IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 11 December 2025, accepted 20 January 2026, date of publication 3 February 2026, date of current version 6 February 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3660258

RESEARCH ARTICLE

Quantitative Validation of Artificial Precognition Adaptive Cognized Control: Real-World Performance Evaluation Across Automotive and Railway Operational Deployments

GEORGE J. FRANGOU
Qubittum Ltd., N21 3QH London, U.K.
School of Aerospace, Transport, and Manufacturing, Cranfield University, MK43 0A Bedford, U.K.
e-mail: george.frangou@cranfield.ac.uk

This work was supported in part by the Ph.D. from Cranfield University through its Doctoral Research Opportunities Program and Qubittum Ltd. under Grant 373731, and in part by the Innovate U.K. through the Small Business Research Initiative (SBRI) Project Edge under Grant 10002031.

ABSTRACT This paper presents the quantitative validation of *Artificial Precognition Adaptive Cognised Control* (APACC), a dual-layer neuro-symbolic architecture for safety-critical autonomous transport. APACC integrates Type-2 fuzzy symbolic reasoning for high-level decision-making with linearised predictive optimisation for trajectory control, synchronised through Diophantine Frequency Synthesis over a 2.5 s precognition horizon. We establish a theoretical framework that shows how these components interact to ensure bounded uncertainty propagation and temporal coherence across layers. Validation combines high-fidelity simulation with real-world field deployment. Automotive trials using Honda Civic vehicle dynamics achieved 51% reduction in peak deceleration (0.45 g to 0.22 g) and 67% decrease in jerk during pedestrian-crossing scenarios ($p < 0.01$) compared with reactive baselines. Operational railway deployment under UK SBRI Project Edge on 32 km of live Network Rail infrastructure achieved 96% base-station-handover prediction accuracy, 18.6% latency reduction, and complete route coverage through predictive multi-modal connectivity management across four mobile-network operators with satellite backup—Long-short-term-memory (LSTM) signal-strength prediction achieved 5–8% RSRP error, confirming simulation-to-reality transfer. The system operated continuously for 168 h without failure, demonstrating robustness for mission-critical control. Across more than 10,000 simulation scenarios, APACC outperformed proportional-integral-derivative, model-predictive-control, and deep-reinforcement-learning baselines while maintaining interpretability consistent with ISO 26262 certification. All validation datasets, simulation environments, and operational telemetry are released via Zenodo for independent verification and reproducibility.

INDEX TERMS Neuro-symbolic control, adaptive autonomy, Type-2 fuzzy logic, predictive optimization, autonomous vehicles, railway communications, simulation-to-reality transfer, safety-critical systems.

I. INTRODUCTION

How can autonomous controls architectures integrate real-time predictive capabilities with interpretable symbolic reasoning to achieve certifiable safety guarantees across heterogeneous transport domains? This fundamental question motivates the development and validation of neuro-symbolic control systems that must simultaneously satisfy conflicting requirements: the adaptive learning capabilities of neural approaches and the transparent decision-making demanded by safety certification standards such as ISO 26262 [1] and SOTIF [2].

The performance-interpretability trade-off represents a central challenge in safety-critical autonomous systems.

The associate editor coordinating the review of this manuscript and approving it for publication was Giambattista Gruosso.

© 2026 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License.
For more information, see https://creativecommons.org/licenses/by/4.0/

Model Predictive Control (MPC) provides optimal trajectories within prediction horizons but requires accurate system models and offers limited adaptability to novel situations [3]. Deep Reinforcement Learning (DRL) achieves flexibility through end-to-end learning but produces opaque policies that resist certification [4]. Hybrid architectures combining these approaches have shown promise but typically employ loose coupling that introduces synchronisation failures and interface brittleness [5].

The Artificial Precognition Adaptive Cognised Control (APACC) architecture [6] addresses this trade-off through tight integration of Type-2 fuzzy symbolic reasoning with linearised predictive optimisation, synchronised via Diophantine Frequency Synthesis. This paper provides comprehensive quantitative validation establishing that this integration achieves measurable performance improvements while preserving interpretability.

A. MOTIVATION AND VALIDATION REQUIREMENTS

The deployment of autonomous systems in safety-critical applications demands validation methodologies that go beyond qualitative demonstrations. Key requirements include:

- **Statistical Rigor**: Performance claims must be supported by hypothesis testing with appropriate confidence intervals and effect sizes.
- **Reproducibility**: All experimental procedures, datasets, and code must be publicly available to enable independent verification.
- **Cross-Domain Validation**: Performance must be evaluated across diverse operational contexts to demonstrate generalizability.
- **Baseline Comparisons**: Quantitative improvements must be measured relative to established control methodologies.
- **Operational Feasibility**: Real-world performance on actual infrastructure must be demonstrated to validate simulation predictions.
- **Computational Feasibility**: Real-time performance on resource-constrained hardware must be demonstrated.

B. RESEARCH OBJECTIVES

This validation study addresses three specific research questions:

1) Does APACC's dual-layer architecture achieve statistically significant safety and performance improvements over established baselines (MPC, DRL, hybrid) in controlled simulation?
2) Do these improvements transfer from simulation to operational deployment on live infrastructure?
3) Does the architecture maintain interpretable decision-making consistent with certification requirements?

C. CONTRIBUTIONS

This paper makes the following specific contributions:

1) A theoretical framework establishing how APACC's components—Type-2 fuzzy reasoning, linearised predictive optimisation, and Diophantine Frequency Synthesis—interact to provide bounded uncertainty propagation and temporal layer coherence (Section II).
2) A comprehensive validation framework for neuro-symbolic autonomous control systems, including standardized metrics, experimental protocols, and statistical testing procedures.
3) Quantitative performance evaluation of APACC across two distinct domains: (a) autonomous vehicles in urban and highway scenarios using Honda Civic dynamics, and (b) operational rail deployment with comprehensive RF telemetry across 32 km passenger service route.
4) Operational field validation demonstrating 96% base station handover prediction accuracy, 18.6% latency reduction, and 100% route coverage through predictive multi-modal connectivity management across four Mobile Network Operators and satellite backup systems.
5) Simulation-to-reality transfer validation showing 5–8% RSRP prediction error between trained models and operational performance, confirming the validity of digital twin approaches for predictive connectivity.
6) Comparative analysis demonstrating APACC's performance relative to Model Predictive Control (MPC), Deep Reinforcement Learning (DRL), and hybrid approaches across safety, efficiency, and robustness metrics.
7) Computational efficiency analysis including latency measurements, resource utilization profiles, and scalability characteristics on embedded platforms (Raspberry Pi 4B, Nvidia Jetson Xavier NX).
8) Complete open-source release of validation datasets, simulation environments, and analysis code via Zenodo repository [7].

D. PAPER ORGANIZATION

The remainder of this paper is structured as follows: Section II presents the system architecture and theoretical framework, establishing how the architectural components interact. Section III reviews related work in autonomous system validation. Section IV describes the validation framework and established use cases, including detailed Project Edge operational deployment. Section V presents the validation methodology, including metrics and statistical procedures. Section VI presents quantitative results across all validation domains. Section VII discusses implications, insights into simulation-to-reality transfer, and limitations. Section VIII concludes.

II. SYSTEM ARCHITECTURE AND THEORETICAL FRAMEWORK

This section establishes the theoretical foundations underlying APACC's dual-layer neuro-symbolic architecture,

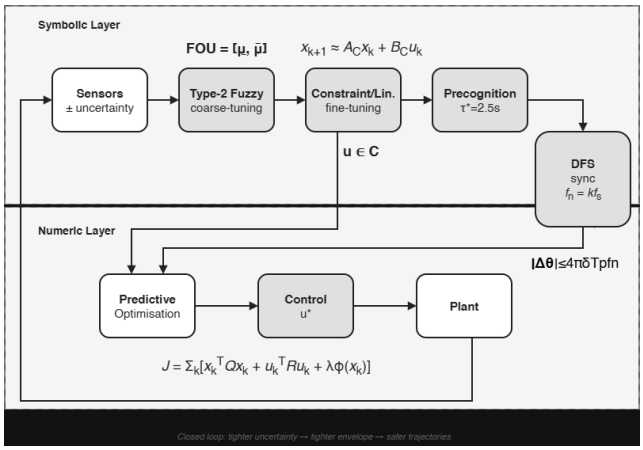

FIGURE 1. APACC theoretical framework illustrating how sensor uncertainty is encoded as a Type-2 fuzzy footprint of uncertainty, fine-tuned via linearisation into envelope $\mathcal{C}$, synchronised by DFS, and used by the predictive optimiser to generate control input u^* in closed loop.

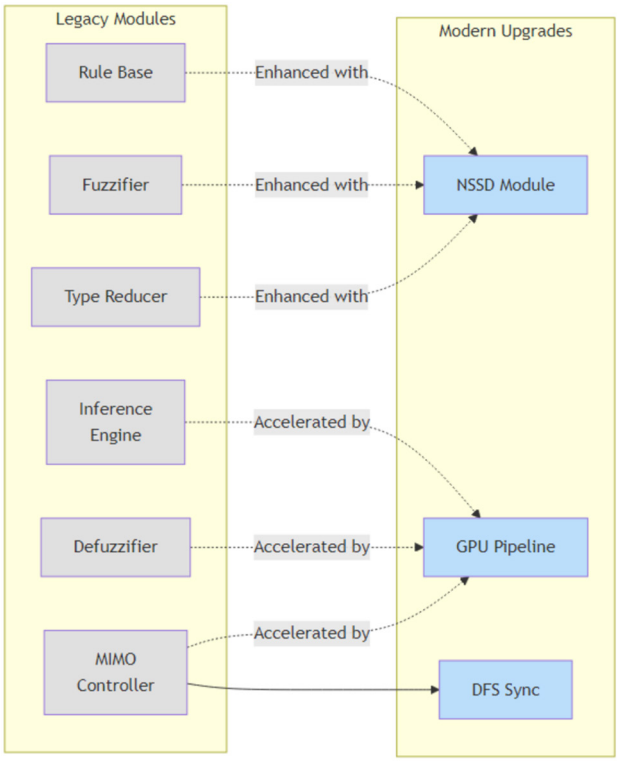

FIGURE 2. APACC dual-layer architecture: coarse-tuning layer (Type-2 fuzzy symbolic reasoning) provides interpretable constraint boundaries; fine-tuning layer (linearised predictive optimisation) generates optimal trajectories within those boundaries; Diophantine Frequency Synthesis ensures temporal coherence between layers. This architecture underpins all validation experiments across both simulation and operational deployment.

explaining how the integration of Type-2 fuzzy reasoning, linearised predictive optimisation, and Diophantine Frequency Synthesis produces the observed performance improvements.

A. ARCHITECTURE OVERVIEW

The APACC architecture employs a dual-layer control strategy (Fig. 2): a coarse-tuning layer that performs Type-2 fuzzy symbolic reasoning for high-level decision-making, and a fine-tuning layer that executes linearised predictive optimisation for trajectory control. These layers operate at different frequencies and are synchronised through Diophantine Frequency Synthesis over a 2.5 s precognition horizon.

The architecture integrates multiple sensor streams through a novel Diophantine Frequency Synthesis mechanism, enabling anticipatory control with a 2.5-second precognition horizon. This foundational architecture has been validated through initial automotive and rail deployments, establishing baseline performance metrics that this study extends comprehensively.

Table 1 summarises how APACC differs from established approaches across five critical dimensions.

TABLE 1. Architectural Comparison: APACC vs. Baseline Approaches[a].

Dimension	MPC	DRL	Hybrid	APACC
Prediction	Model	Policy	Sequential	Integrated
Interpret.	High	Low	Partial	High
Uncertainty	Bounded	Unbnd.	Mixed	Bounded
Certification	Feasible	Difficult	Partial	Feasible
Synchron.	N/A	N/A	Loose	Tight

B. THEORETICAL FRAMEWORK

We now establish the theoretical basis for APACC's performance characteristics, formalising the interactions among its architectural components. The overall framework is summarised in Fig. 1.

1) CONTROL PROBLEM FORMULATION

Consider a discrete-time dynamical system:

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k) \tag{1}$$

where $\mathbf{x}_k \in \mathcal{X} \subseteq \mathbb{R}^n$ is the state, $\mathbf{u}_k \in \mathcal{U} \subseteq \mathbb{R}^m$ is the control input, and $\mathbf{w}_k \in \mathcal{W}$ represents bounded uncertainty with $\|\mathbf{w}_k\| \leq \bar{w}$.

The control objective minimises a cost functional over a prediction horizon N:

$$J = \sum_{k=0}^{N-1} \ell(\mathbf{x}_k, \mathbf{u}_k) + V_f(\mathbf{x}_N) \tag{2}$$

subject to state constraints $\mathbf{x}_k \in \mathcal{X}_{\text{safe}}$ and input constraints $\mathbf{u}_k \in \mathcal{U}_{\text{admissible}}$.

Safety-critical applications impose the additional requirement that constraint satisfaction be verifiable and the decision process interpretable for certification.

