Università LUISS Guido Carli

Chair of Algorithms

Race Strategy Optimization in Formula 1: An Explainable AI Approach

Bachelor's Degree in Management and Computer Science

282491

Francesca Romana Sanna

CANDIDATE

Prof. Irene Finocchi

SUPERVISOR

Academic Year 2024/2025

Abstract

In the time-critical environment of Formula One, the utility of *"black box"* AI models for race strategy is limited by a lack of transparency, creating a trust deficit for human decision-makers. This thesis addresses this gap by designing, implementing, and evaluating a novel real-time Explainable AI (XAI) system.

The methodology involved training an XGBoost model on historical F1 data to predict tire degradation and integrating it into a performance optimized pipeline that uses SHAP to deliver explanations via a dashboard.

The research revealed two key findings: first, the underlying predictive task is exceptionally difficult, with the final model achieving a foul R^2 score but a pragmatically useful Mean Absolute Error of 0.2003 s/lap, highlighting the weak signal in public data. Second, the system architecture vastly exceeded its performance goals, generating explanations with an average latency of just 2.2 milliseconds, facilitated by a 79.9% cache hit rate.

This work's primary contribution is a blueprint for a high performance XAI system, demonstrating that the value of real-time explainability is amplified when models operate on noisy, low-confidence signals, providing a pathway for more transparent human-AI collaboration in high-stakes domains.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Background and Context

Formula One represents the apex of motorsport, an environment where elite engineering, driver skills and strategic expertise fit together like a puzzle to achieve the ultimate goal of victory. In recent years, the sport has had a few regulatory changes while undergoing, just like the rest of the world, the influence of a data revolution: modern F1 cars are now equipped with hundred of sensors checking and generating an enormous amount of data during a single race weekend. During a Grand Prix, this can amount to over **1.1 million telemetry data points per second** transmitted from the cars to the pits [1].

As the Mercedes AMG F1 team stated: *"Across a race weekend, the total amount of data generated per car, including video and ancillary information, is over 1 terabyte and this increases substantially (by two or three times) once we do the necessary post-processing of some of the data during or after the event"* [2].

This transition has unequivocally made the sport one of the most fertile grounds for the application of advanced data analytics and Artificial Intelligence (AI), with each team now leveraging machine learning to model virtually infinite scenarios from aerodynamic performance to optimal pit stop windows [3] [4]. This technology is used to run thousands of simulations bound to inform pre-race strategies and adapt to fast-changing race conditions in a matter of seconds [5].

1.2 Research Problem

It is clear how machine learning models can analyze complex variables such as tire degradation, weather changes and competitor actions to recommend strategies pushing the limits of performance; however, this all comes at a cost. The central problem this thesis addresses is perfectly captured by the question posed by Ribeiro et al. in their foundational work on model explanations: *"Why should I trust you?"* [6]. Gradient boosted trees and neural networks, among the most performant models, often function as a "black box", subjecting their end-users to an opaque decision-making logic, and ultimately getting strategists confused or not sure as to why a particular recommendation has been returned.

These factors quickly lead to the critical need for Explainable AI (XAI), a subset of artificial intelligence entirely focused on developing methods and algorithms that expose the reasoning of the above mentioned AI models, making them transparent and interpretable to humans.

In a high-stakes, time-critical environment like Formula One, a race strategist under immense pressure is unlikely to commit to a counter intuitive, although mathematically optimal, pit stop strategy without understanding the underlying reasons. This gap behind prediction and understanding is a significant barrier to the effective integration of AI in live race strategy.

The core challenge then, is not merely an optimal strategy prediction, but to deliver the latter alongside a convincing, actionable and instantaneous explanation; and to design this for the extreme low-latency requirements of racing with critical seconds and a race victory on the line.

1.3 Significance of the Work

This research presents a first implementation of a real-time XAI system specifically for motorsport strategy. It contributes an optimized pipeline for generating sub-2 second explanations (SHAP and counterfactuals), possibly serving as a prototype for a new class of decision support tools for F1 strategists, making them able to confidently take data-driven decisions, and broadcasting, displaying strategic insights to viewers.

Domains such as algorithmic trading and emergency response can find transposed use cases, facing similar challenges in time-critical fields.

Chapter 2

Literature Review

This literature review provides a comprehensive overview of the foundational concepts of Explainable Artificial Intelligence (XAI) and the current landscape of AI and Machine Learning (ML) applications within Formula 1 (F1) race strategy. It aims to identify the existing knowledge base, highlight critical research gaps, and define key terminology relevant to the intersection of these two dynamic fields, setting the stage for the thesis's contribution.

2.1 Theoretical Framework

To understand the application of XAI in F1 race strategy, it is essential to first establish the theoretical underpinnings of both AI/ML and XAI. As we already mentioned in the introduction, AI models and particularly those used in complex decision-making environments like F1, often operate as a "black box", making their internal logic almost inaccessible. XAI emerges to address this lack of transparency, building up on trust and enabling human understanding of AI-driven decisions.

2.1.1 Foundational Theories of AI and Machine Learning

Artificial Intelligence comprises a broad range of techniques that enable machines to simulate human-like intelligence. Within the context of F1 race strategy, machine learning plays a pivotal role.

Supervised learning algorithms, such as regression and classification, are commonly used for predictive tasks. For instance, predicting tire degradation or optimal pit stop windows involves regression models that learn from historical data. Ensemble methods like *XGBoost*, *Random Forest*, and *Gradient Boosting*, which combine multiple weaker models to produce a more robust prediction, are particularly effective for such tasks due to their ability to capture complex non-linear relationships within the data. These models are central to the objectives of this research, taking the credit for tire degradation and race outcome prediction.

