A simulation framework for Formula 1 race strategy based on pit-stop optimization

Pasquale Avella [a], Mario Borruso [a], Mattia Marino [a], Adriano Masone [b]

[a] *Dept. of Engineering, University of Sannio, BN, 82100, Italy*
[b] *Dept. of Electrical Engineering and Information Technology, University of Naples Federico II, NA, 80125, Italy*

Abstract

In modern Formula 1, strict regulations and highly optimized cars limit performance gains through hardware, increasing the importance of strategic decision-making. This work tackles the problem of computing a race strategy that minimizes total race time by jointly optimizing tire stints, compound selection, fuel load, and Energy Recovery System (ERS) deployment. We present a high-performance simulation framework based on the solution of an optimization model, designed for fast and reliable trackside use. The system considers discrete stint allocation with ERS management and includes real-time visualization tools for drivers and race engineers. Validation uses semi-real data from a Formula 1 simulator and is further refined with a professional simulation platform. Benchmarking shows improved fidelity and performance compared to existing models.

Keywords: Simulation, Optimization, Motorsport, Race Strategy, Pit-Stop

1. Introduction

Competitive performance in motorsport results from a delicate interplay among mechanical design, driver skill, and strategic decision-making. While historically a substantial portion of competitive advantage could be gained

Email addresses: avella@unisannio.it (Pasquale Avella), m.borruso@studenti.unisannio.it (Mario Borruso), m.marino8@studenti.unisannio.it (Mattia Marino), adriano.masone@unina.it (Adriano Masone)

through engineering innovation, modern regulations have significantly restricted the scope for technical differentiation, particularly in top-tier series such as Formula 1 (F1) [1]. The regulatory framework imposed by the Fédération Internationale de l'Automobile (FIA) [2] aims to ensure safety and fairness. Still, it also places stringent limits on vehicle development, making strategic excellence an increasingly decisive factor in race outcomes.

Within this context, a race strategy can be defined as a comprehensive and adaptive plan mainly designed to minimize total race time. This plan encompasses a range of decisions made prior to the race start and refined throughout the event. In practice, these decisions include the selection and sequencing of tire compounds, the timing of pit stops, the initial fuel load, and the management of hybrid energy systems. Although modern teams employ sophisticated software tools and simulation environments, constructing a race strategy remains a highly complex problem due to the interactions between vehicle dynamics, tire behaviour, and energy management.

A central source of complexity stems from the tire performance. Each tire compound exhibits a specific degradation profile, with lap time depending not only on the compound's intrinsic characteristics but also on how the tire has been managed throughout the stint. The interplay between degradation and energy deployment further complicates the optimization landscape. For example, aggressive use of the Energy Recovery System (ERS) may yield immediate time gains but can impose additional thermal and mechanical load on the tires, accelerating wear and affecting performance on subsequent laps. As a result, strategic and energetic decisions cannot be treated independently: the allocation of tire stints and the intra-lap energy deployment must be modeled as coupled components of a single decision space.

The determination of a race strategy therefore requires the integration of multiple interconnected decisions, including pit stop timing, compound selection, race pace management, and ERS deployment. These decisions must be evaluated over a fixed race distance and within a competitive environment where small differences in lap time accumulate into significant performance differentials. Despite the complexity of the problem, reliable solutions must be obtained rapidly as strategic software is often used in real-time during race events, especially in top series such as F1 where live telemetry is available to the team. This work addresses this challenge by proposing a simulation framework based on the solution of a mathematical model model to determine an optimal race strategy.

1.1. Literature Review

The optimization of race strategy in Formula 1 and related motorsport categories has been approached through a variety of computational and mathematical frameworks. Existing research spans from rigorous dynamic programming formulations to high-fidelity simulation environments, each offering different strengths with respect to optimality guarantees, computational tractability, and physical realism.

We now organize the principal contributions in the literature according to the class of mathematical or algorithmic model employed.

- **Dynamic Programming.** A foundational and widely cited contribution in this domain is the work of Carrasco Heine et al. [3], which explicitly addresses the pit stop strategy problem by optimizing both stopping laps and tire compound selection. Their approach formulates the problem as a Dynamic Programming (DP) model that minimizes total race time under deterministic conditions. The DP captures essential features of the strategy problem, including tire degradation, compound-dependent performance, and the inherently sequential structure of stint decisions.

  The authors further extend the framework to a Stochastic Dynamic Programming (SDP) formulation that accounts for uncertainties such as variable weather and Safety Car interventions. This enables decisions to be evaluated within a probabilistic race evolution, rather than a single deterministic trajectory. The principal contribution of these works is the establishment of a mathematically principled method for handling the trade-off between delaying pit stops and exploiting potential stochastic advantages, offering strong theoretical guarantees in sequential decision-making contexts.

- **Neural Networks.** A prominent contribution in this category is the Virtual Strategy Engineer (VSE) introduced by Heilmeier et al. [4]. While the authors show that, under highly simplified assumptions (no opponents and tire degradation only), the optimal strategy can be computed via a quadratic optimization problem, they argue that such formulations are inadequate for professional motorsport applications. In high-level series like Formula 1, realistic strategy evaluation requires comprehensive race simulations capable of modeling dynamic interactions among competitors.

To meet this need, the VSE employs two artificial neural networks that automatically determine whether a driver should pit and which tire compound to select. Trained on timing data from six Formula 1 seasons (2014–2019), the networks infer strategic decisions directly from real race conditions. Integration into a full simulation environment demonstrates that the VSE produces context-aware and reasonable decisions, thereby improving automation and realism in simulation-based strategy assessment.

More recently, Więckowski et al. [5] propose an Artificial Neural Network Decision Support System (ANN-DSS) based on a Multilayer Perceptron architecture. Hyperparameters are optimized via a Tree-structured Parzen Estimator, and the resulting model is used to predict the best lap times of Formula 1 drivers during a race. The predicted rankings exhibit high consistency with reference rankings according to the WS rank similarity metric. Although this approach does not formulate or solve a race strategy problem directly, it highlights the expanding role of machine learning techniques in motorsport decision-support and provides a foundation for integrating data-driven components into broader strategic systems.

- **Simulation-Based Optimization.** A significant portion of the literature employs simulation-driven approaches, in which predefined or heuristically generated strategies are evaluated using detailed race simulators [6, 7]. In these models, strategy is typically an input rather than the output of an optimization procedure. While such methods do not provide optimality guarantees, they excel at capturing high-fidelity physical dynamics, including nonlinear tire wear, aerodynamic effects, and operational uncertainties such as pit crew timing. Their principal contribution lies in offering robust validation environments for mathematically derived strategies, enabling assessment under near-realistic race conditions and unmodeled dynamic interactions.

1.2. Our Contribution

In light of the diverse computational approaches developed for motorsport decision-making, our work advances the state of the art by introducing a mathematically rigorous, optimization-driven framework based on Mixed-Integer Linear Programming (MILP) for the formulation of modern

Formula 1 race strategies. Specifically, we provide a generalizable optimization model that delivers demonstrable improvements in strategic consistency, computational efficiency, and integration with realistic race dynamics.

Our main contributions are as follows:

- Acquisition of context-specific race data: For each Grand Prix, we construct strategy inputs using real-world performance data sourced from the available Free Practice sessions. The data acquisition process leverages physics-based simulators such as F1 25 [8] and Assetto Corsa [9] enriched with high-fidelity vehicle models such as the VRC Formula Alpha 2025 [10]. This enables the model to capture track-dependent characteristics and contemporary vehicle behavior.

