

Hybrid Physics – Machine Learning Framework for Joint 3-DOF Misalignment Compensation via Matching-Network Retuning in Magnetically Coupled Resonant Wireless Power Transfer

Srilatha Gundapaneni^{1,*}, Siva Chakra Avinash Bikkina², Pendurthy Anthony Sunny Dayal³,
Manumula Srinubabu², Manikonda Venkateswararao⁴, and Rangarao Orugu⁵

¹Ramachandra College of Engineering, Andhra Pradesh, India

²Department of ECE, Aditya University, Surampalem 533437, Andhra Pradesh, India

³Department of ECE, SoCDS, Centurion University of Technology and Management, Andhra Pradesh, India

⁴Department of ECE, PACE Institute of Technology and Sciences, Ongole, Andhra Pradesh, India

⁵Department of ECE, Vishnu Institute of Technology, Bhimavaram, Andhra Pradesh, India

ABSTRACT: Magnetically Coupled Resonant (MCR) Wireless Power Transfer (WPT) systems suffer power-transfer efficiency (PTE) degradation due to misalignment of the transmitter/receiver coil pair across multiple degrees of freedom (DOF) simultaneously. Previous methods track lateral offset or rotational drift individually; none have reported correcting combined lateral, vertical, and angular misalignment in real time. The proposed hybrid physics-machine learning model combines an analytical model for a mutual-inductance circuit with a trained feedforward-network (FFN) residual-correction layer to estimate the state of the matching network that produces a maximum PTE when there is simultaneous 3-DOF misalignment. The method is evaluated for a representative 6.78 MHz, 3-turn, 40 mm-diameter coil pair using a numerical electromagnetic model (Neumann-integral mutual inductance evaluation) and an exact model of a series-series (S-S) resonant circuit. The value of nonlinear cross-coupling term between lateral and tilt that is neglected by a separable analytical model increases by $\sim 18\%$ to 79% of the true value of the mutual inductance as misalignment increases, and the value of a learned correction increases strongly with the size of the disturbance. Similarly, the single-axis baseline has an average efficiency gap from the optimal (1.30 percentage points), which is only 49% closed by using the extended analytical model, and 89% closed with the trained residual layer. The shallow FFN is better for this correction than the degree-2 polynomial ridge regressor in an ablation. It is also observed that carrier-frequency retuning has a very small impact on the efficiency in this topology, with the state of the matching network being the major lever. A warning, however, is that the correction may sometimes make the worst-case deficit worse at corners of the envelope where there are fewer samples, meaning that the samples have to be denser or an extrapolation safeguard needs to be included before deployment.

1. INTRODUCTION

In recent years, Wireless Power Transfer (WPT) using magnetically coupled resonance has emerged as a choice solution for contactless power transfer because of its ability to withstand moderate coil separation and misalignment, in contrast to the traditional inductive coupling approach [1, 2]. The coupling coefficient k between transmitter and receiver coils is very sensitive to the position of relative coils and in real deployments such as electric vehicle charging pads, drone charging stations, rotating machinery, it is rare for misalignment to only take place along one axis. The mutual inductance M and resonant behaviour are modified by coupled, nonlinear terms of lateral drift, vertical separation change, and angular tilt, which are not easily represented by an analytical model with a single axis of movement.

Two paths have been taken in order to overcome the efficiency loss due to misalignment: (1) purely analytical or circuit-based frequency retuning, and (2) mechanical or coil-geometry redesign of misalignment tolerance. Neither of the

methods described in Section 2 provides multi-axis compensation. This paper fills this gap by introducing a joint compensation framework for 3-DOF based on an analytical circuit backbone, a trained residual-correction layer, and a numerical electromagnetic and circuit simulation to validate the proposed design rather than a closed-form approximation.

Relative to the closest prior art (Section 2.6, Table 1), the specific engineering contribution is threefold: (i) joint, simultaneous estimation and compensation of all three misalignment axes, rather than one axis at a time; (ii) a physics-informed residual layer that is trained only on the non-separable cross-axis term left over by an explicit analytical backbone, rather than a black-box neural mapping, so the resulting estimate stays interpretable and degrades gracefully outside the training distribution; and (iii) a circuit-level finding (Section 3.2) that, for this topology, matching-network retuning — not carrier-frequency retuning — is the effective compensation mechanism, which directly motivates why the retuning controller of Figure 1 targets R_L rather than f .

* Corresponding author: Srilatha Gundapaneni (drgsl@rcee.ac.in).

