TyrePIML: Real-Time Physics-Informed Machine Learning and Game Theory for Multi-Agent Motorsport Strategy

Mina R. Narman*, Husnu S. Narman†

*Acıbadem Mehmet Ali Aydınlar University, Turkiye, mina.narman@live.acibadem.edu.tr
†Marshall University, Huntington WV, USA, narman@marshall.edu

***Abstract*—The complexity of modern motorsport strategy requires real-time processing of high-frequency telemetry, predictive tire degradation modeling, and dynamic rival response anticipation. While existing studies have successfully utilized offline machine learning for tire analysis, they frequently fail to account for real-time dynamic multi-agent interactions and the impending 2026 energy constraints. This paper presents TyrePIML, a real-time multi-agent race strategy intelligence system that addresses this gap. The proposed system integrates Physics-Informed Machine Learning to accurately predict tire degradation cliffs by combining historical telemetry with physical parameters such as driver-specific degradation profiles. Furthermore, the architecture introduces a Game Theory module utilizing Nash Equilibrium to anticipate rival pit-wall decisions based on dynamic Kinetic Motor Generator Unit battery State of Charge. Engineered with a low-latency WebSocket infrastructure and in-memory caching for real-time deployment, TyrePIML demonstrates a highly scalable and industry-ready solution for high-speed intelligent transportation systems and competitive motorsport strategy.**

***Index Terms*—Formula 1, Physics-Informed Machine Learning, Game Theory, Nash Equilibrium, Telemetry, Intelligent Transportation Systems**

I. Introduction

Modern motorsport, particularly Formula 1, represents the pinnacle of high-speed data analytics and vehicle telemetry. During a race, hundreds of sensors transmit high-frequency data, transforming race strategy from intuitive decision-making into a rigorous computational problem. However, traditional predictive models often rely entirely on historical data, failing to account for real-time physical constraints, dynamic aerodynamic penalties, and the adversarial nature of multi-agent racing environments.

The 2026 Formula 1 power unit regulations introduce unprecedented complexity by significantly increasing the reliance on electrical power and active aerodynamics. Managing the Kinetic Motor Generator Unit (MGU-K) battery State of Charge (SoC) across various track geometries requires the strategic anticipation of rival actions.

Research Gap: Recent literature has extensively explored vehicle dynamics and strategy. Purely data-driven models, such as standard recurrent neural networks (RNNs), have been used to extrapolate historical lap times [1], [2]. Furthermore, Physics-Informed Machine Learning (PIML) has been introduced by Raissi et al. [4] to embed thermodynamic models directly into neural networks. Separately, Wang et al. [10] applied game-theoretic cooperative models for autonomous vehicles. However, a significant gap remains: existing PIML applications are largely computed offline, and game-theoretic models rarely factor in real-time electrical energy deployment (MGU-K constraints) in highly adversarial environments.

Objective and Key Contributions: The primary objective of this research is to develop a real-time, physics-informed computational framework capable of predicting both tire failure points and multi-agent rival responses under 2026 regulations. The key contributions of this paper are threefold:

- An asynchronous, event-driven WebSocket architecture coupled with an in-memory Least Recently Used (LRU) cache, eliminating I/O bottlenecks to enable real-time telemetry processing at the edge.
- A Physics-Informed Machine Learning (PIML) framework that integrates a novel Driver Management Profile (D_p) and thermodynamic constraints to accurately forecast non-linear tire degradation cliffs.
- A Multi-Agent Game Theory module grounded in Nash Equilibrium, designed to predict rival strategic responses dynamically constrained by the 2026 MGU-K energy regulations.

Paper Organization: The remainder of this paper is organized as follows: Section II details the real-time system architecture and WebSocket infrastructure. Section III introduces the Physics-Informed Machine Learning approach for tire degradation. Section IV explains the multi-agent Game Theory module and energy physics. Section V presents the experimental results and performance analysis. Section VI discusses limitations and future work, and Section VII concludes the paper.