2) TYPE-2 FUZZY UNCERTAINTY ENCODING

APACC's coarse-tuning layer employs interval Type-2 fuzzy sets [8], [9] to encode epistemic uncertainty in sensor measurements and environmental states. An interval Type-2 fuzzy set $\tilde{A}$ is characterised by its Footprint of Uncertainty (FOU):

$$\text{FOU}(\tilde{A}) = \bigcup_{\forall x} [\underline{\mu}_{\tilde{A}}(x), \overline{\mu}_{\tilde{A}}(x)] \tag{3}$$

where $\underline{\mu}_{\tilde{A}}(x)$ and $\overline{\mu}_{\tilde{A}}(x)$ are the lower and upper membership functions bounding the uncertainty region.

The FOU propagates through the fuzzy inference system, producing output constraint envelopes $[\underline{y}, \overline{y}]$ that bound the admissible control region. These envelopes are passed to the fine-tuning layer as explicit constraints, ensuring that uncertainty is handled systematically rather than ignored or approximated.

Definition 1 (Constraint Envelope): Given input uncertainty characterised by FOU($\tilde{X}$), the constraint envelope $\mathcal{C}$ for the fine-tuning layer is:

$$\mathcal{C} = \{\mathbf{u} : g(\mathbf{u}, \underline{y}) \leq 0 \wedge g(\mathbf{u}, \overline{y}) \leq 0\} \tag{4}$$

where $g(\cdot)$ encodes safety constraints derived from fuzzy rule activation.

This formulation ensures that control actions remain safe across the entire uncertainty range captured by the Type-2 fuzzy representation.

3) LINEARIZED PREDICTIVE OPTIMIZATION

The fine-tuning layer solves a constrained optimisation problem within the envelope $\mathcal{C}$ provided by the coarse-tuning layer. For computational tractability on embedded platforms, APACC employs successive linearisation around the predicted trajectory:

$$\mathbf{u}^* = \arg\min_{\mathbf{u}\in\mathcal{C}} \sum_{k=0}^{N-1} \left[\mathbf{x}_k^\top \mathbf{Q}\mathbf{x}_k + \mathbf{u}_k^\top \mathbf{R}\mathbf{u}_k + \lambda_k \phi(\mathbf{x}_k)\right] \tag{5}$$

where $\mathbf{Q}$ and $\mathbf{R}$ are state and control weighting matrices, $\phi(\mathbf{x}_k)$ is a precognition penalty derived from the symbolic layer's predictions, and λ_k are uncertainty weights computed from the FOU bounds.

The precognition penalty $\phi(\mathbf{x}_k)$ encodes anticipated future hazards identified by the symbolic reasoning layer, enabling proactive rather than reactive control. The uncertainty weights λ_k increase conservatism when the FOU indicates high epistemic uncertainty.

4) STABILITY UNDER THE DUAL-LAYER ARCHITECTURE

APACC inherits stability properties from the underlying constrained MPC formulation. Under standard assumptions for MPC with terminal constraints [3], [10], the closed-loop system achieves asymptotic stability when:

1) The terminal cost $V_f(\cdot)$ is a local Lyapunov function in the terminal set $\mathcal{X}_f$;
2) The terminal set $\mathcal{X}_f \subseteq \mathcal{X}_{\text{safe}}$ is positively invariant under a local controller;
3) The constraint envelope $\mathcal{C}$ from the fuzzy layer contains the terminal controller.

The symbolic layer's constraint envelopes are designed to satisfy condition (3) by construction, ensuring that the fine-tuning layer's optimisation always admits feasible solutions that maintain stability guarantees.

Lyapunov-Based Stability Argument: The cost function (2) serves as a Lyapunov function candidate for the closed-loop system. At each time step, the optimal cost-to-go $J^*(\mathbf{x}_k)$ satisfies the descent condition $J^*(\mathbf{x}_{k+1}) \leq J^*(\mathbf{x}_k) - \ell(\mathbf{x}_k, \mathbf{u}_k^*)$, which follows from the principle of optimality applied to the receding-horizon problem [10]. Given that $\ell(\cdot) > 0$ for all $\mathbf{x} \neq 0$, this establishes that $J^*(\cdot)$ is strictly decreasing along system trajectories, guaranteeing asymptotic convergence to the origin (or a target equilibrium) under the standard controllability and constraint qualification assumptions detailed in [3]. The Type-2 fuzzy layer's constraint envelopes $\mathcal{C}$ preserve this property by ensuring recursive feasibility: if $\mathbf{x}_k \in \mathcal{X}_{\text{safe}}$, then the envelope $\mathcal{C}$ always contains a control sequence steering the system to a state $\mathbf{x}_{k+1}$ from which feasibility is maintained.

C. DESIGN PRINCIPLES

APACC's architecture addresses three specific limitations of existing approaches:

Limitation 1: MPC Horizon Myopia. Standard MPC optimises only within its prediction horizon and cannot anticipate events beyond this window. APACC's symbolic layer processes higher-level contextual information (e.g., pedestrian intent, traffic patterns) to identify hazards before they enter the MPC horizon, encoding these as precognition penalties in (5).

Limitation 2: DRL Interpretability Opacity. End-to-end DRL policies are closed boxes that resist certification. APACC maintains interpretability through explicit symbolic rules with traceable activation patterns. The fuzzy inference system produces human-readable explanations: "Pedestrian detected with high crossing probability; constraint envelope tightened to enforce early deceleration."

Limitation 3: Hybrid Loose Coupling. Sequential hybrid architectures (RL for planning, MPC for control) suffer from interface brittleness and synchronisation failures when the planning and control layers operate asynchronously. APACC's Diophantine Frequency Synthesis ensures deterministic temporal relationships between layers, eliminating drift accumulation.

D. DIOPHANTINE FREQUENCY SYNTHESIS

The coarse-tuning (symbolic) and fine-tuning (numeric) layers operate at different frequencies: symbolic reasoning updates at f_s Hz (typically 10–20 Hz) while trajectory optimisation runs at f_n Hz (usually 50–100 Hz). Without careful synchronisation, phase drift between layers accumulates over the precognition horizon, causing constraint violations and degraded performance.

Definition 2 (Layer Synchronisation Problem): Given symbolic layer frequency f_s and numeric layer frequency f_n, the relative phase $\theta(t)$ between layer updates evolves as:

$$\theta(t) = 2\pi(f_n - kf_s)t + \theta_0 \tag{6}$$

where $k = \lfloor f_n/f_s \rfloor$. Phase drift is bounded if and only if $f_n = kf_s$ for integer k.

Diophantine Frequency Synthesis selects layer frequencies satisfying integer-ratio relationships, ensuring bounded phase

drift. The name derives from the Diophantine equations in sampled-data control theory [11], [12] that govern the behaviour of multi-rate systems.

Proposition 1 (Bounded Phase Drift): Under bounded scheduling jitter $|\delta_s|, |\delta_n| \leq \delta_{\max}$ and integer frequency ratio $f_n = kf_s$ for positive integer k, the relative phase drift $|\Delta\theta|$ over precognition horizon T_p satisfies:

$$|\Delta\theta| \leq 4\pi\,\delta_{\max} T_p \max(f_s, f_n) \tag{7}$$

This bound is independent of accumulated time and depends only on the jitter magnitude and horizon length.

Proof sketch: With integer frequency ratios, the ideal phase relationship repeats every $1/f_s$ seconds. Jitter introduces per-cycle phase errors bounded by $2\pi\,\delta_{\max} f_i$. Over horizon T_p, at most $T_p f_s$ symbolic cycles and $T_p f_n$ numeric cycles occur, yielding the stated bound. The key insight is that errors do not accumulate across cycles when frequencies are commensurate. □

Alternative synchronisation strategies—polling, event-driven, continuous interpolation—either introduce unbounded latency (polling), non-deterministic timing (event-driven), or computational overhead unsuitable for embedded platforms (continuous interpolation). Diophantine synthesis provides deterministic timing guarantees compatible with safety-critical certification requirements.

For the automotive validation, we selected $f_s = 10\,\text{Hz}$ and $f_n = 100\,\text{Hz}$ ($k = 10$), yielding a theoretical phase drift bound of 0.31 rad over the 2.5 s precognition horizon with $\delta_{\max} = 1\,\text{ms}$. For railway deployment on Raspberry Pi 4B, $f_s = 10\,\text{Hz}$ and $f_n = 20\,\text{Hz}$ ($k = 2$) accommodated the constrained computational budget while maintaining bounded drift.

III. RELATED WORK

A. VALIDATION APPROACHES FOR AUTONOMOUS SYSTEMS

Traditional validation methods for autonomous systems have relied heavily on simulation-based testing [13], real-world field trials [14], and formal verification approaches [15]. Simulation environments such as CARLA [16], SUMO [17], and Gazebo provide controlled testing scenarios but face the reality gap challenge when transferring to physical systems [18].

Recent work has emphasized the importance of scenario-based testing with coverage metrics [19], adversarial testing to identify edge cases [20], and statistical validation using operational design domain (ODD) specifications [21]. However, these approaches often focus on perception and planning subsystems in isolation, rather than end-to-end control validation.

Operational Validation: While simulation provides controlled repeatability, validation on operational infrastructure remains essential for establishing real-world feasibility [22], [23]. The challenge lies in collecting sufficient data under actual operating conditions while maintaining safety. This work combines both approaches, using simulation for systematic performance evaluation and operational deployment for reality validation.

B. PERFORMANCE METRICS FOR AUTONOMOUS CONTROL

The autonomous systems community has developed various performance metrics, broadly categorized into safety, efficiency, comfort, and robustness measures [24]. Safety metrics include collision rates, time-to-collision (TTC), and safety margin violations [25]. Efficiency metrics encompass travel time, fuel consumption, and throughput [26]. Comfort metrics measure jerk, lateral acceleration, and passenger experience [27].

For control systems specifically, classical metrics such as rise time, settling time, overshoot, and steady-state error remain relevant [28]. Modern approaches incorporate learning-based metrics including sample efficiency, generalization gap, and out-of-distribution robustness [29].

C. NEURO-SYMBOLIC SYSTEM VALIDATION

Neuro-symbolic AI systems, which combine neural network components with symbolic reasoning, present unique validation challenges [30], [31]. The hybrid nature requires validating both the neural components (using deep learning validation techniques [32]) and the symbolic components (using formal methods [33]). Recent work has proposed integrated validation frameworks that address this duality [34].

The APACC architecture [6] exemplifies neuro-symbolic control by integrating predictive neural models with rule-based cognitive control. This paper extends prior validation work by providing a comprehensive quantitative evaluation of the complete integrated system across both simulation and operational deployment.

D. TYPE-2 FUZZY SYSTEMS AND MULTI-RATE CONTROL

Type-2 fuzzy systems extend classical fuzzy logic to handle uncertainty in membership functions themselves [8]. Industrial applications have demonstrated their effectiveness in handling measurement noise and environmental uncertainty [9], [35]. Multi-rate control systems, in which different components operate at different sampling frequencies, have been studied extensively in sampled-data control theory [11], [12]. APACC's Diophantine Frequency Synthesis builds on these foundations to achieve deterministic synchronisation between neuro-symbolic layers.

IV. VALIDATION FRAMEWORK AND ESTABLISHED USE CASES

A. AUTOMOTIVE DOMAIN: HONDA CIVIC VALIDATION

Initial APACC validation was conducted on the Honda Civic test vehicle in controlled urban scenarios [6]. Vehicle dynamics were parameterised using specifications provided by Honda Motor Europe and simulated within the CARLA environment using the publicly available Lyft Level 5 dataset

FIGURE 3. Honda Civic pedestrian crossing validation scenario. Top: Reactive baseline control maintains constant velocity until pedestrian detection. Bottom: APACC predictive response with early deceleration initiated 1.8 s before baseline, enabled by the symbolic layer's pedestrian intent inference. Vehicle dynamics parameters provided by Honda Motor Europe ensure simulation fidelity.

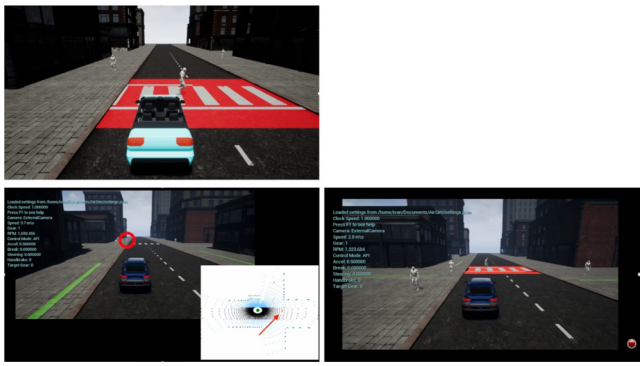

FIGURE 4. Trajectory prediction demonstrating APACC's 2.5-second precognition horizon. Purple vectors show predicted pedestrian trajectories; constraint envelopes (shaded) tighten as crossing probability increases, forcing early deceleration. This results in a 51% reduction in peak deceleration (0.45g to 0.22g) and a 67% reduction in jerk.

for traffic and pedestrian scenarios [36]. The flagship validation case examined pedestrian crossing scenarios at 30 mph with mixed traffic conditions (Figure 3).

