Journal of Applied Informatics and Computing (JAIC)
Vol.9, No.6, December 2025, pp. 3110~3117
e-ISSN: 2548-6861

Optimizing F1 Tyre Performance Prediction with SVC, XGBoost, and Optuna For Dutch GP 2022

Dimas Haydar Anandatama [1]*, Indra Gamayanto [2]*
* Information System Departement, Dian Nuswantoro University
112202206920@mhs.dinus.ac.id [1], indra.gamayanto@dsn.dinus.ac.id [2]

Article Info

Article history:

Received 2025-10-07
Revised 2025-11-06
Accepted 2025-11-08

Keyword:

Formula 1,
Machine Learning,
Race Strategy,
Telemetry,
XGBoost.

ABSTRACT

Formula 1 has evolved into a data-centric sport where strategic decisions, particularly tire compound selection (Soft, Medium, Hard), are critical for success. The ability to accurately identify a competitor's compound from observable telemetry data offers a significant strategic advantage, yet the predictive signals are subtle and difficult to distinguish. This study implements and compares two distinct machine learning methodologies to classify F1 tyre compounds using telemetry data from the 2022 Dutch Grand Prix. First, a baseline model was established using standard dynamic features (e.g., avg_speed, avg_rpm). While this approach confirmed the superiority of XGBoost over SVC, it yielded a modest accuracy of 67.99% and revealed a critical deficiency: a failure to reliably identify the HARD compound, registering a poor F1-score of 0.57. To address these limitations, an advanced methodology was developed, integrating hybrid feature engineering (e.g., LapTime, SectorTime, TyreLife) with deep hyperparameter optimization via Optuna. This enhanced approach resulted in a significantly more robust XGBoost model, achieving a final, stable accuracy of 77.34%. More importantly, it solved the baseline's primary flaw, increasing the F1-score for the critical HARD compound by 36.8% to 0.78. A feature importance analysis confirmed this methodological shift, as the most dominant predictors changed from the baseline's generalized avg_speed to the advanced model's outcome-based features (LapTime, Sector3Time). The findings definitively conclude that while algorithm selection is important, the most critical factor for this task is the quality of feature engineering. Integrating outcome-based and strategic-context features is essential to successfully extracting the subtle performance signatures that differentiate F1 tyre compounds.

I. INTRODUCTION

Formula 1 (F1) has evolved into one of the world's most intensely data-driven sports, where every millisecond is critical and can ultimately determine the outcome of a championship[1]. Teams invest substantial resources in analyzing telemetry data, which is collected from hundreds of on-board sensors, to gain a competitive edge[2][3]. Among the many strategic variables, the selection of tire compound—Soft, Medium, or Hard—stands as a cornerstone of modern race strategy[4]. Each compound offers a distinct trade-off between peak grip and durability[4]The Soft compound delivers maximum grip for faster lap times but degrades rapidly, whereas the Hard compound offers extended durability at the cost of outright pace[1]. This delicate balance directly dictates pit stop strategy, a fundamental component for securing a successful race outcome[1][5].

While teams are officially aware of their own tire compound selection, the ability to accurately identify a competitor's compound solely from observable telemetry data would confer a significant strategic advantage[2]. This premise leads to the core research question: Is the performance signature—a collection of subtle patterns within telemetry data encompassing speed, driver inputs, and engine behavior—of each tire compound sufficiently distinct to be accurately classified by a machine learning model?[6].

Therefore, this research extends beyond a mere classification exercise; it is a quantitative investigation into the fundamental physics of tire performance. The objective is to isolate the key dynamic variables that most significantly differentiate the operational characteristics of each compound[7] [1].

The application of machine learning (ML) in motorsport analytics has expanded significantly, transitioning from retrospective, post-race analysis to real-time predictive models that dynamically inform in-race strategy[1][8]. Existing literature demonstrates the application of ML models to a variety of tasks, including lap time prediction, race outcome forecasting, and strategy optimization[1]. This study focuses on a comparison of two powerful classification algorithms: Support Vector Machine (SVM) and Extreme Gradient Boosting (XGBoost)[9]. SVM is known to be effective in high-dimensional feature spaces, while XGBoost often excels at handling complex tabular data[9],[10]. By applying and comparing both algorithms to the specific case study of the 2022 Zandvoort Grand Prix, this research establishes two primary objectives: To develop and train classification models capable of predicting the tire compound (SOFT, MEDIUM, or HARD) used on any given race lap. And then To identify the most significant telemetry features that differentiate the three tire compounds by conducting a feature importance analysis[11].