II. System Architecture for Real-Time Telemetry

High-speed autonomous racing and intelligent transportation networks require decision-making architectures capable of processing massive data streams with near-zero latency. Traditional data acquisition frameworks often rely on periodic client-side polling, which introduces significant communication overhead and synchronization delays. In competitive motorsport, a latency of merely a few seconds can render a strategic pit-wall decision obsolete.

To overcome these communication bottlenecks, TyrePIML utilizes an asynchronous, event-driven WebSocket architecture. By establishing a persistent, bidirectional communication tunnel between the predictive backend and the strategic dashboard, the system bypasses standard HTTP request-response cycles. Telemetry updates, such as sector times and battery State of Charge (SoC), are actively pushed from the server the exact millisecond a lap is completed. This approach mirrors the ultra-low latency requirements of Vehicle-to-Infrastructure (V2I) and Vehicle-to-Vehicle (V2V) communications in modern ITS deployments [5].

Furthermore, evaluating multi-agent game theory matrices and executing machine learning inferences dynamically requires rapid access to historical track geometries and tire compound datasets. Repeatedly reading these massive datasets from physical storage creates severe disk input/output (I/O) bottlenecks. To achieve real-time operational standards, the TyrePIML backend is optimized using a Least Recently Used (LRU) in-memory caching mechanism.

This architectural optimization treats the strategy engine as a highly efficient edge-computing node [6]. By loading active race subsets directly into the Random Access Memory (RAM), the system eliminates disk retrieval times, effectively reducing inference latency from seconds to milliseconds. This guarantees that the Nash Equilibrium and Physics-Informed Machine Learning (PIML) modules can continuously compute the optimal strategy without halting the live telemetry feed. These architectural choices ensure the system is highly scalable and robust against data pipeline congestion, serving as a practical blueprint for edge-assisted autonomous driving networks. The complete data flow, illustrating the integration of the LRU cache and WebSocket tunnel between the telemetry source and the predictive engines, is detailed in Fig. 1.

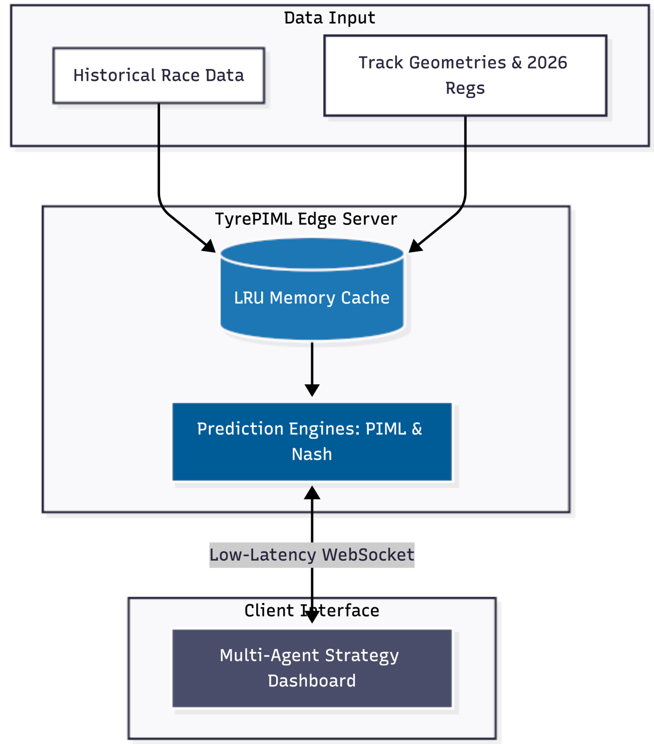

Fig. 1. Proposed system architecture of TyrePIML, detailing the asynchronous WebSocket data pipeline and in-memory LRU caching mechanism for real-time edge computing.

III. Physics-Informed Machine Learning for Tire Degradation

Accurately predicting tire degradation is critical for strategic multi-agent race planning. Traditional data-driven models, which primarily extrapolate historical lap times, frequently fail to predict the catastrophic loss of mechanical grip—commonly referred to in motorsport as the "cliff." This failure occurs because tire degradation is not purely a function of time or distance, but a non-linear thermodynamic process highly sensitive to external constraints [7].