2. RELATED WORK

2.1. Analytical Frequency Tracking under Horizontal Offset

Under horizontal coil offset, Li et al. [3] derive a 6th-order equation for resonant frequencies that maximize output power (Freq-MPP) and efficiency (Freq-MEP), switching between them as a function of operating priority, validated by simulation and experiment.

Assumption: horizontal offset is the only perturbation variable; vertical separation and tilt are not independent variables in the model, so the method cannot distinguish lateral drift, vertical drift, or rotation — all of which co-occur in real deployments.

2.2. Coil-Geometry and Circuit Approaches to Misalignment Tolerance

A separate body of work pursues misalignment tolerance through coil geometry or reconfigurable circuit elements rather than retuning. Wan et al. [4] propose a dual-layer vertical double-D (DVDD) coil structure with a more constant coupling coefficient across combined offsets; Dai et al. [5] apply active magnetic-shielding adjustment to stabilize mutual inductance; Zhang et al. [6] and Yenil [7] use V-I control at the receiver to maintain constant current/voltage charging; and Song et al. [8] add variable capacitance.

In each case, a single physical or circuit degree of freedom is optimized against a coupling coefficient treated as a scalar function of overall misalignment severity, rather than jointly modelling and compensating the three nonlinearly coupled axes addressed here.

2.3. Robustness to Coupling Variation via Circuit Symmetry

A parity-time (PT) symmetric circuit design is used to obtain approximate constant efficiency over a range of coupling strength. A nonlinear PT-symmetric circuit can be shown to attain close to constant transfer efficiency with varying distance between coils without retuning the circuit's frequency and impedance [9], and an extended version of a similar nonlinear PT-symmetric model is extended to constant-efficiency WPT for a drone in-flight charging platform [10]. Some of these results indicate that for some circuit topologies, the coupling can be absorbed via passive means, in addition to the active approach used here; it is possible that this would lead to a lighter burden on the machine learning (ML) layer proposed in Section 3, which we also remark for future work.

2.4. Machine Learning for WPT Tuning

The impedance matching in WPT by means of neural network was done before. Jeong et al. [11] train a feedforward neural network to predict capacitor values in a tunable matching network directly from measured impedance and successfully obtain PTE of approximately 90% over a range of transmitter-receiver distances. Li et al. [12] similarly use a feedforward-backpropagation network for automatic impedance matching. The works both take into account distance variation along one effective axis, rather than misalignment in 3-DOF — and, cru-

cially, neither relies on an analytical physics model underneath and a learned residual on top of it, with the ML model being the full mapping rather than the addition to a physics-based estimate of that mapping. Previous non-learned adaptive matching methods like Beh et al. [13] and Park and Lee [14] demonstrate the general feasibility of magnetic resonance coupling for adaptive matching but do not include any learning component.

2.5. The Gap

The reported approaches (i) do not model and compensate for lateral, vertical, and angular misalignment simultaneously and in real time, as the assumptions of purely analytical or single axis-ML models are violated by the nonlinear cross-coupling between the misalignment axes, (ii) use a learned correction layer to capture the nonlinear cross-coupling between misalignment axes, or (iii) provide predictive retuning with low latency by leveraging a physics-informed hybrid architecture. A physics-informed residual learning network that only learns the residual to a physics-based estimate, not the solution itself, is an effective and general modelling approach for 2D electromagnetic forward modelling, as demonstrated by Shan et al. [15] — directly motivating the residual-correction modelling approach used in Section 3.

2.6. Comparative Positioning Relative to Prior Work

Table 1 positions the proposed method against the studies in Sections 2.1–2.4 by DOF handled, compensation mechanism, and learning component. Because the compared studies use different coil geometries and misalignment ranges, a directly comparable numeric PTE across rows is not meaningful; the table compares mechanism and scope, using this paper's own gap-closure metric (Table 4, Section 6.1) as its comparative anchor. Two points of novelty follow: no prior approach jointly compensates all three axes simultaneously, and the proposed residual layer learns only the non-separable cross-axis term left over by the physics backbone (Eq. (1)), rather than replacing the physical model outright as in [11, 12] — keeping the estimate interpretable outside the training envelope (Section 6.6).

A fully matched, apples-to-apples numeric benchmark against every listed method under an identical coil geometry and misalignment grid is left as future work, since it would require re-implementing each cited controller under our simulation conditions.

3. PROPOSED METHOD

3.1. System Overview

The proposed framework comprises three functional blocks that operate in a closed loop: (i) an analytical physics block to calculate a baseline impedance-matching target, based on an estimated misalignment state; (ii) a trained residual-correction block to predict the systematic error between the analytical backbone mutual-inductance estimate and the actual value, including nonlinear cross-axis coupling; and (iii) a retuning controller that takes both of the above and drives a tunable matching network.