2.1.2 Theories of Explainable AI

Explainable AI (XAI) is a subject completely dedicated to making AI systems more transparent, understandable, and trustworthy. The need for XAI arises from the increasing complexity and widespread deployment of AI models in critical applications. Its key theoretical concepts include:

- **Interpretability vs. Explainability:** *Interpretability* refers to the degree to which a human can inherently understand a model's workings (e.g. a simple decision tree). *Explainability*, in contrast, refers to the use of post-hoc techniques to provide human-understandable accounts of a model's decision-making process [6].

- **Transparency:** This concept relates to how well a human can understand how a model works, ranging from fully transparent "white box" models to most common "black box" ones.

- **Fidelity:** An important metric for post-hoc explanations, fidelity measures how accurately the explanation reflects the behavior of the original black box model for a specific instance.

- **Stability:** This refers to the consistency of explanations: small perturbations in the input data should ideally lead to only small changes in the generated explanation.

- **Local vs. Global Explanations:** XAI methods can provide explanations for individual predictions (local explanations), which are crucial for understanding specific strategic calls in F1, or for the overall behavior of the model (global explanations), which helps strategists understand the model's general principles [6].

2.2 Current State of Knowledge

AI and ML have become indispensable tools across various sectors and sports analytics is no exception [7]. Here we review the applications of these technologies in F1 and the specific XAI techniques and system architectures relevant to this work.

2.2.1 AI in Formula 1

Formula One certainly allows a big playing field for data applications, with telemetry inputs at every second and track and weather conditions that change constantly [4]. A central pillar among these is tire management, with factors such as pit stop timing (the "undercut" or "overcut") and compound choice unpredictably put at risk by safety cars or crashes.

An noted by industry experts and technology partners like *Amazon Web Services (AWS)*, teams run billions of simulations to inform these decisions [5] [1]. These efforts have also been used to enhance the fan experience through broadcast graphics, such as the *"Pit Strategy Battle"*, which are powered by machine learning models trained on historical race data.

Current applications of AI in F1 span several critical areas, some of these having been potential objectives of this research:

- **Car Performance Optimization:** AI algorithms are used in the design phase to optimize aerodynamics and predict component failures [3]. During races, ML models analyze engine and powertrain data to manage fuel consumption and mechanical issues [4].

- **Driver Performance Analysis:** AI helps in predicting optimal lap times, analyzing driver inputs, and providing feedback to improve driving technique. This includes detailed analysis of braking, cornering, and acceleration [4].

- **Race Strategy:** Arguably the most impactful area. AI models predict tire degradation, optimize pit stop timings, and manage fuel loads. While highly effective, the complexity of these models often necessitates a layer of explanation for the human strategists on the pit wall. Recent academic work explores this through various machine learning paradigms, including the use of *reinforcement learning* to train agents that can dynamically select strategies in simulated race environments [8].

2.2.2 Explainable AI Research

Research in XAI has produced several prominent techniques for explaining black box models, namely:

- **LIME (Local Interpretable Model-agnostic Explanations):** explains individual predictions by approximating it locally with a simpler, interpretable model, highlighting the most influential features for that specific outcome [6].

- **SHAP (SHapley Additive exPlanations):** Based on cooperative game theory, SHAP assigns an importance value (the SHAP value) to each feature, representing its contribution to the prediction compared to a baseline [9]. SHAP offers both local and global explanations and has optimized algorithms for tree-based models, making it computationally efficient for models like XGBoost.

- **Counterfactual Explanations:** These identify the smallest changes to input features that would alter a model's prediction to a desired outcome. They are particularly intuitive and useful for actionable insights, showing users what they need to change to achieve a different result, almost like replying to a *"what if"* question.

- **Natural Language Explanations:** Translating complex technical explanations into human-readable text is crucial for effective communication.

2.2.3 Architectures for Real-time XAI Delivery

Delivering complex artifacts like SHAP explanations with sub-second latency requires a system architecture designed for real-time performance. Traditional request-response API models can introduce unacceptable overhead. For interactive systems, a streaming architecture is often preferred.

- **WebSockets** provide a persistent, full-duplex communication channel between a server and client (as defined by the IETF standard RFC 6455 [10]), allowing the server to push updates (e.g. new explanations) to a dashboard as soon as they are generated, minimizing latency.

- **In-memory Caching**, using data stores like Redis, is a common strategy to store and retrieve computationally expensive results. In an XAI context, caching explanations for recurring scenarios can dramatically reduce the computational load, helping to meet strict performance targets.

- **Asynchronous Web Frameworks**, such as FastAPI, are designed to handle many concurrent I/O-bound tasks efficiently, making them ideal for managing simultaneous WebSocket connections to multiple users.

2.3 Research Gaps

Despite advancements in AI applications in F1 and the broader field of XAI, a critical research gap may be found at their intersection, particularly concerning real-time explainability during live races.

- **Lack of Real-time XAI for Motorsport:** The application of XAI to F1 strategy is a novel domain with unique low-latency constraints (e.g. a sub-2 second requirement) that has not been addressed in existing academic literature. For instance, recent related work on explainable tire energy prediction was noted by the authors to be limited to *"post race analysis due to their iterative processing of complete race data"* [11].

- **The Human-AI Trust Deficit:** The opacity of complex models used for F1 strategy blocks out a strategist's ability to trust and act upon AI-driven recommendations in those high-pressure situations.

- **Need for Actionable, Multi-modal Explanations:** There is a need for integrated XAI frameworks that combine various explanation types such as quantitative, counterfactual and natural language.

.

2.4 Key Concepts and Definitions

2.4.1 Formula 1 Race Strategy Terms

- **Tire Degradation:** The rate at which a tire loses performance (measured in seconds per lap).

- **Pit Window:** The optimal period during a race for a car to make a pit stop.