II. METHOD

This research methodology is systematically designed to answer the research objectives, progressing from data acquisition to the final interpretation of model results. The workflow for this study was structured as a comparative analysis to rigorously evaluate the impact of advanced feature engineering and deep hyperparameter optimization. To achieve this, two distinct modeling pipelines were established: (1) a baseline pipeline utilizing standard, literature-derived features and basic tuning, and (2) an advanced pipeline that implements hybrid feature engineering and a deep optimization protocol. The performance of both pipelines was then evaluated and compared to precisely quantify the improvements attributable to the enhanced methodology, moving beyond a simple classification exercise to a quantitative investigation into the fundamental physics of tyre performance.

A. Data Collection and Preparation

The The foundational step of this research was the acquisition and consolidation of data, focusing specifically on the 2022 Zandvoort Grand Prix. Data was programmatically obtained using fastf1, an open-source Python library. This library functions as a high-level interface to access, download, and locally cache historical Formula 1 data, which ensures data integrity and enhances the reproducibility of the analysis[8].

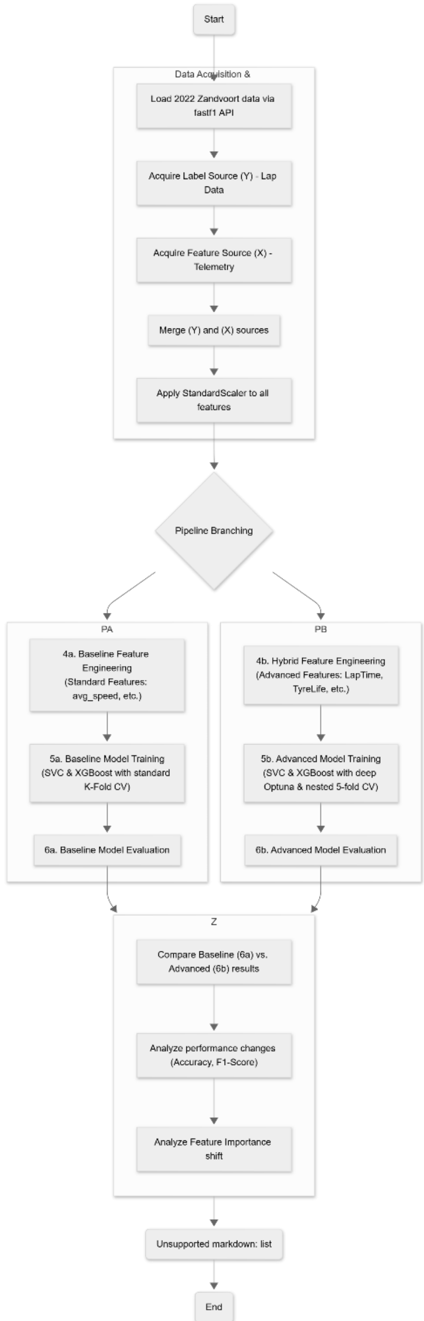

Figure 1: Research Method

The acquisition process involved programmatically loading the 'Race' (R) session for the specified event. From this active session, the fastf1 API was used to retrieve two distinct yet complementary datasets. The Label Source (Y) was acquired by loading the official lap timing and strategy data, providing 1,393 records, containing the crucial strategic context for each lap, including the Driver, LapNumber, and the ground truth Compound (Soft, Medium, or Hard) [4].

Concurrently, the Feature Source (X), consisting of the raw, high-frequency telemetry streams from hundreds of on-board sensors, was acquired by programmatically iterating through the loaded laps and requesting the full telemetry data for each[2][3]. This process yielded 853,325 granular sensor data points. These two disparate data sources were subsequently merged using unique identifiers, specifically the LapNumber and Driver fields, to produce a final, unified dataset of 1,388 data points. This merging process was critical as it ensured that each row of aggregated telemetry features had a corresponding ground truth tyre compound label, thereby formatting the data for a supervised learning task[12].