TABLE 1. Qualitative comparison of the proposed method against representative prior approaches.

Approach	Misalignment DOF handled	Compensation mechanism	Learning component	Reported outcome (as stated by source)
Li et al. [3]	Horizontal offset only (1-DOF)	Frequency retuning (Freq-MPP ↔ Freq-MEP switching)	None	Efficiency/power maximized under horizontal offset only; vertical and tilt not modeled
Wan et al. [4]; Dai et al. [5]	Combined offset, treated as a single severity scalar	Coil-geometry redesign (DVDD)/active magnetic shielding	None	More constant coupling coefficient across offset range; axes not separated or jointly modeled
Zhang et al. [6]; Yenil [7]; Song et al. [8]	Combined offset, treated as a single severity scalar	Variable inductor/capacitor or V-I control at receiver	None	Constant current/voltage charging maintained under misalignment
Jeong et al. [11]; Li et al. [12]	One effective distance axis	Neural network predicts matching-network capacitor values directly (black box)	Full black-box ML, no physics backbone	≈ 90% PTE reported (Jeong et al.) over a one-axis distance range
Proposed method	Simultaneous 3-DOF (lateral Δx , vertical Δz , tilt $\Delta\theta$)	Physics backbone (Eqs. (1)–(2)) + learned residual → matching-network (R_L) retuning	Physics-informed residual only (FFN learns non-separable cross-term)	89% of the single-axis-baseline efficiency-deficit gap closed on the 3-DOF grid; mean tracking deficit reduced from 1.30 pp to 0.14 pp (Table 4)

3.2. Analytical Backbone: Multi-Axis Mutual Inductance Model

The misalignment is represented by lateral offset Δx , vertical change in separation Δz , and angular tilt $\Delta\theta$. In the baseline single-axis model of [3], the resonant frequencies of maximum efficiency and maximum power are expressed as a function of the horizontal offset only. The analytical approximation for mutual inductance is extended to a separable approximation along each of the three axes:

$$M_{\text{approx}}(\Delta x, \Delta z, \Delta\theta) = M_0 \cdot \frac{M(\Delta x, 0, 0)}{M_0} \cdot \frac{M(0, \Delta z, 0)}{M_0} \cdot \cos(\Delta\theta) \quad (1)$$

where M_0 is the mutual inductance at perfect alignment. This approximation omits the cross-coupling term $g(\Delta x, \Delta\theta)$ that represents how lateral offset and angular tilt jointly distort the flux-linkage path in a way not separable into independent single-axis contributions — the component the residual-correction layer is designed to absorb (Section 3.3).

Physically, Eq. (1) treats each misalignment axis as if it acted alone on the flux linkage: the two bracketed ratios rescale M_0 by the fractional coupling loss measured independently under pure lateral offset and pure vertical separation change, while $\cos(\Delta\theta)$ captures the projected-area loss under tilt. This product form is exact only when the three perturbations do not interact, which is why Figure 1 pairs it with the residual layer of Eq. (3), trained specifically to recover the $g(\Delta x, \Delta\theta)$ interaction term. Eq. (1) cannot represent — giving a deployed controller a physically grounded fallback if the learned correction is judged unreliable outside the training distribution, rather than a black-box mapping with no fallback. This fallback is a design

affordance of the architecture; the safeguard logic itself is not yet implemented (Section 6.6).

The series-series (S-S) resonant circuit gives the input impedance:

$$Z_{\text{in}} = R_1 + j\omega L_1 + \frac{1}{j\omega C_1} + \frac{\omega^2 M^2}{R_2 + j\omega L_2 + \frac{1}{j\omega C_2} + R_L} \quad (2)$$

Physically, Eq. (2) is the standard S-S reflected-impedance model: $R_1 + j\omega L_1 + 1/(j\omega C_1)$ is the transmitter's series-resonant branch, and the second term is the receiver's impedance reflected back through M , scaled by $\omega^2 M^2$ and divided by the receiver's own loop impedance (including R_L). Since $\omega^2 M^2$ multiplies the reflected term, misalignment-driven changes in M (Eq. (1) plus the residual of Eq. (3)) change the effective loading, so restoring efficiency requires re-optimizing some circuit element. Eq. (1) alone does not say whether frequency or R_L is the better lever — that is settled by the reactance-cancellation argument immediately below, which is why R_L , not carrier frequency, is identified as dominant for this topology.

