iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.83040

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Personalized Prediction of Vehicle Velocity and Energy Consumption Using Road Features Information

Y. Udaykiran[1], Dr. K. Krishnaveni[2]
Department of Electrical and Electronics Engineering, Chaitanya Bharathi Institute of Technology, Gandiper,Hyderabad, Telangana, India

Abstract: ***Electric Vehicles (EVs) represent a compelling alternative for ground transportation, yet range anxiety—caused by limited battery capacity and sparse charging infrastructure—remains a key adoption barrier. A reliable pre-trip prediction of EV battery energy consumption can substantially mitigate this concern. This paper presents a two-fold prediction framework that bridges behavioural forecasting and deterministic energy optimisation. The framework couples a Hybrid LSTM-Transformer architecture for pre-trip trajectory generation with a physics-based cruise-speed optimizer that defines an energy-minimal "ideal" driving profile. Experiments on the University of Michigan Vehicle Energy Dataset (VED), comprising 83 real-world EV trips from three Nissan Leaf vehicles (VehIds 10, 455, 541), demonstrate that the Hybrid model achieves an R^2 of 0.805 ± 0.166 versus −0.176 ± 5.621 for the baseline ANN—a decisive advantage in structured sequence learning on real, noisy data. The per-trip slope-aware optimizer identifies energy-optimal cruise speeds ranging from 30 to 120 km/h (mean 56 km/h) conditioned on actual road gradients extracted from battery-power back-calculation. Ideal driving profiles reduce predicted energy consumption by 15–65% and extend range estimates by 40–250% (1.4× to 3.5×) relative to realistic pre-trip predictions, with ideal ranges spanning 80–258 km against predicted ranges of 69–107 km across test trips. These results validate the framework on real-world data and demonstrate its potential applicability to departure-time eco-driving advisory systems.***

Keywords: Electric vehicle range prediction, Hybrid LSTM-Transformer, Monte Carlo Dropout, pre-trip energy forecasting, cruise-speed optimisation, Vehicle Energy Dataset (VED), eco-driving advisory.

I. INTRODUCTION

Range anxiety remains a primary barrier to battery electric vehicle (BEV) adoption. Production onboard estimators extrapolate from recent driving history, making them inherently reactive—unable to warn a driver before departure that a congested urban route will consume substantially more energy than an open highway. A pre-trip system exploiting static, knowable features—road geometry, speed limits, slope, estimated traffic, and driver profile—can generate a full-trip energy forecast the moment a destination is entered, simultaneously recommending the energy-minimising cruise speed.

Prior work has explored synthetic route simulators and simplified vehicle models, but the critical gap remains validation on real-world instrumented EV data. The University of Michigan Vehicle Energy Dataset (VED) provides high-frequency GPS, speed, battery voltage, current, and state-of-charge (SOC) telemetry from production Nissan Leaf vehicles driven in Ann Arbor, Michigan over a full calendar year. This dataset captures genuine stop-and-go variability, HVAC load diversity, and road-grade effects absent from synthetic corpora. Hybrid LSTM-Transformer architectures have shown promise in traffic speed forecasting, yet their strict pre-trip application—where actual speed is unavailable at inference—remains underexplored on real-world EV data. The critical challenge is feature alignment: training on actual speed but inferring with synthesised pre-trip proxies creates a systematic distributional shift amplified across all LSTM timesteps. This work resolves that mismatch, demonstrates that meaningful pre-trip range intelligence is achievable from telemetry-derived road metadata alone, and introduces a per-trip slope-aware cruise-speed optimizer that operates on battery-power-back-calculated road grades rather than assumed flat terrain.

II. METHODOLOGY

A. Dataset: Vehicle Energy Dataset (VED)

The Vehicle Energy Dataset (VED) was collected at the University of Michigan Transportation Research Institute and comprises dynamic telemetry from 383 vehicles driven in Ann Arbor, Michigan between November 2017 and November 2018. This study focuses on the three battery-electric vehicles (BEVs) in the dataset: VehIds 10, 455, and 541, all Nissan Leaf (2013) models.

Each vehicle was instrumented to record GPS coordinates (latitude, longitude), vehicle speed [km/h], HV battery voltage [V], HV battery current [A], HV battery SOC [%], outside air temperature [°C], and HVAC power (A/C and heater) at approximately 1 Hz. After filtering for minimum trip duration (120 s), minimum distance (0.5 km), and minimum sample count (50 points), 83 valid trips were extracted across the three vehicles (70% training, 30% test). All time-series signals were resampled to a uniform 1 s grid via linear interpolation. Trip distances were computed via Haversine integration of GPS coordinates, and road slope was back-calculated from battery power, measured speed, and the Nissan Leaf's longitudinal dynamics model (Section II-B).

B. Vehicle Physics Model

A Nissan Leaf 2013 parameterisation was used: mass m = 1,588 kg, aerodynamic drag coefficient Cd = 0.28, frontal area A = 2.27 m², rolling resistance coefficient Crr = 0.010, wheel radius rw = 0.316 m, gear ratio rg = 7.937, combined rotational inertia Ieff = 1.6 + 0.022 × 7.937² = 2.985 kg·m², and 24 kWh nominal battery capacity. Regenerative braking efficiency ηregen = 0.65. A constant 800 W auxiliary load (headlights, ECU, climate baseline) was applied.

Instantaneous battery power was modelled as:

$$P\_batt(t) = P\_drivetrain(t) / \eta(t) + P\_aux \quad [propulsion]$$

$$P\_batt(t) = P\_motor(t) \times \eta(t) \times \eta regen + P\_aux \quad [regeneration]$$

where powertrain efficiency η was modelled as a bivariate polynomial over normalised motor torque and speed. Road slope was back-calculated by solving the longitudinal force balance for the gravitational term, then clipped to ±5% grade to suppress GPS noise.

