Camera–magnetic fusion for lateral coil alignment in dynamic wireless charging on roads with degraded lane markings

A simulation-based methods study

[1]Pravat Satpathy, [2]Sonia Mohapatra, [3]Saswat Swain
[1]Researcher, [2]Researcher, [3]Ph.D Scholar
[1,2]Research and Development Division,
Tekstein Solutions Private Limited, Bhubaneswar, Odisha, India
[3]KSOM — KIIT Deemed to be University,
Bhubaneswar, Odisha, India
[1]analytics.pravat@gmail.com, [2]smctc98@gmail.com, [3]saswat.mck@gmail.com

Abstract*—Dynamic wireless charging of electric vehicles loses coupling rapidly as the receiver drifts sideways over the road pads, so the transmitter needs a reliable estimate of lateral offset. Camera methods depend on visible road markings, which are often worn, interrupted or absent on Indian roads, while magnetic methods work only over an energised pad. This paper develops and evaluates, in simulation, an estimator that combines both. A forward camera reads the lane centre, a two-axis coil at the receiver measures the in-plane angle of the pad field, and a gyro bridges the gaps, all within a four-state extended Kalman filter. The field angle is shown to vary monotonically with offset over the full 70 cm range examined and to be insensitive to transmitter current. The filter also estimates the unknown offset between painted lane and buried pads, which is unobservable from the camera alone. The estimate drives feedforward inverter control with a slow feedback trim. Camera measurements come from rendered road images processed by a real lane detector under five conditions, from fresh markings to near-unmarked surfaces, over 150 runs at 30 to 90 km per hour. The fused estimator reached 2.8 cm root-mean-square lateral error overall and 0.9 cm over powered sections, against 5.8 cm for magnetic sensing alone and 17.0 cm for the camera alone. Power-tracking error fell from 12.0 percent without sensing to 4.6 percent. Fusion gave no benefit on unmarked roads, outlier gating caused occasional lock-out on patchy roads, and the sensitivity analysis attributes most of the remaining power error to the unknown air gap. Hardware validation and air-gap estimation are the next steps.

***Index Terms*—Dynamic wireless power transfer, coil misalignment, sensor fusion, extended Kalman filter, lane detection, magnetic field sensing, electric vehicles, power control.**

I. Introduction

Dynamic wireless power transfer (DWPT) supplies energy to an electric vehicle (EV) while it moves over a sequence of transmitter pads buried in the carriageway. The appeal is a smaller on-board battery and less time spent at chargers [1], [2]. The practical difficulty is that the magnetic coupling between a buried pad and the receiver under the vehicle depends strongly on the lateral offset between the two, and a moving vehicle never holds that offset constant. A driver following a lane wanders laterally by tens of centimetres, while typical light-duty pads are only about half a metre wide. As the offset grows, coupling falls, inverter current rises for the same delivered power, and both efficiency and component stress move in the wrong direction [3], [4].

Two engineering responses exist. The first widens the tolerance of the magnetic structure or compensation network, which costs ferrite, copper and pad area [4], [5]. The second measures the misalignment and uses that measurement, either to regulate power at the transmitter or to guide the vehicle laterally. This paper is concerned with the second route. Published sensing options include camera detection of a ground guideline painted over the pads [6], phase-based radio-frequency identification (RFID) ranging with sub-decimetre accuracy [7], auxiliary sensing coils used to close a steering loop [8], and arrays of magnetoresistive sensors reading the transmitter field directly [9].

Each option leans on an assumption that is fragile on many Indian roads. Vision methods assume a visible, well-maintained reference line; Indian Roads Congress practice specifies such markings [10], but on arterial and urban roads paint is often worn, interrupted by repairs, overlaid by older restriping, or hidden by dust and water. Magnetic methods assume the pad is energised, which in a segmented system happens only over powered sections and, inside them, only near each pad's centre. RFID requires dedicated tags and a known vehicle trajectory to resolve phase ambiguity [7]. No single sensor therefore covers the whole route, including the stretch before a powered section, where alignment matters most because the first pads must be energised and power-limited before any field is available.

Commercial DWPT has begun to leave the test track. A 1.5 km motorway section in France has been operated in live traffic with heavy and light vehicles [11], and readiness reviews of wireless electric road systems describe pilot programmes in several countries [12]. India has no comparable deployment, and Indian road conditions differ from those pilots in exactly the way that weakens vision-only alignment. For that reason the present work is framed as a methods study: it develops and evaluates, in a physically grounded simulation, an alignment

estimator intended for roads where lane markings cannot be trusted. It does not claim field readiness.

The approach fuses a forward camera, which reads lane paint ahead of the vehicle, with a two-axis magnetic sensor at the receiver centre, which reads the angle of the active pad's field. A yaw-rate gyro bridges the gaps between them. The resulting estimate sets the inverter voltage of each pad by feedforward, and a slow feedback trim removes the residual error that lateral sensing cannot see. The contributions are as follows.

1) A magnetic lateral observable, the in-plane field angle $\theta = \mathrm{atan2}(B_x, B_z)$ at the receiver centre, shown by Biot–Savart analysis to be monotone over ±0.70 m of offset, independent of transmitter current, and insensitive to the receiver's own field by symmetry. A two-sensor difference ratio, the more obvious design, is shown to saturate within ±0.19 m and is rejected.
2) A four-state extended Kalman filter (EKF) that carries the unknown offset between painted lane centre and pad centreline as an explicit state. Proposition 2 shows this offset is unobservable from the camera alone and becomes observable whenever a magnetic update occurs, so the system calibrates its own visual reference while passing over powered sections and carries that calibration through unpowered ones.
3) An end-to-end simulation in which camera measurements come from rendered images processed by a real lane detector across five road conditions, rather than from an assumed noise model, combined with closed-form coupling and field calculations validated against direct numerical integration to 10^{-5} relative error.
4) A sensitivity result (Proposition 1) explaining why lateral feedforward alone cannot hold transferred power within ±10%: an unknown ±2 cm air-gap change moves coupling by about 10%, so a slow feedback trim is still required. The controller that combines both is evaluated against feedforward-only and feedback-only alternatives.
5) Ablations and negative results reported in full, including conditions where fusion gives no benefit over a single sensor.

Section II reviews related work. Section III states the problem and assumptions. Sections IV and V give the system architecture and mathematical model, and Section VI the algorithm. Sections VII to IX present the simulation set-up, results and comparison with prior work. Sections X and XI summarise outcomes and limitations.

II. Related Work

Table I summarises the work closest to this study. Quantitative claims are given only where the cited source states them; where a source reports qualitative results, the table says so rather than inferring a number.

TABLE I
Prior Alignment-Sensing Approaches for DWPT

Ref., year	Sensing principle	Reported outcome	Main limitation for degraded roads
[6], 2020	Camera detects a ground guideline over the pads; misalignment mapped to output-power regulation	Abstract reports accurate misalignment acquisition and faster regulation than discrete PI; no figure stated	Needs a dedicated, visible guideline; no field-based cross-check
[7], 2023	Passive UHF RFID phase; maximum-likelihood search along a known trajectory	3.5 cm (laboratory) and about 5 cm (field), single runs; sub-10 cm stated as potential	Tags per section; trajectory and tag-side must be known
[8], 2025	Auxiliary coils sense misalignment; steering controller corrects it	Alignment held within ±5 (unit not stated in abstract); efficiency drop below 5%; robot-scale test	Requires an energised pad; preliminary
[9], 2018	Array of 15 uniaxial TMR sensors reading the transmitter field	Validated by 3-D FEM and experiment over a large range; no figure in abstract	Requires an energised pad; sensor-array cost
This work	Camera lane reading, 2-axis field-angle sensor and gyro fused in a 4-state EKF	Simulation only (Section VIII)	Not yet validated in hardware; see Section XI

PI: proportional–integral. TMR: tunnelling magnetoresistive. FEM: finite-element method. Entries for [6], [8] and [9] are taken from the published abstracts; the full texts were not available to the authors and should be checked before submission.

Lane-marking detection itself is mature. Classical pipelines use inverse perspective mapping (IPM), a filter matched to marking width, and robust line fitting such as random sample consensus (RANSAC) [13]–[15]; learned detectors improve robustness on benchmark data [16]. Neither class is designed to report a reference that is accurate to a few centimetres relative to a buried coil, and both fail when the paint is absent. On the estimation side, delayed camera measurements, outlier gating and covariance adaptation are standard tools [17]–[19]. The gap this paper addresses is therefore not any single component. It is the combination: a reference that remains usable when paint fails, a magnetic observable that does not depend on transmitter current, and an explicit treatment of the offset between what the camera sees and where the coil actually is.