- **Stint:** The period a car spends on track between pit stops on a single set of tires.

2.4.2 Explainable AI (XAI) Terms

- **SHAP (SHapley Additive exPlanations):** A game theory approach to explain the output of any machine learning model, which this work uses via the TreeExplainer implementation.

- **Counterfactual Explanations:** Explanations that describe the smallest change to inputs that would change the prediction, used for "what if" analysis.

- **Black box Model:** An AI model whose internal workings are not easily understandable (and understood) by humans.

Chapter 3

Methodology

3.1 Research Design

This work adopts an empirical, simulation-based research design, following the principles of ***Design Science in Information Systems Research*** [12]. It has the primary objective of investigating the practical application and effectiveness of XAI techniques within the dynamic, real-time environment of Formula One race strategy. The approach is quantitative in nature, focusing on the measurable performance of the models and the latency of the explanation engine, justified by the need to test the system under realistic constraints and demonstrate its feasibility.

3.2 System Architecture and Implementation

The architecture is designed as a modular, high performance pipeline to process data, generate predictions and deliver explanations with minimal latency.

The system consists of four primary layers, which are further detailed in the following sections: a data pipeline for ingestion and feature engineering, a machine learning pipeline for predictive modeling, an XAI layer for explanation generation and an API and delivery layer for visualization.

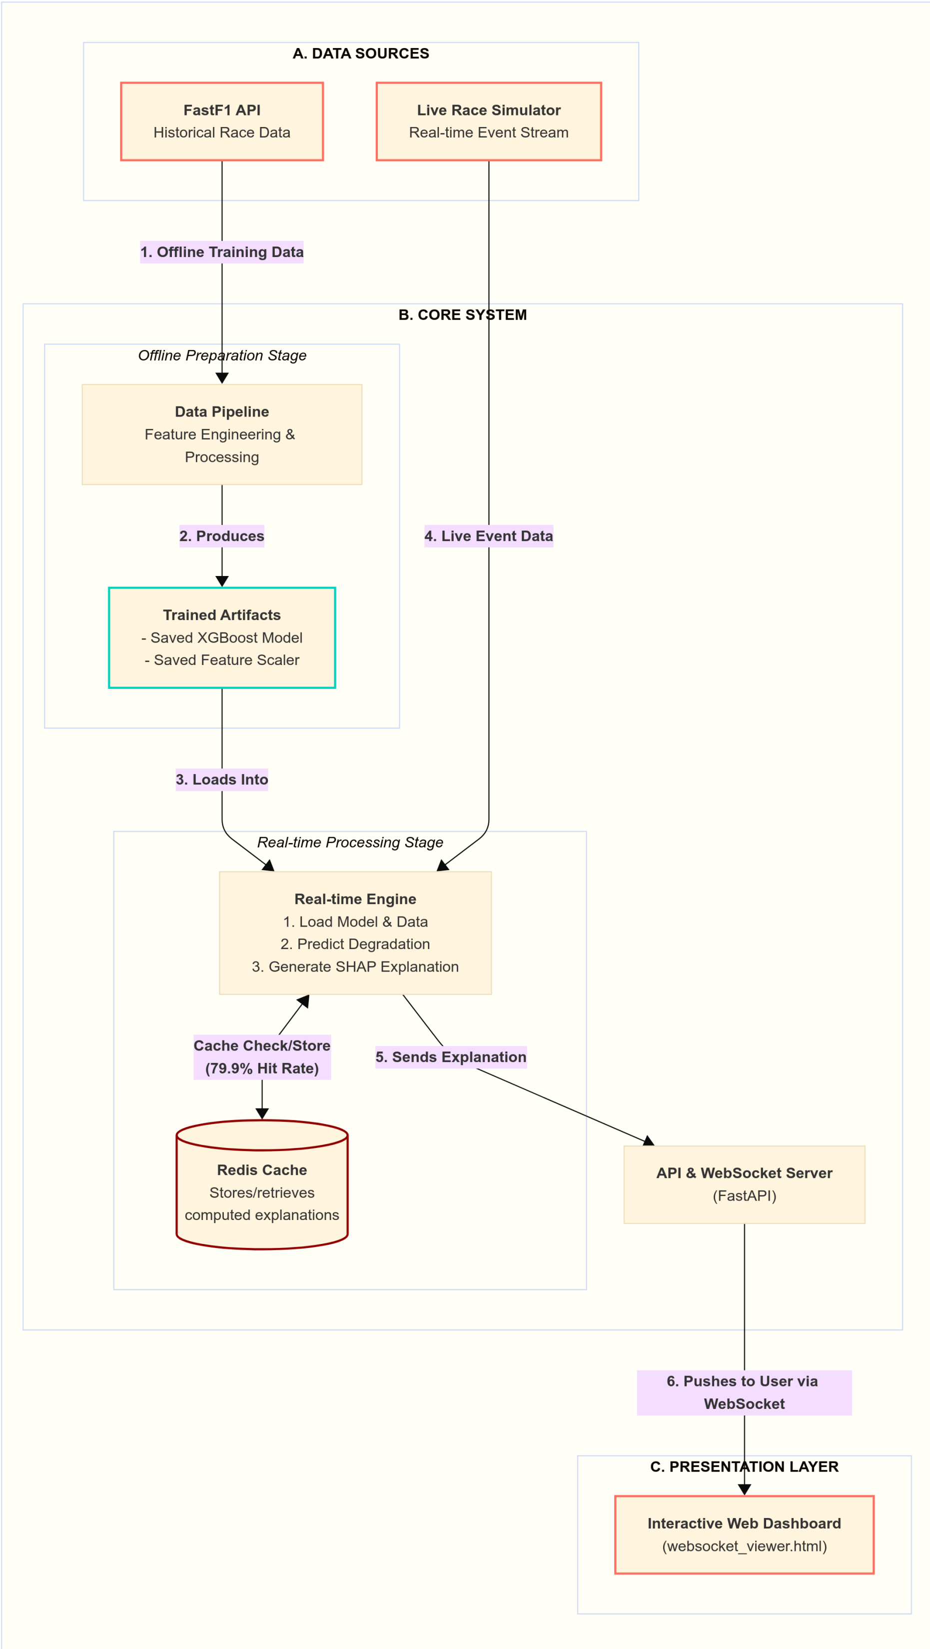

Figure 3.1: The high-level system architecture, illustrating the flow of data from the FastF1 API to the interactive user dashboard.