Circuit-level finding: when the coil and component values used in this study (Section 5) are considered, sweeping the carrier frequency f alone (while keeping fixed the nominal resonance values of C_1 and C_2) gives a very small variation in the achievable PTE for a fixed value of M ; the reactance-cancellation condition is essentially frequency-independent of M for a symmetric resonant pair. The load/matching-network state R_L , on the other hand, has an efficiency-optimal value, which varies significantly as a function of M (as confirmed numerically in Section 5). The retuning problem, therefore, in this paper is not treated as a frequency search problem but rather as




FIGURE 1. Closed-loop control architecture. The misalignment-state estimate feeds the analytical backbone and the ML residual-correction layer in parallel; both outputs combine at the retuning controller, which drives the tunable matching network; impedance sensing closes the loop back to the state estimate.

an estimation of the $M(\Delta x, \Delta z, \Delta \theta)$ as closely as possible to the best possible matching-network state.

3.3. Residual Correction Layer

We do not use an end-to-end black-box approach but instead use a residual-learning approach, where a shallow feedforward network (FFN) is trained to learn only the systematic error between the analytical backbone’s estimate of the mutual-inductance and the true value:

$$\hat{M}(\Delta x, \Delta z, \Delta \theta) = M_{\text{approx}}(\Delta x, \Delta z, \Delta \theta) + \Delta f_{\text{ML}}(\Delta x, \Delta z, \Delta \theta) \quad (3)$$

Two architectures were compared as an ablation (Section 6.3): a degree-2 polynomial ridge regressor and a shallow FFN (one hidden layer, 16 units, tanh activation, L-BFGS solver). FFN model was chosen according to held-out test error (Section 6.3) as the correction layer that was deployed. Training data are created from a numerical sweep of the mutual inductance, as described in Section 5, and not from a physical prototype, which is the first validation of the architecture in a controlled numerical ground truth.

A corrected measured mutual-inductance value M_1 is used to reevaluate the matching-network state R_L that maximizes the prediction of PTE in accordance with the circuit relation used to derive the analytical backbone (Eq. (2)).

3.4. Retuning Controller and Sensing

The matched network is designed to be tunable due to the combined estimate, as in a varactor-tuned coil-link network, similar to that shown in [16] for a WPT system operating at 6.78 MHz. In this study, we do not measure the misalignment state $(\Delta x, \Delta z, \Delta \theta)$, but in a deployed system it would be estimated from input impedance and reflection-coefficient measurements, as in the primary-side, communication-free parameter-estimation approach described in [17] and without the need for auxiliary

position sensors. No hardware control loop has been evaluated in this document, and therefore no discussion of retuning latency will be included in this section.

4. PROBLEM STATEMENT

Given a transmitter-receiver coil pair operating at nominal resonant frequency ω_0 with baseline coupling coefficient k_0 , subject to simultaneous misalignment $\{\Delta x, \Delta z, \Delta \theta\}$, the goal is to determine the matching-network state $R_L^*(\Delta x, \Delta z, \Delta \theta)$ that maximizes instantaneous power transfer efficiency, where the mapping from misalignment to R_L^* is not expressible in closed form for arbitrary combinations of offset and tilt owing to the nonlinear cross-axis term in $M(\Delta x, \Delta z, \Delta \theta)$.

5. SIMULATION SETUP

This is the numerical simulation setup (Figure 2) used throughout the paper — no hardware was developed for this study (see Section 6.4).

As in semi-analytical models of lateral and angular coil misalignment shown to be valid in [18] and [19], the main results in this paper (Sections 5.1–6) are calculated by simply numerically evaluating Neumann’s double-line-integral formula for mutual inductance $M(\Delta x, \Delta z, \Delta \theta)$ between two circular coil loops, with the circuit-level PTE calculated from exact S-S resonant-circuit equations. This is a real numerical electromagnetic and circuit simulation, not a closed-form solution, and ignores second-order physical effects (proximity effect, current distribution, winding parasitics, shielding). A full-wave cross-validation in ANSYS High Frequency Structure Simulator (HFSS) was attempted (Section 6.7) and is reported honestly there for the time being — it succeeded in part for the mutual-coupling term but did not entirely give limitations of this version of the software, and full-wave validation of the complete




FIGURE 2. Simulation configuration. (a) the coil pair and the three misalignment axes (Δx , Δz , $\Delta \theta$) evaluated in this study. (b) the S-S resonant circuit used for all PTE calculations (Eq. (2)), including the mutual inductance M and the tunable load R_L^* driven by the retuning controller of Figure 1.

misalignment grid below would be a priority for follow-up work with a full (non-Student) licence.