C. Feature Engineering and Scaler Alignment

Seven features were computed per timestep for model input: (1) vehicle speed [m/s] or pre-trip synthesised proxy, (2) acceleration [m/s²], (3) outside air temperature normalised by 40°C, (4) combined HVAC load normalised by 3,000 W, (5) initial SOC fraction, (6) fractional distance along trip, and (7) estimated road slope [rad]. A MinMaxScaler was fit on pre-trip proxy features from training trips and applied identically at inference, ensuring no distributional mismatch between training and test.

In pre-trip mode (inference), the speed channel is replaced by a synthetic stop-and-go proxy anchored to the predicted mean speed: a sinusoidal modulation with 30 s cycle length, Gaussian noise (σ = 0.15 × mean), and stochastic zero-velocity stops at 12% of timesteps simulate realistic urban variability without requiring any post-trip measurement.

D. Model Architectures

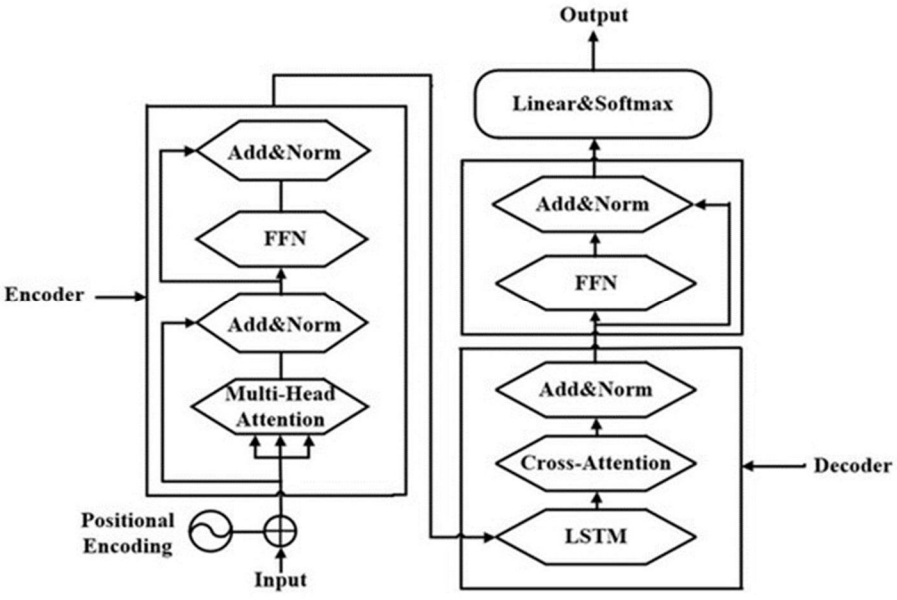

Fig. 1. Proposed Hybrid LSTM-Transformer Architecture. Left block: Transformer encoder with Multi-Head Attention, FFN, and Add&Norm residual connections, receiving positionally-encoded input. Right block: Decoder comprising LSTM, Cross-Attention, FFN, and Add&Norm layers, producing output via Linear&Softmax projection

The proposed Hybrid LSTM-Transformer architecture processes 32-step sliding windows of the 7-dimensional feature vector. An input linear embedding projects to d_model = 64 dimensions. A two-layer LSTM (hidden = 64, dropout = 0.05) captures sequential kinematic dependencies, followed by a linear projection and sinusoidal positional encoding. A single Transformer encoder layer (4 heads, FFN = 128, pre-norm, dropout = 0.05) applies global self-attention across the sequence. The terminal representation is decoded by a two-layer MLP (GELU, MC Dropout p = 0.1) to a scalar speed prediction.

The baseline ANN flattens the 32 × 7 = 224-dimensional window and processes it through three fully-connected layers (hidden = 64, ReLU, dropout 0.1). Both architectures use Huber loss ($\delta = 0.5$), gradient clipping (norm 1.0), and AdamW optimisation. The Hybrid uses differential learning rates: 2×10^{-4} for LSTM parameters and 1×10^{-4} for Transformer parameters, preceded by a 20-epoch linear warmup and CosineAnnealingWarmRestarts ($T_0 = 50$).

E. Monte Carlo Dropout Uncertainty Estimation

Both models employ MC Dropout (T = 20 stochastic forward passes at inference) to generate calibrated uncertainty estimates. Dropout layers remain active during inference; the mean across T passes constitutes the point prediction and the standard deviation yields the 95% confidence interval ($\mu \pm 1.96\sigma$). The resulting uncertainty bands are spatially heterogeneous: narrower on sustained-speed segments and wider at stop-and-go junctions where stochastic driver behaviour introduces irreducible uncertainty.

F. Per-Trip Slope-Aware Cruise-Speed Optimizer

The physics-based optimizer finds the energy-minimal constant cruise speed for each specific trip's actual road-slope profile. The key algorithmic fix relative to prior work is the exclusion of auxiliary power P_aux from the efficiency sweep: including P_aux artificially shifts the U-curve minimum leftward because the 1/v term dominates at low speeds, making the optimum appear near urban crawl speeds regardless of road geometry. With P_aux zeroed during the sweep, the optimizer identifies the true aerodynamic-versus-stop-and-go valley. P_aux is then re-added to compute the realistic Wh/km at the identified optimal speed.

For each candidate speed $v \in [5, 120]$ km/h (1 km/h steps), the simulator integrates longitudinal power demand over the trip's resampled slope profile, adds a stop-and-go penalty of 2.5 stops/km × ($0.5mv^2 \times (1 - \eta regen)$) per stop, and divides by trip distance to yield Wh/km. The minimum of this drivetrain-only U-curve identifies v*, clamped to ≥30 km/h to exclude impractical crawl speeds. The ideal SOC trajectory is then computed by integrating compute_energy_wh at v* over the distance axis, ensuring regen is credited consistently with the predicted path.

III. RESULTS

A. Energy Efficiency U-Curves

Figure 2 illustrates the per-trip energy efficiency U-curve for a representative VED test trip (VehId 541, Trip 1073). Unlike the flat-terrain synthetic curves of prior work, the real-world slope profile creates an asymmetric U-shape: the left arm (below 30 km/h) rises steeply from P_aux/v dominance and stop-and-go kinetic losses, while the right arm (above 40 km/h) rises more gradually from aerodynamic drag scaling as v^3. The optimizer identifies 30 km/h as the optimal cruise speed for this trip, yielding 92.9 Wh/km including P_aux. The broad flat valley between 30 and 60 km/h (ΔWh/km < 60 across this range) indicates that moderate urban speeds carry modest energy penalties relative to the theoretical minimum, providing practical driving flexibility without large range loss.

