024-54178-1
- [9] Emir Ali Karahan, Zheng Liu, and Kaushik Sengupta. 2023. Deep-Learning-Based Inverse-Designed Millimeter-Wave Passives and Power Amplifiers. *IEEE Journal of Solid-State Circuits* 58, 11 (2023), 3074–3088. doi:10.1109/JSSC.2023.3276315
- [10] Thomas H. Lee. 2004. *The Design of CMOS Radio-Frequency Integrated Circuits* (2nd ed.). Cambridge University Press. doi:10.1017/CBO9780511817281
- [11] R. M. Martins, N. Lourenço, F. Passos, R. P. Póvoa, A. Canelas, E. Roca, R. Castro-López, J. S. Sieiro, F. Fernández, and N. Horta. 2019. Two-Step RF IC Block Synthesis With Preoptimized Inductors and Full Layout Generation In-the-Loop. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* 38, 6 (2019), 989–1002. doi:10.1109/TCAD.2018.2834394
- [12] Sunderarajan S. Mohan, Maria del Mar Hershenson, Stephen P. Boyd, and Thomas H. Lee. 1999. Simple Accurate Expressions for Planar Spiral Inductances. *IEEE Journal of Solid-State Circuits* 34, 10 (1999), 1419–1424. doi:10.1109/4.792620
- [13] David Munzer, Siawpeng Er, Minshuo Chen, Yan Li, Naga S. Mannem, Tuo Zhao, and Hua Wang. 2020. Residual Network Based Direct Synthesis of EM Structures: A Study on One-to-One Transformers. In *IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*. 143–146. doi:10.1109/RFIC49505.2020.9218278
- [14] Parinaz Naseri and Sean V. Hum. 2021. A Generative Machine Learning-Based Approach for Inverse Design of Multilayer Metasurfaces. *IEEE Transactions on Antennas and Propagation (TAP)* 69, 9 (2021), 5725–5739. doi:10.1109/TAP.2021.3060142
- [15] Ali M. Niknejad and Robert G. Meyer. 1998. Analysis, Design, and Optimization of Spiral Inductors and Transformers for Si RF IC's. *IEEE Journal of Solid-State Circuits* 33, 10 (1998), 1470–1481. doi:10.1109/4.720393
- [16] Ali M. Niknejad and Robert G. Meyer. 2000. *Design, Simulation and Applications of Inductors and Transformers for Si RF IC's*. Springer. doi:10.1007/b116979
- [17] Sudip Shekhar, Jeffrey S. Walling, and David J. Allstot. 2006. Bandwidth Extension Techniques for CMOS Amplifiers. *IEEE Journal of Solid-State Circuits* 41, 11 (2006), 2424–2437. doi:10.1109/JSSC.2006.883336
- [18] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Annual Conference on Neural Information Processing Systems (NIPS)*. 5998–6008.
- [19] Jiahao Wei, Weiqi Chen, Yajie Gong, Qi Wu, Guangyi Lu, Wei Gao, Lihui Wang, Mei Li, and Haiming Wang. 2023. Highly Efficient Automatic Synthesis of a Millimeter-Wave On-Chip Deformable Spiral Inductor Using a Hybrid Knowledge-Guided and Data-Driven Technique. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* 42, 12 (2023), 4413–4422. doi:10.1109/TCAD.2023.3294449
- [20] Jiahao Wei, Weiqi Chen, Qi Wu, Guangyi Lu, Wei Gao, Lihui Wang, Mei Li, and Haiming Wang. 2024. Microwave Network-Assisted Analysis and Machine Learning-Assisted Synthesis of Arbitrarily Tapped Coils and Its Application to On-Chip Ultrawideband ESD Protection Circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* 43, 12 (2024), 4386–4397. doi:10.1109/TCAD.2024.3416251
- [21] Li-Ye Xiao, Wei Shao, Fu-Long Jin, Bing-Zhong Wang, and Qing Huo Liu. 2021. Inverse Artificial Neural Network for Multi-Objective Antenna Design. *IEEE Transactions on Antennas and Propagation (TAP)* 69, 10 (2021), 6651–6659.
- [22] Yizhou Xu, Chenhao Chu, and Hua Wang. 2026. AI-Assisted Automatic Ultra-Broadband Passive-Active Co-Design Workflow: A T-Coil Peaked Wireline Driver Example. GitHub repository, accessed 2026-04-02. <https://github.com/ETH-IDEAS/AI4TCOIL>
- [23] Huan Yu, Hakki Mert Torun, Mutee Ur Rehman, and Madhavan Swaminathan. 2020. Design of SIW Filters in D-Band Using Invertible Neural Nets. In *IEEE/MTT-S International Microwave Symposium (IMS)*. 72–75.
- [24] C. Patrick Yue and S. Simon Wong. 1998. On-Chip Spiral Inductors With Patterned Ground Shields for Si-Based RF IC's. *IEEE Journal of Solid-State Circuits* 33, 5 (1998), 743–751. doi:10.1109/4.668989
- [25] C. Patrick Yue and S. Simon Wong. 1999. Design Strategy of On-Chip Inductors for Highly Integrated RF Systems. In *ACM/IEEE Design Automation Conference (DAC)*. 982–987. doi:10.1109/DAC.1999.782239
- [26] Chao Zhang, Jing Jin, Weicong Na, Qi-Jun Zhang, and Ming Yu. 2018. Multivalued Neural Network Inverse Modeling and Applications to Microwave Filters. *IEEE Transactions on Microwave Theory and Techniques* 66, 8 (2018), 3781–3797. doi:10.1109/TMTT.2018.2841889
- [27] Jingwei Zhang, Zhun Wei, Kai Kang, and Wen-Yan Yin. 2024. Intelligent Inverse Designs of Impedance Matching Circuits With Generative Adversarial Network. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* 43, 10 (2024), 3171–3183. doi:10.1109/TCAD.2024.3385368