

Mitigating Predictive Bias in Formula 1 Pit Stop Data Using Disaggregated Machine Learning Models

Report By: Neha Gurnani, Participant A, Iman Ibrahim, Salma Ahmed, Momena Amin

Conducted as part of Femetronics

Date: 23/12/2025

Abstract

Machine learning models trained on uneven real-world datasets are vulnerable to systematic predictive bias. In Formula 1, top-tier constructor teams generate cleaner and more consistent pit stop data due to greater resources, while other-tier teams produce noisier observations. Training a single unified Decision Tree Regressor on this imbalanced data resulted in biased performance, with consistently higher Mean Absolute Error (MAE) for other-tier teams across all implementations. To mitigate this effect, a disaggregation-based fairness intervention was applied across three independently implemented projects, in which datasets were split by constructor tier and specialised Decision Tree models were trained separately. This approach reduced MAE disparities from initial differences of up to 0.00305 seconds to as low as 0.00024 seconds, corresponding to bias reductions of up to 92.17%. These findings demonstrate that disaggregated, specialised models substantially improve predictive parity, enhancing the fairness, reliability, and ethical validity of pit stop time predictions across Formula 1 teams.

1 Introduction

Research has highlighted substantial financial disparities between top Formula 1 constructor teams and lower-resourced teams, with budget differences exceeding £150 million prior to the 2021 cost cap¹. Although the cap was introduced to reduce competitive imbalance and has increased to approximately £161 million for the 2026 season, inequalities persist due to expenditures and resources outside its scope. These disparities affect infrastructure, personnel, and operational efficiency, and consequently influence on-track performance and data quality.

Top-tier teams typically generate cleaner and more consistent pit stop data, while lower-tier teams produce noisier observations. When a single unified model is trained on this uneven dataset, it may prioritise higher-quality data patterns, leading to biased predictions that disadvantage lower-resourced teams. Using Mean Absolute Error (MAE) as a fairness metric², this study hypothesises that disaggregation through specialised models will significantly reduce the MAE differential between top-tier and lower-tier Formula 1 teams compared to a unified approach.

2 Background

2.1 Mean Absolute Error (MAE)

Mean Absolute Error (MAE) measures regression accuracy by averaging the absolute difference between true and predicted values⁴:

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

where n represents the number of observations, y_i denotes the true value, and $\hat{y}_i$ represents the predicted value³. Absolute values prevent positive and negative residuals from cancelling, making MAE directly interpretable in the same units as the target variable (seconds)⁴. This is appropriate for pit stop and lap time variation⁷, where typical prediction errors should be summarised without allowing rare anomalies to dominate evaluation⁴. For fairness assessment, MAE can be computed per subgroup, and the absolute difference between group MAEs (accuracy parity gap) quantifies predictive disparity, with values closer to zero indicating stronger predictive parity⁵.

2.2 Disaggregation and Fairness

Disaggregation is a fairness-aware evaluation and mitigation approach that partitions data by a protected attribute to expose and reduce subgroup performance disparities^{5,6}. Here, the constructor tier (top-tier vs. lower-tier) is treated as the protected attribute. Barocas et al. show that strong aggregate performance can conceal severe subgroup underperformance, particularly when disadvantaged groups are underrepresented⁵.

Bynum et al. argue that disaggregated interventions can actively reduce inequality by tailoring modelling decisions to subgroup-specific behaviour⁵. In Formula 1, where data-driven models inform pit strategy and performance analysis, disaggregation helps ensure that predictive systems do not systematically disadvantage lower-resourced teams, improving reliability and ethical validity in motorsport analytics⁷.

3 Implementation

3.1 Shared Implementation Approach

All three projects implemented a fairness-aware machine learning pipeline using a Decision Tree Regressor to predict Lap Time Variation from Formula 1 pit stop data spanning the 2018–2024 seasons. The shared objective was to identify and mitigate predictive bias arising from structural disparities between top-tier and lower-tier constructor teams. While minor implementation details varied, the overall modelling logic was consistent across projects.

The pipeline involved data cleaning and preprocessing, followed by training a unified model as a diagnostic step to identify baseline predictive bias. A disaggregation-based mitigation strategy was then applied by training specialised models for each team tier. Constructor team tier was treated as the protected attribute and Mean Absolute Error (MAE) was used to evaluate predictive parity between groups. Project-specific variations are discussed in Section 3.4.