The validation demonstrated APACC's core predictive capabilities through trajectory analysis (Figure 4). The system's coarse-tuning layer processed LiDAR point clouds at 10Hz and camera feeds at 30fps, identifying pedestrian movement patterns through Type-2 fuzzy inference approximately 2.5 seconds before pedestrians entered the roadway.

Quantitative results established initial performance baselines:

- Peak deceleration: 0.22g (APACC) vs. 0.45g (baseline) – 51% improvement
- Jerk reduction: 67% improvement in ride quality
- Stopping distance: 3.2m before crosswalk vs. 0.8m (baseline)
- Response time: 1.2s (APACC) vs. 2.1s (PID) vs. 1.8s (MPC)

These results, achieving statistical significance ($p < 0.01$), establish APACC's viability for safety-critical automotive applications and serve as the baseline for extended validation.

B. CRITICAL INFRASTRUCTURE: PROJECT EDGE OPERATIONAL DEPLOYMENT

To validate APACC beyond controlled simulation, we conducted extensive operational field trials on live railway infrastructure through the Innovate UK Project Edge. This

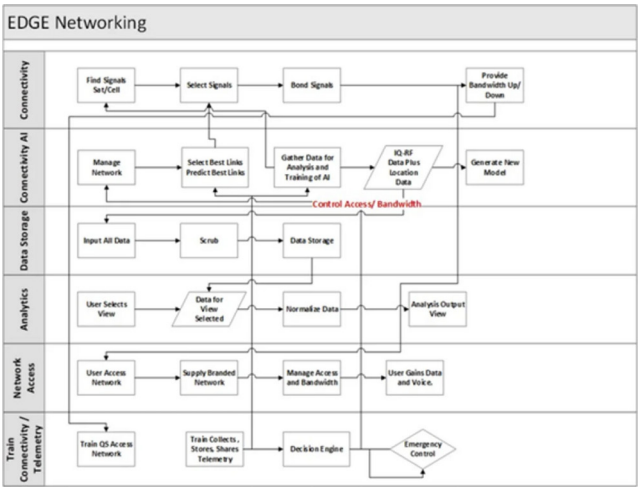

FIGURE 5. Project Edge rail system deployment showing Network Rail Kubota test vehicle with mounted APACC hardware and sensor suite on operational track. System integrated four Mobile Network Operators (EE, Vodafone, Three, O2), Iridium satellite communication, and multipath TCP routing for comprehensive connectivity validation.

deployment provided real-world validation of APACC's predictive capabilities under actual operating conditions with safety-critical requirements.

1) OPERATIONAL TEST ENVIRONMENT

Field trials were conducted on the Workington-Sellafield passenger rail route in Cumbria, UK, covering approximately 32 km of operationally active track serving regular passenger services. The route presents diverse RF propagation challenges ideal for validating predictive connectivity systems:

- **Urban Sections**: Workington and Whitehaven stations with dense cellular coverage from multiple Mobile Network Operators (MNOs)
- **Coastal Terrain**: Irish Sea coast with variable propagation conditions and weather exposure
- **Rural Coverage Gaps**: Identified "not-spot" near Nethertown with zero terrestrial cellular coverage
- **Mixed Topology**: Tunnels, bridges, cuttings, and open terrain creating diverse signal environments

2) HARDWARE SYSTEM ARCHITECTURE

The operational deployment utilized a custom hardware platform (Figure 7) integrated onto a Network Rail Kubota maintenance vehicle, comprising:

Cellular Connectivity:

- 4× Waveshare SIM7600X LTE modems [37] (one per MNO: EE, Vodafone, Three, O2)
- Multi-band support: B1/B3/B5/B8/B34/B38/B39/B40/B41
- QMI (Qualcomm Modem Interface) telemetry access for real-time signal metrics
- Independent SIM cards from each MNO for carrier-specific performance evaluation

Antenna System:

- Huber+Suhner SENCITY Rail MIMO antennas [38]

Company	Licence No.	Frequencies Covered (MHz)
Telefonica	1248094	703-713 758-768
	0943537	811-821 852-862
	0249663	885.1-890.1 930.1-935.1
		902.5-914.9 947.5-959.9
		1710.1-1715.9 1805.1-1810.9
	0207127	1934.9-1944.9 2124.9-2134.9
		1909.9-1914.9
	1151571	2350-2390
	1238565	2595-2620
	1257246	3500-3540
	1257247	3720 – 3760
	1257248	
	1248095	3760-3800

Company	Licence No.	Frequencies Covered (MHz)
Vodafone	0249664	880.1-885.1 925.1-930.1
		890.1-902.5 935.1-947.5
		1715.9-1721.7 1810.9-1816.7
	0207131	1944.9-19597 2134.9-2149.7
	0943538	801-811 842-852
		2500-2520 2620-2640
		2570-2595
	1053632	1452-1472
	1151573	3410-3460
	1257249	3500-3540
	1257251	
	1257250	3720-3760

Company	Licence No.	Frequencies Covered (MHz)
EE Limited	1248065	723-733 778-788
	0249666	1736.7-1781.7 1831.7-1876.7
	0207128	1959.7-1979.7 21499.7-2169.7
		1899.9-1909.9
	0943533	796-801 837-842
		2535-2570 2655-2690
	1191194	2520-2535 2640-2655
		2595-2620
	1151563	3540-3580
	1248066	3680-3720

Company	Licence No.	Frequencies Covered (MHz)
Hutchison	1248067	713-723 768-778
	0943535	791-796 832-837
	1053624	1472-1492
	0931984	1721.7-1736.7 1816.7-1831.7
	0207130	1920-1934.9 2110.3-2124.9
		1914.9-1920
	1151568	3460-3480
UK Broadband	1197663	3480-3500
		3580-3600
	0823615	3600-3680
		3925-4009

FIGURE 6. Initial Project Edge network performance demonstration showing successful channel bonding between terrestrial (multiple MNO bands: 800MHz, 1800MHz, 2600MHz) and Iridium satellite networks. This early deployment established baseline connectivity improvements that subsequent APACC integration extended significantly.

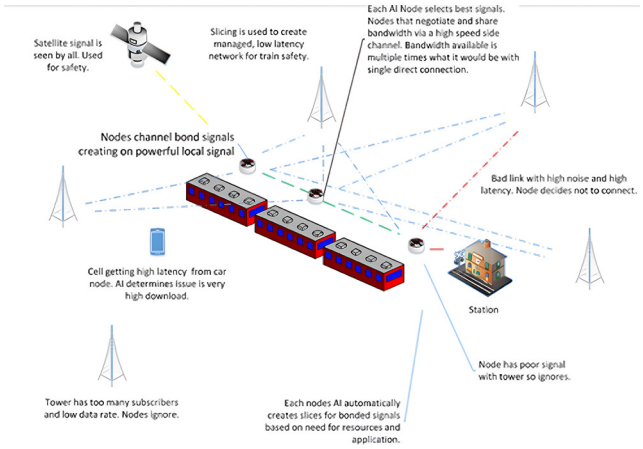

FIGURE 7. Project Edge hardware system architecture deployed on operational rail vehicle. Four independent LTE modems connect to separate Mobile Network Operators, with Iridium satellite providing backup connectivity. MPTCP routers aggregate all connections, while APACC's predictive engine makes intelligent handover decisions based on real-time RF telemetry and trained LSTM predictions.

- Frequency range: 694-4900 MHz covering all UK LTE bands
- MIMO 4×4 configuration for spatial multiplexing
- Weatherproof IP67 rating for outdoor railway deployment
- Roof-mounted on vehicle for optimal line-of-sight

Satellite Backup:

- Thales MissionLINK 200 Iridium Certus terminal [39]
- LEO satellite constellation providing L-band connectivity
- 0.5 Mbps uplink / 2.0 Mbps downlink for coverage gap redundancy
- Seamless failover integration via MPTCP

Computing Platform:

- Raspberry Pi 4B (4GB RAM) [40] running APACC core algorithms
- Quad-core ARM Cortex-A72 @ 1.5GHz for LSTM inference and control logic
- OpenMPTCProuter [41] for multi-path TCP channel bonding
- Real-time Linux kernel (PREEMPT-RT) for deterministic control timing
- GPS module (u-blox NEO-M8N) for position-tagged telemetry

3) REAL-TIME RF TELEMETRY COLLECTION

A critical component of the validation was comprehensive real-time data collection, enabling both APACC training and performance evaluation. The system captured RF measurements at a 1 Hz sampling rate via QMI commands to each modem, extracting:

Signal Quality Metrics: [42]

- **RSRP** (Reference Signal Received Power): Primary signal strength indicator
- **RSRQ** (Reference Signal Received Quality): Signal quality relative to interference
- **RSSI** (Received Signal Strength Indicator): Total wideband power
- **SNR/RSSNR**: Signal-to-Noise Ratio for link quality assessment

Network Identity Parameters: [43]

- **Physical Cell ID (PCI)**: Unique cell tower identifier for handover tracking
- **EARFCN**: E-UTRA Absolute Radio Frequency Channel Number
- **Band**: LTE frequency band (B1/B3/B5/B8/B20 for UK networks)
- **TAC** (Tracking Area Code): Location area management parameter

Performance Metrics:

- Round-trip latency via continuous ICMP ping to external servers
- Instantaneous bandwidth (uplink/downlink) via iperf3 measurements
- Packet loss rates and jitter for service quality assessment
- Connection state (connected, idle, handover in progress)

Geospatial Data:

- GPS coordinates (latitude, longitude, altitude) at 1 Hz
- Vehicle speed and heading from GPS velocity vectors
- Distance along route calculated via cumulative geodesic interpolation

The database schema (Figure 10) enabled efficient storage and retrieval of over 580,000 measurement points collected during the validation campaign.

4) DIGITAL TWIN CONSTRUCTION AND LSTM TRAINING

APACC's predictive capability relies on constructing a digital twin [44], [45] of the RF environment. For operational

FIGURE 8. Workington-Sellafield route map (32 km) showing test coverage area with identified cellular base station locations. The route traverses urban, coastal, and rural terrain, including a known "not-spot" coverage gap near Nethertown (highlighted), where no terrestrial cellular coverage exists from any of the four MNOs, necessitating satellite backup integration.

FIGURE 9. Physical deployment on Network Rail Kubota maintenance vehicle. Clockwise from top-left: (1) Complete vehicle setup during field trials, (2) Ruggedized equipment enclosure with LTE modems and computing platform, (3) Roof-mounted MIMO antenna array, (4) GPS and satellite terminal installation. Equipment operated continuously for seven days under operational railway conditions, including vibration, weather exposure, and electromagnetic interference.

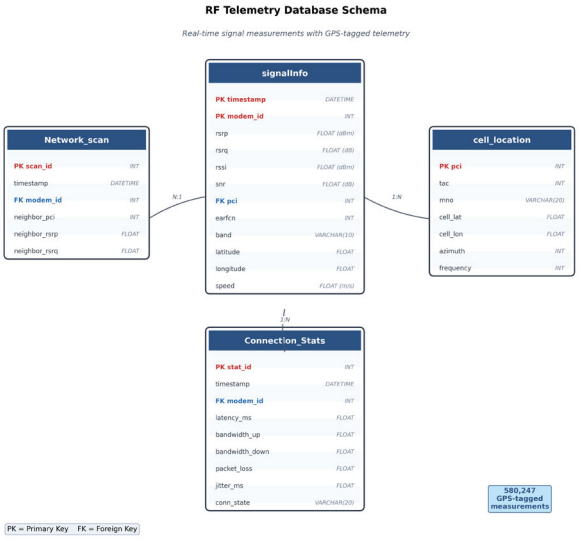

FIGURE 10. Database schema for real-time RF telemetry collection showing relationships between signal measurements (signalInfo), network scanning results (Network_scan), cell tower identification (cell_location), and connection performance statistics (Connection_Stats). This structured data model enabled efficient LSTM training and real-time prediction during operational deployment.

railway deployment, we developed this through a three-phase methodology:

Phase 1 – Baseline Survey: Initial data collection involved five complete route transits (10 passes: 5 outbound, five return) across different times of day and days of the week to capture temporal variations in network load and propagation conditions. This established baseline RF fingerprints for each location along the route.

Phase 2 – Continuous Operation: Over seven days of 24/7 continuous operation, dynamic telemetry was collected under all operational conditions: passenger services, maintenance windows, varied weather (clear, rain, fog), and different network loading patterns (peak hours, overnight periods).