III. Problem Statement

Let $e(t)$ denote the lateral offset of the receiver-coil centre from the centreline of the pad row, positive to the right, and $\psi(t)$ the heading of the vehicle relative to the road. The task is to estimate $e(t)$ continuously along a route that alternates between powered and unpowered sections, and to use the estimate to set the voltage of whichever pad is active so that the transferred power tracks a reference P^*. Four constraints make the problem non-trivial.

- **Unreliable visual reference.** The camera sees lane paint, not coils. Paint may be worn, missing for metres at a time, accompanied by older restriping, or hidden at night by glare. Even when paint is clean, its centre need not coincide with the pad centreline; the offset c between them is unknown to the vehicle.
- **Intermittent magnetic reference.** The field exists only over powered sections and is usable only near the centre of each pad (Section V-B). Before the first pad of a section, no field is available, yet that pad must be energised with a safe voltage.
- **Unmeasured air gap.** Vehicle load and suspension change the air gap by a few centimetres. Coupling depends on gap

as strongly as on lateral offset at small offsets (Proposition 1), and neither sensor measures gap directly in this design.

- **Timing.** At 90 km/h the receiver crosses a 1.25 m pad pitch in 50 ms. A feedback loop that relies on the vehicle reporting received power over a wireless link with 20 ms latency gets about two samples per pad, and about three at 60 km/h.

The following assumptions hold throughout. The road is straight within each analysed section. Pads are identical, equally spaced and energised one at a time. The longitudinal position of the receiver relative to the active pad is known to within a few centimetres from odometry and pad-entry detection. The camera is calibrated at nominal height and pitch, but suffers vibration and load-dependent height change during operation. All results are simulation results; no hardware validation is claimed.

IV. System Architecture

Fig. 1 shows the arrangement. On the vehicle, a forward camera feeds a lane detector that reports the lane-centre line in the vehicle frame together with a quality flag and a detection score. A two-axis sensing coil at the receiver centre measures the horizontal and vertical components of the active pad's field; quadrature demodulation and a pad-core gate turn these into an angle measurement. A gyro and wheel odometry provide yaw rate and speed. All three streams enter an EKF that outputs the lateral offset and its standard deviation. A coupling map converts the estimate into a predicted mutual inductance, which is sent to the roadside, where each pad inverter applies a feedforward voltage and a slow trim based on the power the vehicle reports.

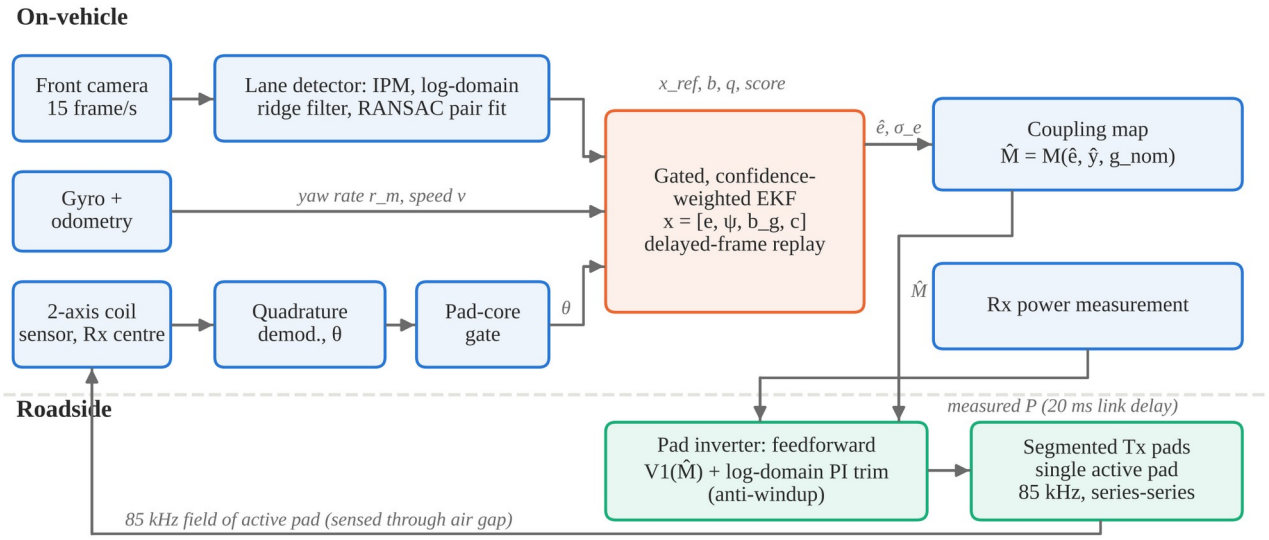


Fig. 1. System architecture. The vehicle fuses camera, magnetic and inertial measurements to estimate lateral offset e; the roadside inverter uses the predicted coupling for feedforward voltage and the reported power for a slow trim. The magnetic channel exists only while a pad is energised and the receiver is over its core region.

A. Transmitter and Receiver

The transmitter pads are rectangular, 0.50 m wide (lateral) and 1.00 m long, with eight turns, repeated at a 1.25 m pitch. The receiver is 0.50 m by 0.45 m with ten turns, at a nominal air gap of 0.15 m. Both sides use series–series (SS) compensation at 85 kHz, the nominal frequency of SAE J2954 [20]. The computed self-inductances are 85.2 µH and 55.7 µH, and the aligned mutual inductance is 12.98 µH (coupling factor 0.188). For an 11 kW target into an 8 Ω equivalent load, the aligned operating point requires 264 V and 43.4 A on the transmitter, with a link efficiency of 96.1% (coil resistances only). The inverter is limited to 600 V and 60 A.

B. Magnetic Sensor Placement

The obvious magnetic design places two vertical-axis sensors symmetrically about the receiver centre and forms the normalised difference of their readings. Table II shows why that design was abandoned. Once one sensor crosses the transmitter conductor, the vertical field it sees changes sign, and the normalised ratio saturates at ±1. Across the six spacings tested (10–35 cm) the monotone range never exceeded ±19 cm. The field angle at the receiver centre, by contrast, is monotone over the full ±70 cm examined, at every gap and longitudinal position tested, with a central slope of 2.3°/cm. Because the receiver coil is symmetric about its own centre, its field there has no horizontal component, which removes the largest interference source from B_x by geometry rather than by calibration.

TABLE II
Candidate Magnetic Observables (Gap 15 cm, y = 0)

Observable	Monotone range (cm)	Central slope
Difference ratio, sensors ±10 cm	±19	2.5 per m
Difference ratio, sensors ±15 cm	±14	6.4 per m
Difference ratio, sensors ±20 cm	±9	14.0 per m
Difference ratio, sensors ±25 cm	±4	35.4 per m
Difference ratio, sensors ±30 cm	±2	52.1 per m
Difference ratio, sensors ±35 cm	±7	6.0 per m
Field angle at receiver centre	±70 (whole grid)	2.3°/cm

Monotone range: offset interval around zero over which the observable changes strictly monotonically; for the difference ratio, slopes below 5% of the central slope are also excluded. The angle was evaluated only to ±70 cm, the edge of the grid.

V. Mathematical Model

A. Coupling

Each coil is represented by concentric rectangular filaments, one per turn, stepping inward by 12 mm. For two pads in

parallel planes, every side is either parallel or perpendicular to every other, and perpendicular pairs contribute nothing to the Neumann integral [21]. For a parallel pair of lengths l and m, perpendicular separation d, and start offset s of the second filament along the first,

$$M_{ij} = (\mu_0/4\pi)\,[G(s+m) - G(s+m-l) - G(s) + G(s-l)] \quad (1)$$

where $G(u) = u\,\sinh^{-1}(u/d) - (u^2 + d^2)^{1/2}$ and the sign of M_{ij} follows the relative current direction. The total mutual inductance $M(x, y, g)$ is the sum over all parallel pairs, where x, y and g are the lateral offset, longitudinal offset and gap. This closed form agreed with direct numerical integration of the Neumann integral to a relative error below 9×10^{-6} at three test points, so the lookup tables built from it carry no meaningful discretisation error. Fig. 2(a) shows the result. Coupling falls to half its peak near ±21 cm and crosses zero near ±38 cm, after which the flux returning outside the transmitter reverses its sign. The curves for different gaps intersect near ±27 cm, so the sign of the gap sensitivity reverses there.