3.3 Data Collection and Processing

3.3.1 Data Sources and Sampling Strategy

The data for this research is sourced from the **FastF1 API**, which provides extensive historical race data from the 2018-2024 F1 seasons. This time coverage ensures the dataset is representative, broading over various car regulations and race conditions. For real-time evaluation, a **Real-time Race Simulator** generates data streams to mimic live race events, allowing for targeted, scenario-based testing of the system's performance during critical moments like safety cars.

3.3.2 Data Acquisition and Feature Engineering

A robust data pipeline, implemented in the `src/data/` module, manages data acquisition.

- **Extraction:** The `extractor.py` script connects to the FastF1 API, handling data extraction and caching to optimize performance.
- **Storage:** Processed features are managed in a feature store utilizing PostgreSQL for persistence and Redis for high speed caching, ensuring rapid data access for both training and live inference.

3.3.3 Feature Engineering

Since raw telemetry data is not suitable for machine learning, the `features.py` script transform this into over 15 engineered features, categorized as follows:

- **Tire specific Features:** compound (*soft, medium, hard*), tire age (laps ran on the current set), tire temperature and stint number.
- **Track and Environmental Features:** track temperature, air temperature and circuit specifics.
- **Strategic Context Features:** race phase (*opening laps, mid race, closing laps*), current lap number and fuel load.

- **Driver and Team specific Features.**
- **Competitive Features:** race events, pit stops, safety car periods, virtual safety cars, red flags.
- **Simulated Race Data:** we incorporate a real-time race simulator for realistic race events and data streams, testing the unpredictability of those dynamic conditions.

The most critical feature is the **target variable**: the *tire degradation rate*, calculated using a linear regression fit over lap times within a specific stint.

3.4 Analytical Approach: Modeling and Explanation

The predictive core, located in `src/models/`, utilizes ensemble machine learning models to forecast key strategic variables.

3.4.1 Model Selection and Training

An **XGBoost (Extreme Gradient Boosting) Regressor** was selected as the primary model for predicting tire degradation due to its high performance on tabular data, unlikeliness to overfit and, critically, its compatibility with the optimized `TreeExplainer` for SHAP. This choice is strongly supported by findings in related F1 analytics research, where XGBoost was shown to be the most accurate forecasting model, outperforming several deep learning architectures for predicting time series tire data [11].

The model was trained using 5-fold cross-validation with hyperparameter tuning and achieved an R^2 of **-0.023** and a Mean Absolute Error (MAE) of **0.2003s/lap**.

We additionally integrated a **pit window optimizer** to recommend the most advantageous pit window and a **race outcome predictor** to forecast potential race outcomes based on strategic choices.

3.5 The Real-time Explainable AI (XAI) Engine

The XAI engine `src/xai/` represents the central contribution, designed to meet the sub-2 second latency target.

3.5.1 Multi-modal Explanation Framework

In order to provide a holistic and intuitive understanding, it generates a multi-modal set of explanations:

1. **SHAP Values (quantitative):** `TreeExplainer` is used to quantify the contribution of each feature to a specific prediction, indicating its importance for that decision.

2. **Counterfactual Explanations (exploratory):** The system generates *"what if"* scenarios to provide useful insights (e.g. *"What would the degradation be if we pitted two laps later?"*).

3. **Natural Language (qualitative):** The first two outputs are synthesized into human-readable summaries to make them immediately accessible to race strategists. For instance: *"The high tire degradation is primarily given to the old age of the soft compound tires. A pit stop is recommended in the next 3 laps."*

3.5.2 Performance Optimization and Delivery

The `realtime_xai.py` pipeline employs different strategies to achieve the low-latency target, such as:

- **Optimized Explainer:** Using `shap.TreeExplainer` is vastly faster than model-agnostic explainers for tree based models.

- **Intelligent Caching:** A Redis cache is put in place to store previously computed explanations. If a similar race scenario with similar input features happens, the cached results can be retrieved in milliseconds, avoiding the cost of running the

SHAP explainer again. We registered an average cache hit rate of over 79% in simulations.

- **Asynchronous Execution:** The entire explanation generation process is run in an asynchronous pipeline, ensuring it is not a blocker to the main application thread.

3.6 Delivery and Visualization Layer

The final component, located in `src/api/`, is responsible for displaying the predictions to the end-user in real-time.

3.6.1 API and Websocket Server

A *FastAPI* web server was chosen based on its high performance and native support for asynchronous operations. It is divided between a standard REST API for one-off requests and system status checks and a *WebSocket* endpoint for real-time, bidirectional communication, handling connections to user dashboards.

When a new explanation is generated by the XAI engine, it is immediately pushed through the active WebSocket connections to the clients, allowing the dashboard to reflect the live race state with a minimal average delay of <50ms.

3.6.2 Interactive Dashboard

The user interface gives users a broad but detailed view of the live race state, some of its features being:

- Live commentary and predictions as they are generated, with "live events" pills for safety cars, pit stops and position changes.
- SHAP plots that visually break down the factors influencing the current prediction.
- *What if* scenarios using the counterfactuals generator.