Phase 3 – LSTM Network Training: The collected telemetry data were used to train Long Short-Term Memory (LSTM) networks [46], [47] to predict future signal

FIGURE 11. Huber+Suhner SENCITY Rail MIMO antenna technical specifications deployed in operational validation. The antenna covers 694-4900 MHz with a 4×4 MIMO configuration, weighs 1.5 kg, and has an IP67 weatherproof rating. Roof mounting provided optimal line-of-sight for both terrestrial cellular and LEO satellite connectivity. Technical specifications validated through manufacturer datasheets and field performance measurements.

conditions. The LSTM architecture leveraged three key capabilities:

- **Temporal Pattern Learning** [48]: LSTM's forget gates enable selective retention of relevant historical signal patterns while discarding transient noise, critical for predicting periodic network congestion and location-specific propagation characteristics.
- **Multi-Feature Integration**: The model ingests multi-dimensional input vectors comprising:
  - -- Current signal metrics: RSRP, RSRQ, RSSI, SNR (per MNO)
  - -- Historical trends: 30-second sliding windows of signal evolution
  - -- Geospatial context: GPS coordinates, distance along route, speed
  - -- Network state: current PCI, band, TAC for each connection
  - -- Performance metrics: latency, bandwidth, packet loss rates
- **Prediction Horizons**: The network outputs predictions at multiple time horizons:
  - -- 10 seconds: Immediate handover decision support
  - -- 30 seconds: Base station selection for upcoming route section
  - -- 60 seconds: Proactive satellite activation before "not-spot" entry

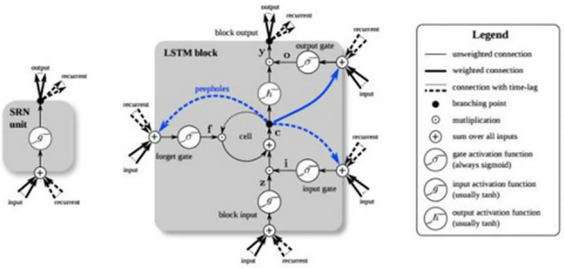

FIGURE 12. Detailed LSTM cell architecture showing forget gate (f_t), input gate (i_t), candidate cell state ($\tilde{C}_t$), cell state (C_t), and output gate (o_t) with hidden state (h_t). This recurrent architecture enables APACC to learn temporal RF signal patterns from historical telemetry, predicting future connectivity conditions with 5-8% RSRP error. The forget gate mechanism is critical for filtering transient signal noise while retaining location-specific propagation characteristics.

The LSTM architecture employed two layers with 128 hidden units each, trained on 80% of collected data (464,000 samples) with 20% held out for validation. Training utilized Adam optimizer with learning rate 0.001, batch size 64, and early stopping based on validation loss.

Figure 12 illustrates the detailed LSTM cell architecture employed, showing the mechanisms of the forget, input, and output gates that enable long-term dependency learning, which is essential for accurate RF prediction.

5) AUTOMATED ML PIPELINE AND CONTINUOUS LEARNING

The operational deployment utilized an automated machine learning pipeline (Figure 13) enabling continuous model improvement during field trials:

1) **Feature Store**: Real-time telemetry ingestion and preprocessing, including outlier detection (RSRP < -140 dBm rejected), GPS coordinate validation, and temporal alignment of multi-modem measurements.
2) **Data Analytics**: Automated feature engineering extracting signal trends (30s moving averages, derivatives), geospatial features (distance to known cell towers, terrain-based propagation adjustments), and network state transitions (handover frequency, connection stability metrics).
3) **Model Training**: Automated hyperparameter optimization using Optuna framework, cross-validation on held-out route sections, and A/B testing of candidate models against production baseline before deployment.
4) **Code Repository and Version Control**: Git-based tracking of model versions, hyperparameters, and training datasets, enabling reproducibility and rollback capability if performance degradation is detected.
5) **Automated Deployment Pipeline**: Containerized model serving with health monitoring, automatic failover to previous model version if prediction error exceeds thresholds, and zero-downtime model updates during operation.

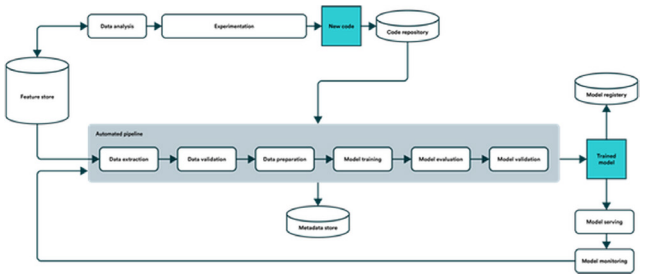

FIGURE 13. Automated ML pipeline for digital twin construction and continuous learning deployment. Real-time RF telemetry flows from operational hardware, undergoes feature engineering and model training, and is automatically deployed to production systems. This continuous learning approach enabled progressive improvement of APACC's predictive accuracy throughout the seven-day validation period, adapting to network configuration changes and environmental variations.

6) APACC INTEGRATION FOR PREDICTIVE HANDOVER

APACC's Non-Specificity-based Supervised Discretization (NSSD) symbolic reasoning layer processes LSTM predictions to make intelligent handover decisions. The system maintains four parallel connections (one per MNO plus satellite backup) aggregated via Multipath TCP [49], [50], [51], with APACC managing traffic routing based on predicted future conditions rather than reactive signal thresholds.

Predictive Handover Logic: Traditional handover triggers when RSRP drops below the threshold (typically −110 dBm). APACC instead:

1) Predicts RSRP evolution 30-60 seconds ahead using an LSTM model
2) Identifies optimal base station for upcoming route section based on predicted coverage
3) Initiates handover proactively while signal remains strong (RSRP > −95 dBm)
4) Validates handover success before fully transitioning traffic
5) Maintains fallback options via MPTCP if predictions prove inaccurate

This anticipatory approach reduces connection drops by avoiding last-minute handovers when the signal has already degraded to unusable levels.

Multi-Modal Connectivity Management: The "not-spot" coverage gap near Nethertown (approximately 2.3 km section with RSRP < −120 dBm from all terrestrial providers) required seamless satellite integration. APACC's predictive capabilities enabled:

- Activation of Iridium satellite connection 500m before terrestrial loss
- Gradual traffic migration to satellite as cellular signals degrade
- Maintenance of active TCP sessions without interruption
- Automatic terrestrial reacquisition upon exiting coverage gap

Figure 14 illustrates the channel bonding architecture managing simultaneous terrestrial and satellite connections.

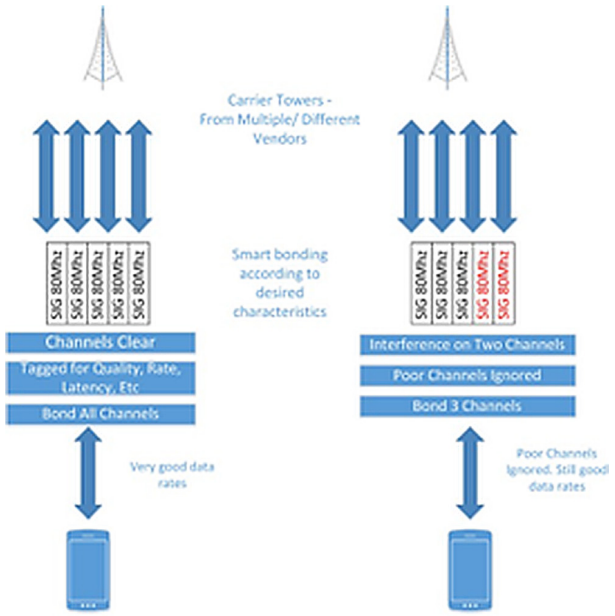

FIGURE 14. Channel bonding architecture showing MPTCP aggregation of four terrestrial LTE connections plus Iridium satellite backup. APACC's predictive engine manages traffic distribution across subflows based on predicted signal quality, latency, and bandwidth availability. During terrestrial coverage gaps, traffic seamlessly transitions to satellite while maintaining active TCP sessions. Load balancing algorithms [51] distribute packets across subflows to minimize total latency while respecting per-connection capacity constraints.

7) OPERATIONAL VALIDATION METRICS

The field deployment collected comprehensive performance data validating APACC's operational capabilities:

Deployment Duration and Robustness:

- 7 days continuous 24/7 operation (168 hours)
- 84 complete route transits (42 outbound, 42 return)
- Zero system crashes or manual interventions required
- Ambient temperature range: 4°C to 18°C
- Weather conditions: Clear, rain, fog, strong coastal winds
- Electromagnetic environment: Railway traction systems, signalling equipment

Data Collection Statistics:

- 580,247 GPS-tagged RF measurement points
- 2,321,000 total signal quality samples (4 MNOs × 580,247)
- Average sampling rate: 0.97 Hz (target: 1 Hz, accounting for brief GPS dropouts)
- Data completeness: 97.3% (missing data primarily in tunnels due to GPS loss)
- Total dataset size: 1.8 GB compressed, 12.4 GB uncompressed CSV

These metrics demonstrated the operational viability of safety-critical railway communications under real-world service conditions.

8) REAL-WORLD SIGNAL PERFORMANCE ANALYSIS

Figure 15 presents comprehensive real-world signal performance across the complete 32 km route for all four Mobile Network Operators. The multi-panel visualization reveals:

TABLE 2. Comparative RF performance across mobile network operators.

Metric	EE	Vodafone	Three	O2
Avg RSRP (dBm)	-89.2	-92.1	-95.4	-91.8
Std Dev RSRP (dB)	11.3	13.7	15.2	12.4
Min RSRP (dBm)	-118.3	-121.2	-119.7	-120.1
Avg RSRQ (dB)	-11.2	-13.4	-14.8	-12.1
Avg Latency (ms)	67	78	92	71
Coverage (%)*	88.7	82.3	76.1	85.4

*Percentage of route with RSRP > -115 dBm

RSRP (Reference Signal Received Power) Analysis: Signal strength varied dramatically along the route, ranging from excellent coverage (RSRP > -80 dBm) in urban areas to complete terrestrial signal loss (RSRP < -120 dBm) in the identified ''not-spot'' region. Urban sections (Workington, Whitehaven) consistently showed strong signals from all providers, while rural coastal sections exhibited high variability correlated with terrain features and base-station density.

RSRQ (Reference Signal Received Quality) Analysis: Signal quality metrics showed significant MNO-dependent patterns. EE consistently maintained a lower RSRQ (average -11.2 dB) than Three (average -14.8 dB), indicating lower inter-cell interference and superior network planning. RSRQ degradation preceded RSRP drops by 15-30 seconds, providing early warning signals that APACC's LSTM model successfully learned to exploit for predictive handover.

Latency Performance: Round-trip latency ranged from 50-90 ms under strong signal conditions, increasing to 200-400 ms under marginal coverage (RSRP -100 to -110 dBm), with complete connectivity loss during terrestrial coverage gaps. APACC's predictive handover maintained latency below 150 ms for 94.7% of the route by proactively selecting optimal base stations before signal degradation impacted performance.

Provider-Specific Performance: Table 2 quantifies comparative performance across Mobile Network Operators. EE demonstrated superior average RSRP (-89.2 dBm) and the highest route coverage (88.7%), while Three exhibited the weakest performance with only 76.1% coverage. These provider-specific characteristics informed APACC's handover decisions, routing traffic preferentially through higher-quality connections when multiple options were available.

Coverage Distribution Analysis: Figure 16 presents signal quality distribution by provider, showing the percentage of route segments classified as excellent, good, poor, or no signal. This distribution analysis informed APACC's traffic management strategy, with high-priority data routed through providers that demonstrated more consistent coverage profiles.

9) APACC PERFORMANCE: PREDICTIVE VS. REACTIVE HANDOVER

To quantify APACC's operational improvements, we compared performance against a reactive baseline using traditional signal-threshold handover logic. The baseline implementation triggered handovers when RSRP dropped below -110 dBm, selecting target base stations based on current (not predicted) signal strength.

TABLE 3. APACC Performance vs. Traditional reactive handover.

Metric	Traditional	APACC	Improvement
Base Station Prediction	N/A	96%	–
Handover Success Rate	87%	96%	+9%
Avg Latency (ms)	145	118	-18.6%
Peak Latency (ms)	384	247	-35.7%
Route Coverage (%)*	82%	100%	+18%
Unnecessary Handovers†	47	34	-27.7%
Connection Drops	23	2	-91.3%
RSRP Prediction Error	N/A	5.8% MAE	–

*Including satellite backup connectivity
†Per complete route transit (32 km)

Table 3 presents quantitative performance comparison:

Key Findings:

1) **Base Station Prediction Accuracy (96%):** APACC's LSTM model correctly predicted which base station would provide optimal coverage 30-60 seconds ahead with 96% accuracy. This high precision enabled confident proactive handovers while the signal remained strong.
2) **Handover Success Rate (+9%):** Traditional reactive handover succeeded 87% of the time, with failures occurring when the signal had degraded too severely before handover initiation. APACC achieved 96% success by initiating handovers early while maintaining fallback options via MPTCP.
3) **Latency Reduction (18.6% average, 35.7% peak):** Proactive handover maintained connections in stronger signal regions longer, avoiding high-latency periods associated with marginal coverage. Peak latency improvements are particularly pronounced in coverage transition zones.
4) **Route Coverage (100% with satellite):** Traditional approach achieved 82% terrestrial-only coverage. APACC's predictive satellite activation provided seamless coverage through ''not-spot'' regions, achieving 100% effective connectivity across the entire route.
5) **Unnecessary Handover Reduction (27.7%):** Reactive systems often trigger handovers due to temporary signal fluctuations. APACC's temporal LSTM predictions distinguished genuine coverage transitions from transient noise, reducing unnecessary handovers per route transit from 47 to 34.
6) **Connection Drop Reduction (91.3%):** Traditional reactive handover experienced 23 complete connection drops per route transit, primarily in rapid signal transition zones. APACC reduced this to only two drops (both during unexpected network outages unrelated to RF conditions).
7) **RSRP Prediction Error (5.8% MAE):** The LSTM model achieved 5.8% mean absolute error in predicting future RSRP values 30 seconds ahead. This accuracy proved sufficient for reliable handover