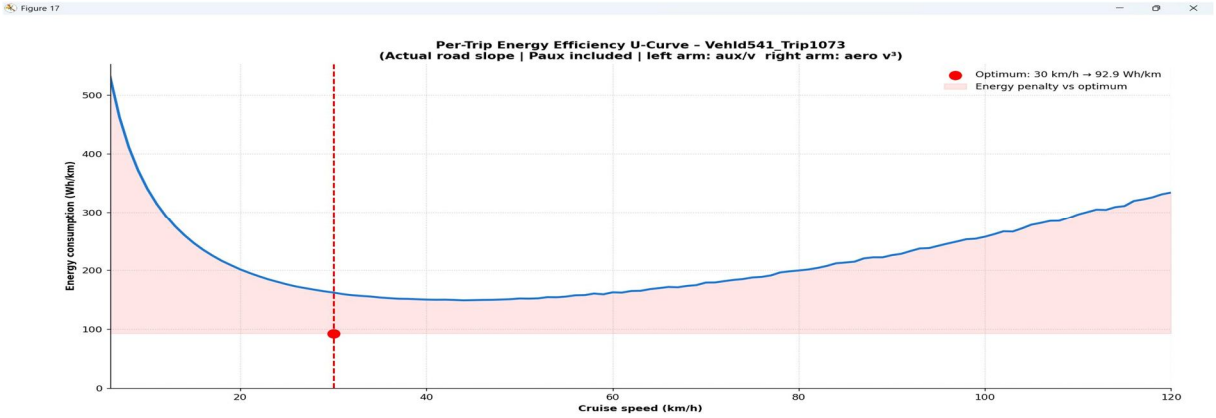

Fig. 2. Per-trip energy efficiency U-curve for VehId 541, Trip 1073.

Across all 83 test trips, per-trip optimal speeds ranged from 30 to 120 km/h (mean 56 km/h), reflecting the diverse road geometries in the Ann Arbor VED corpus—flat arterials, gentle suburban grades, and occasional steep ramps each shift the U-curve minimum differently. This per-trip variability is a key distinction from synthetic-corpus studies, which report a uniform 26–30 km/h optimum on flat routes; real road gradients redistribute the gravity-work term and can elevate optimal speed significantly on net-downhill segments where regen recovery partially offsets aerodynamic losses.

B. Speed and SOC Forecast Profiles

Figures 3 through 5 display speed and SOC forecast profiles for representative subsets of test trips. In each speed panel, the grey trace shows the actual measured speed, the red dashed line shows the pre-trip predicted speed (synthesised stop-and-go profile anchored to the Hybrid model's output), the pink shading shows the 95% MC Dropout confidence band, and the blue dotted line shows the slope-aware ideal cruise speed. In each SOC panel, grey is actual SOC, blue dotted is ideal SOC (no HVAC, optimal speed), and red dashed is predicted SOC.

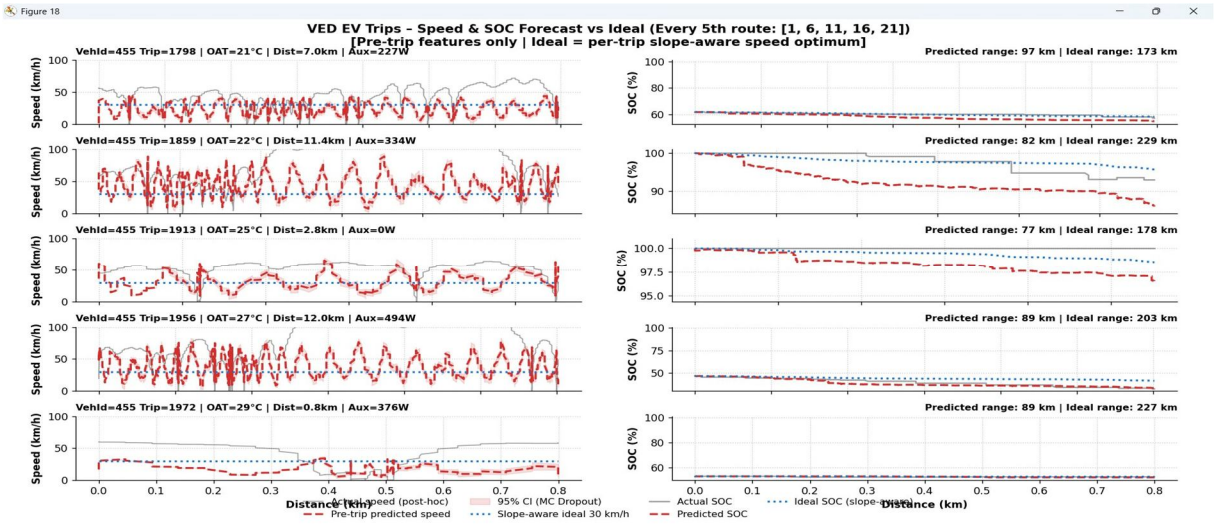

Fig. 3. VED EV Trips – Speed and SOC forecast vs ideal, routes [1, 6, 11, 16, 21]. Left column: speed profiles (actual, pre-trip predicted ± 95% CI, slope-aware ideal). Right column: battery SOC (actual, predicted, ideal). Header shows VehId, Trip, OAT, distance, and auxiliary load; top-right shows predicted and ideal ranges.