5.1. Validation of the Numerical Mutual-Inductance Method

We tested the Neumann double-line-integral calculation for a full 3-DOF sweep using an independently derived, peer-reviewed closed form for the test. In the special case of zero angular tilt (lateral offset Δx and axial separation change Δz only), the mutual inductance between two parallel filamentary circular loops was derived semi-analytically by Babic et al. [18] (their Eq. (3), which is equivalent to Grover's classical result). We used this formula independently and tested it at 10 test points over the misalignment envelope employed in this paper (Δx 0–50 mm, Δz –20 to +20 mm), and calculated each point using both methods.

Wherever a closed-form solution exists, the numerical double-line integral agrees with the closed-form value [18] to within numerical-quadrature tolerance (differences less than $10^{-8}\%$) at all ten test points, confirming that the numerical implementation of the double-line integral is correct. This confirms the part of the mutual-inductance calculation in the paper that used Δx - Δz (lateral-axial). It does not, by itself, validate the angular-tilt dimension ($\Delta \theta$) or the combined 3-DOF case, for which, full-wave EM or hardware validation (Section 6.7) remains the next appropriate step, and this closed-form check is not a substitute for it.

On the near-zero differences in Table 3: these columns are not independent measurements that happen to agree — Babic et al.'s closed form [18] is a semi-analytical, elliptic-integral evaluation of the same Neumann double-line integral we evaluate numerically, a different computational route to the identical physical model, not a different model. The two routes should therefore agree to within the numerical quadrature's convergence tolerance, not to some fundamentally different level of accuracy. The observed $\sim 10^{-8}\%$ ($\sim 10^{-10}$ relative) difference matches a typical adaptive-quadrature tolerance, not pure floating-point round-off, which for double precision is several

orders smaller (machine epsilon $\approx 2.2 \times 10^{-16}$ relative); a percent-level difference, by contrast, would indicate a genuine implementation bug. This check therefore validates implementation correctness only — it is deliberately not a comparison against an independently derived alternative model, and Section 6 draws no accuracy claims from it beyond that scope.

5.2. Independent Implementation (MATLAB)

As a reproducibility check, the full pipeline (mutual inductance, S-S circuit PTE, matching-network search, residual-correction training, polynomial-ridge ablation) was independently re-implemented in base MATLAB — with the FFN and ridge regression coded from first principles (manual back-propagation with momentum; normal-equation ridge) — so the pipeline can be run and audited without any paid toolboxes.

The FFN hyperparameters (16 units, tanh, lr 0.05, momentum 0.9, 4000 epochs) were cross-checked by rerunning the identical algorithm in Python on the earlier, smaller calibration grid (± 30 mm/ ± 15 mm/ $\pm 20^\circ$): held-out root-mean-square error (RMSE) was ~ 0.28 nH in Python versus 0.40 nH for scikit-learn on the same grid, confirming that the architecture and training procedure are sound. Section 6's results use the larger grid (Table 2) and come from the Python implementation; the MATLAB scripts are provided and can be rerun on the larger grid to reproduce the same figures.

6. RESULTS

6.1. Main Results

Table 4 shows the mean PTE on the 133-point held-out test grid for four operating conditions: the single-axis reactive baseline (matching-network state derived from a misalignment estimate using only Δx , but not Δz and $\Delta \theta$ at all); the analytical backbone alone (without any residual correction); the proposed method (analytical backbone with FFN residual correction); and the true optimum (analytical backbone with numerical search of R_L against the ground-truth M at each grid point).

TABLE 2. Simulation configuration.

Parameter	Value
Coil geometry	3-turn planar spiral, 20 mm radius (40 mm diameter)
Nominal separation D_0	50 mm
Operating frequency f_0	6.78 MHz (ISM band, consistent with [16, 20])
Coil series resistance $R_1 = R_2$	0.10Ω ($Q \approx 363$)
Coupling coefficient k_0	0.0181
Figure of merit $k_0 Q$	6.57
Misalignment grid	$\Delta x \in [-50, 50]$ mm (9 pts), $\Delta z \in [-20, 20]$ mm (7 pts), $\Delta \theta \in [-35, 35]^\circ$ (7 pts); 441 points total
Train/test split	308/133 points (70/30, random)

TABLE 3. Neumann-integral vs. closed-form (Babic et al. [18]) mutual inductance.

Δx (mm)	Δz (mm)	M , Neumann integral (nH)	M , closed form (nH)	Diff.
0	0	15.423856	15.423856	0%
0	15	8.076142	8.076142	0%
0	-15	33.739669	33.739669	0%
30	0	8.210335	8.210335	0%
30	15	5.191844	5.191844	0%
50	0	2.690415	2.690415	0%
50	-20	0.621158	0.621158	0%

TABLE 4. Mean PTE and tracking-deficit statistics across the extended 3-DOF test grid ($n = 133$).