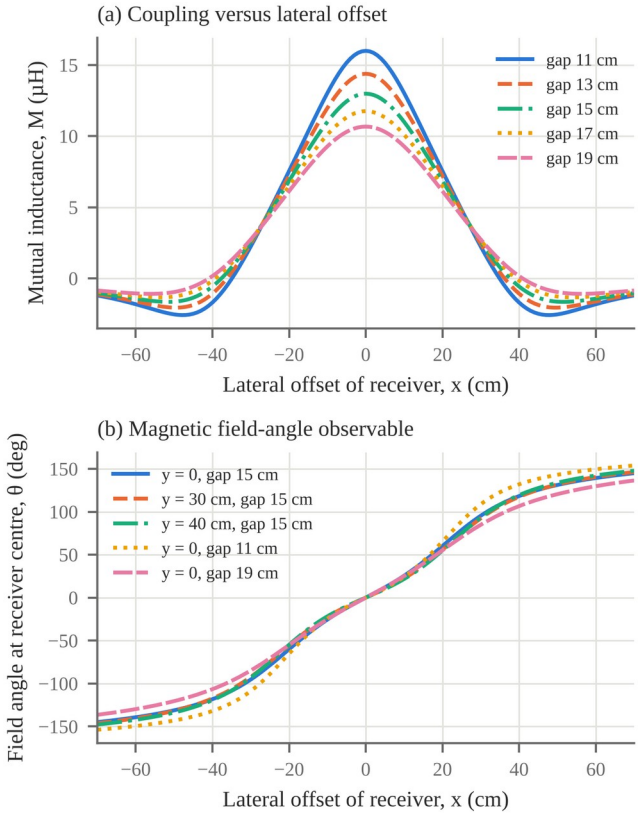


Fig. 2. (a) Mutual inductance versus lateral offset for five air gaps, from the closed-form Neumann model. (b) In-plane field angle at the receiver centre versus lateral offset for different longitudinal positions y and gaps. The angle is monotone over the full range; gap and longitudinal position perturb it by a few degrees.

B. Field-Angle Observable

The field of a straight filament segment at a point at perpendicular distance ρ, with end-point angles α_1 and α_2, is

$$\mathbf{B} = (\mu_0 I / 4\pi\rho)(\cos\alpha_1 - \cos\alpha_2)\,\hat{\boldsymbol{\varphi}} \quad (2)$$

and the pad field is the vector sum over its segments. The measured angle is

$$\theta = \operatorname{atan2}(B_x, B_z'), \quad B_z' = B_z + \sin(\delta\varphi)\,B_{rx}\,I_2/I_1 \quad (3)$$

The ratio removes the transmitter current. The second term models leakage of the receiver's own vertical field B_{rx} through a demodulation phase error $\delta\varphi$. In an SS link at resonance the receiver current I_2 lags the transmitter current by 90°, so quadrature demodulation rejects it ideally; a residual phase error of 1.5° (one standard deviation) is assumed. Because $I_2/I_1 = \omega M/(R_2+R_L)$, the leakage term depends on coupling but not on inverter voltage.

Beyond about 0.45 m from the pad centre, both in-plane components fall towards zero and the angle loses meaning. Magnetic updates are therefore accepted only when the estimated longitudinal position satisfies $|\hat{y}| < 0.42$ m, which leaves the magnetic channel available for 67% of the distance over a powered section. The error of the angle measurement, expressed as lateral error, follows from first-order propagation

$$\sigma_x \approx \sigma_\theta / |\partial h/\partial x| \quad (4)$$

where $h(x, y)$ is the nominal angle map at 15 cm gap. Table III lists the error budget obtained from 4000 random cases with gap varying uniformly by ±2 cm, odometry error of 3 cm, gain mismatch of 1% between axes, leakage as above, field noise of 0.5 µT, and a per-pad field-map bias of 1°. Near the centre the error is a few millimetres. It grows to 4.5 cm beyond 50 cm offset, where the angle curve flattens. Unknown air gap is the largest single contributor (1.9 cm RMS over the range), followed by receiver-field leakage (1.3 cm).

TABLE III
Magnetic Lateral-Error Budget (Pad Core, 4000 Cases)

Offset band (cm)	σ_θ (deg)	σ_x (cm)	Bias (cm)
0–10	0.84	0.33	-0.06
10–20	1.93	0.56	-0.02
20–30	3.16	0.90	0.02
30–40	3.64	1.60	0.04
40–50	3.76	2.87	0.24
50–60	3.65	4.47	-0.02

Single-source RMS lateral error over $|x| < 60$ cm: gap 1.91 cm; leakage 1.27 cm; odometry 0.30 cm; noise 0.26 cm; axis mismatch 0.18 cm. The per-pad 1° bias is added in the time-domain simulation, not in this table.

C. Power Transfer and Feedforward Control

For an SS link at resonance with coil resistances R_1, R_2 and load R_L, the reflected impedance and transferred power are

$$Z_r = \omega^2 M^2/(R_2+R_L), \quad P = R_L[\omega M\,V_1 / ((R_1+Z_r)(R_2+R_L))]^2 \quad (5)$$

Inverting (5) for a target P^* with an estimated coupling $\hat{M}$ gives the feedforward voltage

$$V_1 = (P^*/R_L)^{1/2}\,(R_1+\hat{Z}_r)(R_2+R_L)\,/\,(\omega\hat{M}) \quad (6)$$

Proposition 1 (*power sensitivity to coupling error*). Under (6), the delivered power is $P/P^* = (M/\hat{M})^2[(R_1+\hat{Z}_r)/(R_1+Z_r)]^2$. When $\omega^2M^2 \gg R_1(R_2+R_L)$, this reduces to $P/P^* \approx (\hat{M}/M)^2$, so that

$$\delta P/P \approx 2(\partial \ln M/\partial x)\,\delta x + 2(\partial \ln M/\partial g)\,\delta g \quad (7)$$

Here $\delta x = \hat{x} - x$ and $\delta g = \hat{g} - g$ are the estimation errors behind $\hat{M}$. *Proof.* Substitute (6) into (5); the factors (R_2+R_L) and R_L cancel. In the stated regime $R_1+Z_r \approx Z_r \propto M^2$, giving the reduced form, and (7) is its first-order expansion. ∎

Two consequences matter. First, in the stated regime an over-estimate of coupling gives over-delivery, which in an SS link means excess receiver current, so under-estimating is the safe direction. Outside that regime the factor $(M/\hat{M})^2$ dominates and the sign reverses, so the rule must not be applied at large offsets. Second, the gap term is not small. At the design point $\partial \ln M/\partial g = -5.0\ \mathrm{m}^{-1}$, so a 2 cm gap error alone moves power by about 20%, while $\partial \ln M/\partial x$ is zero at the centre and $-6.5\ \mathrm{m}^{-1}$ at 20 cm.

The proposition was checked against the exact expression: at x = 0.2 m with a 1 cm lateral error the exact ratio is 0.896 against 0.875 from the approximation. The approximation fails at large offset; at x = 0.3 m with a 2 cm lateral error, where $\omega^2 M^2$ is comparable to $R_1(R_2+R_L)$, the exact value is 0.968 against 0.495. In that regime the transmitter current limit binds in any case: at 20 cm offset only 5.8 kW can be delivered within 60 A.

The gap term motivates a slow trim. Every 10 ms the inverter receives the power measured on the vehicle 20 ms earlier and updates a log-domain correction

$$u_{k+1} = u_k + (\kappa/2)[\ln P^* - \ln P_{k-d}],\ \ V_1 \leftarrow \exp(u)\, V_1 \qquad (8)$$

with $\kappa = 0.3$ and $|u| \le 0.3$. The factor ½ follows from $P \propto V_1^2$. Integration is suspended while the inverter is at its current or voltage limit and the error would push further into it (conditional-integration anti-windup [22]). Without that rule the first version of the trim over-delivered by about 20% after every limited interval; the fix is reported here because it changed the controller ranking.