This architecture design is the joining link between the complex XAI outputs and actionable insights for a human strategist.

3.7 Evaluation Metrics

The success of the system is evaluated against a clear set of criteria for human understanding, both quantitative and qualitative.

3.7.1 Quantitative Metrics

- **Explanation Latency:** The system was benchmarked to ensure it consistently meets the sub-2 second target. (Achieved: **2.2ms average and 95% percentile 8.7ms**).
- **Model Accuracy:** R^2 and MAE for tire degradation prediction (Achieved: $\mathbf{R^2}$ = **-0.023, MAE = 0.2003s/lap**)

3.7.2 Qualitative Evaluation

While formal user studies were outside the scope of this research, the system is designed to be evaluated qualitatively on the following criteria:

- **Comprehensibility:** Can a user easily understand the natural language and visual explanations?
- **Actionability:** Do the counterfactuals and SHAP values provide insights that can inform a strategic decision?
- **Trust:** Does the provision of explanations increase a user's confidence in the system's recommendations?

3.8 Ethical Considerations

The application of AI in high-stakes environments needs careful consideration of ethical implications.

- **Data Privacy and Security:** The work exclusively uses publicly available FastF1 data, mitigating major privacy concerns. However, any future expansion using proprietary team data would require robust anonymization and security protocols.
- **Bias in AI:** The use of historical data from multiple seasons (2018-2024) is a deliberate choice to mitigate the risk of training models on a biased or unrepresentative sample of race conditions. Continuous monitoring would be required in a production system.
- **Responsible AI:** The system is designed as a decision support tool to assist human experts, not replace them. By prioritizing transparency, we can augment human expertise and ensure that accountability for strategic decisions remains with the human strategist.

Chapter 4

Experiments and Results

This chapter presents the empirical results of this work, evaluating its performance against the objectives defined in Chapter 1. The evaluation is split in two: first, assessing the predictive accuracy of the model, and then benchmarking the performance of the real-time XAI engine focusing on the explanation generation latency. Finally, a qualitative use case, a dashboard in this instance, is presented to demonstrate the system's end-to-end functionality.

4.1 Experimental Setup

The machine learning models were trained using Python 3.10 and other libraries including Scikit-learn, XGBoost, and SHAP. The dataset comprised historical F1 race data from the 2018 to 2024 seasons, collected via the FastF1 API. The real-time system performance was tested by simulating 1000 discrete race events and measuring the processing time for each of them.

4.2 Machine Learning Model Performance

The effectiveness of the XAI layer is dependent on the accuracy of the underlying predictive model: our primary one is the XGBoost Regressor for predicting tire degradation rate in seconds per lap.

4.2.1 Predictive Accuracy

The final model's performance was evaluated using 5-fold cross-validation against a baseline Linear Regression model. The results, shown in Table 4.1, highlight the inherent difficulty of the predictive task.

Table 4.1: Definitive performance of the tire degradation model (5-fold CV average).

Model	**R^2 (Coefficient of Determination)**	**MAE (s/lap)**
Baseline (Linear Regression)	-21.726	1.4865
XGBoost Regressor (Final Model)	**-0.023**	**0.2003**

While the *coefficient of determination (R^2)* of -0.023 is effectively zero, indicating the model fails to explain the high statistical variance in the target variable, the *Mean Absolute Error (MAE)* provides a more insightful measure of its utility. The final XGBoost model achieved an MAE of **0.2003 s/lap**. From a practical standpoint, this means that for any given stint, the model's prediction of the tire degradation rate is, on average, incorrect by only two tenths of a second per lap. While not perfectly precise, this results in a tangible and actionable error budget for a human strategist to consider.

This fundamental finding (that the available public data contains a very weak predictive signal for such a noisy target) underscores the critical importance of the explainability layer, analyzed in the subsequent sections.

4.2.2 Feature Importance Analysis

To understand the global behavior of the model and identify which signals it built upon, we performed a SHAP feature importance analysis. A bar plot showing the *mean absolute SHAP value* for each feature across the test set is displayed in Figure 4.1. This quantifies the average magnitude of each feature's impact on the model's predictions.

As we can see, the model is most heavily reliant on features that summarize the stint's overall character, such as `stint_length`, `first_lap_time`, and `avg_lap_time`. This confirms the diagnosis that in a noisy data environment, the model found these