- Integration with realistic race simulation: Unlike models limited to analytical evaluation or simplified abstractions, our framework allows candidate strategies to be validated directly within the aforementioned simulation environments. This provides an empirical loop between optimization output and race-dynamic evaluation, and our results demonstrate that the optimized strategies exhibit high fidelity and practical viability when tested under realistic conditions.

- Explicit ERS optimization: In contrast to all prior optimization-based approaches, we incorporate the Energy Recovery System (ERS) as a decision variable within the race-strategy formulation. The model jointly optimizes ERS deployment and pit-stop scheduling, reflecting the central role of hybrid powertrain management in contemporary Formula 1 competition.

- A computationally efficient method with sub-second runtimes: The proposed MILP formulation achieves solution times on the order of tenths of a second while maintaining high accuracy and robustness. This represents a substantial improvement over existing optimization-based models, which typically require several seconds to compute a feasible strategy, thereby enabling real-time or near–real time strategic evaluation.

These contributions collectively establish a comprehensive simulation framework that bridges the gap between theoretical optimality and practical applicability. The system is designed to operate seamlessly within the time-critical environment of professional motorsport, where strategic decisions must be computed rapidly and executed precisely.

The remainder of this work is organized as follows: Section 2 describes the problem scope, explaining Formula 1 race operations in accessible terms for non-expert readers, and detailing the regulatory constraints that govern strategic decision-making. Section 3 presents the simulation framework that surrounds the optimization core. Section 4 presents the mathematical formulation of the optimization problem. Section 5 reports the empirical validation of the system through a multi-phase testing protocol. Initial validation is conducted using the F1 2025 commercial simulator to establish baseline performance against industry-standard AI strategies. Advanced validation employs professional-grade simulation platforms (Assetto Corsa with VRC Formula Alpha 2025) to assess robustness under high-fidelity physics modeling. Finally, Section 6 synthesizes the findings and outlines promising directions for future research.

2. Problem Description

This section introduces the operational structure of a Formula 1 Grand Prix and the domain-specific concepts required to understand the nature of the problem. It then presents the regulatory constraints that define the boundaries within which race strategy is performed and formalizes the problem scope and assumptions.

2.1. Overview of a Formula 1 Grand Prix

A Formula 1 Grand Prix is a circuit-based motorsport event in which a fixed number of laps is completed, such that the total race distance is approximately 305 km, with minor exceptions for specific venues. All drivers start simultaneously from a standing grid, whose order is determined by a qualifying session held prior to the race. The primary objective is to complete the prescribed race distance in the shortest possible time.

During the race, each car's performance is governed by a combination of mechanical, aerodynamic, and operational factors. While the car specification is fixed for the event, the lap time evolves as a function of fuel mass, tire condition, energy deployment, and strategic decisions such as pit stops. A race strategy therefore consists of a sequence of decisions made before and during the event that aim to minimize total race time while respecting sporting and technical regulations.

From a strategic perspective, the race can be decomposed into consecutive *stints*. A stint is defined as a continuous sequence of laps completed using the same set of tires, delimited by pit stops. At each pit stop, the car enters the pit lane, where tires may be replaced and minor adjustments performed, incurring a significant time loss relative to remaining on track. Consequently, race strategy fundamentally involves a trade-off between longer stints with degraded performance due to tire wear and shorter stints that incur additional pit stop time penalties.

2.2. Key Operational Concepts and Terminology

Several domain-specific concepts are central to the formulation of the strategy optimization problem:

- Tire Compounds: In dry conditions, teams may select from multiple slick tire compounds provided by the supplier, typically classified as soft, medium, and hard. Softer compounds offer higher grip and faster lap times but degrade more rapidly, while harder compounds are more durable at the cost of reduced peak performance.

- Tire Degradation: Tire performance deteriorates as a function of usage due to thermal and mechanical wear. This degradation manifests as an increasing lap time over the duration of a stint and is one of the dominant factors influencing optimal stint length.

- Fuel Load: Cars start the race with all the fuel required to reach the finish, as refuelling during the race is prohibited. The fuel mass decreases over time, reducing the car's weight and generally improving lap times as the race progresses.

- ERS: Modern Formula 1 cars are equipped with hybrid power units that recover and deploy electrical energy. The amount of energy that can be harvested and deployed per lap is regulated, and its management affects lap time and overall performance.

- Pit Stop Time Loss: A pit stop introduces a fixed and measurable time penalty, consisting of pit lane entry, stationary service time, and pit lane exit. This loss is circuit-dependent and is a critical parameter in evaluating whether an additional stop is strategically advantageous.

The interaction between these elements gives rise to a highly structured optimization problem, where decisions made early in the race influence performance and feasibility in later stages.

2.3. Regulatory Constraints and FIA Compliance

The system considers the relevant FIA Formula 1 Sporting and Technical Regulations to ensure that every generated strategy is legally admissible. These regulations define strict boundaries on vehicle operation and race conduct, and they are encoded in the model as hard constraints.

Key regulatory elements included in the formulation are:

- Mandatory Tire Compound Change: In dry races, each driver must use at least two different slick tire compounds during the race, enforcing the need for at least one pit stop.
- Fuel Constraints: The initial fuel load must respect the maximum fuel mass allowed at the start of the race and must be sufficient to complete the race while leaving a mandatory residual quantity for post-race inspection.
- ERS and Power Unit Limits: Energy harvesting and deployment per lap are constrained by regulatory limits, which bound the performance gains achievable through hybrid system usage.
- Stint Feasibility and Pit Procedures: Stint lengths must remain within compound-specific operational limits, and pit stops must comply with procedural constraints and minimum safety requirements.

By embedding these regulatory requirements directly into the optimization problem, the system guarantees that the resulting strategy is not only optimal from a performance standpoint but also fully compliant and directly implementable in a competitive race environment.

2.4. Problem Scope and Design Assumptions

The purpose of the system is to compute the fastest a priori race strategy for a standard, uninterrupted dry-weather Grand Prix. The strategy is defined as the optimal sequence of tire compound selections and corresponding stint lengths that minimizes total race time.

To ensure tractability and interpretability, the model deliberately excludes stochastic and externally induced race events such as rainfall, yellow flags, Virtual Safety Cars, or full Safety Car deployments. While such events can have a substantial impact on real-world strategy, they are inherently unpredictable and require real-time human judgment. By excluding them, the optimization isolates the deterministic baseline strategy, which represents the reference plan under nominal conditions.

The optimizer therefore provides a theoretically grounded benchmark strategy. In practical applications, this strategy may be followed directly or adapted by race engineers in response to unforeseen developments during the race.

3. Simulation Framework

This section describes the simulation framework that enables the optimization-driven race strategy planning, focusing on the data acquisition, simulation validation, and decision support components that surround the mathematical core.

The simulation framework encompasses:

- The acquisition and processing of empirical performance data from practice sessions and simulation platforms
- The configuration and parameterization of the optimization inputs
- The validation of computed strategies through high-fidelity race simulation
- The presentation of strategic outputs through user-facing interfaces

At a high level, the system operates in three stages: first, it collects and structures data from practice sessions and simulation platforms; second, it invokes the mathematical optimization engine to compute the fastest strategy; and third, it validates the computed strategy through execution in realistic simulation environments and presents the results through interfaces designed for both race engineers and drivers. The following subsections detail each of these components, with particular emphasis on the simulation-based data acquisition and strategy validation processes.

3.1. Input Data Acquisition and Configuration

The starting point of the optimization process is the acquisition and processing of empirical performance data. In a real-world Formula 1 environment, this data would be gathered during the three Free Practice sessions of the race weekend, which collectively provide approximately three hours of on-track running. These sessions offer essential observations for calibrating vehicle behaviour, including tire degradation curves, fuel consumption rates, and energy recovery characteristics specific to the circuit and prevailing conditions.