3.2 Data and Feature Engineering

All projects used the same CSV dataset containing Formula 1 pit stop data. Rows with missing target values were removed during preprocessing to ensure reliable evaluation. The target variable, Lap Time Variation, was measured in seconds.

Feature selection prioritised physical race conditions rather than organisational identity. Input features included stint length, air temperature, tyre compound, and constructor information. Categorical variables were one-hot encoded for compatibility with the Decision Tree model.

A derived categorical feature, `Team_Category`, was used solely for fairness analysis and data splitting and was excluded from the specialised model inputs. Including this feature would cause label leakage, allowing the model to rely on team identity rather than learning relationships from race conditions, which would artificially lower error for dominant teams and compromise fairness and

generalisability. Excluding Team_Category ensured that any bias reduction resulted from disaggregation alone, not explicit access to team status.

3.3 Unified and Specialised Models

Each project first trained a single unified Decision Tree model on the full dataset to diagnose baseline predictive bias. Model performance was evaluated separately for top-tier and lower-tier teams using MAE, and bias was quantified as the absolute difference in MAE between groups.

To mitigate the identified bias, a disaggregation strategy was applied by splitting the training data according to Team_Category and training separate specialised models for each team tier. Removing the team category feature from these models ensured predictions were learned from race conditions rather than team identity, improving fairness and generalisation.

3.4 Project-Specific Implementation

Project A (Iman and Participant A)

Project A extended the shared pipeline by analysing directional bias through signed prediction errors, identifying minor overestimation for both team tiers, with slightly higher bias for lower-tier teams. These effects were small given dataset size and variability.

Project B (Salma and Momena)

Project B closely followed the shared pipeline but applied a stricter definition of top-tier teams, resulting in a larger baseline MAE disparity under the unified model. This led to a more pronounced bias reduction after disaggregation while maintaining comparability with other projects. For clarity, we included simple interpretive checks that display the magnitude of bias, as well as the differences in MAE and percentages of bias reduction.

Project C (Neha)

Project C strictly separated bias diagnosis and mitigation phases, ensuring that bias measurement and reduction were evaluated independently. This strengthened methodological clarity and supported a clearer causal interpretation of the fairness intervention.

4 Results

This section presents the quantitative results of the unified and specialised models, evaluated using MAE to assess predictive parity between team tiers.

4.1 Unified Model Results: Baseline Bias

The unified model results provide an initial assessment of predictive bias prior to any fairness intervention. For each project, Mean Absolute Error was computed separately for top-tier and lower-tier teams using a single Decision Tree model.

Table 1: Unified Model Mean Absolute Error by Project

Project	UM Top-Tier MAE	UM Other-Tier MAE	UM Bias Difference
A	0.0323	0.0337	0.0014
B	0.0319	0.0350	0.00305
C	0.0071	0.0097	0.0025

4.2 Specialised Model Results: Bias Mitigation

Following the identification of baseline bias, each project applied a disaggregation strategy by training specialised models for top-tier and lower-tier teams. Model performance was again evaluated using MAE.

Table 2: Specialised Model Mean Absolute Error by Project

Project	SM Top-Tier MAE	SM Other-Tier MAE	SM Bias Difference
A	0.0331	0.0336	0.0005
B	0.0362	0.0360	0.00024
C	0.0115	0.0111	0.0004

Across all projects, specialised models reduced MAE disparities between team tiers, indicating improved predictive parity.

4.3 Bias Reduction Comparison Across Projects

To compare the effectiveness of the mitigation strategy, bias reduction was calculated as the percentage decrease in the MAE difference between top-tier and lower-tier teams.

Table 3: Bias Reduction Summary

Project	Initial Bias	Final Bias	Bias Reduction
A	0.0014	0.0005	65.3%
B	0.00305	0.00024	92.17%
C	0.0025	0.0004	83.8%

These results demonstrate that the disaggregation approach consistently reduced predictive bias across all projects, confirming its effectiveness as a fairness intervention.

4.4 Project-Specific Results and Visualisation

Each project produced a bar chart comparing unified and specialised model performance for top-tier and lower-tier teams.

Project A (Iman and Partipant A)

Figure 1: Project A Bar Chart




Figure 1 shows a modest bias reduction after disaggregation, likely due to smaller group sizes.