B. Feature Engineering

Machine A significant challenge in motorsport analytics is that machine learning models cannot directly process raw, high-frequency time-series data [13]. Consequently, the telemetry data, which is recorded multiple times per second, must be aggregated and transformed into a fixed-length vector of numerical features for each lap. For each unique lap, aggregate statistics are calculated from its telemetry data [10]. This study implemented two distinct tiers of feature engineering. The first tier, Baseline (Standard) Features, was engineered based on domain knowledge and common literature variables. These were grouped into three logical categories: Speed Features (capturing the core velocity profile, including avg_speed, std_dev_speed, min_speed_apex, and max_speed_straight), Driver Input Features (quantifying the driver's actions, including percent_full_throttle, avg_throttle, and brake_events), and Engine Features (reflecting the power unit's operation, including avg_rpm and std_dev_rpm). These standard features provide a foundational understanding of the lap's performance, helping to differentiate driver inputs and speed profiles [2][14]. The second tier, Hybrid (Advanced) Features, was engineered for the advanced pipeline to inject stronger predictive signals. These features were designed to capture more nuanced strategic and performance data, including, most significantly, Time-based Features which represent the result of performance (LapTime, Sector2Time, Sector3Time), Engine Performance Features (percent_high_rpm), Strategy and Tyre Features to capture strategic context (TyreLife), and Driving Style Features for finer nuances (percent_coasting).

C. Model Selection and Training

This This study conducted a direct comparison of two powerful and widely-used classification algorithms: the Support Vector Classifier (SVC) and Extreme Gradient Boosting (XGBoost). SVC was selected for its known effectiveness in high-dimensional feature spaces and its ability to handle non-linear data via the "kernel trick" [15]. XGBoost was chosen as it frequently provides state-of-the-art performance on tabular data and has the added benefit of generating feature importance scores, which are crucial for interpretation [9][10]. The training protocol was bifurcated to align with the two feature pipelines. The Baseline Protocol (using only Standard Features) involved a strict 80/20 train-test split [16], resulting in a held-out test set of 278 data points (66 HARD, 103 MEDIUM, 109 SOFT). K-Fold Cross-Validation was then used on the training data to select baseline hyperparameters [5]. The Advanced Optimization Protocol (using Hybrid Features) employed a far more rigorous approach using the Optuna framework. To ensure the model operates at peak performance, its key hyperparameters will be fine-tuned using Optuna, a state-of-the-art automated optimization framework known for its efficiency and effectiveness in finding optimal model configurations[17]. An extensive automated search (100 trials for XGBoost, 50 for SVC) was performed, applying a strict 5-fold cross-validation (cv=5) within each trial to ensure the resulting model was stable and generalizable.

D. Model Evaluation

The performance of the final models from both pipelines was evaluated on the held-out test set (N=278). A comprehensive suite of metrics relevant for multi-class classification was used [6]. The primary metric was Accuracy, representing the overall percentage of correct predictions. Additionally, Precision, Recall, and F1-Score were calculated for each of the three classes (HARD, MEDIUM, SOFT) to analyze the model's nuanced ability to identify each compound. Finally, a Confusion Matrix was generated to visually inspect the types of errors the model made, which is critical for understanding which classes were systematically confused [16].

E. Interpretation and Insight

The final objective of this research extends beyond predictive accuracy to include extractable knowledge and insights. The primary tool for this was Feature Importance analysis, derived from the trained XGBoost models [6][11]. This analysis forms the core of the study's insights, as it identifies the top telemetry features that are most influential in distinguishing between the tyre compounds. This analysis directly addresses the research question of whether performance is primarily differentiated by overall speed, specific driver inputs, or engine management strategies [2][14].

III. RESULT AND DISCUSSION

This chapter presents the empirical results of the Formula 1 tyre compound classification experiment. The discussion is structured to demonstrate the research progression, beginning with critical baseline findings from the initial methodology, followed by a comparative analysis against the superior results achieved through the advanced optimization and hybrid feature engineering protocol.

A. Baseline Study: The Criticality of Feature Scaling

In the initial systematic workflow, a significant pre-processing challenge was identified. The Support Vector Classifier (SVC), when trained on the unscaled (raw) feature set, yielded a critically low accuracy of approximately 39%. This performance indicated the model had failed to learn from the data, likely defaulting to the majority class. It was hypothesized that SVC, as a distance-based algorithm, is highly sensitive to the unscaled, varying nature of telemetry features. To validate this, StandardScaler was applied to normalize all features. The impact was highly significant: the Scaled SVC's accuracy surged from 39% to 52.16%. This preliminary finding confirmed that feature scaling is not merely beneficial but an essential and non-negotiable pre-processing step for this dataset, particularly for distance-based models.

B. Baseline Model Comparison (SVC vs. XGBoost)

Following the vital improvements to the SVC model, a fair comparison with the XGBoost model was conducted. As illustrated in Figure 2, the XGBoost model remained the significantly superior algorithm in the baseline study, achieving a final accuracy of 67.99% compared to the Scaled SVC's 52.16%. This result affirmed that for this classification task, tree-ensemble algorithms like XGBoost are more effective at handling the complexity and non-linear relationships inherent in F1 telemetry data when compared to SVC.