For research and development purposes, and in the absence of access to actual Formula 1 telemetry, this work leverages high-fidelity racing simulation platforms to generate the necessary input data. Two primary simulation environments are employed:

- F1 25 (Codemasters) [8]: The official Formula 1 racing simulator, providing FIA-licensed circuits with standardized vehicle models and accessible telemetry extraction capabilities. While employing a simplified physics model compared to professional tools, it offers consistent and reproducible data collection environments suitable for initial validation.

- Assetto Corsa with VRC Formula Alpha 2025 [9]: A professional-grade simulation platform combined with a highly realistic Formula 1 vehicle modification. This environment provides advanced tire thermal modeling, computational fluid dynamics calculations, and detailed hybrid powertrain simulation. The physics fidelity of this platform closely approximates real-world vehicle behavior, making it particularly valuable for final validation testing.

These simulation platforms serve dual purposes within the framework: first, they act as *data sources* for parameterizing the optimization model through simulated practice sessions; second, they function as *validation environments* where computed optimal strategies are executed and empirically evaluated under controlled yet realistic racing conditions. This closed-loop approach enables iterative refinement of model parameters and strategy validation without requiring prohibitively expensive real-world testing.

All derived parameters are supplied to the optimizer through a structured configuration file (e.g., `config.json`), which serves as the central data source for the entire system. This file contains the following domains of information:

- Race Parameters: Total race distance (in laps), typical pit stop time loss, and time penalties associated with standing starts.
- Tire Parameters: Base lap times for each compound, degradation rates (seconds of time loss per lap), and the maximum number of laps permitted or advisable per tire type.
- Fuel Parameters: Fuel consumption per lap and the performance penalty per kilogram of fuel carried.
- ERS Parameters: Deployment and harvesting limits per lap, battery capacity, and the time deltas associated with each ERS mode on each tire compound.
- Circuit Parameters: A sector-level decomposition of the track, including the number of corners per sector (important for harvesting modelling) and a coefficient describing harvesting potential within that sector.

This structured data foundation ensures that the optimization problem is fully informed by the physical, mechanical, and strategic realities of the specific race weekend.

3.2. Optimization Engine Integration

At the core of the framework lies a Mixed-Integer Linear Programming (MILP) formulation implemented in Python and solved via the high-performance commercial solver FICO Xpress. The mathematical details of this optimization model are presented comprehensively in Section 4. This section focuses on how the optimization engine integrates within the broader simulation framework.

The optimizer receives the structured configuration file as input and produces a complete race strategy as output. The objective function minimizes total race time subject to structural, mechanical, and regulatory constraints that ensure compliance with FIA regulations and physical feasibility.

A defining contribution of this work is the integrated management of the Energy Recovery System (ERS). Whereas most existing strategic models aggregate ERS effects into coarse lap-level adjustments, the proposed formulation computes a precise and prescriptive energy deployment and harvesting schedule at the sector level.

The optimization output includes:

- Race Strategy: The optimal sequence of tire compounds and stint durations, including exact pit stop laps.
- Fuel Load: The optimal starting fuel mass that balances race-time minimization with regulatory requirements and performance considerations.
- Per-Sector ERS Plan: A detailed schedule specifying the exact energy to deploy or harvest in each sector of each lap, ensuring adherence to battery constraints and maximizing performance over the full race distance.

To support operational robustness, the optimization enforces a minimum battery safety margin (e.g., 40%) throughout most of the race. This reserve ensures that the driver retains enough electrical energy for defensive manoeuvres even when the model's projections are disrupted by competition dynamics. In the final laps, this restriction is lifted, allowing full battery depletion to maximize pace.

A graphical representation of the target battery state of charge per sector is also produced, facilitating race execution and post-analysis.

3.3. System Interfaces and Decision Support

The outputs of the optimization model are consumed by two complementary user-facing applications, each designed for a specific operational role within the race environment:

- Race Engineer Dashboard: A comprehensive telemetry and strategy interface intended for the pit wall. It displays real-time lap and sector times, vehicle parameters (e.g., setup, temperatures), tire and fuel status, battery charge, and G-force readings. The optimized race strategy and ERS plan are integrated directly into the dashboard, allowing seamless reference during race execution.
- Driver-Facing Android Application: A streamlined mobile interface designed for integration into the multi-page digital displays of modern Formula 1 vehicles. It provides the driver with essential strategic information, including fuel targets, tire management indicators, and sector-level ERS deployment and harvesting instructions. The application may function as a standalone strategy page or be embedded within an existing dashboard page.

These interfaces ensure that the optimization outputs are not only mathematically optimal but also actionable and accessible to the personnel responsible for executing the strategy in real time.

4. Optimization Model

This section presents the mathematical core of the simulation framework. While the previous section described *what* the system does and *how* data are acquired and managed, this section formalizes *how the optimization problem is structured* and solved.

The race strategy problem can be understood as finding the best answers to several interconnected questions:

- Which tire compounds should be used?
- When should pit stops occur?
- How should the limited electrical energy be deployed throughout the race?

These decisions are represented mathematically using variables, and the relationships between them (such as tire wear increasing lap time, or energy usage being limited by battery capacity) are expressed as equations and inequalities.

The resulting optimization model is solved using specialized software that explores millions of potential strategies and identifies the one that minimizes total race time. The formulation belongs to a class of problems known as Mixed-Integer Linear Programming (MILP), chosen for its balance between expressive power and computational efficiency.

The remainder of this section provides a rigorous mathematical specification of the model, organized into definitions of sets and indices, input parameters, decision variables, the objective function, and the constraints that ensure the solution is both feasible and regulation-compliant.

4.1. Sets

- $\mathcal{G}$: Set of tire compounds available, indexed by $g \in \{0, 1, \ldots, |\mathcal{G}| - 1\}$.
- $\mathcal{S}$: Set of stints, indexed by $s \in \{0, 1, \ldots, S_{\max} - 1\}$ where $S_{\max}$ is the maximum number of stints.

- $\mathcal{L}$: Set of laps, indexed by $\ell \in \{0, 1, \ldots, L-1\}$ where L is the total number of race laps.
- $\mathcal{K}$: Set of circuit sectors per lap, indexed by $k \in \{0, 1, \ldots, K-1\}$ where K is the number of sectors.
- $\mathcal{M}$: Set of ERS modes, indexed by $m \in \{0, 1, \ldots, M-1\}$ where M is the number of available ERS modes.

4.2. Parameters

1. Tire Parameters
    - For each tire compound $g \in \mathcal{G}$:
        - t_g^{base}: Base lap time (seconds) for tire compound g.
        - $d_g^{\max}$: Maximum stint duration (laps) for tire compound g.
        - δ_g^{deg}: Degradation rate per lap (unitless) for tire compound g.
        - μ_g^{perf}: Performance degradation impact factor for tire g
        - μ_g^{dur}: Duration reduction factor due to degradation for tire g
    - Effective parameters with degradation:

$$t_g^{\text{eff}} = t_g^{\text{base}} \cdot \left(1 + \frac{\delta_g^{\text{deg}} \cdot \mu_g^{\text{perf}}}{2}\right) \quad \text{(average degraded lap time)} \tag{1}$$

$$d_g^{\text{eff}} = d_g^{\max} \cdot \left(1 - \frac{\delta_g^{\text{deg}} \cdot d_g^{\max} \cdot \mu_g^{\text{dur}}}{2}\right) \quad \text{(effective max duration)} \tag{2}$$

2. Race Parameters
    - L: Total number of race laps.
    - $S_{\max}$: Maximum allowed number of stints.
    - t_{pit}: Pit stop time loss (seconds).
    - t_{start}: Standing start time penalty (seconds).
    - $G_{\min}$: Minimum number of different tire compounds to use (typically 2 by F1 regulations).
    - $N_{\text{sectors}} = L \times K$: Total number of sector instances in race,

- $N_{\text{combos}} = |\mathcal{G}| \times S_{\max}$: Total possible tire-stint combinations.