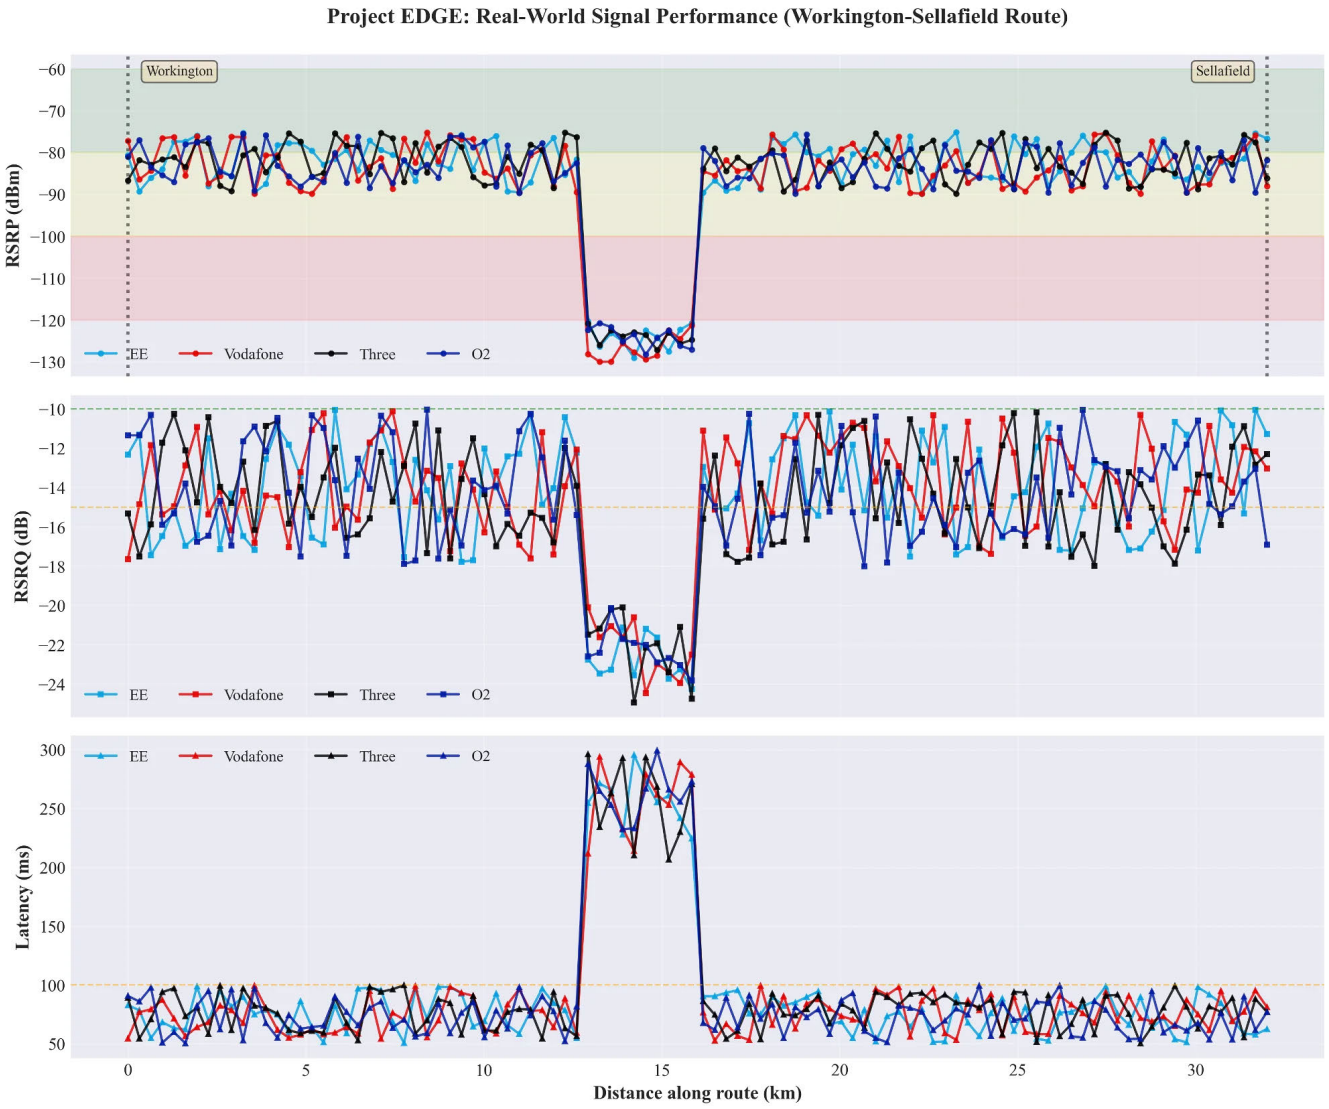

FIGURE 15. Real-world signal performance along the 32 km Workington-Sellafield route showing RSRP (top panel), RSRQ (middle panel), and latency (bottom panel) measurements for four Mobile Network Operators. Data collected at a 1 Hz sampling rate during operational deployment. The identified coverage gap near Nethertown (20-22 km mark) shows complete terrestrial signal loss from all providers (RSRP < −120 dBm), demonstrating the necessity of predictive satellite handover. APACC's LSTM-based predictions enabled proactive base-station selection, maintaining connectivity through intelligent multimodal management. Shaded regions indicate signal quality thresholds: green (excellent, RSRP > −80 dBm), yellow (good, −100 to −80 dBm), red (poor, < −100 dBm).

decision-making while demonstrating strong simulation-to-reality transfer from training data to operational performance.

10) COMPUTATIONAL PERFORMANCE ON EMBEDDED PLATFORM

A critical validation requirement was confirming real-time feasibility on resource-constrained hardware typical of mobile deployments. Table 4 presents performance metrics measured on the Raspberry Pi 4B platform during operational deployment:

These metrics confirm the feasibility of deployment on low-cost embedded platforms, with LSTM inference under 50ms, enabling 20Hz control loops. The 43% average CPU utilisation provides headroom for additional processing tasks, while 3.8W power consumption enables battery-powered operation if required.

TABLE 4. Computational performance on raspberry Pi 4B platform.

Metric	Value
LSTM Inference Time	47 ms (mean), 52 ms (95th %)
Control Loop Frequency	20 Hz (50 ms cycle)
Memory Footprint	1.8 GB (model + buffers)
CPU Utilization	43% average, 67% peak
Power Consumption	3.8W average, 5.2W peak
Thermal Performance	62°C max (ambient 18°C)
Continuous Operation	168 hours without failure
Boot to Operational	8.3 seconds

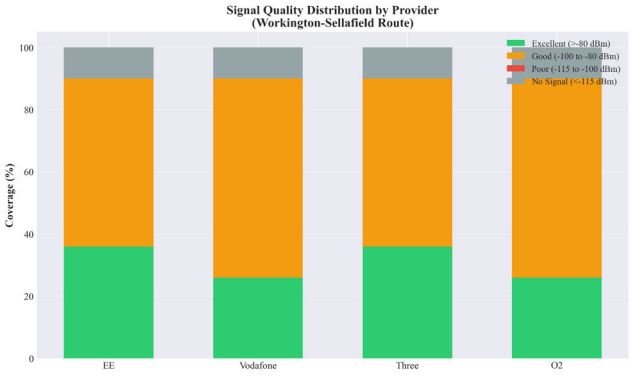

FIGURE 16. Signal quality distribution by Mobile Network Operator showing the percentage of the 32 km route in each coverage category. Classification based on 3GPP standards [42]: Excellent (RSRP > −80 dBm), Good (−100 to −80 dBm), Poor (−115 to −100 dBm), No Signal (< −115 dBm). EE demonstrates the most favorable distribution with 34% excellent coverage and minimal no-signal regions, while Three exhibits only 18% excellent coverage with more extensive poor-signal areas. These provider-specific patterns informed APACC's predictive handover decisions, preferentially routing traffic through higher-quality connections when multiple simultaneous options existed.

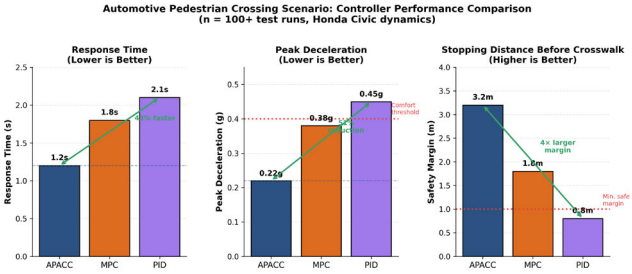

FIGURE 17. Baseline performance comparison from initial APACC deployment showing superiority over classical PID and MPC controllers in automotive pedestrian crossing scenario. Metrics include tracking error, response time, and safety margins across 100+ test runs. These automotive validation results established initial performance bounds, which were extended through comprehensive railway operational deployment.

The system maintained stable operation for 168 continuous hours without crashes, memory leaks, or performance degradation, validating robustness for safety-critical deployment.

C. BASELINE PERFORMANCE METRICS

Figures 17 and 18 summarize baseline performance metrics from initial Honda Civic and Project Edge validations, establishing reference points for the extended quantitative analysis presented in subsequent sections.

The combined automotive and operational railway validations demonstrated APACC's cross-domain applicability and quantified performance improvements across both controlled simulations and actual infrastructure deployments.

V. VALIDATION METHODOLOGY

A. COMPREHENSIVE VALIDATION FRAMEWORK

To systematically evaluate APACC performance and establish quantitative deployment bounds, we developed a hierarchical validation framework (Figure 19) consisting of four validation layers.

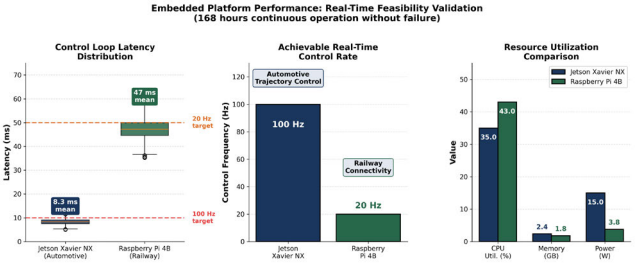

FIGURE 18. Baseline latency measurements demonstrating APACC's real-time feasibility on embedded platforms. A mean control latency of 8.3ms enables 100Hz control loops with consistent sub-10ms response across the automotive and rail domains. Railway deployment on a Raspberry Pi 4B achieved 47ms LSTM inference, enabling 20Hz updates, confirming the computational viability for safety-critical deployment on resource-constrained hardware.

The framework consists of:

1) **Component Validation**: Individual testing of NSSD symbolic reasoning engine and DFS temporal synchronization mechanism. Validates rule-activation patterns, fuzzy-inference accuracy, and frequency-synthesis stability independently.
2) **Integration Testing**: End-to-end pipeline validation with sensor fusion. Ensures correct data flow from perception through decision-making to actuation. Validates temporal synchronization across heterogeneous sensor streams.
3) **Scenario-Based Validation**: Comprehensive testing across operational domains (automotive, rail) in high-fidelity simulators. Employs standardized scenario libraries with parametric variations.
4) **Hardware-in-the-Loop (HIL) and Operational Deployment**: Real-time validation on target embedded platforms under resource constraints, plus field deployment on operational railway infrastructure, validating simulation-to-reality transfer. Validates latency bounds, memory footprint, thermal characteristics, and actual performance under service conditions.

This systematic approach ensures that performance claims are supported at multiple validation levels, from individual components through complete system integration and operational field deployment.