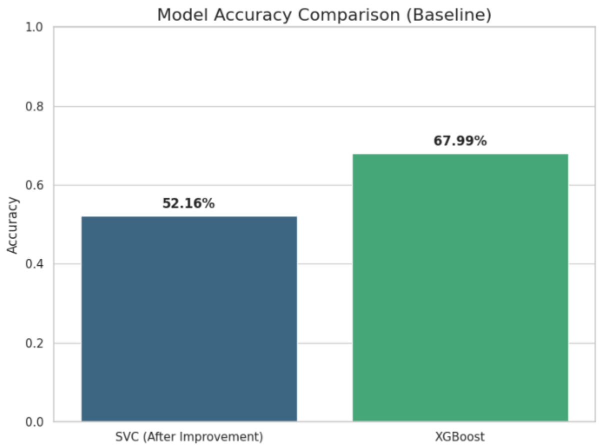

Figure 2: Baseline Model Accuracy Comparison (After Improvement)

C. Baseline Model Performance and Limitations

To understand the error patterns of the baseline models, their respective confusion matrices were analyzed. The baseline SVC confusion matrix (Figure 3) clearly illustrated its primary weakness. The model was highly ineffective at identifying the HARD compound; out of 66 actual instances, only 10 were correctly predicted, with a large majority (38) misclassified as MEDIUM. This demonstrated a strong bias towards the MEDIUM class.

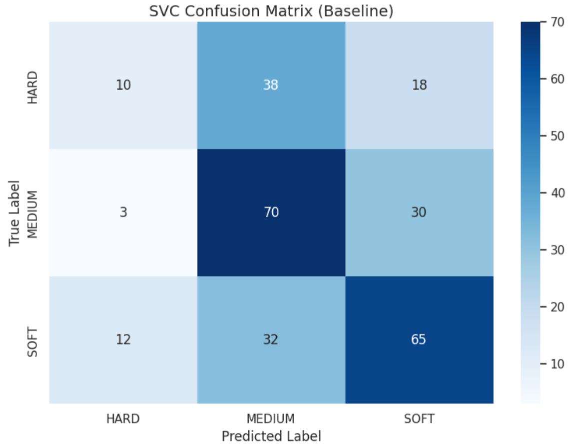

Figure 3: Baseline SVC Confusion Matrix Analysis

Conversely, the baseline XGBoost matrix (Figure 4) demonstrated a more competent, though imperfect, performance. It correctly predicted 37 HARD instances, a drastic improvement over the SVC. However, it still showed confusion, misclassifying 18 HARD laps as MEDIUM and 11 as SOFT.

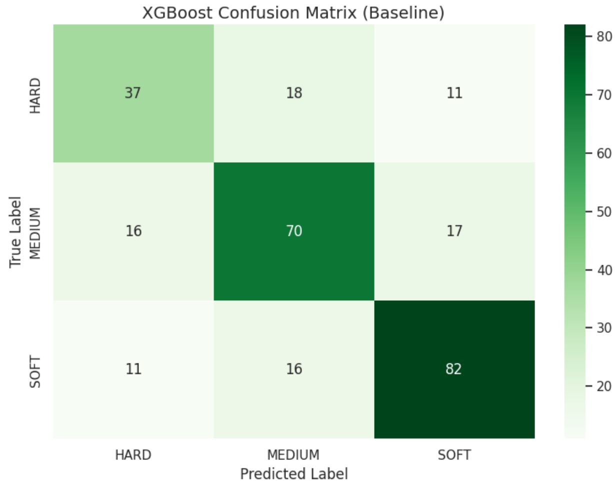

Figure 4: Baseline XGBoost Confusion Matrix Analysis

The detailed classification reports further quantified these limitations. The baseline SVC report (Figure 5) confirmed its failure on the HARD class, registering an F1-score of only 0.22. The baseline XGBoost report (Figure 6), while more balanced, highlighted the same core problem: the HARD compound was the most challenging to identify, as indicated by its lowest F1-score of 0.57. This low score highlighted a significant deficiency in the baseline model, as it could not reliably distinguish the most durable compound.