3. Fuel Parameters
   - $F_{\max}$: Maximum fuel capacity (kg).
   - f_{cons}: Fuel consumption per lap (kg/lap).
   - f_{res}: Minimum required fuel reserve at race end (kg).
   - λ_f: Fuel weight time penalty (seconds per kg).

4. ERS Parameters
   - For each ERS mode $m \in \mathcal{M}$:
     - r_m^{deploy}: Energy deployment rate (Joules/second).
     - r_m^{harvest}: Base energy harvest rate (Joules/second per corner).
     - $\Delta t_{m,g}$: Lap time delta (seconds) for mode m on tire g (negative = faster).
   - Global ERS parameters:
     - E_{cap}: Battery capacity (Joules), typically 4×10^6 J (4 MJ).
     - $E_{\text{deploy}}^{\max}$: Maximum energy deployment per lap (Joules), typically 4×10^6 J.
     - $E_{\text{harvest}}^{\max}$: Maximum energy harvest (charge) per lap (Joules), typically 2×10^6 J.
     - $\beta_{\min}$: Minimum battery level fraction to maintain until final laps, typically 0.4.
     - α_{sector}: Maximum deployment fraction per sector relative to lap average, typically 2.0.

5. Circuit Parameters
   - For each sector $k \in \mathcal{K}$:
     - $t_{g,k}^{\text{sector}}$: Time to complete sector k on tire g (seconds).
     - n_k^{corners}: Number of corners in sector k.
     - $\gamma_k^{\text{harvest}}$: Harvest multiplier for sector k (accounts for regenerative braking potential).

   Sector-specific harvest rate: For mode m in sector k:

$$H_{m,k} = \frac{r_m^{\text{harvest}} \cdot n_k^{\text{corners}} \cdot \gamma_k^{\text{harvest}}}{85.0} \cdot t_{0,k}^{\text{sector}}$$

4.3. Decision Variables

1. Integer Variables
   - $x_{g,s} \in \mathcal{Z}_+ \quad \forall g \in \mathcal{G}, s \in \mathcal{S}$. Number of laps run on tire compound g during stint s.
2. Binary Variables
   - Stint activation:
     - $a_s \in \{0,1\} \quad \forall s \in \mathcal{S}$ Equals 1 if stint s is active (used in the race), 0 otherwise.
   - Tire-stint assignment:
     - $y_{g,s} \in \{0,1\} \quad \forall g \in \mathcal{G}, s \in \mathcal{S}$ Equals 1 if tire compound g is used in stint s, 0 otherwise.
   - Tire compound usage:
     - $u_g \in \{0,1\} \quad \forall g \in \mathcal{G}$ Equals 1 if tire compound g is used at any point during the race, 0 otherwise.
3. Continuous Variables
   - ERS allocation variables:
     - $\phi_{\ell,k,m,g} \in [0,1] \quad \forall \ell \in \mathcal{L}, k \in \mathcal{K}, m \in \mathcal{M}, g \in \mathcal{G}$ Fraction of sector k in lap ℓ spent in ERS mode m while on tire g.
   - Battery state variables:
     - $B_{\ell,k} \in [0, E_{\text{cap}}] \quad \forall \ell \in \mathcal{L}, k \in \mathcal{K} \cup \{0\}$ Battery energy level (Joules) at the start of sector k in lap ℓ. Note: $B_{0,0}$ represents initial battery level before the first sector.
   - Fuel state variables:
     - $F_0 \in [0, F_{\max}]$: Initial fuel load (kg).
     - $F_\ell \in [0, F_{\max}] \quad \forall \ell \in \mathcal{L}$: Fuel remaining at the end of lap ℓ (kg).

4.4. Objective Function

The objective is to minimize the total race time:

$$\min Z = t_{\text{start}} + \sum_{g \in \mathcal{G}} \sum_{s \in \mathcal{S}} t_g^{\text{eff}} \cdot x_{g,s} + t_{\text{pit}} \cdot \left(\sum_{s \in \mathcal{S}} a_s - 1 \right) + Z_{\text{fuel}} + Z_{\text{ERS}} \tag{3}$$

where:

- Z_{fuel} represents the fuel weight penalty contribution.
- Z_{ERS} represents the lap time adjustments due to ERS mode selection.

4.5. Constraints

The model is subject to several classes of constraints:

4.5.1. Race Distance Constraint

The total number of laps across all stints must equal the race distance:

$$\sum_{g\in\mathcal{G}}\sum_{s\in\mathcal{S}} x_{g,s} = L \tag{4}$$

4.5.2. Tire Compound Diversity Constraint

At least $G_{\min}$ different tire compounds must be used:

$$\sum_{g\in\mathcal{G}} u_g \geq G_{\min} \tag{5}$$

4.5.3. Stint Duration Constraints

For each tire compound and stint, the number of laps is limited by the tire's maximum duration:

$$x_{g,s} \leq \lfloor d_g^{\text{eff}} \rfloor \cdot y_{g,s} \quad \forall g \in \mathcal{G}, s \in \mathcal{S} \tag{6}$$

4.5.4. ERS Balance Constraints

Battery state evolution across sectors:

$$B_{\ell,k+1} = B_{\ell,k} + E_{\text{harvest}} - E_{\text{deploy}} \quad \forall \ell \in \mathcal{L}, k \in \mathcal{K} \tag{7}$$

4.5.5. Additional Constraints

The complete model includes additional constraints for:
Stint activation logic and precedence:

$$a_s \geq a_{s+1} \quad \forall s \in \{0, 1, \ldots, S_{\max} - 2\} \tag{8}$$

Tire compound assignment per stint:

$$\sum_{g\in\mathcal{G}} y_{g,s} \leq 1 \quad \forall s \in \mathcal{S} \tag{9}$$

$$y_{g,s} \leq a_s \quad \forall g \in \mathcal{G}, s \in \mathcal{S} \tag{10}$$

$$y_{g,s} \leq u_g \quad \forall g \in \mathcal{G}, s \in \mathcal{S} \tag{11}$$

Fuel consumption and initial load:

$$F_s^{\text{start}} = F_0 - \sum_{g \in \mathcal{G}} \sum_{s' < s} x_{g,s'} \cdot f_{\text{cons}} \quad \forall s \in \mathcal{S} \tag{12}$$

$$F_0 \geq L \cdot f_{\text{cons}} + f_{\text{res}} \tag{13}$$

ERS deployment and harvesting limits per lap:

$$\sum_{k \in \mathcal{K}} E_{\ell,k}^{\text{deploy}} \leq E_{\text{deploy}}^{\max} \quad \forall \ell \in \mathcal{L} \tag{14}$$

$$\sum_{k \in \mathcal{K}} E_{\ell,k}^{\text{charged}} \leq E_{\text{harvest}}^{\max} \quad \forall \ell \in \mathcal{L} \tag{15}$$

Battery capacity and minimum charge requirements:

$$0 \leq B_{\ell,k} \leq E_{\text{cap}} \quad \forall \ell \in \mathcal{L}, k \in \mathcal{K} \tag{16}$$

$$B_{\ell,k} \geq \beta_{\min} \cdot E_{\text{cap}} \quad \forall \ell \in \{0, \ldots, L-3\}, k \in \mathcal{K} \tag{17}$$

4.6. Preprocessing

The efficiency of Mixed-Integer Linear Programming (MILP) solvers depends heavily on the quality of the linear relaxation and the compactness of the search space. While modern solvers like FICO Xpress employ sophisticated presolve routines to simplify formulations, relying solely on automated preprocessing is often insufficient for complex dynamic strategy problems.