Metric	Single-axis baseline	Analytical only	Proposed	True optimum
Mean PTE, aligned	73.84%	—	—	73.84%
Mean PTE, 3-DOF grid	44.89%	45.52%	46.05%	46.19%
Min PTE, test grid	0.16%	—	0.17%	—
Mean tracking deficit from optimum	1.30 pp	0.67 pp	0.14 pp	0
Max tracking deficit	6.95 pp	—	9.11 pp	0
Std. dev. of tracking deficit	1.57 pp	—	0.84 pp	0

The raw PTE improvement across the full grid is modest (44.89% to 46.05%, +1.16 pp) since this coil geometry’s achievable efficiency is naturally lower toward the envelope edges. Tracking accuracy relative to the achievable ceiling is more informative: extending the analytical backbone to all three axes with no ML closes only 49% of the single-axis baseline’s efficiency-deficit gap, while adding the trained residual-correction layer closes 89% of that gap (Figure 3). Section 6.5 explains why the ML layer contributes so much more here than for a small, near-aligned envelope: the nonlinear lateral-tilt cross-term that the analytical model omits grows sharply, in relative terms, as misalignment increases.

6.2. Ablation: Residual-Correction Architecture

As required by the model-selection question raised in Section 3.3, two architectures were compared for the residual-correction layer, evaluated on held-out mutual-inductance prediction error, with results summarized in Table 5.

The residual signal is, as expected, larger in absolute value and much larger as a percentage of the analytical estimate of the backbone than on the original, smaller grid, and has a standard deviation of 3.91 nH (37% of the mean analytical-backbone estimate), definitely a significant signal and not numerical noise. The FFN is able to capture more of this variance than the polynomial baseline, which is consistent with the residual being a smooth but non-polynomial function of the $(\Delta x, \Delta z, \Delta \theta)$ parameters involving both $\cos(\Delta \theta)$ and elliptic-integral-like terms from the underlying calculation of the mutual inductance.

6.3. Retuning Latency (Engineering Estimate, Not Simulated)

This study did not construct or test a hardware control loop. Based on representative component datasheets, an order-of-magnitude estimate for the time needed to settle the voltage into the range of accuracy required for impedance sensing is: impedance-sensing settling ~ 200 ns; FFN forward pass (one hidden layer, 16 units) ~ 100 ns at typical embedded clock




FIGURE 3. (a) mean PTE across the 3-DOF test grid for each method and (b) mean tracking deficit from the true optimum, single-axis baseline vs. proposed method.

TABLE 5. Residual-correction model ablation.

Model	Test RMSE (nH)	5-fold CV RMSE (nH)
Degree-2 polynomial ridge	2.393	3.810 ± 0.691
Shallow FFN (16 units, tanh, L-BFGS) — selected	1.072	—

rates; varactor or PIN-diode matching-network settling ~ 300 – 500 ns per component; giving an estimated settling time of 0.6 – 0.8 μ s, which is significantly less than the required 1 μ s. If a paper includes a latency claim for the hardware implementation of a specific control loop, then this value should be replaced by a measured value in the paper.

6.4. Why the ML Correction Matters More at Larger Misalignment

A direct calculation from the mutual-inductance model explains why widening the misalignment envelope beyond the smaller ± 30 mm/ ± 15 mm/ $\pm 20^\circ$ development grid sharpened the case for the residual-correction layer. The nonlinear lateral-tilt cross-term omitted by the separable analytical model grows from $\sim 18\%$ ($\Delta x = 30$ mm, $\Delta \theta = 20^\circ$) to $\sim 79\%$ of the true coupling at the edge of the extended envelope ($\Delta x = 60$ mm, $\Delta \theta = 50^\circ$): the further the coil pair drifts from alignment, the more the separable approximation of Eq. (1) misrepresents the true coupling, and the more a learned correction has to contribute. This matches Table 4: the analytical extension alone closes a smaller share of the gap on the larger grid (49%) than on the smaller development grid (86%), while the ML-corrected method still closes most of the gap (89%) despite the harder envelope.

6.5. Caveat: Worst-Case Behaviour at the Envelope Edge

The mean-deficit improvement is not a worst-case guarantee. At one sparsely-sampled corner of the extended test grid, the ML-corrected method's tracking deficit (9.11 pp) exceeded the single-axis baseline's worst case (6.95 pp) — the trained FFN evidently extrapolated poorly there despite improving the mean and closing most of the gap overall. This is a legitimate limita-

tion, not a reason to discard the approach: it indicates that the training grid should be denser near deployment-relevant envelope edges and that a deployed controller should include a simple safeguard (e.g., reverting to the analytical-only estimate, or clipping the FFN correction) when the input state falls outside the well-covered training region.