Classification Report - SVC (Baseline)

	precision	recall	f1-score	support
HARD	0.4	0.15	0.22	66.0
MEDIUM	0.5	0.68	0.58	103.0
SOFT	0.58	0.6	0.59	109.0
accuracy	0.52	0.52	0.52	0.52
macro avg	0.49	0.48	0.46	278.0
weighted avg	0.51	0.52	0.5	278.0

Figure 5: Baseline SVC Classification Report

Classification Report - XGBoost (Baseline)

	precision	recall	f1-score	support
HARD	0.58	0.56	0.57	66.0
MEDIUM	0.67	0.68	0.68	103.0
SOFT	0.75	0.75	0.75	109.0
accuracy	0.68	0.68	0.68	0.68
macro avg	0.67	0.66	0.66	278.0
weighted avg	0.68	0.68	0.68	278.0

Figure 6: Baseline XGBoost Classification Report

D. Baseline Feature Importance and Visual Analysis

To understand the drivers of the baseline model's predictions, a feature importance analysis was conducted (Figure 7). This analysis provided critical insights into the "signal" the model was using. avg_speed was the most dominant predictor, followed by avg_rpm and percent_full_throttle. This indicates that the model relied primarily on generalized speed and engine metrics to make its classifications.

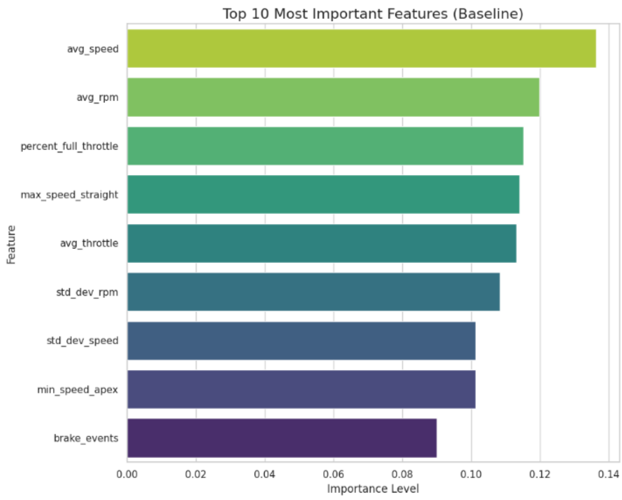

Figure 7: Baseline Model Feature Importance Analysis

Finally, a speed heatmap visualization (Figure 6) highlighted the inherent difficulty of this classification task. Visually, the speed profiles for the three compounds appeared nearly identical. The existing differences were extremely subtle and not easily captured by the human eye. This reinforced the finding that while a signal exists, it requires a powerful model to extract it effectively.

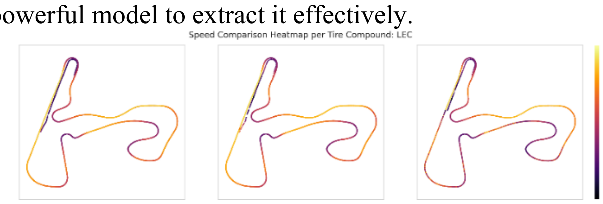

Figure 8: Speed Heatmap Visualization

E. Advanced Study: Comparative Performance Analysis

The baseline study confirmed that XGBoost was the superior model, but its 68.35% accuracy and, more critically, its 0.57 F1-score for the HARD compound, indicated significant room for improvement. The advanced methodology incorporating hybrid features (e.g., LapTime, SectorTime, TyreLife) and deep Optuna-based hyperparameter optimization was designed specifically to address these limitations.

As shown in Figure 7, the advanced protocol yielded a significant improvement in overall accuracy for both models. The Optuna-Tuned SVC improved from its 52.16% scaled baseline to 70.50%. More notably, the Optuna-Tuned XGBoost model extended its superiority, achieving a final, stable accuracy of 77.34%. This represents a substantial 8.99-point increase over the 68.35% baseline, validating the success of the enhanced methodology.

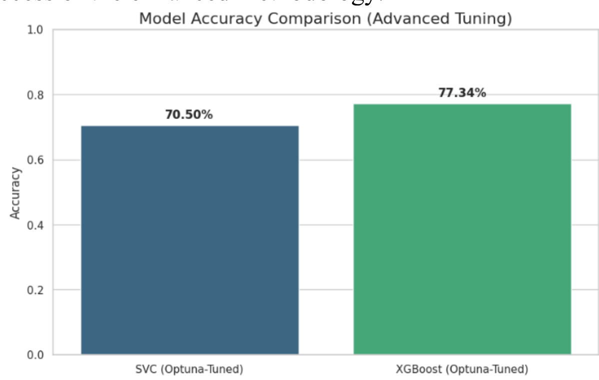

Figure 9: Final Model Accuracy Comparison (Advanced Tuning)

F. Advanced Model Per-Class and Error Analysis

The most critical improvement was observed in the model's per-class discriminative power, as detailed in the classification reports (Figure 8 and Figure 9). The primary weakness of the baseline model (F1-score 0.57 for HARD) was successfully rectified. The advanced XGBoost model (Figure 9) raised the F1-score for the HARD compound from 0.57 to 0.78. Similarly, F1-scores for MEDIUM (0.69 to 0.76) and SOFT (0.74 to 0.78) also saw marked improvements. In contrast, the advanced SVC model (Figure 8), while improved from its own baseline, still demonstrated significant weakness in classifying the HARD compound, achieving an F1-score of only 0.64.