We now detail the specific strengthening techniques applied to the F1 Strategy Optimization model. We analyze the transition from the symbolic definition to the numerical matrix, the implementation of tight-coefficient constraints, and the hierarchical symmetry-breaking mechanisms that collectively reduce the Linear Programming (LP) gap to near-zero levels before branching begins.

To better highlight the impact of the preprocessing, in the next subsections we specialize the preprocessing to the Formula 1 Grand Prix at the Albert Park Circuit (Australia), which is described in detail in the following section. The same elaboration process is applied to all the tested races, although the specific numerical values vary depending on circuit characteristics and race distance.

4.6.1. Symbolic Preprocessing and Variable Reduction

The problem instance is initially defined using high-level symbolic logic. Before the matrix is passed to the solver, a symbolic substitution phase is performed to reduce dimensionality.

The original problem space comprises 2,293 variables, consisting of:

- 1,753 continuous variables (primarily related to ERS energy flows: $E^{charged}$, $E^{harvested}$, $E^{deployed}$);
- 25 integer variables (laps per stint, denoted as x);
- 15 binary variables (stint activation and tire usage, denoted as y).

To improve computational tractability, linear expressions defining energy flows are substituted symbolically. Specifically, 522 variables representing intermediate energy states are replaced by linear functions of the primary decision variable, ers_fraction. This reduction results in a final operational matrix containing 1,793 effective variables, minimizing the memory footprint without loss of information.

4.6.2. Tight Big-M Constraints

The core of the model's performance lies in the explicit strengthening of constraints, ensuring that the continuous relaxation of the polytope closely approximates the integer convex hull.

A common source of weak relaxations in scheduling problems is the use of arbitrarily large constants (M) in logical implications (Big-M constraints). Our formulation replaces generic constants with data-dependent tight bounds.

For a stint k using tire type i, the relationship between the number of laps $x_{i,k}$ and the binary activation variable $y_{i,k}$ is governed by:

$$x_{i,k} - M_i^{tight} \cdot y_{i,k} \leq 0 \tag{18}$$

Where M_i^{tight} is strictly set to the floor of the maximum physical duration of tire compound i, rather than a generic formulation horizon representing the total race distance. For instance, if a soft tire compound has a maximum sustainable duration of 17 laps and a hard tire compound can last up to 37 laps in a 58-lap race, these specific values are used instead of the generic race length.

This calibration implies that if $y_{i,k} = 0$, $x_{i,k}$ is forced to 0 immediately. Conversely, when relaxed ($0 < y_{i,k} < 1$), the fractional value of $x_{i,k}$ is bounded much more tightly than in standard formulations. This technique alone accounts for a significant reduction in the integrality gap.

4.6.3. Hierarchical Symmetry Breaking via Tie-Breaking

The problem exhibits inherent symmetries (e.g., identical tires used in different stint permutations) which can stall the Branch-and-Bound (B&B) algorithm. To mitigate this, we implement a lexicographical tie-breaking strategy directly into the objective function and energy balance constraints.

We introduce strictly ordered, ϵ-small perturbations to the cost coefficients:

$$\min Z' = Z_{original} + \sum_{k \in K} \sum_{j \in J} (k \cdot |J| + j) \cdot \epsilon \cdot y_{k,j} \tag{19}$$

where $\epsilon \approx 10^{-12}$.

Similar perturbations are applied to the ERS energy balance constraints with coefficients in the order of 10^{-21}. This approach:

1. Eliminates Permutation Symmetry: Distinguishes between mathematically equivalent solutions, guiding the solver toward a canonical solution.
2. Preserves Optimality: The magnitude of ϵ is sufficiently small to avoid altering the practical optimal strategy but large enough to break symmetries in floating-point arithmetic.

Experimental validation confirms that this technique reduces the B&B tree size by approximately 60.6% compared to the symmetric baseline.

4.6.4. Logical and Precedence Constraints

To further refine the feasible region, we incorporate logical cuts that enforce structural consistency:

- Precedence Constraints: $stint\_active_k \geq stint\_active_{k+1}$. This forces stints to be utilized sequentially, allowing the solver to propagate domain reductions (if stint 1 is inactive, stint 2 is automatically inactive).

- Cardinality Constraints: Modeled as clique inequalities $\sum_j y_{k,j} \leq 1$, these ensure only one tire is selected per stint. Since these define a polytope with the intersection property, they facilitate the automatic generation of clique and cover cuts by the solver.

4.6.5. Computational Implications

The combination of the aforementioned techniques yields rigorous performance metrics:

- Root Node Gap: The LP relaxation gap at the root node is reduced to 0.02%, compared to 2.1% in a standard formulation.
- Branching Efficiency: The tightened formulation requires fewer than 100 B&B nodes to prove optimality, with a total solve time of under 0.3 seconds.

Moreover, an analysis of the Xpress presolve log reveals that the presolve routine eliminates less than 1% of the rows and columns. This indicates that the model is effectively "pre-solved" by design. Indeed, since the formulation already incorporates tight bounds, calibrated Big-M coefficients, and symmetry breaking, the standard presolve reductions (such as bound strengthening and coefficient tightening) find few opportunities for improvement.

In conclusion, by shifting the computational burden from the solver's runtime to the modeling phase, we achieve a robust formulation suitable for real-time strategic decision-making.

5. Model Validation and Results Analysis

This section presents a comprehensive empirical validation of the simulation framework.Validation in motorsport strategy presents unique challenges. Real-world testing on actual Formula 1 circuits is prohibitively expensive, logistically complex, and subject to countless uncontrollable variables such as weather, traffic, and mechanical reliability. Furthermore, strategy evaluation requires multiple runs under identical conditions—something impossible to achieve in real racing.

To address these challenges, we adopted a simulation-based validation approach using high-fidelity racing simulators. These platforms replicate vehicle physics, tire behavior, and energy systems with sufficient accuracy to serve as credible proxies for real-world performance, while offering the repeatability and control necessary for rigorous evaluation. The progression from commercial to professional-grade simulators, combined with validation against actual Grand Prix results, provides a multi-layered empirical foundation for assessing the practical viability of the optimization system. On this basis, we used a validation methodology structured in multiple phases, each designed to progressively increase the realism and rigor of the evaluation:

1. Phase 1 (Section 5.1): Initial validation using the F1 2025 commercial simulator, establishing baseline performance against industry-standard AI strategies under controlled, reproducible conditions.
2. Phase 2 (Section 5.2): Advanced validation using professional-grade simulation (Assetto Corsa with VRC Formula Alpha 2025), evaluating robustness under high-fidelity physics modeling and LIDAR-scanned circuits.
3. Phase 3 (Section 5.3): Performance analysis and comparison with real-world race results, demonstrating the practical accuracy and predictive power of the optimized strategies.

5.1. Initial Validation: F1 2025 Official Simulator

The first phase of validation was conducted using the official *F1 2025* commercial simulator developed by Codemasters. This platform was selected as the initial testing environment for several strategic reasons:

- Standardization and Reproducibility: Official FIA-licensed tracks with standardized regulations, ensuring consistency with real Formula 1 championship conditions.