6.6. Attempted Full-Wave Cross-Validation (Partial, Inconclusive)

A full-wave cross-validation of the aligned-case mutual inductance was attempted in ANSYS HFSS (Student Edition), modelling transmitter and receiver as single-turn circular wire loops with lumped-port excitation (Figure 4). The extracted mutual-coupling term ($\text{Im}(Z_{12}) \approx 0.66$ – 0.72 Ω) was stable and of the correct sign but roughly an order of magnitude larger than the analytical value ($\text{Im}(Z_{12}) \approx 0.073$ Ω). The self-impedance terms ($\text{Im}(Z_{11})$, $\text{Im}(Z_{22})$) never converged to physically plausible values under any port, meshing, or excitation configuration tried — consistently returning large negative reactances, the signature of an unintended capacitive-gap artifact rather than the small positive loop reactance expected — so no $L_1/L_2/M$ values are reported from this simulation. We report the attempt rather than omit it: it points to lumped-port de-embedding sensitivity for electrically small, solid 3D wire loops as a concrete obstacle for follow-up work with a full (non-Student) license, a wave-port or coaxial-feed excitation better suited to closed loops, and denser port-gap meshing than the Student license allows.

7. DISCUSSION

The results support four claims at different confidence levels. With high confidence: jointly estimating the 3-DOF misalign-




FIGURE 4. The ANSYS HFSS coil-pair model built for the full-wave cross-validation attempt: single-turn Tx_coil and Rx_coil (20 mm radius, 50mm nominal separation, aligned case shown), each with a lumped-port excitation bridging a small gap cut into the loop.

ment state and re-deriving the matching-network state closes most of the tracking-error gap left by single-axis compensation (Table 4), with the improvement growing as the misalignment envelope widens (Section 6.5). With moderate confidence: the residual layer meaningfully outperforms a low-order polynomial baseline, indicating the cross-axis coupling term is genuinely nonpolynomial rather than a simple quadratic correction. As an honest caveat, not a claim: the mean-case improvement does not guarantee the worst case, and a deployed system needs a safeguard against FFN extrapolation error at sparsely-sampled envelope edges (Section 6.6). With lower confidence, pending full-wave EM or hardware validation: the retuning-latency figure should be treated as provisional.

8. CONCLUSION

This paper has presented a hybrid physics-machine learning framework for joint 3-DOF misalignment compensation in MCR-WPT, combining an analytical mutual-inductance circuit backbone with a trained residual-correction layer. Over a realistic ± 50 mm lateral/ ± 20 mm vertical/ $\pm 35^\circ$ tilt envelope, the method tracks the true efficiency optimum far more closely than a single-axis baseline, closing 89% of the mean tracking-error gap versus 49% for the analytical extension alone — because the nonlinear lateral-tilt cross-term that the analytical model omits grows from $\sim 18\%$ to $\sim 79\%$ of the true coupling as misalignment increases, so the learned correction's value grows with disturbance size. We also report a circuit-level finding relevant to related work: matching-network retuning, not carrier-frequency retuning, is the dominant lever in this topology. A genuine limitation is reported rather than omitted: the correction's worst-case deficit at one sparsely-sampled envelope corner exceeded the baseline's, indicating a need for denser edge-of-envelope training data or an extrapolation safeguard in any

deployed controller. A toolbox-free MATLAB port of the full pipeline is provided alongside the Python implementation (Section 5.2) for independent verification. The main remaining limitations are incomplete full-wave EM validation (Section 6.7) and absence of hardware validation, both identified as immediate next steps.