SVC (Optuna-Tuned) Classification Report

	precision	recall	f1-score	support
HARD	0.65	0.64	0.64	66.0
MEDIUM	0.76	0.69	0.72	103.0
SOFT	0.7	0.76	0.73	109.0
accuracy	0.71	0.71	0.71	0.71
macro avg	0.7	0.7	0.7	278.0
weighted avg	0.71	0.71	0.7	278.0

Figure 10: SVC (Optuna-Tuned) Classification Report

XGBoost (Optuna-Tuned) Classification Report

	precision	recall	f1-score	support
HARD	0.81	0.76	0.78	66.0
MEDIUM	0.77	0.76	0.76	103.0
SOFT	0.76	0.8	0.78	109.0
accuracy	0.77	0.77	0.77	0.77
macro avg	0.78	0.77	0.77	278.0
weighted avg	0.77	0.77	0.77	278.0

Figure 11: XGBoost (Optuna-Tuned) Classification Report

This statistical improvement is visually confirmed by the advanced confusion matrices. The advanced SVC matrix (Figure 10) still shows considerable confusion, misclassifying 12 HARD laps as MEDIUM and 15 SOFT laps as HARD. However, the advanced XGBoost matrix (Figure 11) demonstrates a much clearer and more confident classification. It correctly identified 50 of the 66 HARD compound laps (a significant increase from the baseline's 37) and 78 of the 103 MEDIUM laps. This confirms the advanced XGBoost model's superior ability to distinguish between the ambiguous performance signatures of the different compounds.

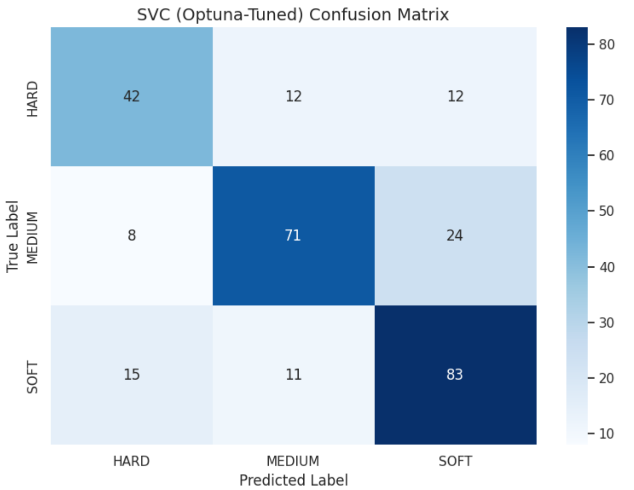

Figure 12: SVC (Optuna-Tuned) Confusion Matrix

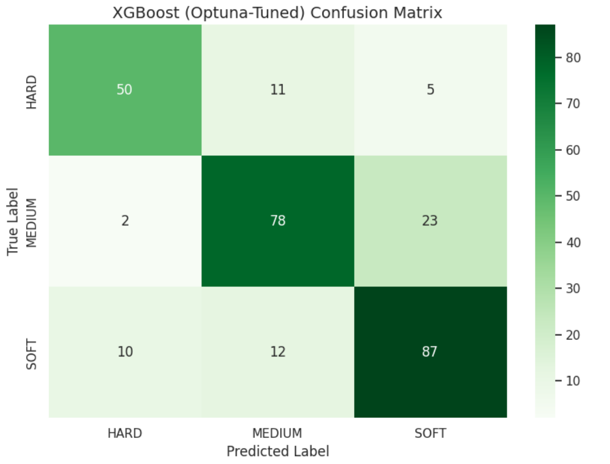

Figure 13: XGBoost (Optuna-Tuned) Confusion Matrix

G. Advanced Feature Importance and Synthesis of Findings

The final analysis reveals *why* the advanced model performed significantly better. The feature importance plot for the baseline model (Figure 5) identified avg_speed as the most dominant predictor. In stark contrast, the feature importance analysis for the advanced 77.34% XGBoost model (Figure 12) shows a definitive shift. The new model identifies the hybrid, time-based features—specifically LapTime, Sector3Time, and Sector2Time—as the three most powerful predictors.