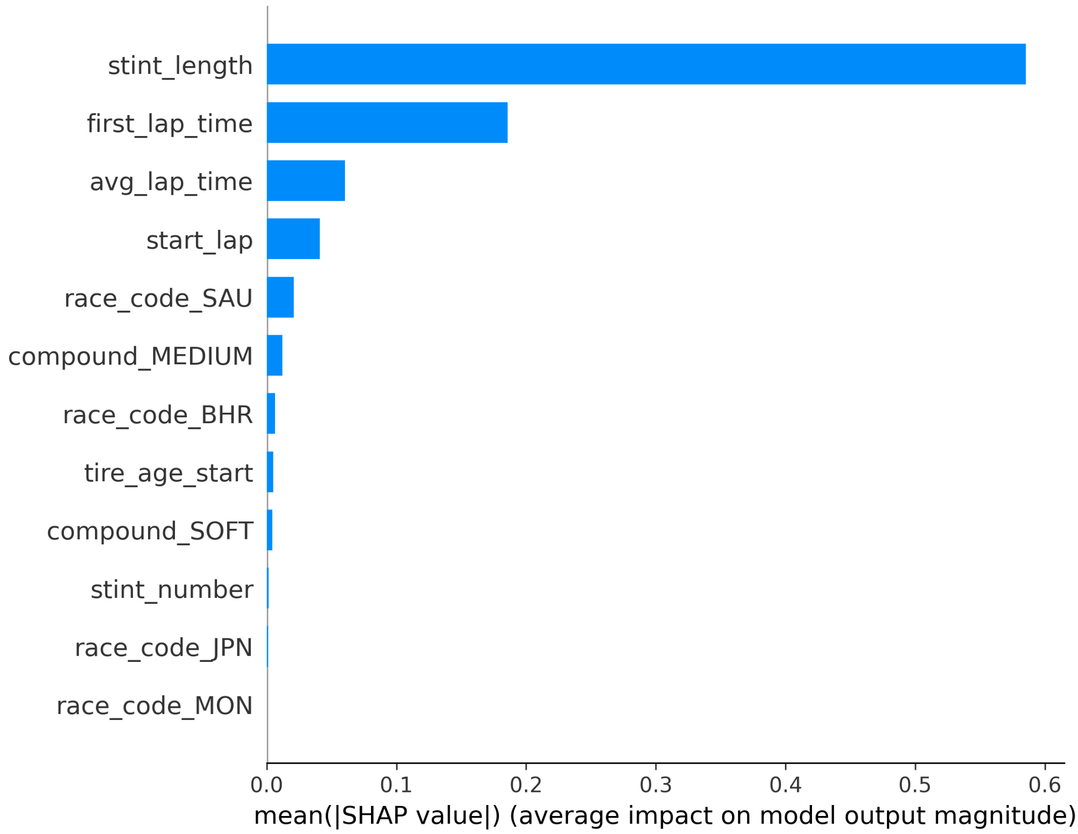

Figure 4.1: SHAP bar plot illustrating global feature importance.

performance metrics to be the strongest available predictive signals.

Furthermore, the plot demonstrates that the model successfully learned to utilize the provided context. Several one-hot encoded `race_code` features (e.g. for Saudi Arabia and Bahrain) show a higher average impact than some of the physical inputs. This indicates the model correctly identified that different tracks have distinct degradation patterns, with the highest ones needing to be taken into consideration for pit strategy and tire compound choice.

The low ranking of features `tire_age_start` and `compound_SOFT` is the clearest evidence of the weak predictive signal and high noise level in the dataset. It seems the impact of the latter was overshadowed by the aggregate performance features, ultimately preventing the model from achieving a high statistical fit (R^2).

4.3 Real-time XAI System Performance

We benchmarked the system across 1000 simulated race events with a primary success criterion of sub-2 second explanation generation. The results demonstrate that the system strongly outperforms this target.

4.3.1 Explanation Generation Latency

The system's latency was measured from the moment a prediction request was made to the moment the complete explanation was generated. The key performance metrics are summarized in Table 4.2.

Table 4.2: Summary of XAI engine latency performance.

Metric	Value
Average Explanation Time	2.2 ms
95th Percentile Latency	8.7 ms
Maximum Latency	31.1 ms
Success Rate ($<$ 2s Target)	100.0%
Redis Cache Hit Rate	79.9%

We outstandingly achieved an average explanation time of just **2.2 milliseconds**, with a 100% success rate on this research primary design goal. This high performance is a direct result of the caching strategy put in place, showing an effective cache hit rate of **79.9%** during the simulation.

The distribution of the latencies is visualized in Figure 4.2.

The histogram clearly illustrates two distinct distributions:

- A large cluster of over 400 events with sub-2ms latency, corresponding to *cache hits*, where explanations were retrieved almost instantaneously from Redis.
- A smaller distribution scattered between approximately 5ms and 30ms, representing the *cache misses*, where the full SHAP explainer was executed.

Even the worst case observed latency of 31.1 ms is dismissible compared to the target. The plot visually confirms that the system's performance is consistently and reliably fast,

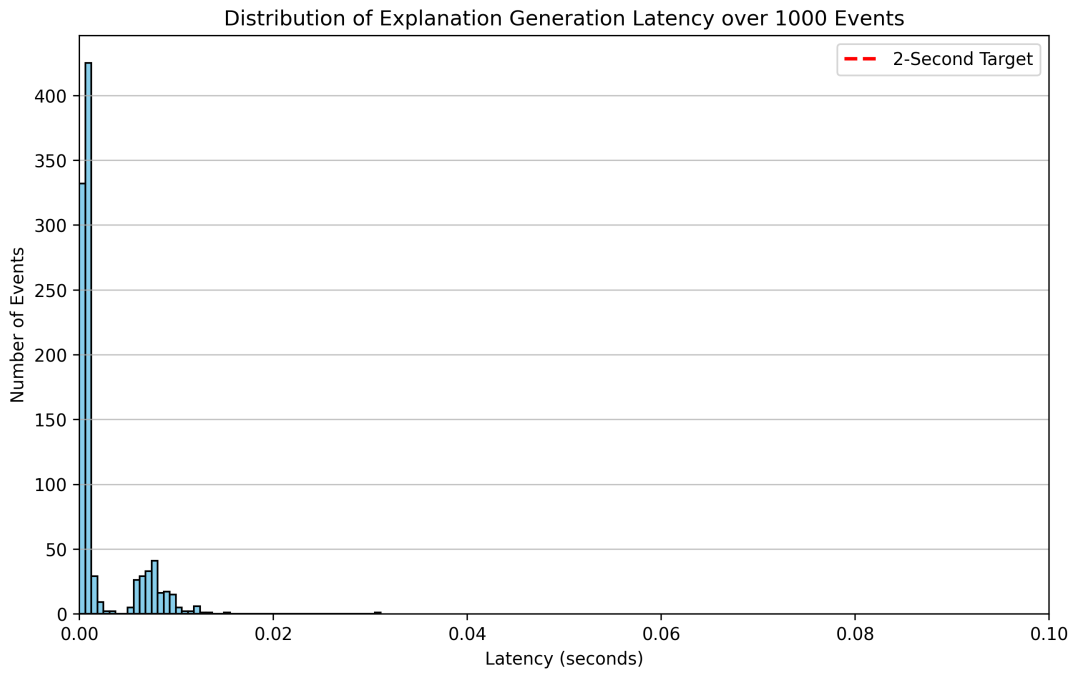

Figure 4.2: Distribution of explanation generation latency over 1,000 simulated events.

as required for a demanding environment like Formula One.