REFERENCES

- [1] Kurs, A., A. Karalis, R. Moffatt, J. D. Joannopoulos, P. Fisher, and M. Soljačić, "Wireless power transfer via strongly coupled magnetic resonances," *Science*, Vol. 317, No. 5834, 83–86, 2007.
- [2] Karalis, A., J. D. Joannopoulos, and M. Soljačić, "Efficient wireless non-radiative mid-range energy transfer," *Annals of Physics*, Vol. 323, No. 1, 34–48, 2008.
- [3] Li, J., J. Zou, Y. Zou, and Z. Li, "Study on frequency splitting and segmented tracking for high-performance wpt under horizontal offset," *Progress In Electromagnetics Research C*, Vol. 170, 57–65, 2026.
- [4] Wan, X., J. Jiang, and B. Liu, "Optimization study on misalignment tolerance of a wireless power transfer system with double-layer vertical DD coils," *Wireless Power Transfer*, Vol. 12, No. 1, e025, 2025.
- [5] Dai, Z., M. Li, H. Xu, M. Ji, and L. Zhang, "Strong misalignment tolerance wireless power transfer with active adjustment of magnetic shielding," *AIP Advances*, Vol. 14, No. 1, 015320, 2024.
- [6] Zhang, H., Y. Chen, D.-H. Kim, Z. Li, M. Zhang, and G. Li, "Variable inductor control for misalignment tolerance and constant current/voltage charging in inductive power transfer system," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, Vol. 11, No. 4, 4563–4573, 2023.
- [7] Yenil, V., "A wireless power transfer system with high misalignment tolerance based on dual side variable inductor control," *IEEE Access*, Vol. 12, 134 937–134 947, 2024.
- [8] Song, K., Y. Lan, R. Wei, G. Yang, F. Yang, W. Li, J. Jiang, C. Zhu, and Y. Li, "A control strategy for wireless EV charging system to improve weak coupling output based on variable in-

- ductor and capacitor,” *IEEE Transactions on Power Electronics*, Vol. 37, No. 10, 12 853–12 864, 2022.
- [9] Assawaworrarit, S., X. Yu, and S. Fan, “Robust wireless power transfer using a nonlinear parity-time-symmetric circuit,” *Nature*, Vol. 546, No. 7658, 387–390, 2017.
- [10] Zhou, J., B. Zhang, W. Xiao, D. Qiu, and Y. Chen, “Nonlinear parity-time-symmetric model for constant efficiency wireless power transfer: Application to a drone-in-flight wireless charging platform,” *IEEE Transactions on Industrial Electronics*, Vol. 66, No. 5, 4097–4107, 2019.
- [11] Jeong, S., T.-H. Lin, and M. M. Tentzeris, “Range-adaptive impedance matching of wireless power transfer system using a machine learning strategy based on neural networks,” in *2019 IEEE MTT-S International Microwave Symposium (IMS)*, 1423–1425, Boston, MA, USA, 2019.
- [12] Li, Y., W. Dong, Q. Yang, J. Zhao, L. Liu, and S. Feng, “An automatic impedance matching method based on the feedforward-backpropagation neural network for a WPT system,” *IEEE Transactions on Industrial Electronics*, Vol. 66, No. 5, 3963–3972, 2019.
- [13] Beh, T. C., M. Kato, T. Imura, S. Oh, and Y. Hori, “Automated impedance matching system for robust wireless power transfer via magnetic resonance coupling,” *IEEE Transactions on Industrial Electronics*, Vol. 60, No. 9, 3689–3698, 2013.
- [14] Park, B.-C. and J.-H. Lee, “Adaptive impedance matching of wireless power transmission using multi-loop feed with single operating frequency,” *IEEE Transactions on Antennas and Propagation*, Vol. 62, No. 5, 2851–2856, 2014.
- [15] Shan, T., J. Zeng, X. Song, R. Guo, M. Li, F. Yang, and S. Xu, “Physics-informed supervised residual learning for electromagnetic modeling,” *IEEE Transactions on Antennas and Propagation*, Vol. 71, No. 4, 3393–3407, 2023.
- [16] Lim, T., J. Lee, and Y. Lee, “Tunable coil-links for multiple-receiver wireless power transfer system with arbitrary position and power division,” *IEEE Access*, Vol. 10, 96 862–96 872, 2022.
- [17] Xiao, W., R. Shen, B. Zhang, D. Qiu, Y. Chen, and F. Xie, “Multiple parameters estimation based on transmitter side information in wireless power transfer system,” *IEEE Access*, Vol. 7, 164 835–164 843, 2019.
- [18] Babic, S., F. Sirois, and C. Akyel, “Validity check of mutual inductance formulas for circular filaments with lateral and angular misalignments,” *Progress In Electromagnetics Research M*, Vol. 8, 15–26, 2009.
- [19] Zhang, X., H. Meng, B. Wei, S. Wang, and Q. Yang, “Mutual inductance calculation for coils with misalignment in wireless power transfer,” *The Journal of Engineering*, Vol. 2019, No. 16, 1041–1044, 2019.
- [20] Liu, Y., N. Ha-Van, P. Jayathurathnage, J. Kyryä, and S. A. Tretyakov, “Design and performance analysis of a repeater-assisted multireceiver wireless power transfer system operating at MHz frequencies,” *IEEE Open Journal of the Industrial Electronics Society*, Vol. 4, 773–785, 2024.