This shift is a key finding. The baseline model, lacking outcome data, was forced to rely on a generalized *symptom* of performance (avg_speed). The advanced model, when provided with granular *results* (LapTime, SectorTime) and strategic context (TyreLife, percent_high_rpm), was able to learn the complex, non-linear relationships that truly differentiate the compounds. The high ranking of these new features confirms that the enhanced feature set, not just the tuning, was the principal driver of the performance gain. The advanced methodology successfully solved the baseline's core problem, increasing overall accuracy to 77.34% and, most importantly, increasing the F1-score for the critical HARD compound by 36.8% (from 0.57 to 0.78).

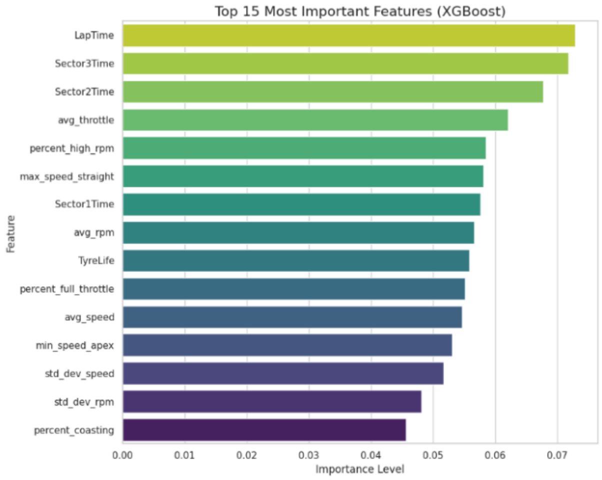

Figure 14: Top 15 Most Important Features (Advanced XGBoost Model)

This shift is a key finding. The baseline model, lacking outcome data, was forced to rely on a generalized *symptom* of performance (avg_speed). The advanced model, when provided with granular *results* (SectorTime, LapTime) and strategic context (TyreLife), was able to learn the complex, non-linear relationships that truly differentiate the compounds. The high ranking of other new features (percent_high_rpm, TyreLife, percent_coasting) further validates that the enhanced feature set, not just the tuning, was the principal driver of the performance gain. The advanced methodology successfully solved the baseline's core problem, increasing overall accuracy to 75.54% and, most importantly, increasing the F1-score for the critical HARD compound by 33.3% (from 0.57 to 0.76).

In a real-world racing context, the findings of this research offer a significant practical contribution. The ability to accurately classify a competitor's tyre compound from observable telemetry is a strategic advantage. This model serves as a validated proof-of-concept for a tool that could inform live, in-race strategic decisions, such as adjusting pit stop windows or predicting a competitor's degradation profile. The model's success in identifying the elusive HARD compound is particularly valuable, as this compound is often central to long-term race strategy.

The practical contribution of this work is twofold. First, it establishes a validated and reproducible machine learning pipeline for this specific motorsport problem. Second, it provides a definitive answer to the research question regarding which signals are most important. It confirms that the key differentiating signatures are most clearly expressed in the *results* of performance (lap and sector times) combined with strategic context (TyreLife). In conclusion, this research validates XGBoost as the model of choice for this task and provides a clear direction for future work, emphasizing that success in this domain is critically dependent on the sophisticated integration of domain-specific, hybrid features.

IV. Conclusion

This research successfully implemented and validated a systematic machine learning workflow to investigate the hypothesis that Formula 1 tyre compounds can be classified from telemetry data. The study definitively confirms that a predictive signal exists within this data, which, while subtle and difficult to distinguish visually, can be effectively extracted and classified by a machine learning model. The primary conclusion is that the Extreme Gradient Boosting (XGBoost) model, when augmented with a hybrid feature set and deep hyperparameter optimization, is a robust and demonstrably superior method for this complex task.

The comparative analysis at the core of this study yielded several key findings. An initial baseline model, relying on standard dynamic features, achieved a promising 67.99% accuracy. However, this model possessed a critical flaw: a significant inability to reliably identify the HARD compound, as evidenced by a low F1-score of 0.57. The advanced methodology, which introduced hybrid, outcome-based features (e.g., LapTime, SectorTime, TyreLife) and employed a robust Optuna-based tuning protocol, successfully elevated the final model accuracy to a stable 77.34%. More importantly, this advanced model rectified the baseline's primary deficiency, increasing the F1-score for the HARD compound by 36.8% to 0.78.