To address this, TyrePIML implements a Physics-Informed Machine Learning (PIML) architecture built upon Long Short-Term Memory (LSTM) neural networks. The sequential nature of the LSTM is uniquely suited for time-series telemetry data [8]. However, unlike standard LSTMs, the network is constrained by real-time physical parameters. At any given lap t, the input feature vector X_t is formulated as:

$$X_t = [L_t, T_{track}, C_{type}, D_p] \tag{1}$$

In this formulation, L_t is the elapsed lap count, T_{track} represents the dynamic track temperature, C_{type} represents the categorical compound coefficient (Soft, Medium, or Hard), and D_p is the dynamically calculated Driver Management Profile.

The Driver Management Profile (D_p) is a core contribution of this system, functioning as a continuous penalty or reward scalar. Rather than assuming uniform degradation across the grid, TyrePIML calculates a specific degradation slope for each agent. The profile is mathematically derived by comparing the ego-vehicle's performance degradation rate against the field average:

$$D_p^{(i)} = \frac{\nabla\tau^{(i)}}{\frac{1}{N}\sum_{j=1}^{N}\nabla\tau^{(j)}} \tag{2}$$

In this context, $\nabla\tau^{(i)}$ denotes the lap-time drop-off rate (seconds lost per lap) for driver i, and N is the total number of active drivers on the same compound. A $D_p^{(i)}$ value greater than 1.0 indicates poor tire management, accelerating the predicted approach to the cliff.

The PIML module continuously outputs the Remaining Useful Life (RUL) of the tire before the grip coefficient enters the non-linear failure phase. By embedding the D_p scalar and thermodynamic inputs directly into the LSTM structure, TyrePIML achieves high-fidelity predictions of the cliff point. This physical grounding prevents the neural network from generating mathematically valid but physically impossible degradation curves, ensuring reliable deployment in edge-case racing scenarios (Fig. 2).

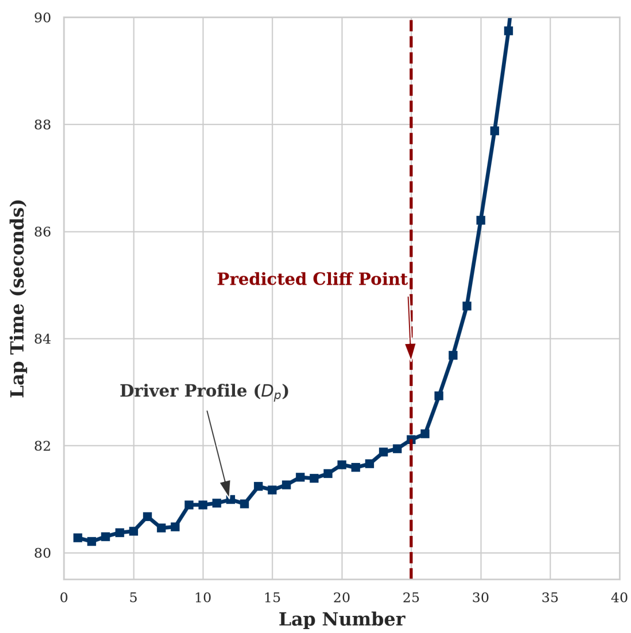

Fig. 2. PIML-predicted tire degradation curve. The model accurately forecasts the non-linear "cliff" point by integrating the driver-specific management profile (D_p) and physical compound constraints.

IV. Multi-Agent Game Theory and 2026 Energy Physics

A fundamental limitation of conventional race strategy software is the assumption of a static environment. In reality, both motorsport and intelligent transportation networks are adversarial, multi-agent ecosystems. When an ego-vehicle executes a strategic maneuver, such as an early pit stop (undercut), rival agents will dynamically alter their trajectories and energy deployment to neutralize the advantage.

To model this adversarial interaction, TyrePIML implements a Game Theory module grounded in Nash Equilibrium [9]. The decision-making process is formulated as a non-cooperative game where each vehicle seeks to maximize its track position utility. However, the action space for any given rival is strictly constrained by the 2026 FIA power unit regulations, which drastically increase the reliance on the Kinetic Motor Generator Unit (MGU-K) and introduce active aerodynamic wake penalties.