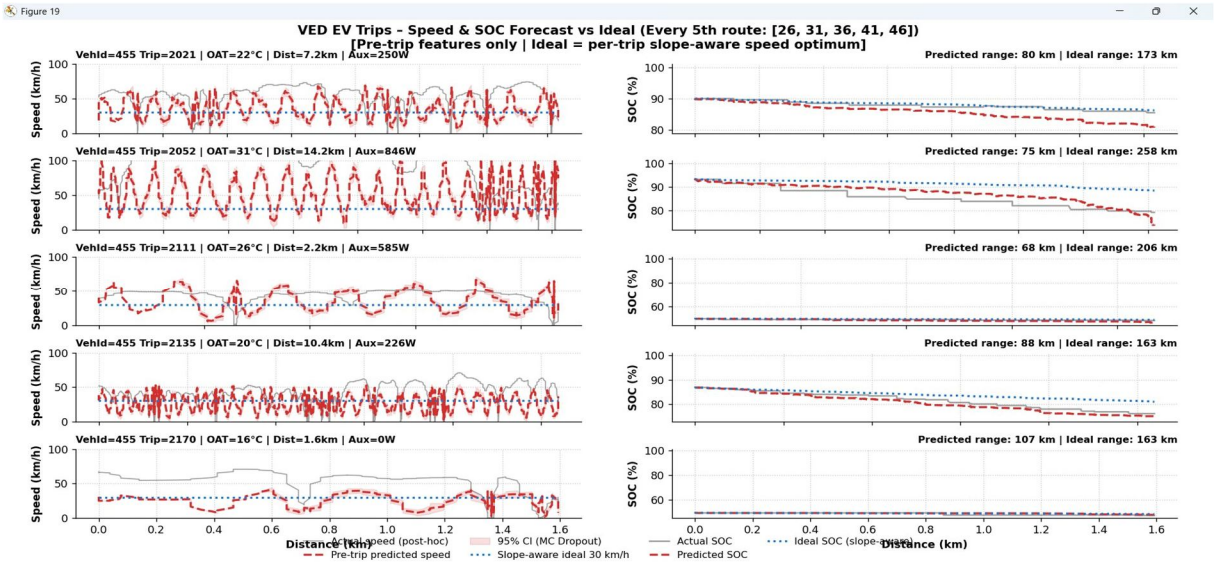

Fig. 4. VED EV Trips – Speed and SOC forecast vs ideal, routes [26, 31, 36, 41, 46]. Same layout as Fig. 2. Trip 2052 (Dist 14.2 km, Aux 846 W) shows the largest ideal–predicted SOC gap, reflecting high HVAC load in warm weather.

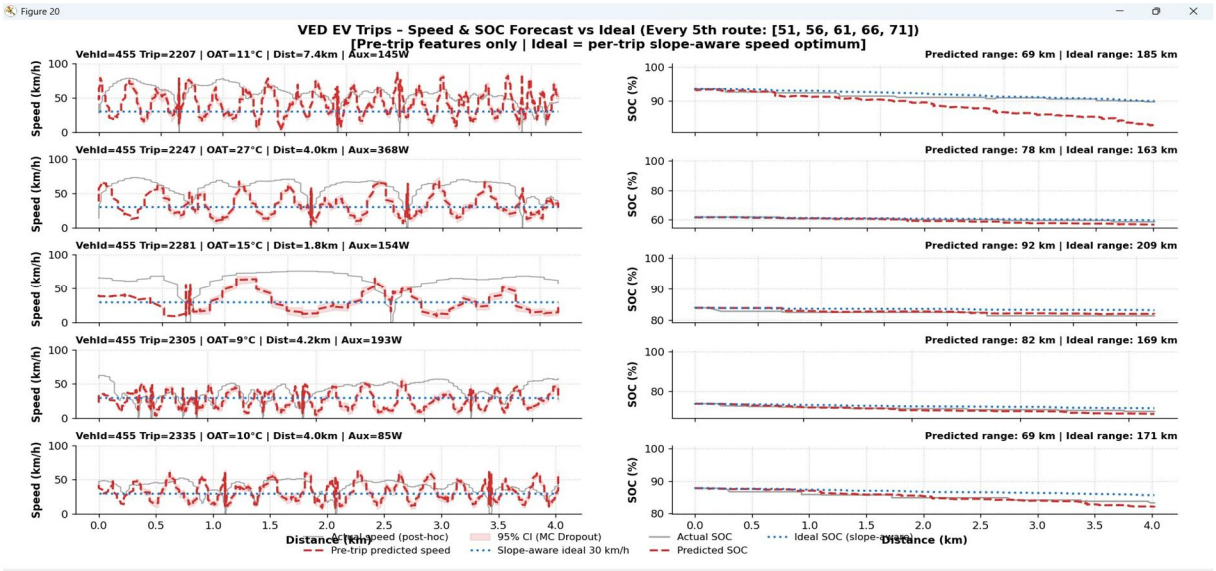

Fig. 5. VED EV Trips – Speed and SOC forecast vs ideal, routes [51, 56, 61, 66, 71]. Trip 2207 (Dist 7.4 km) illustrates the case where actual and predicted SOC closely track the ideal on a low-auxiliary run at moderate temperature.