Project B (Salma and Momena)

Figure 2: Project B Bar Chart




Figure 2 demonstrates that the unified model is biased toward top-tier teams (MAE 0.032 vs 0.035). The specialised model reduces this bias by equalising error (0.036), yet accuracy for lower-tier teams marginally worsens.

Project C (Neha)

Figure 3: Project C Bar Chart




Figure 3 shows near-equal MAE values after disaggregation, indicating achieved predictive parity despite a modest increase in overall error.

4.5 Cross-Project Comparison

Across all projects, unified models disadvantaged lower-tier teams, while disaggregated models consistently reduced bias.

5 Conclusion

This study examined whether predictive bias arising from structural inequalities in Formula 1 data could be mitigated using disaggregation-based machine learning. Using Mean Absolute Error as a fairness metric, the results support the hypothesis that specialised models trained on disaggregated team-category data reduce predictive disparity between top-tier and lower-tier teams. Across the three projects, MAE differentials were reduced by 65.3%-92.17%, demonstrating that disaggregation is an effective and robust bias mitigation strategy, even when overall prediction accuracy does not improve.

Each project contributed complementary insights: Project A analysed directional bias to improve interpretability of over- and underestimation across team tiers; Project B applied a stricter definition of top-tier teams, producing a larger baseline bias and a more pronounced reduction after mitigation; and Project C maintained a strict separation between bias diagnosis and mitigation, strengthening methodological clarity and causal interpretation.

From an ethical and professional perspective, this work aligns with responsible AI principles by addressing systematic performance disparities. The observed bias reduction supports compliance with emerging AI fairness regulations, such as the EU AI Act⁸, and aligns with IEEE and ACM professional standards. By prioritising fairness alongside technical performance, the approach promotes trust among teams, sponsors, and fans^{9 10}. Future work could extend this methodology to other motorsports, such as MotoGP or the World Rally Championship, and explore longer historical datasets or ensemble models to further improve accuracy while maintaining fairness.

References

1. GoToTheGrid (2025) *The truth behind F1's budget cap*. Available at: <https://www.gotothegrid.com/en/blog/the-truth-behind-f1s-budget-cap> (Accessed: 12 December 2025).
2. ScienceDirect (n.d.) *Mean absolute error*. Available at: <https://www.sciencedirect.com/topics/engineering/mean-absolute-error> (Accessed: 17 December 2025).
3. Burch, D. (2023) *Mean absolute error in machine learning: what you need to know*. Arize AI. Available at: <https://arize.com/blog-course/mean-absolute-error-in-machine-learning-what-you-need-to-know/> (Accessed: 17 December 2025).
4. Hyndman, R.J. and Athanasopoulos, G. (2021) 'Evaluating point forecast accuracy (Section 5.8)', in *Forecasting: Principles and Practice*. 3rd edn. Melbourne: OTexts. Available at: <https://otexts.com/fpp3/accuracy.html> (Accessed: 20 December 2025).
5. Barocas, S., Guo, A., Kamar, E., Krones, J., Morris, M.R., Vaughan, J.W., Wadsworth, D. and Wallach, H. (2021) *Designing disaggregated evaluations of AI systems: choices, considerations, and tradeoffs*. arXiv. Available at: <https://arxiv.org/abs/2103.06076> (Accessed: 19 December 2025).
6. Bynum, L., Loftus, J.R. and Stoyanovich, J. (2021) *Disaggregated interventions to reduce inequality*. arXiv. doi: <https://doi.org/10.1145/3465416.3483286> (Accessed: 19 December 2025).
7. Ambler, W. (2024) *How data analysis transforms F1 race performance*. Catapult. Available at: <https://www.catapult.com/blog/f1-data-analysis-transforming-performance> (Accessed: 19 December 2025).
8. European Union (2025) *The EU Artificial Intelligence Act*. Available at: <https://artificialintelligenceact.eu/> (Accessed: 29 December 2025).
9. Logan, N. (2020) *Why are there no British riders in MotoGP?* Blackcircles.com. Available at: <https://www.blackcircles.com/news/moto-gp-no-british-riders> (Accessed: 29 December 2025).
10. Gray, A. (2024) *World Rally Championship: Rally Ireland bid falls through over funding*. BBC Sport, 16 April. Available at: <https://www.bbc.co.uk/sport/motorsport/68828597> (Accessed: 29 December 2025).