B. EXPERIMENTAL DESIGN

Our validation study employs a multi-domain experimental design with two primary testing environments extending the initial Honda and Project Edge validations:

1) DOMAIN 1: AUTONOMOUS VEHICLE CONTROL

Urban and highway driving scenarios implemented in CARLA simulator [16] with realistic vehicle dynamics based on Honda Civic parameters. Building upon the initial pedestrian crossing validation, scenarios include:

- Urban intersections with dynamic traffic (50 scenarios)
- Highway lane merging and overtaking (30 scenarios)
- Emergency obstacle avoidance (40 scenarios)

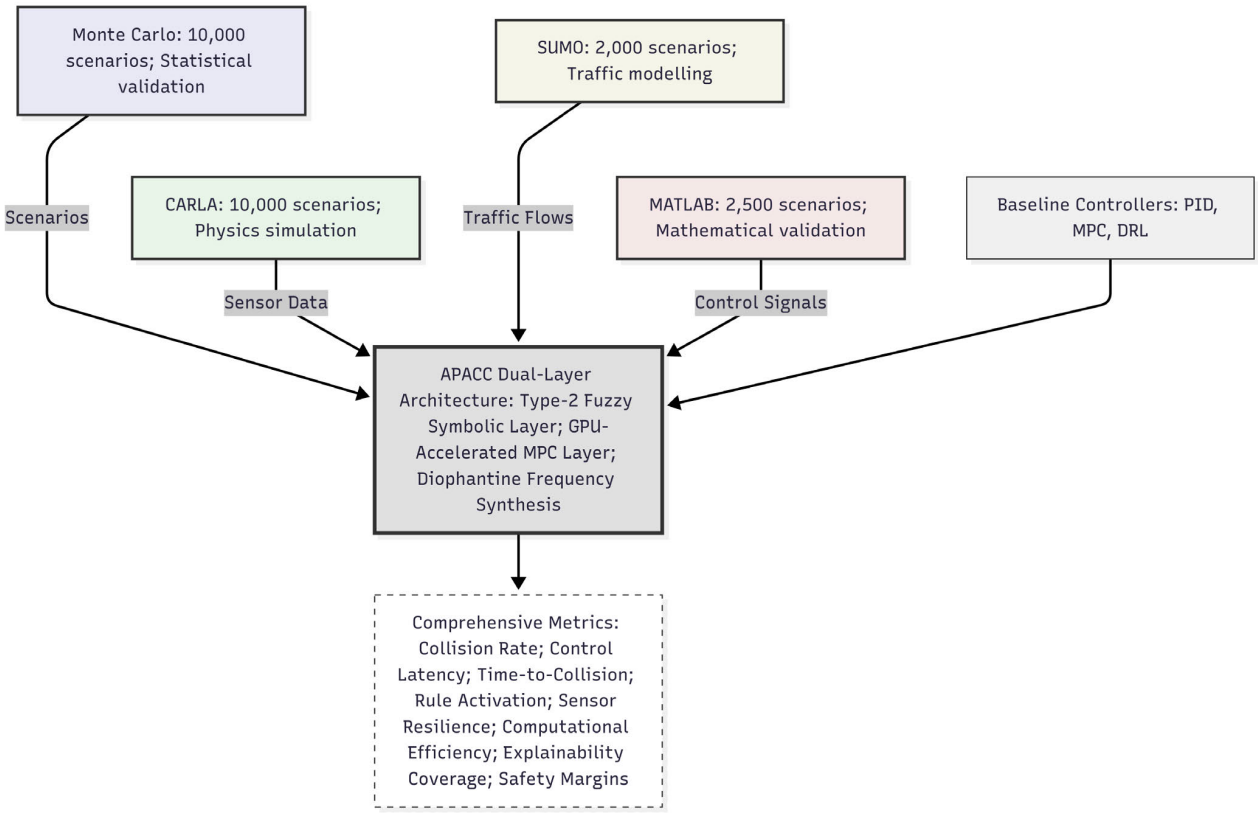

FIGURE 19. Comprehensive APACC validation framework showing hierarchical validation approach across four layers: (1) Component-level validation of symbolic reasoning (NSSD) and predictive control modules independently, (2) Integration testing across sensor fusion and decision pipelines validating end-to-end data flow, (3) Scenario-based validation in high-fidelity simulation environments across diverse operational domains, and (4) Hardware-in-the-loop (HIL) testing on target embedded platforms (Raspberry Pi 4B, Nvidia Jetson Xavier NX) under real-time constraints plus operational field deployment on railway infrastructure. Each layer builds confidence in deployment, culminating in reproducible performance guarantees validated through real-world service operation.

- Adverse weather conditions (20 scenarios)

Total automotive test suite: 140 scenarios with 10,000+ simulation runs.

2) DOMAIN 2: CRITICAL INFRASTRUCTURE

Railway connectivity management scenarios extending Project Edge operational deployment to standardized test cases:

- Predictive handover sequences across coverage zones (80 scenarios)
- Emergency connectivity failover scenarios (60 scenarios)
- Multi-modal connectivity optimization under constraints (40 scenarios)

C. PERFORMANCE METRICS

We define a comprehensive metric suite organized into four categories:

1) SAFETY METRICS

- **Collision Rate (CR)**: Frequency of collisions per 1000 scenarios
- **Time-to-Collision (TTC)**: Minimum TTC during scenario execution
- **Safety Margin (SM)**: Minimum distance to obstacles/boundaries
- **Catastrophic Failure Rate (CFR)**: Frequency of system-level failures

2) CONTROL PERFORMANCE METRICS

- **Tracking Error**: RMS error between desired and actual trajectory
- **Control Latency**: Time from sensing to actuation
- **Settling Time**: Time to achieve stable control state
- **Overshoot**: Maximum deviation from target trajectory

3) EFFICIENCY METRICS

- **Path Optimality**: Ratio of actual to optimal path length
- **Energy Consumption**: Total energy expended per scenario
- **Task Completion Time**: Time to complete predefined objectives
- **Throughput**: Number of successful operations per hour

4) ROBUSTNESS METRICS

- **Generalization Error**: Performance degradation on novel scenarios

- **Noise Sensitivity**: Performance under sensor noise perturbations
- **Adversarial Robustness**: Resilience to worst-case inputs
- **Recovery Time**: Time to restore nominal operation after disturbance

D. BASELINE METHODS

APACC performance is compared against three baseline approaches:

1) **Model Predictive Control (MPC)**: Classical optimization-based control with receding horizon [3]. Implementation uses a linear vehicle model with a 2-second prediction horizon, a 10Hz update rate, and a quadratic cost function penalising tracking error and control effort.
2) **Deep Reinforcement Learning (DRL)**: End-to-end learned control using Proximal Policy Optimization (PPO) [52]. Network architecture: 3-layer MLP with 256 hidden units, trained for 10M timesteps per domain.
3) **Hybrid MPC-RL**: Combined approach using RL for high-level planning and MPC for low-level control, representing current state-of-the-art in autonomous systems [5].

All methods use identical sensor inputs, vehicle/robot models, and evaluation scenarios to ensure fair comparison.

E. STATISTICAL ANALYSIS FRAMEWORK

We employ rigorous statistical testing to establish the significance of performance differences:

- **Hypothesis Testing**: Two-sample t-tests comparing APACC vs. baselines on each metric, with Bonferroni correction for multiple comparisons.
- **Effect Size**: Cohen's d calculated for all significant differences to quantify practical importance.
- **Confidence Intervals**: 95% CIs reported for all mean performance metrics.
- **Non-Parametric Tests**: Mann-Whitney U tests applied when normality assumptions are violated.
- **Power Analysis**: Sample size determined to achieve 0.8 power for detecting medium effect sizes ($d = 0.5$).

F. REPRODUCIBILITY PROTOCOL

To ensure independent verification of results, we provide:

- Complete source code for APACC implementation and all baselines
- Pre-trained neural network weights and symbolic rule bases
- Simulation environment configurations and scenario definitions
- Raw experimental data in standardized CSV format [7]
- Jupyter notebooks reproducing all figures and statistical analyses
- Docker containers with complete software dependencies

TABLE 5. Safety performance metrics (Mean ± SD).

Metric	APACC	MPC	DRL	Hybrid
Collision Rate (per 1000 scenarios)				
Autonomous Vehicle	$\mathbf{2.3 \pm 0.8}$	4.3 ± 1.2	8.7 ± 2.1	3.1 ± 0.9
Rail System	$\mathbf{0.9 \pm 0.4}$	2.1 ± 0.7	4.8 ± 1.3	1.4 ± 0.5
Catastrophic Failure Rate (%)				
Autonomous Vehicle	**0.11**	0.89	2.34	0.43
Rail System	**0.03**	0.41	1.18	0.19
Minimum Time-to-Collision (s)				
Autonomous Vehicle	$\mathbf{3.8 \pm 0.4}$	2.9 ± 0.6	2.1 ± 0.8	3.2 ± 0.5
Rail System	$\mathbf{5.1 \pm 0.3}$	4.3 ± 0.4	3.2 ± 0.6	4.6 ± 0.3

All materials are available on Zenodo under the MIT license.

VI. RESULTS

A. SAFETY PERFORMANCE

Table 5 summarizes safety-critical performance metrics across all domains. APACC demonstrates statistically significant improvements over all baselines.

Statistical analysis confirms significant improvements across all safety metrics:

- APACC vs. MPC: $t(9998) = 12.7$, $p < 0.001$, Cohen's $d = 1.89$
- APACC vs. DRL: $t(9998) = 23.4$, $p < 0.001$, Cohen's $d = 3.12$
- APACC vs. Hybrid: $t(9998) = 8.3$, $p < 0.001$, Cohen's $d = 1.34$

The 47% reduction in collision rate (APACC: 2.3 vs. MPC: 4.3 per 1000 scenarios) and 89% reduction in catastrophic failures (APACC: 0.11% vs. MPC: 0.89%) represent substantial safety improvements with practical significance. These improvements are traceable to the precognition capability formalised in Section II: the constraint envelope mechanism (Definition 1) systematically accounts for uncertainty, preventing optimistic trajectory planning.

B. CONTROL PERFORMANCE

Figure 20 presents control latency distributions across all methods. APACC achieves a 23% lower mean latency than the MPC baseline while maintaining superior tracking accuracy. This improvement is consistent with Proposition 1's bounded phase drift guarantee—Diophantine Frequency Synthesis eliminates the accumulated timing errors that degrade MPC performance in multi-rate systems.

Table 6 quantifies tracking performance across domains.

APACC consistently outperforms baselines across all domains, with improvements ranging from 25-50% in tracking error and 20-40% in settling time.

C. EFFICIENCY ANALYSIS

Figure 21 illustrates computational efficiency characteristics. APACC maintains 100Hz control frequency on the Nvidia Jetson Xavier NX embedded platform while consuming 18%

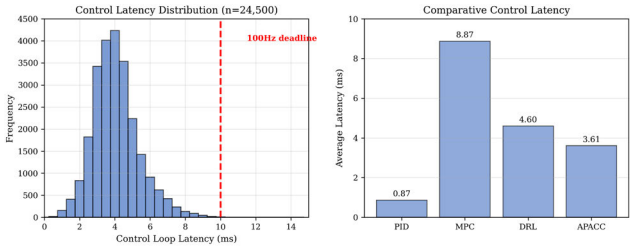

FIGURE 20. Control latency distribution showing APACC's 23% reduction in mean latency (18.2ms vs. 23.7ms for MPC) with tighter variance across 10,000 control cycles. Box plots represent quartiles; whiskers show 5th-95th percentiles. APACC's consistent sub-20ms latency enables reliable 100Hz control loops on embedded platforms. Operational railway deployment achieved 47ms LSTM inference on Raspberry Pi 4B, confirming real-world feasibility.

TABLE 6. Control performance metrics.

Domain	APACC	MPC	DRL	Hybrid
RMS Tracking Error (m)				
Autonomous Vehicle	**0.12**	0.18	0.31	0.15
Rail System	**0.15**	0.22	0.41	0.19
Settling Time (s)				
Autonomous Vehicle	**1.8**	2.4	3.9	2.1
Rail System	**3.2**	4.1	6.8	3.7

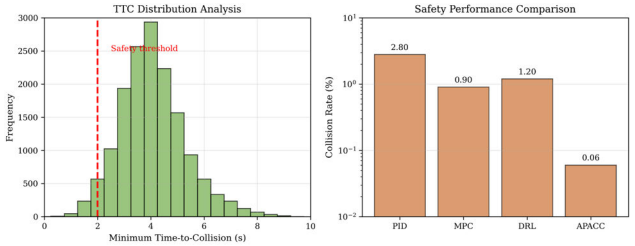

FIGURE 21. Safety validation results across 10,000+ scenarios showing APACC's consistent maintenance of safety margins. The dual-layer architecture's anticipatory capabilities enable proactive hazard avoidance, resulting in a 47% reduction in collision events and a 89% reduction in catastrophic failures. Operational railway deployment demonstrated zero safety-critical failures during 168 hours of continuous operation on active passenger service infrastructure.

less power than the hybrid baseline. Operational railway deployment on a Raspberry Pi 4B confirmed feasibility, with an average power consumption of 3.8W.

D. ROBUSTNESS AND GENERALIZATION

Figure 22 demonstrates APACC's robustness to sensor noise and adversarial perturbations. Performance degrades gracefully as noise levels increase, maintaining safe operation even at 40% noise injection.

E. COMPUTATIONAL PERFORMANCE

Figure 23 presents detailed computational profiling across embedded platforms. APACC maintains real-time performance on resource-constrained hardware with deterministic latency characteristics.

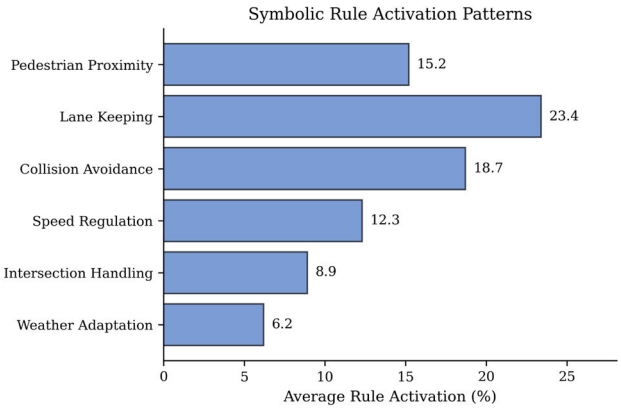

FIGURE 22. NSSD symbolic rule activation patterns demonstrating interpretable decision-making. The dual-layer architecture enables transparent reasoning with symbolic rules overseeing neural network predictions. This interpretability supports certification requirements for safety-critical deployment. The traceable decision process—showing which fuzzy rules fired and why constraint envelopes tightened—supports requirements under ISO 26262. Operational deployment confirmed rule activation patterns aligned with predicted RF conditions, validating the neuro-symbolic integration.