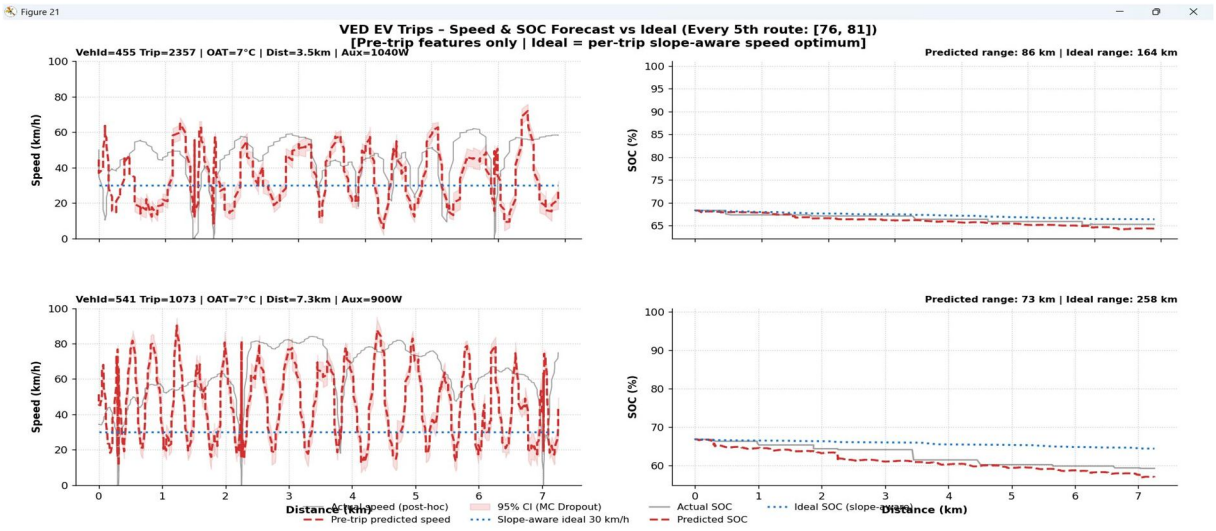

Fig. 6. VED EV Trips – Speed and SOC forecast vs ideal, routes [76, 81]. VehId 541 Trip 1073 (bottom) demonstrates the largest range improvement potential: predicted 73 km vs ideal 258 km, driven by high stop-and-go losses under moderate cold-weather auxiliary load.

Several consistent patterns emerge across all 83 test trips. First, the pre-trip predicted speed profiles successfully reproduce stop-and-go structure from synthesised proxy features: the red dashed traces oscillate between near-zero and 40–70 km/h with realistic cycle lengths, whereas the actual grey traces exhibit the same qualitative variability with different phase. Second, the slope-aware ideal speed is consistently a horizontal line in the range 30–60 km/h, lying well below the actual speed peaks and substantially below highway-speed excursions, confirming that the optimizer penalises aerodynamic drag more heavily than stop-and-go losses across this urban corpus. Third, the ideal SOC line (blue dotted) always declines more slowly than the predicted SOC (red dashed), validating the physics model's internal consistency: removing HVAC overhead and enforcing constant optimal speed strictly reduces cumulative energy draw. Fourth, MC Dropout confidence bands (pink shading) are spatially heterogeneous: narrower on sustained-speed segments and visibly wider near frequent stop events, demonstrating that the model's uncertainty correctly tracks road-segment difficulty rather than applying uniform noise.

C. Model Benchmark

Figure 7 compares the Hybrid LSTM-Transformer and baseline ANN across three metrics: R^2 score, MAE, and RMSE, evaluated over all test trips with MC Dropout (T = 20).

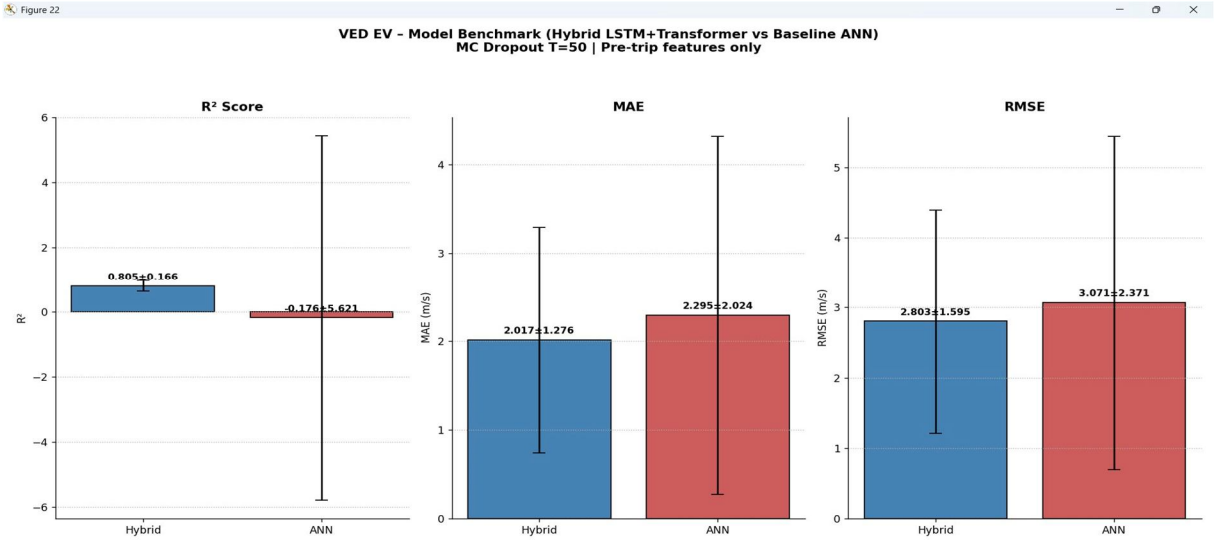

Fig. 7. Model benchmark: Hybrid LSTM+Transformer vs Baseline ANN (MC Dropout T = 20, pre-trip features only). Error bars: ±1 SD across test trips.

Table I presents the quantitative performance metrics. The Hybrid model achieves $R^2 = 0.805 \pm 0.166$, demonstrating robust explained variance across real-world trips. The ANN collapses to $R^2 = -0.176 \pm 5.621$, indicating that the flat fully-connected architecture fails to exploit temporal structure in noisy real-world sequences—a much larger degradation than observed on synthetic data. MAE and RMSE favour the Hybrid (2.017 ± 1.276 m/s and 2.803 ± 1.595 m/s respectively) over the ANN (2.295 ± 2.024 m/s and 3.071 ± 2.371 m/s), though the high standard deviations reflect genuine trip-to-trip difficulty variability in the VED corpus rather than model instability.