D. Lateral Dynamics and Estimator

The EKF state is $\mathbf{x} = [e, \psi, b_g, c]^T$: lateral offset, heading, gyro bias, and the offset of painted lane centre from pad centreline. With speed v and measured yaw rate r_m, the discrete process model over a step Δt = 5 ms is

$$e^+ = e + v\Delta t \sin\psi,\quad \psi^+ = \psi + \Delta t(r_m - b_g) \qquad (9)$$

with b_g and c held constant apart from process noise. The process noise comes from gyro white noise (0.005 rad/s per sample), a small bias random walk, and a random walk on c of 2.5×10^{-6} m² per metre travelled, which corresponds to about 1 cm of drift over 40 m. The lane detector fits the lane-centre line $x = a + b{\cdot}y$ in camera coordinates and reports its lateral position at a look-ahead Y_{ref} = 8 m and its slope. With $L = Y_{ref}$ + 2.6 m (the camera sits 2.6 m ahead of the receiver),

$$\mathbf{z}_c = [(c - e)/\cos\psi - L\tan\psi,\ -\tan\psi]^T + \mathbf{n}_c \qquad (10)$$

Because θ is monotone in x (Fig. 2(b)), the magnetic channel does not need to be linearised. The nominal map $h(x, \hat{y})$ at 15 cm gap is inverted numerically, and the result enters the filter as a direct position measurement

$$z_m = h^{-1}(\theta, \hat{y}) = e + n_m \qquad (11)$$

which is accepted only if it lies inside the calibrated range $|z_m| \le 0.60$ m. The first version of the estimator instead linearised h about the current estimate. On one of the ten test seeds this diverged: when the estimate approached the flat ends of the angle curve, the Kalman gain grew as the inverse of the local slope, and a single update moved the estimate by several metres. The diverged runs inflated the magnetic-only error by an order of magnitude. All results reported here use the inverted form, and every run was repeated after the change.

The camera noise $\mathbf{n}_c$ has nominal standard deviations of 1.2 cm and 0.10°, estimated on development seeds with fresh markings that are disjoint from the test seeds. When only one line is found, the lane centre is inferred from the nominal 3.5 m width and the lateral variance is increased by (6 cm)². The magnetic noise n_m follows Table III at $|z_m|$, inflated by a factor of two in standard deviation because the dominant errors (gap, leakage) are correlated in time and a white-noise model would otherwise overstate the information in each sample.

Proposition 2 (*observability of the paint-to-pad offset*). With camera measurements alone, e and c are not separately observable. With at least one magnetic position update, they are.

Proof. In (10) the first row depends on e and c only through $c - e$, and the process model does not couple them, so the observability matrix has the direction $[1, 0, 0, 1]^T$ in its null space. The slope row of (10) observes ψ directly, and b_g is then observable through the heading dynamics in (9), so this is the only unobservable direction. A magnetic update (11) adds the row [1, 0, 0, 0], which is not orthogonal to it, so the null space vanishes. The argument is for the small-angle linearisation and deterministic observability; it says nothing about how fast c converges, which Section VIII examines. ∎

The practical reading is that a camera-only estimator inherits whatever offset exists between the paint and the pads, however good its detector, while the fused estimator learns that offset over powered sections and keeps it through unpowered ones. The ablation in Section VIII tests whether this matters in practice.

Two robustness rules complete the estimator. A camera measurement with innovation $\mathbf{v}$ is rejected if $\mathbf{v}^T\mathbf{S}_g^{-1}\mathbf{v} > 13.8$ (the 99.9% point of χ^2 with two degrees of freedom), where $\mathbf{S}_g$ is computed with the nominal noise. The covariance used for the update itself is inflated by the ratio of a reference score to the frame's detection score, bounded between 1 and 25. Gating on the nominal rather than the inflated covariance prevents a weak detection from widening its own acceptance region. After eight consecutive camera rejections, the lateral and heading variances are inflated so that the filter can re-acquire rather than lock out a correct but surprising measurement. Magnetic measurements are gated in the same way (one degree of freedom, threshold 10.8), and 20 consecutive rejections inflate the lateral variance. The ungated variants E4 and E7 in Section VII disable both gates. Camera frames arrive 60 ms after capture; the filter rolls back to the capture step, applies the update and re-propagates, including any magnetic updates in between [18].

VI. Algorithm

Algorithm 1 summarises one estimator step and the roadside control step it feeds.

Algorithm 1: Fused alignment estimation and pad power control

for each 5 ms step k **do**
  propagate $\mathbf{x}$, $\mathbf{P}$ with (9) using r_m, v
  if pad energised **and** $|\hat{y}| < 0.42$ m **and** k mod 4 = 0 **then**
    $\theta \leftarrow \text{atan2}(B_x, B_z)$ from quadrature demodulation
    $z_m \leftarrow h^{-1}(\theta, \hat{y})$; **if** $|z_m| \le 0.6$ m: gate (1 dof, 10.8); update
  for each camera frame delivered at k (captured at k_c) **do**
    restore $\mathbf{x}$, $\mathbf{P}$ at k_c
    IPM; log-domain ridge filter; RANSAC per side; pair check
    if line found **then** gate on $\mathbf{S}_g$; update with score-weighted $\mathbf{R}$
    if 8 consecutive rejections **then** inflate P_{ee}, $P_{\psi\psi}$
    re-propagate k_c+1 … k, replaying magnetic updates
  send $\hat{e}$, $\hat{y}$ to roadside; $\hat{M} \leftarrow M(\hat{e}, \hat{y}, g_{nom})$
  $V_1 \leftarrow \exp(u)\cdot(6)$; apply 600 V / 60 A limits
  if k mod 2 = 0 **and** received P_{k-4} valid **then** update u by (8) with anti-windup
end for

Complexity is dominated by the lane detector. The EKF has a fixed four-dimensional state, so each predict or update is $O(1)$

in route length; the delayed-frame replay adds at most 12 predict steps per frame, which is $O(L_d)$ with L_d the latency in steps. The detector is $O(N_p)$ in the number of bird's-eye pixels (63 000 here) plus $O(I \cdot N_c)$ for RANSAC with $I = 120$ iterations over N_c candidate points. Memory is $O(N_p + L_d)$. Measured times appear in Section VIII.

VII. Simulation Set-Up

Everything was implemented in Python 3 with NumPy, SciPy and OpenCV and run on a two-core cloud CPU. Coil and field quantities come from the closed forms of Section V, tabulated on a 1 cm lateral grid, a 2 cm longitudinal grid and five gaps, and interpolated during simulation. Table IV lists the main parameters.

TABLE IV
Simulation Parameters

Quantity	Value
Route	240 m; powered 40–120 m and 160–240 m
Pads	0.50 × 1.00 m, 8 turns, 1.25 m pitch
Receiver	0.50 × 0.45 m, 10 turns, gap 15 cm nominal
Link	85 kHz, SS, R_L = 8 Ω, P^* = 11 kW, 600 V / 60 A
Gap variation	±2 cm per run (load) + 5 mm at 1.5 Hz
Paint-to-pad offset c	N(0, 5 cm) per run; lane width 3.5 m, s.d. 0.1 m
Driver wander	2nd order, 6 s period, ζ = 0.6; s.d. about 17 cm
Camera	640 × 480, f = 700 px, 1.30 m high, 12° down, 15 frame/s, 60 ms latency
Camera vibration	pitch 0.25° + 0.2° at 1.5 Hz; roll 0.15°; height ±1 cm at 1.5 Hz, lowered up to 2 cm by load
Gyro	0.005 rad/s white, bias ±0.002 rad/s
Magnetic	50 Hz updates; Table III error sources; 1° per-pad bias
Power link	10 ms control period, 20 ms reporting delay
Speeds; seeds	30, 60, 90 km/h; test seeds 101–110

Camera measurements are not drawn from a noise model. For every frame the simulator traces each pixel's ray to the ground, samples a procedurally generated road texture at 1 cm lateral and 2 cm longitudinal resolution, and applies vignetting, optical blur, motion blur matched to speed and exposure, signal-dependent noise and 8-bit quantisation. The detector of Section IV then runs on the rendered image. Five road conditions are defined; Fig. 3 shows one frame of each. S1 has fresh markings. S2 has 45% paint wear and short gaps. S3 adds 60% wear, missing paint over 40% of the length, tar-sealed cracks, repair patches, tree shadows, road dust and an intermittent ghost line left by earlier restriping 0.3–0.6 m from the true one. S4 is wet and dark, with headlight glare streaks and an 8 ms exposure. S5 has paint over only about 12% of its length. A sixth, mixed route concatenates S1, S3, S2, S5 and S3 segments.