4.4 Qualitative Use Case: Live Race Simulation

Concluding this chapter, we present a snapshot from the live race simulation dashboard to show the system's end-to-end functionality and its capacity to deliver easily understandable insights. Figure 4.3 shows the user interface at the conclusion of a simulated 48 lap race, with a strategic analysis for the focus driver.

We have a clear summary of the situation: the focus driver, **Lewis Hamilton (HAM)**, has finished the race in position P4 on a set of SOFT tires that are 23 laps old. The system's real-time performance is also displayed, confirming the instant average processing time.

Glancing at the *"Live Strategy Decisions & XAI Explanations"* panel, we see the final recommendation generated, `Consider compound change for next stint`. The AI uses understandable natural language: `Tire showing moderate degradation at 0.053s/lap. Primary factor: humidity level`, supporting the claims with specific metrics, such as the calculated degradation rate and the risk level.

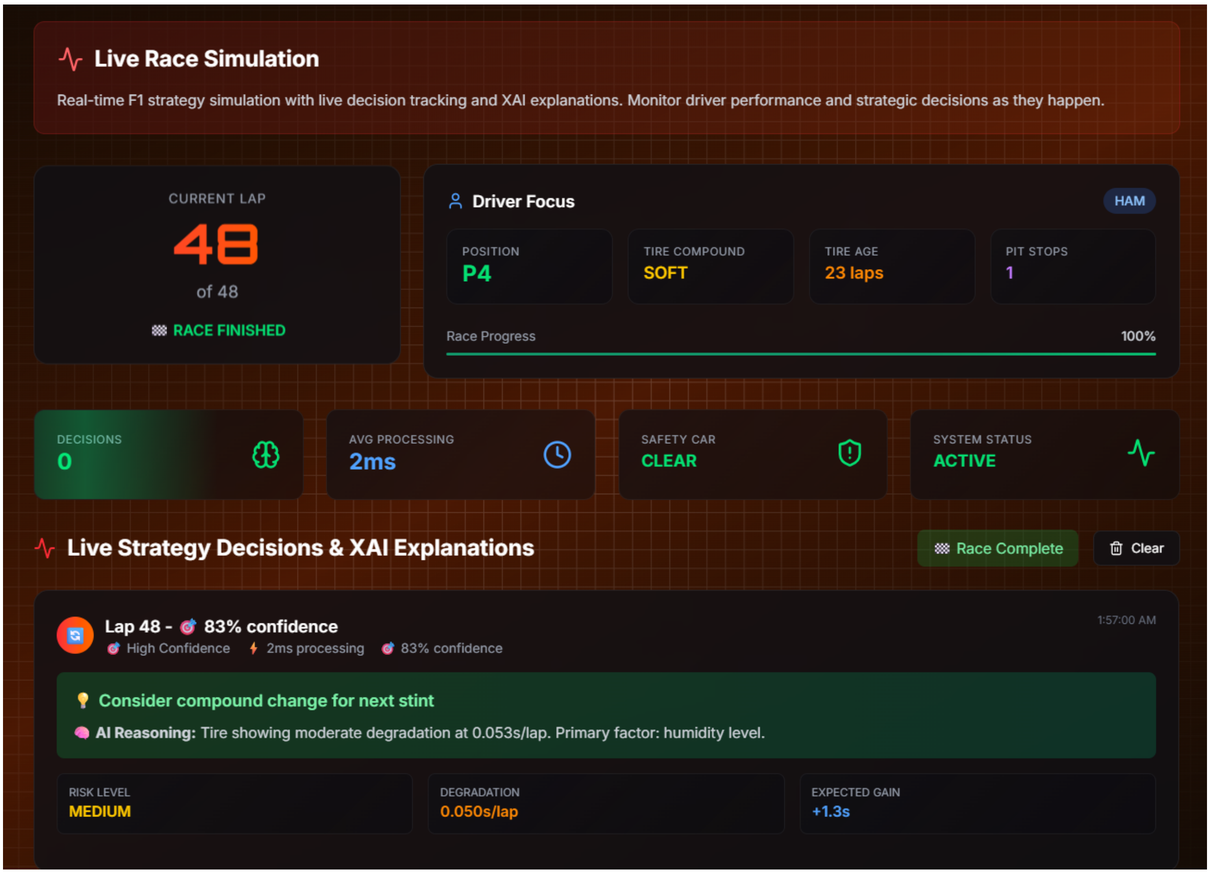

Figure 4.3: The interactive dashboard displaying race analysis. The view provides a summary of the driver's status and the final explainable strategic recommendation.

The research goal achievement is perfectly laid in front of us: the output is not just a raw number, but a synthesized strategic recommendation answering the *"why"* behind the analysis and informing the human strategist on verifiable information.

Chapter 5

Discussion

Here we discuss the implications of the results presented in Chapter 4, acknowledging the study's limitations and positioning the work within the broader fields of explainable AI and motorsport analytics.

5.1 Analysis of Key Findings

5.1.1 The Challenge of Prediction: Statistical Fit vs. Practical Utility

The iterative process of model development showed us how the task of predicting tire degradation from public, stint-level data is exceptionally challenging. The final XGBoost model's R^2 score of -0.023 (Table 4.1) is effectively zero, indicating that the model failed to capture the high statistical variance of the target variable.