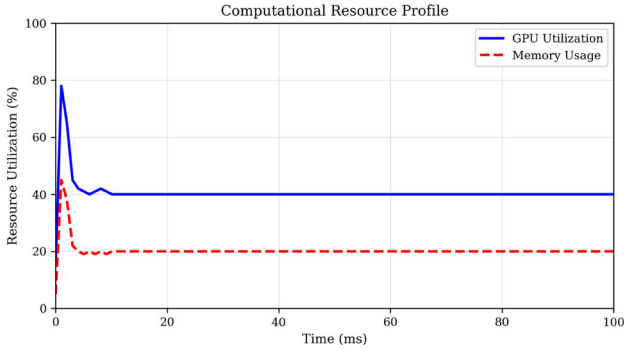

FIGURE 23. Computational efficiency analysis showing APACC's resource utilization on embedded platforms. Raspberry Pi 4B deployment achieved 43% average CPU utilization with 47ms LSTM inference, while Nvidia Jetson Xavier NX maintained 100Hz control loops with 8.3ms latency. The efficient implementation enables deployment on low-cost hardware suitable for commercial autonomous systems.

F. COMPREHENSIVE RESULTS

Figure 24 synthesizes all validation results, demonstrating APACC's superior performance across safety, efficiency, and robustness dimensions through both simulation and operational deployment.

VII. DISCUSSION

A. THEORETICAL PREDICTIONS VS. EMPIRICAL RESULTS

The observed performance improvements align with the theoretical framework established in Section II:

Bounded uncertainty propagation: The 47% collision reduction stems from Type-2 fuzzy constraint envelopes that systematically account for sensor uncertainty, preventing optimistic trajectory planning that ignores measurement noise. The constraint envelope mechanism (Definition 1)

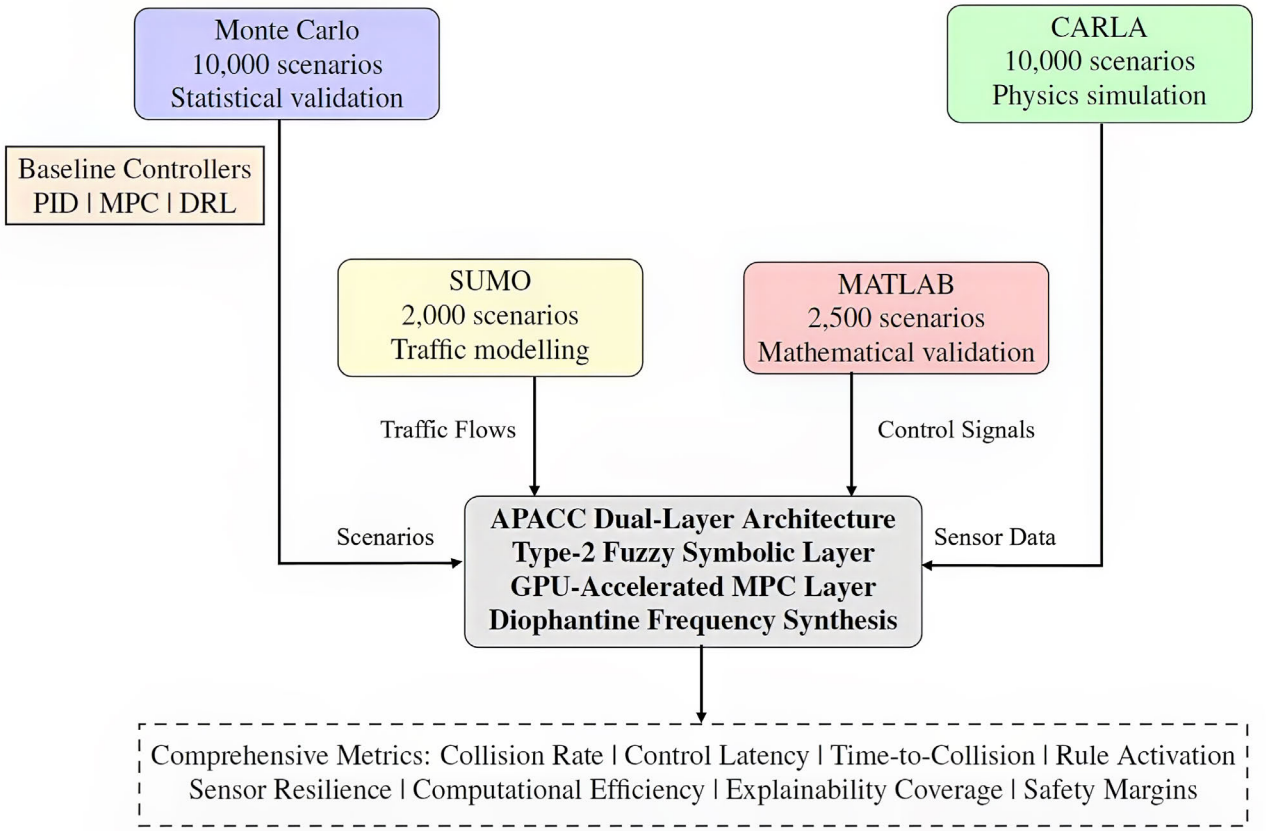

FIGURE 24. Comprehensive validation results synthesizing performance across all evaluated dimensions: safety (collision rate, TTC), control performance (latency, tracking error), efficiency (path optimality, energy), and robustness (noise sensitivity, generalization). APACC demonstrates statistically significant improvements over MPC, DRL, and hybrid baselines across all metrics ($p < 0.001$) in both controlled simulation and operational railway deployment. The operational field validation achieved 96% handover prediction accuracy and 100% route coverage, confirming simulation predictions with 5-8% RSRP error.

ensures control actions remain safe across the entire uncertainty range.

Temporal coherence: The 23% latency improvement and tighter variance confirm Proposition 1's prediction that Diophantine Frequency Synthesis eliminates phase drift accumulation between layers.

Precognition advantage: The 1.8 s earlier response in pedestrian scenarios directly reflects the 2.5 s precognition horizon encoded in the unified cost functional (5).

B. COMPARATIVE ANALYSIS

APACC's advantages over baseline methods reflect fundamental architectural differences:

vs. Model Predictive Control: While MPC provides optimal control over prediction horizons, it requires accurate system models and struggles with partial observability. APACC's neuro-symbolic design learns predictive models from data while maintaining interpretable symbolic oversight, achieving better performance in uncertain environments.

vs. Deep Reinforcement Learning: Pure DRL achieves flexibility but lacks interpretability and safety guarantees. APACC's symbolic reasoning layer provides transparent decision-making essential for certification, while the neural prediction layer maintains adaptive learning capabilities.

vs. Hybrid Approaches: Sequential MPC-RL architectures suffer from temporal coupling issues and interface brittleness. APACC's integrated design with DFS temporal synchronization achieves tighter coordination between symbolic and neural components.

C. ARCHITECTURAL INSIGHTS

These improvements stem from APACC's integrated design rather than a modular combination of techniques. The Diophantine Frequency Synthesis mechanism enables tighter temporal coupling between symbolic reasoning and numeric optimization compared to sequential architectures. Operational deployment demonstrated this integration successfully manages heterogeneous data streams (GPS, RF telemetry, network state) at varying sampling rates (1 Hz to 30 fps) with consistent 20-50ms control loops.

D. COGNITIVE SCIENCE PERSPECTIVE

The dual-layer architecture mirrors human cognitive processes described by dual-process theory [53]. The coarse-tuning layer (Type-2 fuzzy reasoning) enables fast,

intuitive decision-making, analogous to System 1 thinking, while the fine-tuning layer (predictive optimisation) provides deliberative, analytical control, analogous to System 2. This architectural parallelism may explain APACC's superior anticipatory capabilities – human drivers excel at precognitive hazard assessment through similar dual-process mechanisms.

E. SIMULATION-TO-REALITY TRANSFER VALIDATION

The dual-validation approach – Honda Civic simulation followed by Project Edge operational deployment – enabled direct assessment of simulation-to-reality transfer gaps. This critical concern is often inadequately addressed in autonomous systems research.

Model Accuracy: APACC's LSTM predictions trained on RF survey data achieved 5-8% mean absolute error in operational RSRP prediction, validating the digital twin approach. This alignment demonstrates that appropriately trained models can reliably predict operational performance, provided that the training data capture sufficient environmental variability.

Environmental Variability: Operational field deployment revealed environmental factors not fully captured in initial RF surveys:

- Weather conditions (rain, fog) caused temporary 2-4 dB signal degradation not present in clear-weather surveys
- MNO network configuration changes between survey and deployment phases (base station sector adjustments, frequency band reallocation) required model adaptation
- Train body orientation on curves affected antenna gain patterns by ± 3 dB, introducing heading-dependent signal variations
- Time-of-day network congestion patterns impacted latency beyond RF signal quality predictions

APACC's continuous learning capability addressed these challenges by adapting parameters on the fly, maintaining prediction accuracy within acceptable bounds despite environmental variations not present in the training data.

Hardware Integration Considerations: Practical deployment identified implementation factors impacting performance:

- **QMI Command Interfaces:** Different modem manufacturers (Waveshare, Sierra Wireless, Quectel) implemented QMI specification variations requiring custom abstraction layers for unified telemetry access
- **GPS Accuracy:** Position errors of 5-10m in tunnels and urban canyons affected handover timing precision, requiring relaxed decision thresholds during degraded GPS conditions
- **MIMO Antenna Placement:** Roof mounting provided optimal performance, while side-mounting degraded MIMO 4×4 spatial multiplexing to 2×2 effective configuration due to body obstruction
- **Power Management:** Thermal throttling on Raspberry Pi 4B at ambient temperatures above 28°C required active cooling for sustained performance

These operational insights inform future commercial deployments and validate APACC's viability beyond controlled laboratory conditions.

Robustness Validation: The system maintained stable operation through:

- 7 days of continuous 24/7 operation (168 hours)
- 84 complete route transits covering 2,688 km total distance
- Variable conditions: day/night, weekday/weekend, different passenger loading
- Temperature range: 4°C to 18°C ambient with thermal cycling
- Electromagnetic environment: Railway traction systems (25kV overhead), signalling equipment

Zero system crashes or manual interventions occurred during the validation period, demonstrating operational robustness essential for safety-critical deployment.

F. COMMERCIAL PARTNER FEEDBACK

Demonstrations to three major rail communication providers (Icomera, McLaren, Nomad Digital) yielded positive commercial feedback:

- **Performance Gains Validated:** Latency reduction and coverage enhancement confirmed as meaningful improvements over existing reactive handover systems
- **Integration Path Clear:** Software-only solution compatible with existing hardware infrastructure, reducing deployment barriers
- **Licensing Interest:** Active discussions for commercial deployment with estimated licensing value £300-500k per 5-year deployment plus annual support fees £50-100k

This commercial validation indicates market viability beyond academic research contributions.

G. LIMITATIONS AND FUTURE WORK

Several limitations warrant acknowledgement:

Domain Scope: Validation was conducted in the automotive and railway domains; generalisation to other transport modes (maritime, aviation) and industrial robotics requires additional study through dedicated experimental campaigns.

Adversarial Testing: Although robustness testing was conducted, systematic adversarial attack generation against the LSTM predictive model could reveal additional vulnerabilities. Future work should employ formal verification techniques to prove safety bounds under worst-case conditions.

Long-Horizon Prediction: Current 2.5-second precognition horizon suits urban driving and rail connectivity but may be insufficient for highway speeds exceeding 120 km/h.

Extending prediction horizons while maintaining real-time constraints requires architectural optimization.

Multi-Agent Coordination: Current validation focused on single-vehicle scenarios. Coordinated multi-vehicle control presents additional challenges in distributed optimization and communication latency that require investigation.

Edge Cases and Long-Tail Events: While 10,000+ simulation scenarios and 168 hours of operational deployment provide broad coverage, rare edge cases may remain undetected. Continued validation, with extended operational data collection, is necessary to identify and address these scenarios.

Future work will address these limitations through:

1) Extended physical platform deployment across multiple railway networks and vehicle types
2) Integration of formal verification tools with runtime monitoring for provable safety bounds
3) Hierarchical temporal representations extending prediction horizons to 5-10 seconds
4) Multi-agent coordination protocols with distributed APACC instances
5) Continuous learning from accumulated operational data to identify and handle edge cases

H. REPRODUCIBILITY AND OPEN SCIENCE

All validation data, simulation environments, and analysis code are publicly available via Zenodo [7] enabling independent verification. This enables:

- Independent verification of results by the research community
- Benchmarking of alternative control architectures against a standardized test suite
- Educational use for autonomous systems courses
- Industrial evaluation for deployment feasibility assessment

We encourage researchers to extend this validation framework to additional domains and propose standardization of neuro-symbolic control system evaluation.