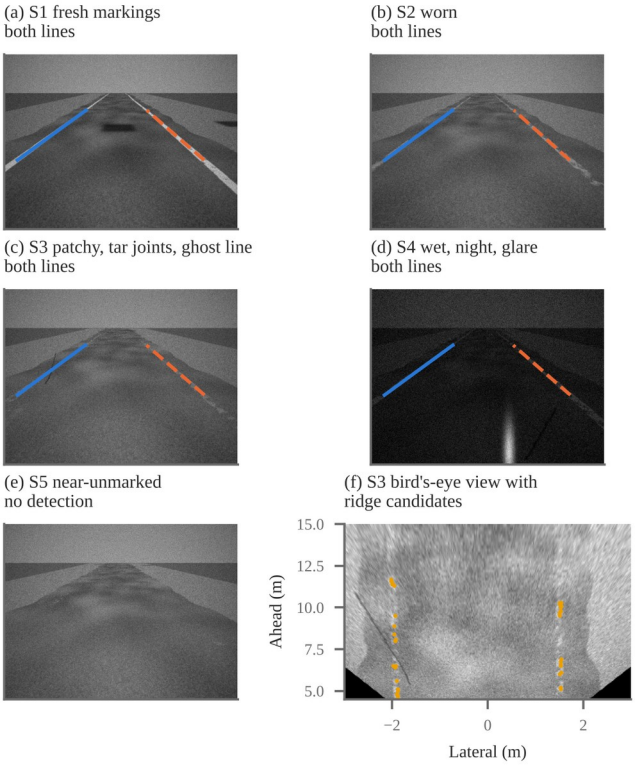


Fig. 3. Rendered camera frames for the five road conditions with the fitted lane lines (solid: left, dashed: right), and (f) the bird's-eye view of the S3 frame with the ridge-filter candidates that survive thresholding.

The detector was developed on separate seeds. Its first version normalised each bird's-eye row by its median and failed completely on S3 because a partly shadowed row has no meaningful median; it was replaced by a ridge filter on log intensity, which turns multiplicative shading into an additive offset that the symmetric ridge response cancels. Rendering parameters for S3 were also revised after the first build placed implausibly large shadows across most frames. Both changes were made before any test seed was run.

Seven estimators are compared on identical trajectories and measurements: camera-only with fixed noise and no gating (E1), camera-only with gating and score weighting (E2), magnetic-only with gyro (E3), naive fusion without gating or the c state (E4), the proposed estimator (E5), and E5 without the c state (E6) or without gating (E7). Truth, road texture and sensor noise use separate random streams, so the same seed gives the same trajectory under every road condition. Nine power controllers are compared: no sensing (feedforward assuming $x = 0$); feedback-only trim reset at each pad; feedback-only trim carried across pads; feedforward from E1, E2, E3 or E5; feedforward from E5 with the carried trim; and an ideal feedforward with true coupling and gap. Power metrics are computed over pad cores where the ideal controller delivers at least 20% of P^*.

VIII. Results and Discussion

The test set comprises 150 runs (five conditions, three speeds, ten seeds). Each run covers 240 m and yields about 144 rendered frames at 90 km/h and 432 at 30 km/h. Lateral wander had a standard deviation of about 17 cm at every speed. Intervals in brackets are 95% bootstrap intervals over runs; paired differences use the same runs for both estimators.

Because a seed gives the same trajectory under every road condition, the five conditions are not independent samples. Intervals pooled over all conditions therefore understate uncertainty, most of all for E3, whose 150 pooled runs contain only 30 distinct cases; the per-condition intervals are the ones to rely on.

A. Lane Detector Across Road Conditions

TABLE V
Lane Detector Performance on Test Seeds

Cond.	Detected (%)	Both lines (%)	Median err. (cm)	90th pct. (cm)	>20 cm (%)
S1	100.0	98.5	1.3	2.8	0.2
S2	100.0	89.1	1.4	3.5	0.3
S3	60.2	15.1	4.5	14.9	6.4
S4	70.1	23.2	4.2	36.6	12.7
S5	0.1	0.0	9.6	28.7	22.2

Errors are in the lane-centre position at 8 m look-ahead, over detected frames only. S5 produced 9 detections in 30 runs, so its error statistics are not meaningful.

Table V and Fig. 4 show the detector. On fresh and worn paint it finds the lane in essentially every frame, with a median error of 1.3–1.4 cm. Degradation shows up in two different ways. On the patchy road (S3) the detection rate falls to 60%, and only 15% of detections contain both lines, so most fall back on the nominal lane width and inherit half of its width variation (about 5 cm standard deviation). Gross errors above 20 cm are 6.4% of detections, mostly the ghost line. Wet night (S4) keeps a higher detection rate (70%) but has the heavier tail: 12.7% of detections are off by more than 20 cm, because glare streaks pass the width-matched ridge filter and the single-line fallback accepts them. The 90th-percentile error there is 37 cm. That distinction matters for fusion: a missing measurement costs information, but a confident wrong one costs accuracy unless it is rejected.

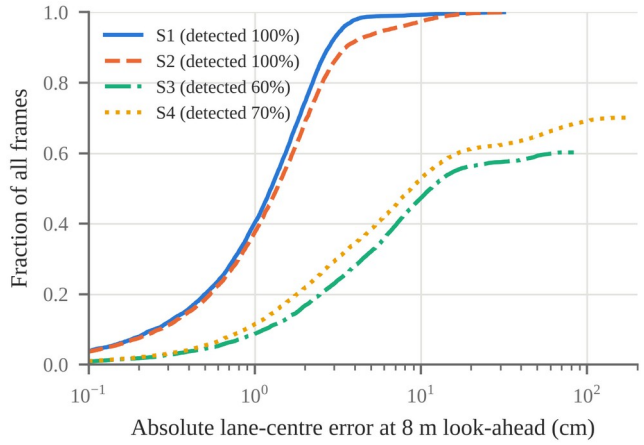


Fig. 4. Cumulative distribution of lane-centre error over all frames, including frames with no detection (the curves saturate at the detection rate). S5 is omitted because it produced almost no detections.

B. Lateral Estimation Accuracy

TABLE VI
Lateral RMSE (cm) by Road Condition, All Speeds

Cond.	E2 camera	E3 magnetic	E5 proposed
S1	3.22 (2.77–3.66)	5.83 (4.44–7.43)	1.49 (1.32–1.67)
S2	3.20 (2.77–3.65)	5.83 (4.41–7.37)	1.50 (1.34–1.66)
S3	6.79 (5.99–7.64)	5.83 (4.40–7.37)	3.01 (2.46–3.70)
S4	4.05 (3.56–4.58)	5.83 (4.44–7.39)	2.20 (1.88–2.54)
S5	67.60 (50.06–87.19)	5.83 (4.46–7.42)	5.83 (4.42–7.34)
All	16.97 (12.06–22.90)	5.83 (5.20–6.51)	2.81 (2.41–3.27)

Root-mean-square error over the full 240 m route, averaged over 30 runs per condition (150 for All); 95% bootstrap interval in parentheses.

Table VI gives the headline comparison and Fig. 5 splits it by route segment. Over all runs the proposed estimator reached 2.81 cm, against 5.83 cm for magnetic sensing alone and 16.97 cm for the camera alone. The paired difference from magnetic-only was -3.02 cm (-3.70 to -2.37), and E5 was better in 73% of runs.

The split by segment explains where the benefit comes from. Inside powered sections the magnetic channel dominates: E5 and E3 achieve 0.92 cm and 1.25 cm respectively, while the camera alone is limited to about 3 cm even on fresh paint by the unknown paint-to-pad offset c. Over unpowered sections the order changes: magnetic-only dead-reckons on the gyro and reaches 9.86 cm, whereas E5 holds 4.56 cm because the camera keeps correcting it. At the entry to each powered section, the point that sets the first pad's voltage, the mean error was 3.6 cm for E5, 9.0 cm for E3 and 11.3 cm for E2.

The benefit depends on the road. With good paint (S1, S2) fusion halves the camera's error, from 3.2 cm to 1.5 cm, with E5 better in every paired run. On the patchy road it reduces 6.8 cm to 3.0 cm. On the near-unmarked road (S5) fusion is exactly equal to magnetic sensing alone (5.83 cm): with almost no detections the camera contributes nothing, and the camera-only estimator degrades to gyro dead-reckoning with an RMSE of 68 cm. This is the expected outcome, and it is reported because it bounds what fusion can do: on a road without paint, the estimator is only as good as its magnetic channel and its gyro.