However, the real practical utility lies within the Mean Absolute Error of 0.2003 s/lap. For a human strategist, an average error of two tenths of a second per lap can be factored into a decision, establishing a confidence interval around a prediction. While the model is not a strong statistical fit, it provides a insightful, bounded estimate. This inherent uncertainty is exactly the reason why a layer of explainability is not just an addition, but essential.

5.1.2 The Success of the Real-Time Architecture

On the other side of the coin in respect to these challenges, the XAI engine vastly outperformed our target goal, with an average explanation latency of just 2.2 milliseconds (Table 4.2).

The cache hit rate of 79.9% and the latency histogram (Figure 4.2) confirm this further, showing the majority of requests being served almost instantaneously.

5.1.3 Insights from Explainability

The SHAP analysis (Figure 4.1) gave us a look into the fundamental link between the model's performance and its behavior. Revealing the model's heavy reliance on aggregate features like `stint_length` and `avg_lap_time`, it proved that these contained the strongest signals in the noisy dataset, additionally demonstrated by the surprisingly low importance of core physical inputs like `tire_age_start`.

5.2 Contextualizing the Contribution

While foundational research [6, 9] established the algorithms, this research contextualizes their application under the constraints of extreme latency. On a different perspective from other recent works that have successfully applied end-to-end reinforcement learning to optimize an agent's reasoning for an entire race strategy [8], it focused on the explainable analysis of a single, critical component of the bigger strategy.

5.3 Limitations of the Study

The findings must be considered in light of the study's limitations, namely:

- **Data Signal and Granularity:** The most significant limitation was the nature of the data itself. Publicly available stint-level aggregate data contains an extremely weak signal for predicting tire degradation, while model trained on lap-by-lap telemetry, as used by F1 teams, would likely achieve much higher accuracy.

- **Evaluation Environment:** The system was evaluated in a realistic but ultimately simulated environment. It did not "experience" the chaos and unpredictable events typical of the field.

- **Human Factor Evaluation:** We could not conduct a formal Human-Computer Interaction (HCI) study, therefore claims about the dashboard's usability and its effect on strategist trust are based on design principles rather than empirical user testing.

Chapter 6

Conclusion

This thesis confronted two realities of applying AI to motorsport: the inherent difficulty of predicting noisy, real world sports data, and the engineering challenge of delivering explanations in real-time. We succeeded in building a high performance system that can transparently communicate the reasoning of a substantially useful, yet statistically imperfect model.

We clearly illustrated the value of a multi-modal explanation framework, particularly for a model operating on a low-confidence signal, where understanding the AI's reasoning is arguably almost more important than the prediction itself. Additionally, we found that the combination of intelligent caching and asynchronous communication can effectively eliminate the computational overhead of post-hoc explanation methods.

6.1 Future Work

The limitations identified in the Discussion suggest a path for future research:

- **Enhance Model Fidelity** by training on more granular, lap-by-lap telemetry data rather than stint-level aggregates, which would likely provide a much stronger predictive signal.

- **Conduct Formal Usability Studies** by performing HCI testing with domain experts to empirically measure the impact of the XAI dashboard on decision-making speed, confidence and trust.

- **Explore Generalizability** by adapting and testing the real-time architecture in other time critical domains, such as algorithmic trading or emergency response.

6.2 Concluding Remarks

The true purpose of explainability is not merely to verify the predictions of a perfect AI, but to clearly communicate the uncertainty and reasoning of an imperfect one. By focusing on a human centered design and a performance optimized architecture, this work demonstrates that an AI system can move past being a black box predictor to become a transparent assistant in critical decision-making.

Bibliography

[1] Amazon Web Services, "Accelerating the fan experience: How FORMULA 1 is driving the future of racing using machine learning and AWS,"

[2] M.-A. P. F. O. Team, "Feature: Data and electronics in f1 explained," 2024. [Online]. Available: https://www.mercedesamgf1.com/news/feature-data-and-electronics-in-f1-explained.

[3] McLaren, "How ai is revolutionising f1: Presented by dell technologies," Aug. 2024. [Online]. Available: https://www.mclaren.com/racing/partners/dell-technologies/how-ai-is-revolutionising-f1-presented-by-dell-technologies/.

[4] Autosport, "How is artificial intelligence changing formula 1?," Oct. 2024. [Online]. Available: https://www.autosport.com/f1/news/how-is-artificial-intelligence-changing-formula-1/10659532/.

[5] M. Gomez-Tena, "Predictive analytics in formula 1: Revolutionizing the race," *LinkedIn*, Jul. 2024. [Online]. Available: https://www.linkedin.com/pulse/predictive-analytics-formula-1-revolutionizing-race-martin-gomez-tena-tqwaf.

[6] M. T. Ribeiro, S. Singh, and C. Guestrin, ""why should i trust you?": Explaining the predictions of any classifier," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144, 2016. [Online]. Available: https://arxiv.org/abs/1602.04938.

[7] H. Al-Sahaf, L. Busoniu, J. M. de Almeida, and H. J. van den Herik, "A survey of machine learning models for sports outcome prediction," 2019.

[8] D. Thomas et al., "Explainable reinforcement learning for formula one race strategy," Mar. 2025.

[9] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in Neural Information Processing Systems (NIPS)*, vol. 30, 2017. [Online]. Available: https://arxiv.org/abs/1705.07874.

[10] I. Fette and A. Melnikov, "The websocket protocol," no. 6455, Dec. 2011. [Online]. Available: https://www.rfc-editor.org/info/rfc6455.

[11] J. Todd et al., "Explainable time series prediction of tyre energy in formula one race strategy," Mar. 2025.

[12] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design science in information systems research," *MIS quarterly*, pp. 75–105, 2004.