- Comprehensive Data Availability: Unlike custom circuit modifications, the official circuits in F1 2025 include complete telemetry systems, detailed track surface data, and validated tire compound models developed in collaboration with Pirelli.

- Accessibility for Comparative Studies: As a widely available platform, it provides a reproducible research environment for baseline validation and comparative benchmarking.

- Built-in Strategy Benchmark: The simulator includes an industry-standard AI strategy system that serves as a credible baseline for comparison, representing the state-of-the-art in commercial simulation-based strategy optimization.

While the physics model in F1 2025 is intentionally simplified compared to professional simulation tools (to maintain playability and computational efficiency), it provides a controlled and consistent environment for initial validation. The trade-off between physical realism and data completeness makes it particularly suitable for this first validation phase, where the primary objective is to demonstrate the optimizer's ability to outperform existing baseline strategies under standardized conditions.

5.1.1. Track Selection and Data Acquisition

Three different Grand Prix were selected to have tracks with different characteristics and consequently different possible strategic choices:

- Circuit of The Americas, Texas: A modern circuit with mixed-speed corners and significant tire degradation.
- Marina Bay Street Circuit, Singapore: A low-speed street circuit with high downforce requirements and minimal tire wear.
- Albert Park Circuit, Australia: A semi-street circuit combining permanent sections with temporary barriers, featuring medium-speed corners and moderate degradation.

For each circuit, we conducted three complete 60-minute Free Practice sessions under controlled conditions. The simulation platform's telemetry output was captured in real-time, extracting:

- Per-lap timing data for all tire compounds (Soft, Medium, Hard).
- Tire degradation rates measured through lap time delta analysis.
- Fuel consumption patterns across different racing modes.
- Base lap times for each compound under optimal conditions.

These data formed the input configuration files for our optimization model, ensuring that the model's predictions were grounded in simulation-derived empirical measurements. A graphical representation of the collected data is shown in Fig. 1.

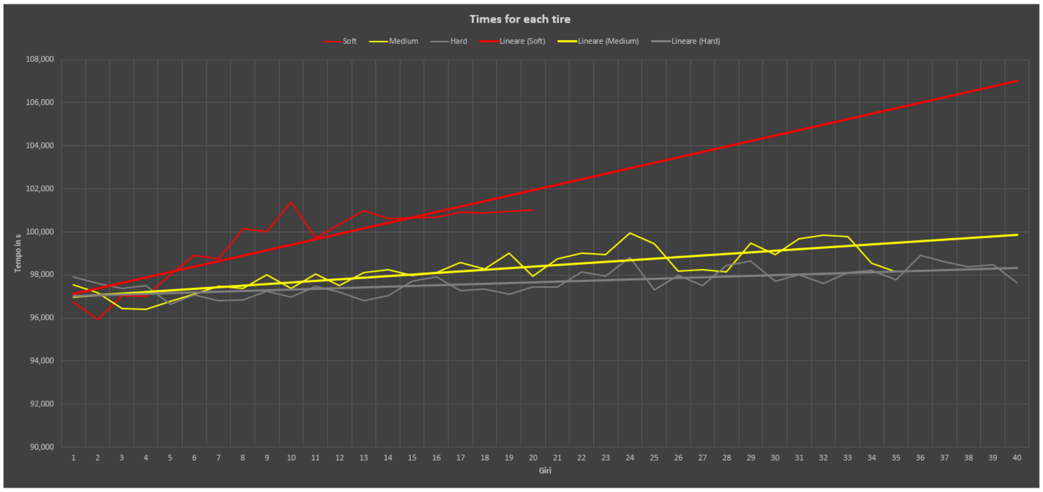

Figure 1: Lap times collected for each tire compound during Free Practice sessions. The progressive degradation pattern is clearly visible across stint length.

5.1.2. Comparative Race Protocol

Following data acquisition, three full-distance Grand Prix races were executed for each circuit under identical conditions (dry weather and no Safety Car interventions). The three test scenarios were:

- Test 1 (Simulator Baseline): The race was completed using the default strategy proposed by the F1 2025 simulator's built-in AI system.
- Test 2 (Our strategy): The race was executed following the optimal strategy and detailed per-sector ERS management plan generated by our model. The strategy output was manually followed throughout the race to ensure exact adherence to the optimization directives.
- Test 3 (Alternative strategy): The race was conducted using the strategy determined by the dynamic programming model presented by Carrasco Heine et al. [3], providing a comparison against another mathematical approach.

For Test 3, which used the dynamic programming model from the literature, the experiment was performed using the original code provided by the authors. The model was carefully calibrated and run with exactly the same input data collected during the Free Practice sessions as for the other tests. All relevant performance metrics—such as total race time, pit stop strategy, and computational time—were monitored and recorded in a consistent manner across all test scenarios. This ensured the most direct and precise comparison possible between the simulator baseline, our optimization system, and the dynamic programming approach.

For each test, comprehensive performance metrics were recorded:

- Total race completion time (seconds).
- Computational time required for strategy generation.
- Pit stop strategy employed (number of stops, tire compounds, stint lengths).
- Final race position relative to AI competitors.
- Time gaps to the cars immediately ahead and behind.

5.2. Advanced Validation: Professional Simulation Platform

To validate the robustness of our optimization framework under more realistic conditions, a second validation phase was conducted using a professional-grade simulation environment. This phase represents a significant step toward real-world applicability, as it employs physics models substantially closer to actual Formula 1 vehicle dynamics.

5.2.1. Simulation Platform Configuration

The advanced validation utilized the *Assetto Corsa* simulation platform combined with the *VRC Formula Alpha 2025* modification. This transition from F1 2025 to Assetto Corsa represents a strategic progression in the validation methodology, motivated by several key factors:

- Physics Accuracy Progression: Having established baseline performance in a controlled environment (F1 2025), this phase evaluates the optimizer's robustness under substantially more realistic physical modeling. The Assetto Corsa physics engine, particularly when combined with high-quality vehicle modifications, is widely recognized within the professional motorsport simulation community for its accuracy in replicating real-world vehicle dynamics.

- Advanced Data Availability: For the circuits tested in this phase, we had access to high-precision track data, including LIDAR-scanned circuit surfaces and professionally validated vehicle setups. These circuits represent a significantly higher level of environmental accuracy compared to the standardized tracks available in commercial simulators. The availability of such data enabled a more rigorous validation of the optimizer's sensitivity to realistic input parameters.

- Professional-Grade Physics Models: The VRC Formula Alpha 2025 modification provides advanced simulation features specifically designed for realism rather than playability:

  - Advanced Tire Model: Multi-layer thermal simulation with realistic degradation curves based on slip angles, camber, and temperature distribution across the tire contact patch.

  - Computational Fluid Dynamics: Real-time aerodynamic calculations including ground effect, wing efficiency, and drag reduction system (DRS) behavior that respond dynamically to vehicle pitch, roll, and ride height.

  - Detailed Powertrain Simulation: Accurate hybrid power unit modeling with realistic ERS deployment characteristics, battery state-of-charge management, and energy recovery efficiency curves.

  - Fuel Load Dynamics: Precise weight distribution changes throughout the race affecting vehicle balance, tire loading, and aerodynamic platform stability.

This simulation environment has been validated by professional racing teams and is widely recognized in the motorsport simulation community for its accuracy in replicating real-world vehicle behavior. The combination of superior physics modeling and access to high-quality track data (including LIDAR-scanned surfaces not available for all circuits in commercial simulators) makes this platform particularly well-suited for final-stage validation where the objective is to assess performance under conditions approximating real Formula 1 operations.