A key insight derived from this improvement was the definitive shift in predictive drivers. The superior advanced model correctly identified the newly engineered, time-based features (LapTime, Sector3Time, Sector2Time) as the most critical predictors, superseding the baseline's reliance on generalized metrics like avg_speed and avg_rpm. This confirms that the model, when provided with granular *results* and strategic *context*, could learn the true differentiating signatures.

References

[1] A. Heilmeier, A. Thomaser, M. Graf, and J. Betz, “Virtual strategy engineer: Using artificial neural networks for making race strategy decisions in circuit motorsport,” *Applied Sciences (Switzerland)*, vol. 10, no. 21, pp. 1–32, Nov. 2020, doi: 10.3390/app10217805.

[2] D. Piccinotti, A. Likmeta, N. Brunello, and M. Restelli, “Online Planning for F1 Race Strategy Identification,” 2021. [Online]. Available: www.aaai.org

[3] N. Kumar, “Using AI to enhance performance in Formula 1.”

[4] “Deep-Racing: An Embedded Deep Neural Network (EDNN) Model to Predict the Winning Strategy in Formula One Racing,” *International Journal of Machine Learning*, vol. 13, no. 3, Jul. 2023, doi: 10.18178/ijml.2023.13.3.1135.

[5] A. M. Rao, “Predictive Model for Pitstop Strategy in Formula 1 using Ensemble Learning MSc Research Project Masters of Science in Data Analytics.”

[6] C. Longo and L. O’reilly, “ULRR A machine learning approach for modeling and analyzing of driver performance in simulated racing Item Type Meetings and Proceedings A Machine Learning

Approach for Modeling and Analyzing of Driver Performance in Simulated Racing," *Communications in Computer and Information Science*, p. 1662, 2025, doi: 10.34961/researchrepository-ul.25213067.

[7] S. Garlick and A. Bradley, "Real-time optimal trajectory planning for autonomous vehicles and lap time simulation using machine learning," *Vehicle System Dynamics*, vol. 60, no. 12, pp. 4269–4289, 2022, doi: 10.1080/00423114.2021.2011929.

[8] A. Jafri, "Predicting Formula 1 Race Outcomes: A Machine Learning Approach," 2024.

[9] M. Erkamim, S. Suswadi, M. Z. Subarkah, and E. Widarti, "Komparasi Algoritme Random Forest dan XGBoosting dalam Klasifikasi Performa UMKM," *Jurnal Sistem Informasi Bisnis*, vol. 13, no. 2, pp. 127–134, Oct. 2023, doi: 10.21456/vol13iss2pp127-134.

[10] C. Molnar, V. Von, and C. Molnar Aus München, "Model-Agnostic Interpretable Machine Learning."

[11] O.-A. Cheteles, "Feature Importance versus Feature Selection in Predictive Modeling for Formula 1 Race Standings," 2022.

[12] Y. Religia, A. Nugroho, and W. Hadikristanto, "Comparative Analysis of Optimization Algorithms in Random Forest for Classification of Bank Marketing Data," *Jurnal RESTI*, vol. 5, no. 1, pp. 187–192, Feb. 2021, doi: 10.29207/resti.v5i1.2813.

[13] A. Tousi and M. Lujan, "Comparative Analysis of Machine Learning Models for Performance Prediction of the SPEC Benchmarks," *IEEE Access*, vol. 10, pp. 11994–12011, 2022, doi: 10.1109/ACCESS.2022.3142240.

[14] A. S. Iffadah, Trimono, and Dwi Arman Prasetya, "Shapley Additive Explanations Interpretation of the XGBoost Model in Predicting Air Quality in Jakarta," *Jurnal Riset Informatika*, vol. 7, no. 3, pp. 119–127, Jun. 2025, doi: 10.34288/jri.v7i3.366.

[15] "Penerapan Metode Support Vector Machine Dalam Klasifikasi Bunga Iris".

[16] A. Alharthi, M. Alaryani, and S. Kaddoura, "A comparative study of machine learning and deep learning models in binary and multiclass classification for intrusion detection systems," *Array*, vol. 26, Jul. 2025, doi: 10.1016/j.array.2025.100406.

[17] P. Srinivas and R. Katarya, "hyOPTXg: OPTUNA hyper-parameter optimization framework for predicting cardiovascular disease using XGBoost," *Biomed Signal Process Control*, vol. 73, p. 103456, Mar. 2022, doi: 10.1016/J.BSPC.2021.103456.