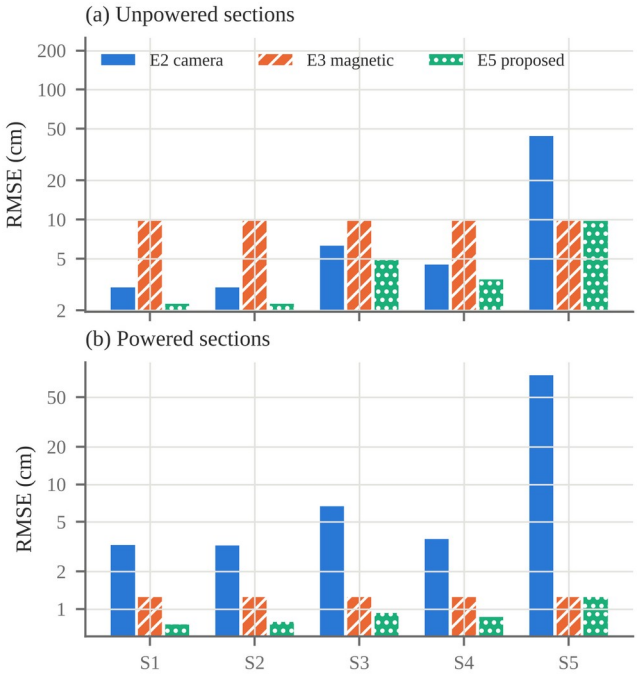


Fig. 5. Lateral RMSE by road condition, split into (a) unpowered and (b) powered sections (logarithmic scale). The camera-only estimator is limited by the paint-to-pad offset in (b) and fails on S5; magnetic-only is limited by gyro drift in (a).

Speed had little effect on the proposed estimator, whose RMSE was 2.85, 2.83 and 2.74 cm at 30, 60 and 90 km/h. Higher speed shortens the time spent in unpowered stretches, which helps the dead-reckoning components, and reduces the

number of frames per metre, which hurts the camera; the two roughly cancel. The filter's own uncertainty was well calibrated: the true error lay inside $\pm 2\sigma_e$ for 93.4% of samples, close to the nominal 95%. The camera-only filter without score weighting covered only 42%, meaning it believed itself several times more accurate than it was.

C. Example Trace and a Failure Mode

Fig. 6 shows the S3 run at 60 km/h with seed 101, which was fixed as the illustration seed before any test run. It turned out to be the worst of three S3 runs in which gating failed. In the first 20 m, before any powered section, the detector returned only single-line, low-score measurements. The gated filters (E2, E5, E6) drifted to an error of 68 cm at 15 m and recovered by about 22 m once the re-acquisition rule inflated their covariance; the ungated E7 stayed within 13 cm. The fault lies in the start-up transient: the gate is computed with nominal camera noise, so while the heading estimate is still poor, a run of weak but correct measurements can be rejected long enough for the error to grow. The episode ended well before the first powered section; inside powered sections the fused estimators then tracked to within about 2 cm in this run. The same weakness appeared twice more on S3, at 30 km/h with seeds 105 and 108, this time in the second unpowered stretch just before the second powered section: E5 reached 36 cm and 20 cm while E7 stayed below 9 cm. In all three cases the gated filter rejected a run of weak but correct single-line detections while its own estimate was drifting on the gyro. Three of the 30 S3 runs is not a rare event. Gating with a fixed nominal covariance is therefore unsuitable on patchy roads as designed here; a gate that widens with time since the last accepted measurement, or an initialisation from several frames, is the obvious remedy, but neither was tested.

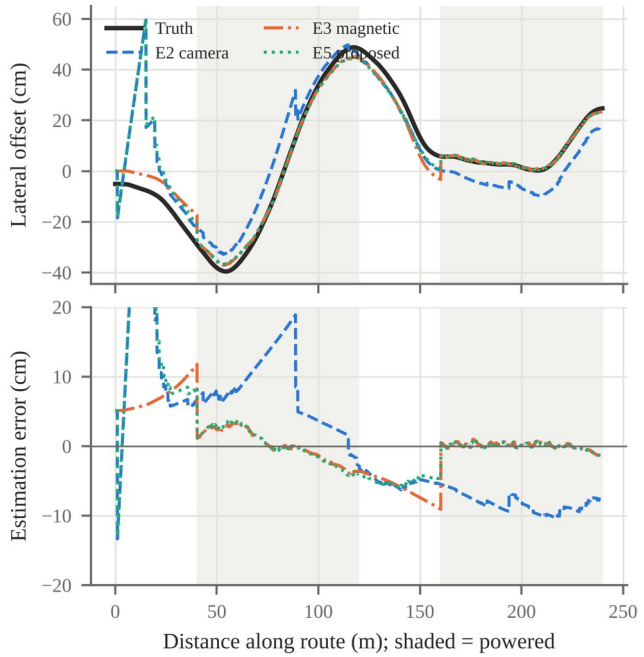


Fig. 6. Example run: S3, 60 km/h, seed 101 (chosen in advance; it is the worst of three S3 runs in which gating failed). (a) True and estimated lateral offset. (b) Estimation error, clipped at ±20 cm (the start-up excursion of E2 and E5 reaches 68 cm at 15 m). Shaded bands are powered sections, where the magnetic channel is available.

D. Ablation

TABLE VII
Ablation: Lateral RMSE (cm)

Cond.	E4 naive	E6 no *c*	E7 no gate	E5 full
S1	1.93	1.78	1.49	1.49
S2	1.99	1.73	1.50	1.50
S3	3.90	3.05	2.69	3.01
S4	6.94	2.35	4.53	2.20
S5	7.02	5.83	5.82	5.83
All	4.36	2.95	3.21	2.81

E4: no gating, no score weighting, no *c* state. E6: E5 without the *c* state. E7: E5 without gating.

Table VII isolates each component. The *c* state gave a small gain overall, 0.14 cm (0.07 to 0.21), larger on clean paint (0.29 cm on S1) where the paint offset is the dominant camera error, and indistinguishable from zero on S3, where single-line fallback errors are larger than *c*. Its clearer effect is on calibration and entry: without it, uncertainty coverage on S1 fell from 93% to 77%, the mean section-entry error rose from 2.0 cm to 3.0 cm, and the camera rejection rate rose from 1.8% to 4.9% because the filter kept seeing a consistent bias it had no state to absorb. The gain is not uniform: E5 beat E6 in only 53% of paired runs, with ties throughout S5.

Gating is the component whose value depends most on the road. On wet night roads, where glare produces confident wrong detections, removing it doubled the error (4.53 cm against 2.20 cm). On S1 and S2 it made no difference. On S3 the ungated filter was slightly better (2.69 against 3.01 cm), although the paired interval (-0.18 to 0.95 cm) includes zero; this is the cost of the gating failures described above. The naive fusion E4, which lacks all three refinements, was worse than E5 on S1–S4 and fell behind the robust camera-only filter on wet roads; on S5 the two could not be separated (paired interval -3.56 to 0.00 cm).

Fig. 7 tests Proposition 2 directly. Panel (a) varies the paint-to-pad offset on S2 at 60 km/h. The camera-only error grows with the offset, from 1.48 cm at zero to 11.04 cm at 10 cm, and fusion without the *c* state reaches 8.87 cm. The proposed estimator also degrades, from 1.04 to 4.46 cm, but the segment breakdown shows where. At the entry to the first powered section, before any magnetic update, its error equals the offset (10.6 cm at 10 cm), as Proposition 2 requires. At the entry to the second section, after one powered section has been driven, its error is 1.3, 1.2 and 1.2 cm for offsets of 0, 5 and 10 cm, whereas without the *c* state it is 11.0 cm at 10 cm. Once learned, the offset no longer matters. With zero offset the *c* state costs nothing measurable (1.04 against 0.99 cm). Panel (b) shows the estimate of *c* along one route: it stays at its prior through the first unpowered stretch, converges within the first powered section, and holds through the second unpowered stretch.