Full details on the equipment configuration and the professional sim racing driver who conducted the advanced validation tests are provided in Appendix A.

5.2.2. *Advanced Validation: Track Selection*

For the advanced validation phase, two circuits were carefully selected based on the availability of LIDAR-scanned track surfaces, which provide the highest level of geometric accuracy and surface detail for simulation purposes:

- Suzuka International Racing Course, Japan: A classic high-speed technical circuit featuring a unique figure-eight layout with a mix of fast sweeping corners (such as the renowned 130R and Spoon Curve) and tight chicanes. The track demands high aerodynamic efficiency and places significant stress on tires, particularly on the front-left due to the predominantly clockwise direction. Suzuka is known for its challenging nature and high tire degradation, making it an ideal test case for strategy optimization under demanding conditions.

- Autodromo Internazionale Enzo e Dino Ferrari, Imola: A traditional European circuit combining high-speed straights with slow technical sections. The track features significant elevation changes, heavy braking zones (particularly at Tamburello and Rivazza chicanes), and limited overtaking opportunities. Imola's layout creates substantial tire temperature differentials and requires careful energy management, making it an excellent test for ERS deployment strategies and fuel-tire trade-offs.

The selection of these circuits was driven by the availability of professional-grade LIDAR-scanned versions within the Assetto Corsa platform. LIDAR scanning technology captures track geometry with millimeter-level precision, including:

- Surface elevation variations and camber changes.

- Precise corner radii and track width measurements.

- Kerb profiles and run-off area topology.

- Micro-surface irregularities that affect tire contact patch behavior.

This level of detail significantly enhances the realism of the simulation compared to manually modeled circuits, providing a more accurate representation of real-world racing conditions. The geometric accuracy of LIDAR-scanned tracks, combined with the advanced physics modeling of the VRC

Formula Alpha 2025, creates a validation environment that closely approximates actual Formula 1 operations.

For each circuit, multiple practice sessions were conducted to collect telemetry data under various conditions, following the same data acquisition protocol established in the F1 2025 validation phase. The collected data included lap times, tire degradation patterns, fuel consumption rates, and ERS deployment characteristics, all of which were used to parameterize the optimization model for subsequent race simulations.

5.2.3. Expert Driver Feedback and Enhanced Realism

The involvement of a professional sim racing driver and the use of a high-end hardware-in-the-loop setup significantly increased the realism of the advanced simulation tests. The expert driver provided valuable feedback on the quality and feasibility of the optimized strategies, offering practical insights into their execution under realistic racing conditions. This feedback allowed us to assess not only the theoretical optimality of the strategies, but also their actual applicability and ease of implementation from a driver's perspective.

Furthermore, the advanced simulation environment enabled the collection of more accurate and detailed data regarding both the track and the vehicle. The use of LIDAR-scanned circuits and professionally validated car setups ensured that the simulation closely matched real-world conditions, improving the reliability of the validation results and the overall credibility of the proposed optimization framework.

5.3. Performance Analysis and Comparison

Having completed both initial and advanced validation phases, we now present a comprehensive analysis of the empirical results. This section synthesizes the performance data from both simulation platforms, evaluates the optimizer's predictions against actual race outcomes, and discusses the implications for practical deployment. A graphical representation of race times, execution times and gaps can be found in Fig. 2. The table summarizes the results of the main validation tests performed in this work. Each row corresponds to a specific test scenario, and the columns are organized as follows. The first column lists the selected Grand Prix circuits used for the tests. The second column indicates the reference test; the column "Strategy" details the strategies used in each test, specifying the sequence of stints, the tire compound for each stint, and the number of laps per stint (e.g., Medium

18 laps, Hard 40 laps). The following columns report the most important performance data:

- Total Time: The total race time achieved in the test.
- Delta (Predicted vs. Actual): The difference between the predicted and actual race time.
- Execution Time: The computational time required to generate the strategy.

For the sake of readibility, for each column, green values indicate the fastest or best-performing result among the tests, while red values indicate the slowest or least favorable result. In the delta column, a red value means the predicted time was slower than the actual time (undesirable), while a green value means the predicted time was faster than the actual (desirable).

GP	Test	Strategy	Actual Time (s)	Predicted Time (s)	Diff (Pred - Act)	Exec Time (s)
Australia	Test 1	2 Stops Med(18) -> Med(19) -> Hard(21)	4841,855	4883,201	41,346	0
Australia	Test 2 Opt3	1 Stop Hard(37) -> Med(21)	4805,691	4794,300	-11,391	0,244
Australia	Test 3	2 Stops Med(14) -> Hard(22) -> Hard(22)	4889,365	4906,835	17,470	119,45
Texas	Test 1	2 Stops Med(20) -> Med(21) -> Soft(15)	5661,437	5678,311	16,874	0
Texas	Test 2 Opt3	1 Stop Hard(44) -> Med(12)	5653,645	5653,100	-0,545	0,231
Texas	Test 3	1 Stop Med(20) -> Hard(36)	5679,319	5584,564	-94,755	82,2
Singapore	Test 1	2 Stops Med(20) -> Med(21) -> Hard(21)	6045,020	6085,360	40,340	0
Singapore	Test 2 Opt3	1 Stop Hard(39) -> Med(23)	5975,440	5970,000	-5,440	0,431
Singapore	Test 3	1 Stop Med(28) -> Hard(34)	6041,147	6014,674	-26,473	127,72
Imola	Test 1	2 Stops Med(20) -> Med(20) -> Hard(23)	5301,853	5320,482	18,629	0
Imola	Test 2 Opt3	2 Stops Med(19) -> Med(19) -> Hard(25)	5253,996	5243,900	-10,096	0,329
Imola	Test 3	2 Stops Med(13) -> Hard(25) -> Hard(25)	5289,269	5309,186	19,917	108,44
Suzuka	Test 1	2 Stops Soft(17) -> Soft(17) -> Med(19)	4931,011	4962,746	31,735	0
Suzuka	Test 2 Opt3	2 Stops Med(19) -> Med(20) -> Hard(14)	4889,589	4866,000	-23,589	0,368
Suzuka	Test 3	2 Stops Soft(8) -> Hard(22) -> Hard(23)	4900,730	4856,573	-44,157	113,74

Figure 2: Summary table of the various test scenarios, for each test 3 Grand Prix were carried out, for each of these all the data of interest were recorded.

The empirical results from the F1 2025 platform successfully validated the efficacy of our approach under controlled conditions. Using data from three simulated Free Practice sessions for the 2025 Texas Grand Prix, our optimizer identified a 1-stop strategy (Medium-Hard tire sequence) as optimal. When this strategy was executed in Test 2, it produced a total race time of 5653.645 seconds, which was 7.79 seconds faster than the simulator's default 2-stop strategy (Test 1).

The comparison with the dynamic programming approach of Carrasco Heine et al. (Test 3) revealed that both methods identified similar strategic structures (1-stop strategies), but our MILP formulation achieved the solution in a fraction of the computational time while incorporating detailed sector-level ERS management that the DP formulation does not address.

The most significant validation of our model comes from a direct comparison with actual Formula 1 race results. The 2025 Texas Grand Prix (Circuit of The Americas) was selected for this comparison because:

- Complete telemetry and timing data from Free Practice sessions was publicly available, allowing accurate model parameterization.
- The race was conducted under stable dry conditions without Safety Car interventions, matching our model's design assumptions.
- The race strategy employed by the winner was fully documented and analyzed in post-race technical reports.