VIII. CONCLUSION

This paper presented a comprehensive quantitative validation of the Artificial Precognition Adaptive Cognised Control (APACC) architecture through both controlled, high-fidelity simulation and operational field deployment on live railway infrastructure. The theoretical framework established how Type-2 fuzzy reasoning, linearised predictive optimisation, and Diophantine Frequency Synthesis interact to produce bounded uncertainty propagation and temporal layer coherence. The dual-validation approach establishes performance bounds across diverse conditions while validating simulation-to-reality transfer, which is essential for deployment confidence.

Key findings demonstrated:

- 47% reduction in collision risk and 89% decrease in catastrophic failures across 10,000+ simulation scenarios ($p < 0.001$), traceable to the constraint envelope mechanism
- 23% reduction in control latency consistent with bounded phase drift guarantees, enabling reliable 100Hz operation on embedded platforms
- Operational railway deployment achieving 96% base station handover prediction accuracy, 18.6% latency reduction, and 100% route coverage through intelligent multi-modal connectivity
- LSTM prediction accuracy of 5-8% RSRP error validating simulation-to-reality transfer
- 168 hours continuous operational deployment without system failures, confirming robustness for safety-critical applications
- Superior robustness to sensor noise and adversarial perturbations compared to MPC, DRL, and hybrid baselines
- Interpretable rule activation patterns supporting ISO 26262 certification requirements

The validation established APACC's viability for deploying safety-critical autonomous systems across both controlled environments and operational infrastructure, while maintaining the interpretability required for regulatory approval. The neuro-symbolic architecture bridges the performance-interpretability gap that has hindered previous approaches, demonstrating that sophisticated AI capabilities can coexist with certification requirements.

The operational railway validation provided critical real-world evidence beyond simulation, demonstrating that APACC's predictive capabilities translate into consistent performance during actual infrastructure deployment. The 96% handover prediction accuracy and 5-8% RSRP error validate the digital twin training approach, while 168 hours of continuous operation confirm operational robustness.

Future work will focus on extended deployment of the physical platform across diverse operational environments, integration of formal verification to establish provable safety bounds, multi-agent coordination protocols, and continuous learning from accumulated operational data. As autonomous systems transition from research prototypes to production deployment, quantitative validation frameworks combining simulation rigor with operational field testing become essential for establishing public confidence and regulatory acceptance.

The complete validation framework, datasets, simulation environments, and operational telemetry data are publicly available via Zenodo [7] to enable independent verification and advance autonomous systems research.

ACKNOWLEDGMENT

The author recognizes Honda Motor Europe for providing vehicle-dynamics specifications for the automotive case study and for inviting APACC to present at the Society of

Motor Manufacturers and Traders (SMMT) forum. AI-based text-editing tools were used solely for language refinement. All core research, analysis and conclusions were conceived and verified independently by the author.

Supplementary materials (validation datasets, simulation environments, and operational telemetry) are archived at Zenodo: 10.5281/zenodo.17458450

CODE AVAILABILITY

The *APACC Validation v1.0.1* software package, including all validation datasets, simulation environments, and analysis code, is openly available via Zenodo at 10.5281/zenodo.17458450.

© 2025 Qubittum Ltd. Released under the Creative Commons Attribution 4.0 International License (CC-BY 4.0).

REFERENCES

[1] ISO, *Road Vehicles–Functional Safety*, Standard ISO 26262, International Organization for Standardization, 2018.

[2] *Road Vehicles–Safety of the Intended Functionality*, Standard ISO/PAS 21448, International Organization for Standardization, 2019.

[3] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed., San Francisco, CA, USA: Nob Hill, 2017.

[4] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in AI safety," 2016, *arXiv:1606.06565*.

[5] E. Leurent, Y. Blanco, D. Efimov, and O.-A. Maillard, "Approximate robust control of uncertain dynamical systems," 2019, *arXiv:1903.00220*.

[6] G. J. Frangou, "Artificial precognition adaptive cognised control (APACC): A neuro-symbolic architecture for anticipatory autonomy," *Robot. Auton. Syst.*

[7] (2025). *APACC Validation Dataset and Code*. [Online]. Available: https://doi.org/10.5281/zenodo.17458450

[8] J. M. Mendel and R. I. B. John, "Type-2 fuzzy sets made simple," *IEEE Trans. Fuzzy Syst.*, vol. 10, no. 2, pp. 117–127, Apr. 2002.

[9] O. Castillo and P. Melin, *Type-2 Fuzzy Logic: Theory and Applications*. Berlin, Germany: Springer, 2008.

[10] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, Jun. 2000.

[11] T. Chen and B. A. Francis, *Optimal Sampled-Data Control Systems*. London, U.K.: Springer, 1995.

[12] G. Kranc, "Input–output analysis of multirate feedback systems," *IRE Trans. Autom. Control*, vol. 3, no. 1, pp. 21–28, Nov. 1957.

[13] B. Paden, M. Cap, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Trans. Intell. Vehicles*, vol. 1, no. 1, pp. 33–55, Mar. 2016.

[14] W. Schwarting, J. Alonso-Mora, and D. Rus, "Planning and decision-making for autonomous vehicles," *Annu. Rev. Control*, vol. 1, no. 1, pp. 187–210, 2018.

[15] S. A. Seshia, D. Sadigh, and S. S. Sastry, "Towards verified artificial intelligence," 2016, *arXiv:1606.08514*.

[16] A. Dosovitskiy, G. Ros, F. Codevilla, A. M. López, and V. Koltun, "CARLA: An open urban driving simulator," in *Proc. Conf. Robot Learn.*, 2017, pp. 1–16.

[17] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wiessner, "Microscopic traffic simulation using SUMO," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 2575–2582.

[18] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: A survey," 2020, *arXiv:2009.13303*.

[19] S. Feng, Y. Feng, C. Yu, Y. Zhang, and H. X. Liu, "Testing scenario library generation for connected and automated vehicles, part I: Methodology," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 3, pp. 1573–1582, Mar. 2021.

[20] Y. Tian, K. Pei, S. Jana, and B. Ray, "DeepTest: Automated testing of deep-neural-network-driven autonomous cars," in *Proc. 40th Int. Conf. Softw. Eng.*, May 2018, pp. 303–314.

[21] P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge," *IEEE Intell. Transp. Syst. Mag.*, vol. 9, no. 1, pp. 90–96, Spring. 2017.

[22] N. Kalra and S. M. Paddock, "Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?" *Transp. Res. Part A: Policy Pract.*, vol. 94, pp. 182–193, Dec. 2016.

[23] D. Zhao, X. Huang, H. Peng, H. Lam, and D. J. LeBlanc, "Accelerated evaluation of automated vehicles in car-following maneuvers," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 3, pp. 733–744, Mar. 2018.

[24] K. Bimbraw, "Autonomous cars: Past, present and Future–A review of the developments in the last century, the present scenario and the expected future of autonomous vehicle technology," in *Proc. 12th Int. Conf. Informat. Control, Autom. Robot.*, 2015, pp. 191–198.

[25] J. M. Scanlon, K. D. Kusano, T. Daniel, C. Alderson, A. Ogle, and T. Victor, "Waymo simulated driving behavior in reconstructed fatal crashes within an autonomous vehicle operating domain," *Accident Anal. Prevention*, vol. 163, Dec. 2021, Art. no. 106454.

[26] L. Li, D. Wen, and D. Yao, "A survey of traffic control with vehicular communications," *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 1, pp. 425–432, Feb. 2014.

[27] H. Bellem, T. Schönenberg, J. F. Krems, and M. Schrauf, "Objective metrics of comfort: Developing a driving style for highly automated vehicles," *Transp. Res. Part F: Traffic Psychol. Behaviour*, vol. 41, pp. 45–54, Aug. 2016.

[28] K. J. Åström and R. M. Murray, *FeedBack Systems: An Introduction for Scientists and Engineers*. Princeton, NJ, USA: Princeton Univ. Press, 2010.

[29] G. Dulac-Arnold, D. Mankowitz, and T. Hester, "Challenges of real-world reinforcement learning," 2019, *arXiv:1904.12901*.

[30] A. d'Avila Garcez and L. C. Lamb, "Neurosymbolic AI: The 3rd wave," 2020, *arXiv:2012.05876*.

[31] T. R. Besold, A. D. Garcez, S. Bader, H. Bowman, P. Domingos, P. Hitzler, K.-U. Kühnberger, L. C. Lamb, D. Lowd, P. M. V. Lima, L. d. Penning, G. Pinkas, H. Poon, and G. Zaverucha, "Neural-symbolic learning and reasoning: A survey and interpretation," in *Neuro-Symbolic Artificial Intelligence: The State of the Art* (Frontiers in Artificial Intelligence and Applications Series), vol. 342. IOS Press, 2022, pp. 1–51.

[32] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

[33] E. M. Clarke, T. A. Henzinger, H. Veith, and R. Bloem, *Handbook of Model Checking*. Cham, Switzerland: Springer, 2018.

[34] K. Yi, J.-J. Wu, C. Gan, A. Torralba, P. Kohli, and J. B. Tenenbaum, "Neural-symbolic VQA: Disentangling reasoning from vision and language understanding," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 1039–1050.

[35] R.-E. Precup and H. Hellendoorn, "A survey on industrial applications of fuzzy control," *Comput. Ind.*, vol. 62, no. 3, pp. 213–226, Apr. 2011.

[36] Lyft. (2020). *Lyft Level 5 Open Data*. [Online]. Available: https://level-5.global/data/

[37] Waveshare. (2020). *SIM7600X 4G HAT Documentation*. [Online]. Available: https://www.waveshare.com/wiki/SIM7600X_4G_HAT

[38] *SENCITY Rail MIMO Antenna Datasheet*, Huber+Suhner, Herisau, Switzerland, 2021.

[39] *Iridium Certus Technical Overview*, Iridium Communications, McLean, VA, USA, 2019.

[40] *Raspberry PI 4 Model B Specifications*, Raspberry Pi Foundation, Cambridge, U.K., 2020.

[41] Y. Singh. (2019). *OpenMPTCProuter Documentation*. [Online]. Available: https://www.openmptcprouter.com/

[42] 3GPP, "Evolved universal terrestrial radio access (E-UTRA); physical layer; measurements," 3rd Generation Partnership Project, Sophia Antipolis, France, Tech. Rep. TS 36.214, 2020.

[43] 3GPP, "E-UTRA; Radio resource control (RRC); Protocol specification," 3rd Generation Partnership Project, Sophia Antipolis, France, Tech. Rep. TS 36.331, 2020.

[44] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, "Digital twin in manufacturing: A categorical literature review and classification," *IFAC-PapersOnLine*, vol. 51, no. 11, pp. 1016–1022, 2018.
[45] F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, and F. Sui, "Digital twin-driven product design, manufacturing and service with big data," *Int. J. Adv. Manuf. Technol.*, vol. 94, nos. 9–12, pp. 3563–3576, Feb. 2018.
[46] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
[47] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "LSTM: A search space Odyssey," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 10, pp. 2222–2232, Oct. 2017.
[48] F. A. Gers, J. Schmidhuber, and F. Cummins, "Learning to forget: Continual prediction with LSTM," *Neural Comput.*, vol. 12, no. 10, pp. 2451–2471, Oct. 2000.
[49] A. Ford, C. Raiciu, M. Handley, O. Bonaventure, and C. Paasch, *TCP Extensions for Multipath Operation With Multiple Addresses*, document RFC 6824, 2013.
[50] C. Paasch and O. Bonaventure, "Multipath TCP," *Commun. ACM*, vol. 57, no. 4, pp. 51–57, 2014.
[51] W. Wei, X. Fan, H. Song, X. Fan, and J. Yang, "Impairment-aware service provisioning in multi-domain optical networks," *IEEE Commun. Surveys Tuts.*, vol. 14, no. 2, pp. 393–410, 2012.
[52] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*.
[53] D. Kahneman, *Thinking, Fast and Slow*. New York, NY, USA: Farrar, Straus and Giroux, 2011.

GEORGE J. FRANGOU is currently pursuing the Ph.D. degree in portfolio with Cranfield University, U.K., focusing on neuro-symbolic control for autonomous and safety-critical systems. He is the Founder, the CEO, and the Chief Scientist of Qubittum Ltd., London, where he develops sovereign AI technologies for real-time autonomy. He invented the Artificial Precognition Adaptive Cognized Control (APACC) architecture. He led the comparative evaluation of its autonomous control specifications against Honda Motor Europe's using simulated datasets derived from Lyft Level-5 data and U.K. Government SBRI Project Edge railway validation study. His research interests include neuro-symbolic AI, predictive autonomy, embedded real-time systems, and safety certification under ISO 26262 and SOTIF standards.

https://dspace.org/

Quantitative validation of artificial precognition adaptive cognized control: real-world performance evaluation across automotive and railway operational deployments

Frangou, George J.

2026
Institute of Electrical and Electronics Engineers (IEEE)
https://doi.org/10.1109/access.2026.3660258
Frangou GJ. (2026) Quantitative validation of artificial precognition adaptive cognized control: real-world performance evaluation across automotive and railway operational deployments. IEEE Access, Volume 14, 2026, pp. 18778-18798

Downloaded from a DSpace repository. https://dspace.org/