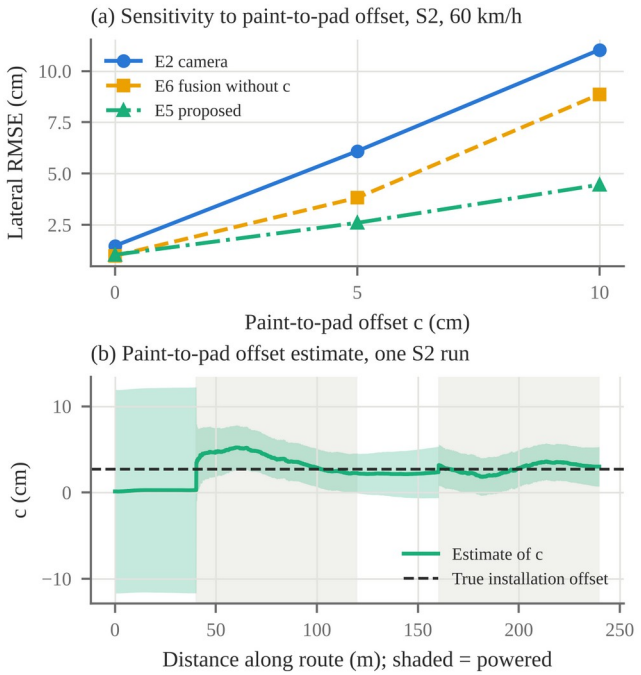


Fig. 7. (a) Lateral RMSE versus paint-to-pad offset c (S2, 60 km/h, ten seeds per point). (b) Estimate of c with $\pm 2\sigma$ band along one S2 route; shaded bands are powered sections, and the dashed line is the true installation offset.

E. Power Control

TABLE VIII
Power-Tracking Performance in Pad Cores, All Runs

Controller	RMS error (% P^*)	>1.15 P^* (%)	<0.85 ideal (%)
No sensing (x = 0)	12.0 (11.3–12.6)	12.2	1.8
Feedback trim, reset per pad	10.7 (10.1–11.2)	10.1	1.7
Feedback trim, carried	5.6 (5.3–6.0)	1.8	1.6
Feedforward, E2	16.4 (14.8–18.1)	4.9	19.0
Feedforward, E3	8.3 (7.6–9.0)	6.1	4.1
Feedforward, E5	8.2 (7.6–8.9)	6.1	4.1
Feedforward E5 + trim	4.6 (4.2–5.0)	0.9	1.5

Error relative to an ideal feedforward that knows true coupling and gap, subject to the same 600 V / 60 A limits; evaluated in pad cores where the ideal controller delivers at least 20% of P^*; the over-power column is computed over all energised samples.

Table VIII and Fig. 8 compare the controllers. Without sensing, assuming perfect alignment, the inverter over-drives every misaligned pass; power exceeded 1.15 P^* for 12.2% of the time and the RMS error was 12.0%. A feedback trim reset at each pad barely helps (10.7%): with a 20 ms reporting delay it gets about two corrections per pad at 90 km/h and three at 60 km/h. Carrying the trim across pads is much better (5.6%), because lateral offset and gap change slowly compared with the pad pitch.

Feedforward from the estimates splits into two groups. Camera-only feedforward was the worst controller (16.4%), and under-delivered more than 15% for 19% of the time, because its paint-referenced errors translate directly into coupling errors through Proposition 1. Magnetic and fused feedforward performed the same (8.3% and 8.2%): inside pad cores both estimators are within about a centimetre, and the remaining error is the air gap, which neither sees. Adding the carried trim to fused feedforward gave the lowest error, 4.6% (4.2–5.0), and the fewest over-power events (0.9%). The paired difference from the carried trim alone was -1.09 percentage points (-1.24 to -0.94), negative in 98% of runs. By speed the margin was 0.9, 0.7 and 1.6 percentage points at 30, 60 and 90 km/h, largest at the highest speed but not monotone.

Two honest qualifications follow. First, the improvement of fused feedforward plus trim over trim alone is about one percentage point of RMS error; it is real, but it is not the main argument for fusion. The main argument is the unpowered-section and section-entry accuracy of Section VIII-B, which is what a lane-guidance aid or a first-pad safety check would use. Second, for power control alone, a magnetic-only estimator is as good as the fused one, so a system that only regulates power within powered sections does not need the camera.

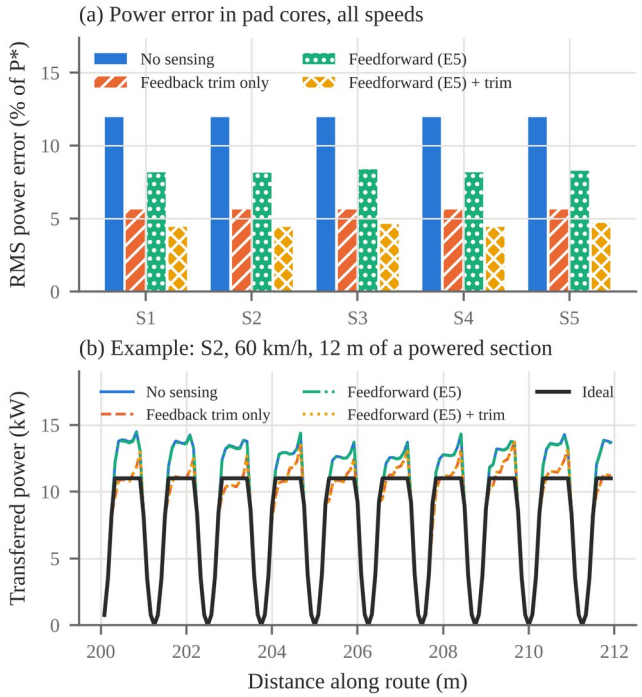


Fig. 8. (a) RMS power-tracking error in pad cores for four controllers, by road condition. (b) Transferred power over 12 m of a powered section (S2, 60 km/h, seed 101). In this run the vehicle is lightly loaded and the air gap is about 2 cm larger than nominal, so the coupling predicted at nominal gap exceeds the true value and feedforward alone over-delivers like the unsensed controller; in this window the trim brings both trimmed controllers to the 11 kW target.

F. Mixed-Quality Route

The mixed route, where paint quality changes every 40–60 m (S1, S3, S2, S5, S3), is closest to an Indian arterial road. Over 10 runs at 60 km/h the proposed estimator reached 1.95 cm (1.68–2.23), against 5.74 cm for magnetic-only and 5.72 cm for camera-only, with a mean section-entry error of 3.0 cm. The fused feedforward with trim held power to 4.5% RMS error.

G. Computational Cost

Rendering plus detection took 35–47 ms per frame in single-threaded Python. Timed separately over 120 frames each on S1, S3 and S4, the detector alone needed a median of 8–11 ms (95th percentile 14–17 ms); the remainder is the simulator's renderer, which a real system does not run. The detector therefore fits the 67 ms budget of 15 frame/s even in Python, and its operations (two remaps, a cumulative-sum box filter, and RANSAC on a few hundred points) map directly onto embedded vision hardware. The four-state EKF with replay is negligible by comparison. A full 240 m run, including all seven estimators and nine controllers, took 24 s on average.

IX. Comparative Analysis

Table IX compares the approach with prior alignment-sensing work on the attributes that matter for deployment. The accuracy column mixes reported experimental values with this paper's simulated values, and the two are not directly comparable: the RFID figures [7] are single hardware runs, while the values here come from a simulation whose fidelity is limited in ways listed in Section XI.

TABLE IX
Qualitative and Quantitative Comparison with Prior Work

Attribute	Vision [6]	RFID [7]	Coils / TMR [8], [9]	This work
Reference needed	Painted guideline	Tags at known positions	Energised pad	Paint where available, energised pad
Works before first pad	Yes, if paint visible	Yes, at tags	No	Yes, degrades with paint
Works without paint	No	Yes	Yes, over pads	Over pads; gyro elsewhere
Lateral accuracy	Not stated in abstract	3.5 cm lab, about 5 cm field (reported)	Not stated in abstract	0.9 cm powered, 4.6 cm unpowered (simulated)
Added road hardware	Paint	Tags	None beyond pads	None beyond pads
Added vehicle hardware	Camera	Reader, antenna	Sensor coils or TMR array	Camera, 2-axis coil, gyro
Computational load	Image processing	Grid-search MLE	Low	Image processing + 4-state EKF
Evidence level	Experiment	Lab and field runs	Robot test; FEM + experiment	Simulation only

MLE: maximum likelihood estimation. Entries for [6], [8] and [9] reflect what their abstracts state; the full papers may report more.

Three points follow from the comparison. The proposed scheme needs no road infrastructure beyond the pads themselves, unlike guideline painting or tag installation, which matters where road authorities restripe without coordinating with a charging operator. It is the only option in the table that both works before the first pad and survives the loss of paint, though in the second case only through the magnetic channel and gyro. Its evidence base is the weakest: until it is tested in hardware, its numbers should be read as what the method can achieve under the modelled error sources, not as a field accuracy.