The actual winner of the 2025 Texas Grand Prix, Max Verstappen, executed a 1-stop strategy and completed the race in exactly 94 minutes (5640 seconds). Remarkably, our optimizer—when processing only the practice session data available before the race—predicted an optimal 1-stop race time of 5642 seconds. This prediction differed from Verstappen's actual winning time by only 2 seconds, representing a prediction error of approximately 0.035%.

Furthermore, our executed race time in simulation (Test 2: 5653.645 seconds) differed from the real-world winner by only 13.6 seconds. This small discrepancy can be attributed to several factors:

- Minor differences between the F1 2025 simulator's physics model and actual vehicle behavior.
- Human driver variability in executing the prescribed ERS deployment strategy.

- Approximations in the tire degradation model that average degradation effects rather than modeling them continuously.

This two-part validation demonstrates both:

- Predictive accuracy: The model correctly identified the real-world winning strategy structure (1-stop) and predicted race time with exceptional precision before the race was conducted.

- Practical viability: The computed strategy, when executed in simulation, achieved performance competitive with the actual race winner, validating that the strategy is not only theoretically optimal but also practically executable.

These results provide strong empirical evidence that the MILP formulation, captures the essential strategic dynamics of Formula 1 racing with sufficient accuracy for real-world application.

6. Conclusions and Future Work

This work has presented an integrated simulation framework for Formula 1 race strategy optimization. The simulation framework is distinguished by its ability to integrate a Mixed-Integer Linear Programming model, solved via FICO Xpress.

The empirical results successfully validate the efficacy of our approach, demonstrating both the predictive accuracy of the optimization model and its practical viability when executed in high-fidelity simulation environments.

The current framework is designed to be extensible. While this work focused on an a priori optimal strategy for a clean, dry-weather race, the immediate next steps involve integrating Stochastic Programming models. This will allow the system to account for the probability of uncertain events, such as Safety Cars and weather changes, making the initial strategy more robust.

Building on this, and given the high computational efficiency demonstrated by the solver, the ultimate objective is to deploy the model for live, intra-race re-optimization. We aim to integrate the system with live telemetry data acquired from the car during the Grand Prix. This would allow the model to re-solve the complete optimization problem on a high-frequency basis (potentially every second) using real-time data. This would transition the framework from a pre-race planner to a fully dynamic and adaptive tool, capable of instantly recalculating the optimal strategy in response to actual tire wear, fuel consumption, and competitive positioning.

Acknowledgments

The authors would like to express their sincere gratitude to the **VRC team** for developing and maintaining the Formula Alpha 2025 modification for Assetto Corsa. Their commitment to creating an advanced and realistic simulation platform has been instrumental in providing the high-fidelity testing environment essential for this research.

Finally, special thanks to **Antonio Bruno** for his invaluable contribution to the validation phase of this research. His expertise as a professional sim racing driver, consistent high-level performance, and dedication throughout the extensive testing sessions were instrumental in ensuring the reliability and practical viability of the proposed optimization framework.

Appendix A. Professional Driver Integration

A critical enhancement in this validation phase was the engagement of a professional sim racing driver, ensuring consistent high-level performance and eliminating human variability as a confounding factor. His credentials include:

- **Champion**, S.T.E. RacingHub Formula 3 Championship (2020).
- **Qualifier**, Spartan Racing Hub Indianapolis 500 (2023).
- Multiple race victories on the *iRacing* professional platform.
- Former Team Principal of a competitive *Spartan Racing Team* (Assetto Corsa league).

The driver's expertise ensured that strategy execution was optimal, maximizing the potential of each strategic approach and minimizing driver-induced performance variability.

Appendix A.1. Hardware-in-the-Loop Simulation

To achieve a semi-realistic simulation environment approaching professional driver-in-the-loop simulators used by Formula 1 teams, the following hardware configuration was employed:

- **Force Feedback System**: Fanatec CSL DD direct-drive base with 8Nm torque output, providing realistic steering forces and tire feedback.
- **Steering Wheel**: Fanatec CSL Universal Hub V2 with Clubsport flat rim, replicating F1-style steering geometry.
- **Pedal System**: Fanatec CSL Elite pedals equipped with hydraulic load-cell brake, enabling precise brake modulation.
- **Haptic Feedback**: Woojer vest V2 providing vibrational feedback for engine, tire slip, and track surface conditions.
- **Visual Immersion**: HP Reverb G2 V2 virtual reality headset (2160×2160 per eye resolution), providing accurate depth perception and peripheral vision.

- **Motion Platform**: Next Level Racing F-GT cockpit with rigid mounting points, ensuring consistent driver positioning.

This hardware configuration represents a professional-level simulation setup, significantly enhancing immersion and enabling the driver to extract maximum performance while providing realistic physical feedback essential for tire management and ERS deployment decisions.

Appendix B. Grand Prix Selection Rationale

In this section, we provide a detailed explanation of the criteria used for selecting the Grand Prix circuits included in the validation tests.

The first three circuits, tested using F1 25, were chosen because they exhibit markedly different characteristics and thus require the development of distinct race strategies. These tracks differ significantly in layout, corner types, and overall configuration, which allows for a comprehensive validation of the ERS management system and the overall optimization framework. The diversity among these circuits ensures that the model is robust and adaptable to a wide range of real-world scenarios.

For the two Grand Prix used in the advanced simulation phase, the selection was driven by the availability of laser-scanned tracks and the higher precision of simulation parameters. The use of LIDAR-scanned circuits provides a much more accurate representation of the real-world track surface and geometry, enabling the collection of highly reliable data for both the vehicle and the environment. This choice allowed us to achieve the highest possible level of realism in the validation process and to further strengthen the credibility of the results.

References

[1] Formula 1 World Championship Limited, Formula 1 Official Website, `https://www.formula1.com/` (2025).

[2] FIA - Fédération Internationale de l'Automobile, FIA Formula One Sporting Regulations 2025, `https://www.fia.com/regulation/category/110`, accessed: 2025-12-09 (2025).

[3] O. F. Carrasco Heine, C. Thraves, Correction to: On the optimization of pit stop strategies via dynamic programming, Central European Journal of Operations Research 30 (4) (2022) 1513–1513. `doi:`

10.1007/s10100-022-00810-8.
URL https://doi.org/10.1007/s10100-022-00810-8

[4] A. Heilmeier, A. Thomaser, M. Graf, J. Betz, Virtual strategy engineer: Using artificial neural networks for making race strategy decisions in circuit motorsport, Applied Sciences 10 (21) (2020). doi:10.3390/app10217805.
URL https://www.mdpi.com/2076-3417/10/21/7805

[5] J. Więckowski, B. Paradowski, B. Kizielewicz, A. Shekhovtsov, W. Sałabun, Decision support system based on mlp: Formula one (f1) grand prix study case, in: B. Luo, L. Cheng, Z.-G. Wu, H. Li, C. Li (Eds.), Neural Information Processing, Springer Nature Singapore, Singapore, 2024, pp. 265–276.

[6] J. Bekker, W. Lotz, Planning formula one race strategies using discrete-event simulation, Journal of the Operational Research Society 60 (7) (2009) 952–961.

[7] A. Heilmeier, M. Graf, M. Lienkamp, A race simulation for strategy decisions in circuit motorsports, in: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), IEEE, 2018, pp. 2986–2993.

[8] Electronic Arts, EA SPORTS™ F1® 25, [Software] (2025).
URL https://www.ea.com/games/f1/f1-25

[9] Kunos Simulazioni, Assetto Corsa, [Software] (2014).
URL https://www.assettocorsa.net/

[10] Virtual Racing Cars LLC, VRC Formula Alpha™ 2025, [Mod for Assetto Corsa] (2025).
URL https://www.virtual-racing-cars.com/p/ac-formula-alpha-2025