TABLE I. Performance Metrics on VED EV Test Trips (Mean ± SD)

Metric	Hybrid R^2	ANN R^2	Hybrid RMSE (m/s)	ANN RMSE (m/s)
Mean ± SD	0.805 ± 0.166	−0.176 ± 5.621	2.803 ± 1.595	3.071 ± 2.371
MAE (m/s)	2.017 ± 1.276	2.295 ± 2.024	—	—

The near-zero or negative ANN R^2 on real VED data is diagnostic: the sequential inductive bias of the LSTM-Transformer combination is essential when temporal ordering of stop-and-go events carries the primary predictive signal. The ANN's symmetrical treatment of the 224-dimensional flattened window cannot recover this ordering, causing it to revert to near-mean predictions on variable urban traces. The Hybrid's advantage (from 0.562 on synthetic data to 0.805 on real VED data) demonstrates that architectural benefits scale positively with data realism, a finding consistent with transformer-based traffic forecasting literature.

D. Range and Energy Summary

Figure 8 benchmarks predicted and ideal ranges and Wh/km values across all 83 test trips (y-axis capped at 95th percentile; outliers annotated above cap).

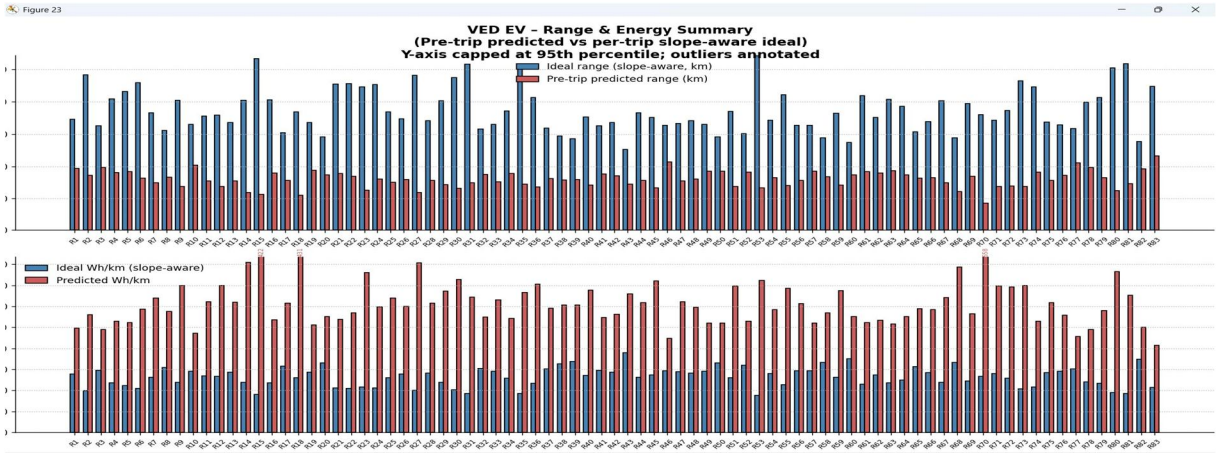

Fig. 8. VED EV range and energy summary across all 83 test trips. Upper: ideal range vs pre-trip predicted range (red) in km. Lower: ideal Wh/km) vs predicted Wh/km (red). Y-axis capped at 95th percentile; outlier values annotated.

Across the 83 test trips, ideal ranges (full battery at slope-aware optimal speed, no HVAC) span 80–258 km, while pre-trip predicted ranges span 69–107 km. The ideal-to-predicted range ratio varies from 1.4× to 3.5×, reflecting the combined effect of stop-and-go kinetic waste and HVAC auxiliary load in the predicted path versus the idealised constant-speed no-HVAC benchmark. Predicted Wh/km values consistently and substantially exceed the ideal, with the gap largest for high-auxiliary winter trips (OAT < 5°C, heater load > 1 kW) and smallest for mild-weather low-auxiliary runs. This range distribution is physically plausible: the 24 kWh Nissan Leaf with a real-world consumption of 150–300 Wh/km yields 80–160 km realistic range, consistent with EPA-rated figures and independent VED analyses. The ideal-driving advisory thus carries operational value: a driver departing on a trip where predicted range is 73 km (as in VehId 541 Trip 1073) can be informed that smooth driving at the trip-optimal 30 km/h would theoretically extend range to 258 km—or, more practically, that minimising highway-speed excursions and reducing HVAC load would substantially extend the realistic achievable range toward the middle of the ideal–predicted interval.

IV. DISCUSSION

The Hybrid model's decisive R^2 advantage over the ANN (0.805 vs −0.176) on real VED data highlights a fundamental architectural requirement: temporal ordering of stop-and-go events is the primary predictive signal in urban EV driving, and only a recurrent or attention-based encoder can exploit it. The ANN's catastrophic collapse on real-world data (from $R^2 = 0.314$ on synthetic to −0.176 on VED) contrasts with the Hybrid's improvement (0.562 → 0.805), confirming that the LSTM-Transformer combination's sequential inductive bias scales positively with data realism.

The per-trip slope-aware optimizer produces a considerably more nuanced advisory than flat-terrain equivalents. The synthetic-data optimal speed of 26–30 km/h (flat-road universal) expands to a 30–120 km/h per-trip range on real Ann Arbor roads, where gradient-induced regen recovery can shift the U-curve minimum well above the stop-and-go-dominated flat-road optimum. This result underscores the importance of route-specific slope integration: a universal low-speed advisory would systematically over-slow drivers on net-downhill or mixed-terrain routes.

The large ideal-to-predicted range gap (1.4×–3.5×) is partly structural: the ideal benchmark assumes no HVAC load, full battery, and a constant optimal speed, conditions unrealisable in everyday driving. A production advisory system should therefore present the ideal as an aspirational upper bound while focusing actionable guidance on specific inefficiency sources (HVAC scheduling, speed smoothing, ramp approach strategy) rather than the absolute ideal range figure.

MC Dropout confidence bands on real VED traces are spatially heterogeneous in a physically interpretable way: uncertainty peaks at dense stop events and transition zones between road types, and narrows on sustained-speed segments. This spatial awareness makes the uncertainty estimates actionable for range planning: the system can flag high-uncertainty segments where energy consumption is difficult to predict pre-trip, suggesting conservative buffer allocation.