On cost and complexity, the added vehicle hardware is modest: a forward camera of the kind already fitted for lane-keeping assistance, a two-axis pickup coil, and the gyro already present in the vehicle's stability control. Scalability along the road is better than for tag-based methods because nothing is added per pad. The sustainability argument is indirect: better alignment estimates reduce over-drive and current stress, and they allow the pad to be energised only when the receiver is confirmed over it, which also limits stray-field exposure to the reference levels of [23].

X. Key Outcomes

- **Method.** A field-angle magnetic observable that is monotone over the full ±70 cm examined, a four-state EKF that learns the paint-to-pad offset, and a feedforward-plus-trim power controller, all evaluated end-to-end with rendered images.
- **Accuracy.** Lateral RMSE of 2.81 cm over all conditions (0.92 cm in powered sections), against 5.83 cm for magnetic-only and 16.97 cm for camera-only, in simulation.
- **Power.** RMS power-tracking error of 4.6% with over-power events at 0.9% of the time, against 12.0% and 12.2% without sensing.
- **Engineering insight.** By Proposition 1 and the error budget, a 2 cm gap error moves power by about 20% while the fused lateral error in pad cores is about 1 cm, so the unknown air gap, not lateral error, is the likely limit on feedforward power control. This was inferred, not isolated by a controller using true gap.
- **Limits found.** No benefit over magnetic sensing on unmarked roads; gating lock-out in 3 of 30 patchy-road runs; no power-control benefit of the camera inside powered sections.
- **Relevance.** The design targets Indian roads, where paint is unreliable and no DWPT deployment exists yet; it requires no road infrastructure beyond the pads.

XI. Conclusion

This paper set out to estimate the lateral alignment of a dynamic wireless charging receiver on roads where lane paint cannot be trusted. Combining a camera that reads the paint with a magnetic sensor that reads the pad field, and carrying the offset between the two references as a filter state, gave a lateral error of a few centimetres along the whole route and below 1 cm over powered sections in simulation. The camera mattered most where the magnetic field is absent, before and between powered sections; the magnetic channel mattered most where paint is poor. Neither alone was adequate everywhere.

The limitations are substantial and are stated plainly. All results are simulated; no hardware was built. The camera model renders a planar textured road with no other vehicles, no road curvature, no rain on the lens, and no rolling-shutter effects, and its scene parameters were set by judgement rather than measured from Indian roads. The coil model uses filament conductors without ferrite or shielding, which changes both coupling and field shape; the angle observable must be re-derived from a finite-element model of a real pad. Field distortion by steel reinforcement in concrete roads, neighbouring vehicles and pad-to-pad manufacturing spread is represented only by a 1° per-pad bias. The driver model has a smaller wander (about 17 cm standard deviation) than pavement-design values, which corresponds to a driver following an alignment cue. Air gap is not estimated, and it dominates the remaining power error. Ten seeds per condition give reasonably tight intervals for mean errors but say little about rare failures; the gating lock-out seen in 3 of 30 patchy-road runs shows that such failures exist and need a larger study. The air-gap attribution of the residual power error rests on the sensitivity analysis rather than on a controlled comparison.

Future work follows directly. The field magnitude, normalised by a transmitter current reported over the existing communication link, should carry information on gap; adding it

as a measurement could remove most of the residual power error. A gate that widens with time since the last accepted measurement, or multi-frame initialisation, should address the lock-out. A bench rig with a single pad, a moving receiver on a linear stage, and a camera over printed road textures would test the magnetic observable and the fusion logic before any road trial. Images from Indian roads would replace the procedural textures, and a learned lane detector could be compared with the classical one on the same footing.

ACKNOWLEDGMENT

[To be completed by the authors. Sponsor acknowledgments go here, not on the first page.]

REFERENCES

[1] S. Li and C. C. Mi, "Wireless power transfer for electric vehicle applications," *IEEE J. Emerg. Sel. Topics Power Electron.*, vol. 3, no. 1, pp. 4–17, Mar. 2015.

[2] S. Y. Choi, B. W. Gu, S. Y. Jeong, and C. T. Rim, "Advances in wireless power transfer systems for roadway-powered electric vehicles," *IEEE J. Emerg. Sel. Topics Power Electron.*, vol. 3, no. 1, pp. 18–36, Mar. 2015.

[3] Y. J. Jang, "Survey of the operation and system study on wireless charging electric vehicle systems," *Transp. Res. Part C, Emerg. Technol.*, vol. 95, pp. 844–866, 2018.

[4] C. Wang, W. Ren, Y. Chen, and X. Li, "Review of high-misalignment tolerance techniques in wireless power transfer systems," *Energies*, vol. 19, no. 3, Art. no. 713, Jan. 2026, doi: 10.3390/en19030713.

[5] W. Zhang and C. C. Mi, "Compensation topologies of high-power wireless power transfer systems," *IEEE Trans. Veh. Technol.*, vol. 65, no. 6, pp. 4768–4778, Jun. 2016.

[6] [AUTHORS TO BE VERIFIED], "Vision-based rapid power control for a dynamic wireless power transfer system of electric vehicles," *IEEE Access*, vol. 8, 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9075994

[7] H. Sun, X. Ma, R. Q. Hu, and R. Christensen, "Precise coil alignment for dynamic wireless charging of electric vehicles with RFID sensing," arXiv:2312.12565, Dec. 2023.

[8] H. Mahdi, A. R. Abdeltwab, M. S. Hassan, A. Osman, and A. M. Benaya, "Misalignment detection and correction for dynamic wireless charging of electric vehicles," in *Proc. Int. Telecommun. Conf. (ITC-Egypt)*, Cairo, Egypt, Jul. 2025, pp. 625–630, doi: 10.1109/ITC-Egypt66095.2025.11186564.

[9] X. Liu, W. Han, C. Liu, and P. W. T. Pong, "Marker-free coil-misalignment detection approach using TMR sensor array for dynamic wireless charging of electric vehicles," *IEEE Trans. Magn.*, vol. 54, no. 11, Art. no. 4002305, Nov. 2018, doi: 10.1109/TMAG.2018.2844863.

[10] Indian Roads Congress, *Code of Practice for Road Markings*, IRC:35-2015, New Delhi, India, 2015.

[11] VINCI Autoroutes, "World's first dynamic wirelessly-charged motorway shows highly promising real-world results," press release, Oct. 22, 2025.

[12] A. Wendt *et al.*, "Wireless electric road systems: Technology readiness and recent developments," in *Proc. IEEE Wireless Power Technol. Conf. Expo (WPTCE)*, Kyoto, Japan, 2024, pp. 177–182, doi: 10.1109/WPTCE59894.2024.10557264.

[13] M. Aly, "Real time detection of lane markers in urban streets," in *Proc. IEEE Intell. Veh. Symp.*, Eindhoven, The Netherlands, 2008, pp. 7–12.

[14] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981.

[15] S. P. Narote, P. N. Bhujbal, A. S. Narote, and D. M. Dhane, "A review of recent advances in lane detection and departure warning system," *Pattern Recognit.*, vol. 73, pp. 216–234, Jan. 2018.

[16] X. Pan, J. Shi, P. Luo, X. Wang, and X. Tang, "Spatial as deep: Spatial CNN for traffic scene understanding," in *Proc. AAAI Conf. Artif. Intell.*, 2018, pp. 7276–7283.

[17] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York, NY, USA: Wiley, 2001.

[18] T. D. Larsen, N. A. Andersen, O. Ravn, and N. K. Poulsen, "Incorporation of time delayed measurements in a discrete-time Kalman filter," in *Proc. 37th IEEE Conf. Decis. Control*, Tampa, FL, USA, 1998, pp. 3972–3977.

[19] R. K. Mehra, "On the identification of variances and adaptive Kalman filtering," *IEEE Trans. Autom. Control*, vol. AC-15, no. 2, pp. 175–184, Apr. 1970.

[20] SAE International, *Wireless Power Transfer for Light-Duty Plug-in/Electric Vehicles and Alignment Methodology*, SAE Standard J2954, Oct. 2020.

[21] F. W. Grover, *Inductance Calculations: Working Formulas and Tables*. New York, NY, USA: Dover, 2004 (reprint of 1946 ed.).

[22] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC, USA: ISA, 2006.

[23] International Commission on Non-Ionizing Radiation Protection, "Guidelines for limiting exposure to time-varying electric and magnetic fields (1 Hz to 100 kHz)," *Health Phys.*, vol. 99, no. 6, pp. 818–836, Dec. 2010.