Let S_{rival} represent the real-time battery State of Charge (SoC) of a competing vehicle. When the ego-vehicle initiates an undercut, the rival evaluates a strategic response set $A = \{\text{Defend}, \text{Cover}, \text{Harvest}\}$. The utility function $U(a_i, a_{-i})$ is highly dependent on S_{rival}.

For instance, the "Defend" action requires deploying a 350kW electrical override to prevent the undercut. If $S_{rival} < \delta_{critical}$ (where $\delta_{critical}$ is the track-specific energy threshold required for a full-lap defense), the "Defend" action is mathematically eliminated from the feasible set. The rival is then forced to choose between "Cover" (immediately matching the pit stop) or "Harvest" (conceding track position to regenerate kinetic energy for a later offensive).

The system continuously solves for the Nash Equilibrium, defined as the state where no agent can improve its payoff by unilaterally changing its strategy:

$$U_i(a_i^*, a_{-i}^*) \geq U_i(a_i, a_{-i}^*) \tag{3}$$

For all available actions $a_i \in A$, the system evaluates the optimal response. By continuously pushing track-specific recovery profiles (e.g., high-braking zones for energy harvesting) through the WebSocket tunnel, TyrePIML dynamically updates the SoC of all agents. This physics-informed game theory allows the engine to predict the exact probability of a rival's response, moving beyond isolated lap-time predictions into the realm of dynamic, multi-agent behavioral forecasting [10], as mapped in Fig. 3.

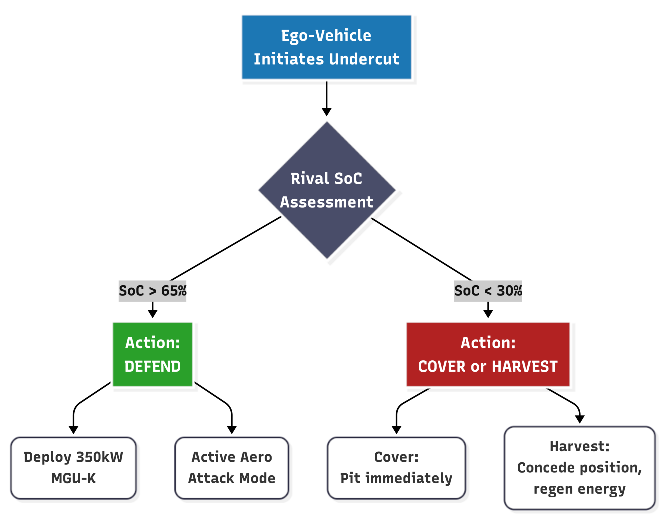

Fig. 3. Strategic decision tree illustrating rival response mechanisms. The action space is strictly constrained by the MGU-K battery State of Charge (SoC), forcing agents into conditional Game Theory behaviors.

V. Experiments and Results

To validate the proposed TyrePIML architecture, simulations were conducted using historical telemetry datasets modified to reflect the anticipated 2026 aerodynamic and energy deployment parameters. The experimental setup evaluated the predictive accuracy of the PIML tire degradation module against a purely data-driven Random Forest (RF) baseline and a standard Long Short-Term Memory (LSTM) network.

As summarized in Table I, the PIML model, incorporating the continuous Driver Management Profile (D_p), successfully predicted the non-linear grip loss (the cliff point) with a Mean Absolute Error (MAE) of 1.2 laps. In contrast, the baseline RF model and standard LSTM, which lacked physical and thermodynamic constraints, exhibited significantly higher error rates, frequently overestimating the remaining useful life of the tire.

TABLE I
Performance Comparison of Predictive Models

Model Architecture	**Cliff Pred. MAE (Laps)[a]**	**RMSE**	**Latency (ms)**
Baseline Random Forest	3.8	0.42	45.0
Standard LSTM	2.4	0.28	22.5
TyrePIML (Proposed)	**1.2**	**0.14**	**14.2**

[a]Mean Absolute Error for predicting the non-linear grip loss.

Furthermore, performance profiling of the WebSocket and LRU caching infrastructure demonstrated a computational latency of 14.2 milliseconds per telemetry update, compared to the 45.0 milliseconds of standard HTTP polling methods. This confirms the system's viability for high-speed edge deployment on a live pit wall.

VI. Limitations and Future Work

While TyrePIML provides a robust strategic framework, certain limitations remain. The current PIML model assumes highly accurate sensor data; however, in real-world scenarios, telemetry pipelines are susceptible to packet loss and sensor noise. Additionally, the Game Theory module currently models rival responses based on deterministic MGU-K thresholds, without accounting for stochastic driver errors or unpredictable safety car deployments. Future work will focus on integrating reinforcement learning (RL) agents to replace the static Nash Equilibrium matrix, allowing the system to adapt to highly volatile, probabilistic race conditions. Research will also expand the model to include V2V communication protocols for cooperative energy harvesting in civilian autonomous fleets.

VII. Conclusion

This paper presented TyrePIML, an advanced, real-time race strategy engine designed to tackle the complexities of the 2026 Formula 1 regulations. By discarding traditional polling mechanisms in favor of an event-driven WebSocket architecture and LRU in-memory caching, the system achieves the ultra-low latency required for high-speed edge computing.

The integration of Physics-Informed Machine Learning (PIML) provides highly accurate predictions of non-linear tire degradation cliffs by accounting for physical track conditions, compound coefficients, and driver-specific management profiles. Furthermore, the application of Nash Equilibrium-based game theory transforms static lap-time predictions into a dynamic, multi-agent behavioral model, strictly constrained by real-time MGU-K energy reserves and active aerodynamic wake penalties.

Beyond competitive motorsport, the architecture of TyrePIML holds significant deployment value for the broader Intelligent Transportation Systems (ITS) industry. The methodologies developed for multi-agent adversarial racing directly translate to cooperative and non-cooperative trajectory planning in connected and autonomous vehicles (CAVs). The real-time energy prediction and game-theoretic decision-making modules can be deployed in commercial hybrid-electric fleets to optimize battery State of Charge (SoC) over dynamic routes. Simultaneously, the low-latency WebSocket data pipeline serves as a robust blueprint for Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communication networks [11]. Ultimately, TyrePIML demonstrates that the extreme engineering demands of motorsport can yield scalable, physics-informed solutions for the future of autonomous mobility.

References

[1] J. Betz et al., "Autonomous Vehicles on the Edge: A Survey on Autonomous Vehicle Racing," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 3, pp. 458–488, 2022.

[2] W. Schwarting, J. Alonso-Mora, and D. Rus, "Planning and Decision-Making for Autonomous Vehicles," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 187–210, May 2018.

[3] A. Sciarretta and L. Guzzella, "Control of hybrid electric vehicles," *IEEE Control Systems Magazine*, vol. 27, no. 2, pp. 60–70, April 2007.

[4] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, Feb. 2019.

[5] M. Chiang and T. Zhang, "Fog and IoT: An Overview of Research Opportunities," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 854–864, Dec. 2016.

[6] K. Zhang et al., "Mobile Edge Computing for Vehicular Networks: A Promising Network Paradigm with Predictive Off-Loading," *IEEE Vehicular Technology Magazine*, vol. 12, no. 2, pp. 36–44, June 2017.

[7] H. B. Pacejka, *Tire and Vehicle Dynamics*, 3rd ed. Oxford, U.K.: Butterworth-Heinemann, 2012.

[8] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

[9] J. F. Nash, "Equilibrium points in n-person games," *Proc. Nat. Acad. Sci. U.S.A.*, vol. 36, pp. 48–49, Jan. 1950.

[10] R. Wang et al., "Game-Theoretic Cooperative Lane Changing under Connected and Automated Environments," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 8, pp. 5448–5459, Aug. 2021.

[11] E. Moradi-Pari et al., "Connected and Automated Vehicles: Real-World Implementations and Challenges," *IEEE Communications Magazine*, vol. 55, no. 12, pp. 119–125, Dec. 2017.