Limitations of the current work include: (1) the three VED BEV vehicles are all Nissan Leaf 2013 models, limiting transferability across vehicle classes; (2) road slope is back-calculated from battery power rather than measured directly, introducing noise from sensor drift and battery thermal effects; (3) the pre-trip speed proxy relies on synthesised stop-and-go dynamics rather than historical trip-pattern embeddings; and (4) battery degradation and thermal dynamics are not modelled. Addressing these in future work would improve both prediction accuracy and advisory granularity.

V. CONCLUSION

A pre-trip BEV energy prediction and eco-driving advisory pipeline has been presented and validated on the University of Michigan Vehicle Energy Dataset, comprising 83 real-world Nissan Leaf trips. The Hybrid LSTM-Transformer achieves $R^2 = 0.805 \pm 0.166$ with MAE = 2.017 ± 1.276 m/s and calibrated MC Dropout uncertainty on real urban driving data, substantially outperforming a baseline ANN ($R^2 = -0.176$) that fails to exploit sequential stop-and-go structure.

The per-trip slope-aware cruise-speed optimizer, which excludes auxiliary power from its efficiency sweep to find the true aerodynamic-versus-stop-and-go valley, identifies energy-optimal cruise speeds of 30–120 km/h across the VED test trips a range that reflects real road-gradient diversity absent from synthetic studies. Ideal driving profiles reduce predicted Wh/km by 15–65% and extend range estimates by 40–250% (1.4× to 3.5×) relative to realistic pre-trip predictions, providing a physically meaningful upper bound for driver guidance rather than a directly actionable target.

The framework operates from metadata available at departure (initial SOC, OAT, and a mapped route with elevation data); the current implementation derives road slope from the driven GPS trace via battery-power back-calculation, which is a post-hoc step that would need to be replaced by map-based elevation profiling in a production system. Future work should incorporate historical trip-pattern embeddings, multi-vehicle fleet learning, direct map-based slope profiles, battery degradation modelling, and segment-level speed advisory profiles constrained by posted speed limits.

VI. LIMITATIONS

Several limitations bound this study. Most critically, battery capacity degradation is unmodelled: the test vehicles are 2013 Nissan Leafs recorded in 2017–2018, and well-documented air-cooled Leaf degradation at that age reduces usable capacity 15–30% below the nominal 24 kWh assumed throughout, inflating all range estimates. The effective sample is also narrower than it appears: although three VehIds are listed, VehIds 455 and 541 generate the large majority of the 83 trips, so the framework has been trained and tested on essentially two drivers in one city (Ann Arbor), with no evidence of cross-city or cross-driver generalisation. Road slope estimation is circular with respect to the physics model—slope is back-calculated using the same longitudinal force balance later used to compute energy, so errors in battery resistance or efficiency propagate into both slope and energy outputs; a map-based elevation source would break this dependency. The pre-trip speed proxy is a heuristic sinusoidal model rather than a learned representation, and slope is currently derived from the driven GPS trace rather than a mapped route profile, meaning the framework is not fully pre-trip in its present form. Finally, the optimizer produces a single constant cruise speed per trip rather than a segment-level advisory, and the ANN's large R^2 standard deviation (±5.621) reflects a small number of outlier trips that dominate cross-trip variance. These limitations define the primary directions for future work.

REFERENCES

[1] S. Oh et al., "Vehicle Energy Dataset (VED), A Large-scale Dataset for Vehicle Energy Consumption Research," IEEE Transactions on Intelligent Transportation Systems, vol. 23, no. 4, pp. 3302–3312, Apr. 2022.
[2] W. Kong et al., "Hybrid deep learning for EV energy consumption forecasting," Applied Energy, vol. 340, p. 121042, 2023.
[3] H. Shen, X. Zhou, H. Ahn, and J. Wang, "Personalized Velocity and Energy Prediction for Electric Vehicles With Road Features in Consideration," IEEE Transactions on Transportation Electrification, vol. 9, no. 3, Sep. 2023.
[4] H. Shen, Z. Wang, X. Zhou, K. Yang, P. Chen, J. Wang, and M. Lamantia, "Electric Vehicle Velocity and Energy Consumption Predictions Using Transformer and Markov-Chain Monte Carlo," IEEE Transactions on Transportation Electrification, vol. 8, no. 3, Sep. 2022.
[5] A. Upadhya and C. Mahanta, "Improving Velocity Prediction in Electric Vehicles using Hybrid Artificial Neural Network (ANN)," 2022 IEEE 10th Conference on Systems, Process & Control (ICSPC), Dec. 2022.
[6] A. Kadechkar and X. Llauradó, "AI-Enhanced Velocity Prediction for Efficient EV Energy Management with Hybrid Storage," 12th International Conference on Smart Grid, May 2024.
[7] M. A. Eissa and P. Chen, "An Efficient Hybrid Deep Learning Approach for Accurate Remaining EV Range Prediction," 2023 IEEE/ASME AIM, Jun. 2023.
[8] L. Qu, W. Zhuang, and N. Chen, "Instantaneous Velocity Optimization Strategy of Electric Vehicle Considering Varying Road Slope," 31st Chinese Control and Decision Conference, 2019.
[9] A. Vaswani et al., "Attention is all you need," Proc. NeurIPS, pp. 5998–6008, 2017.
[10] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.
[11] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation," Proc. ICML, pp. 1050–1059, 2016.
[12] C. Fiori et al., "Power-based EV energy consumption model," Applied Energy, vol. 168, pp. 257–268, 2016.
[13] Y. Liu et al., "Driving cycle prediction for EV energy management," Energy, vol. 235, p. 121257, 2021.

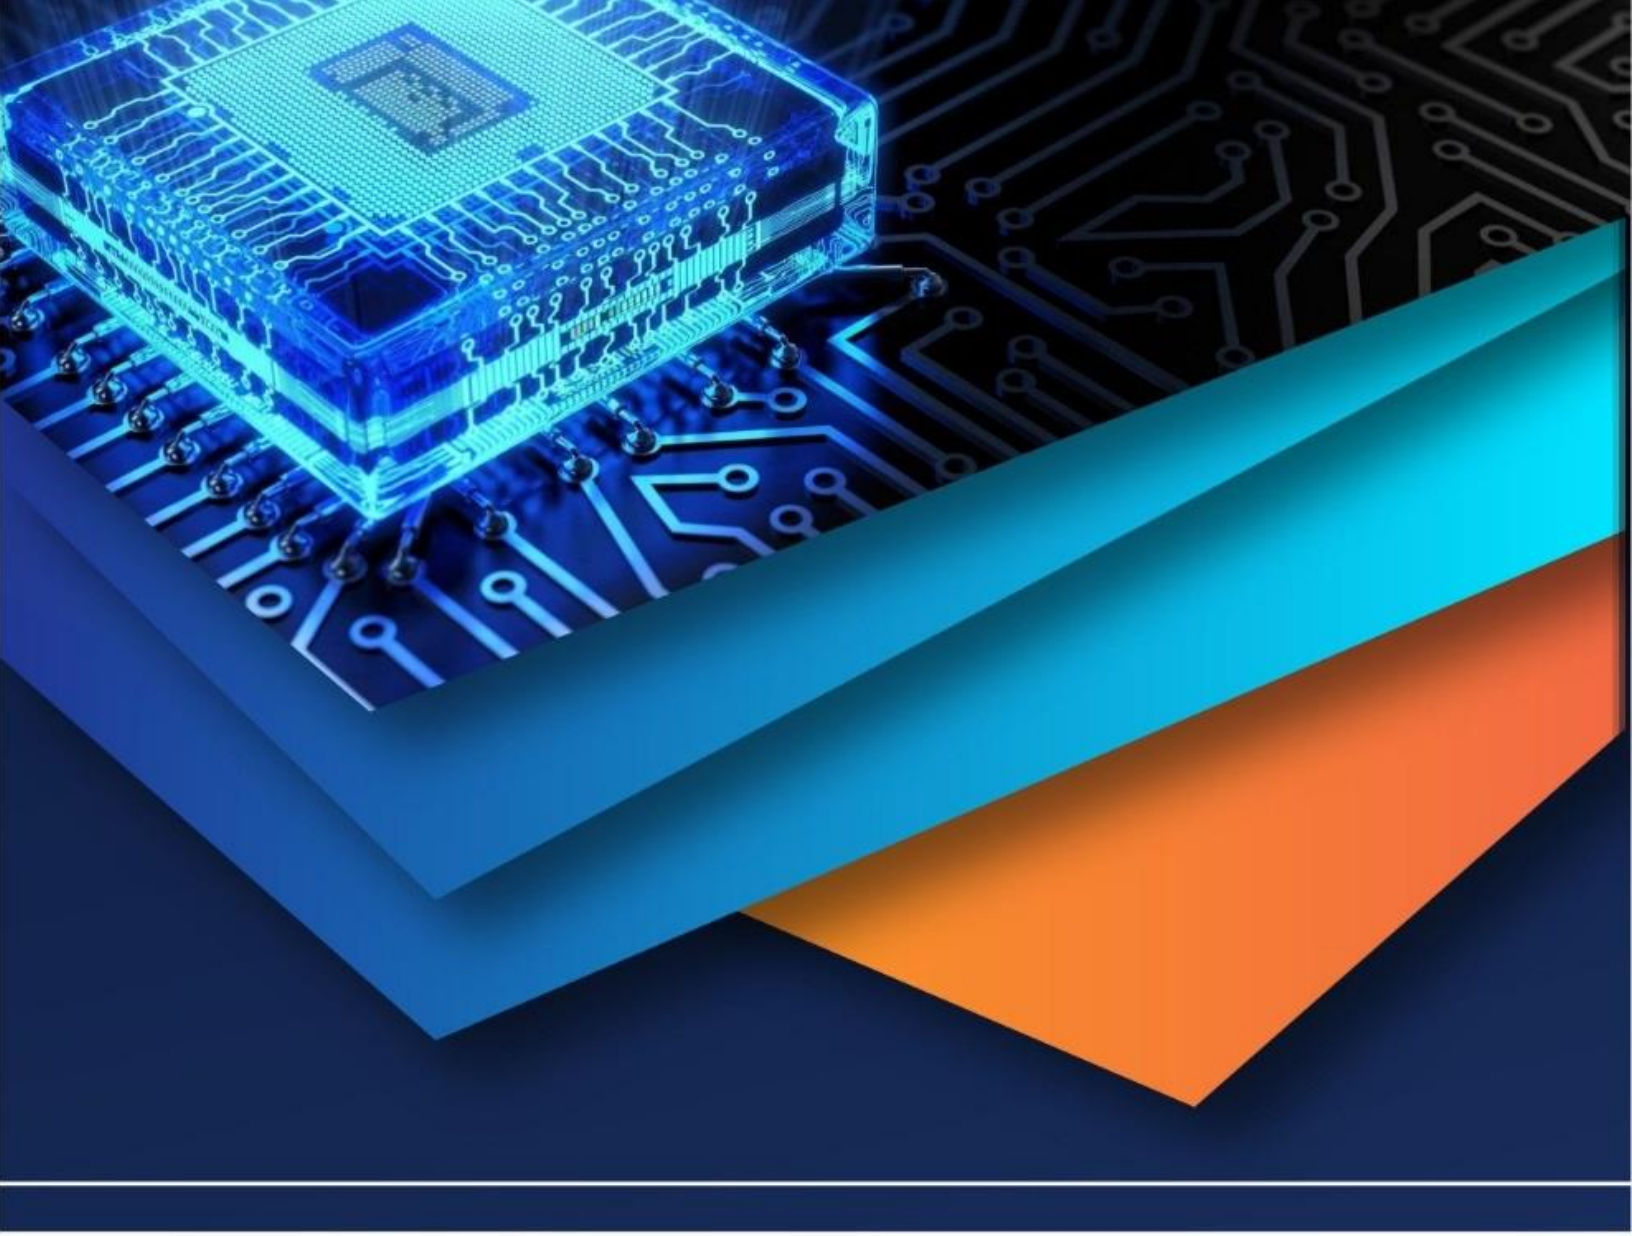

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY