AI-Assisted Substitution Decisions: A Fuzzy Logic Approach to Real-Time Game Management

Pedro Passos Farias[1]

[1]Institute of Computing – Universidade Federal Fluminense (UFF)

Av. Gal. Milton Tavares de Souza – 24210-346 – Niterói – RJ – Brazil

pedropassos@id.uff.br

October 2025

Abstract

With millions on the line every match, top soccer clubs still make critical substitution decisions based largely on intuition. This paper introduces an AI-powered Decision Support System (DSS) that brings data-driven rigor to one of the game's most crucial tactical moments. Using fuzzy logic to model expert coaching knowledge, our system provides real-time substitution priorities by integrating validated performance metrics (`playerankScore`), fatigue (`minutesPlayed`), age, and disciplinary risk (`TemCartaoAmarelo`). A key innovation is its contextual logic, which modulates disciplinary risk based on a player's tactical position (`roleCluster`), reflecting deeper tactical awareness. Validation through case studies confirms the system's ability to balance conflicting factors and escalate priority in high-risk scenarios, providing a tangible competitive advantage for real-time game management when every decision counts.

1 Introduction

The financial stakes in modern soccer have never been higher, creating a vast economic disparity between clubs. In the 2023/24 season, Manchester City earned £175.9 million in prize money for winning the Premier League title ge (2025). In contrast, Fluminense, a Brazilian club with more modest economic power, received a total of US$50.71 million for reaching the semifinals of the 2025 Club World Cup—a sum that could have escalated to US$40 million for winning the finalESPN.com.br (s.d.).

Despite these immense financial consequences, one of the most impactful in-game decisions: player substitutions, often remains guided by intuition and limited visible cues. As top coaches like Guardiola acknowledge, the timing and choice of a substitution can be the difference between winning titles and missing crucial opportunities.

This gap is particularly evident because tactical decision-making during a soccer match is one of the most complex and high-impact processes in team management. Substitutions, changes in tactical posture, or defensive adjustments are frequently made by coaches based on a combination of empirical knowledge and a multifactorial assessment of inherently imprecise variables. Factors like "the player looks tired," "performance has recently dropped," or "the risk of expulsion is high" are not binary, but rather gradual. Performance evaluation is, in itself, a notable challenge. Although the availability of large-scale event data has grown, there is no single, universally accepted metric that captures all facets of a player's performance Pappalardo et al. (2019).

This work addresses the problem of tactical decision-making under uncertainty through the implementation of a fuzzy logic control system. Fuzzy logic is an artificial intelligence tool particularly suited for this domain, as it allows for modeling "expert knowledge" (a coach's rules) and handling vague linguistic concepts, such as "high age," "low performance," or "medium fatigue" (Marliere, 2017). The objective of this paper is to design, implement, and validate a decision support system (DSS) focused specifically on player substitution priority. The system uses four inputs (performance, fatigue, age, and disciplinary risk) to generate a fuzzy output (Substitution Priority), providing a quantitative tool to assist the coaching staff with one of their most critical real-time decisions."

2 Literature Review

The academic literature on data analysis in soccer provides the theoretical foundation for this work, which lies at the intersection of performance evaluation and tactical decision modeling. The first methodological challenge is to quantify a player's performance. The PlayeRank framework, proposed by Pappalardo Pappalardo et al. (2019), addresses this problem by defining a multidimensional and role-aware evaluation metric, academically validated against professional scout assessments. Given its robustness, we adopt the `playerankScore` as the basis for our main performance input, avoiding the subjectivity of creating a new metric. The second challenge is how to interpret imprecise data for decision-making. The literature justifies the use of fuzzy logic as the central methodological tool for this purpose. Works such as that by Huarachi-Macuri et al. LACCEI (2023) demonstrate the effectiveness of fuzzy systems for evaluating player performance by modeling imprecise linguistic concepts like "Stamina" and "Agility". This approach validates the "what" of our system: the transformation of quantitative metrics (like `minutesPlayed` and `IdadeDoJogador`) into the fuzzy concept of *Fatigue*. Similarly, the tactical decision support literature, such as Marliere's work Marliere (2017) in the context of robot soccer, uses fuzzy control systems to arbitrate between tactical actions (e.g., "heavy defense," "light attack"). This approach validates the "why" of our system: the generation of a decision output, in our case, the *Substitution Priority*. Despite these foundations, a gap was identified in the existing literature. While previous works focus on evaluating player performancePappalardo et al. (2019); LACCEI (2023), or modeling general tactical aspects Marliere (2017), there is still no decision support system that integrates validated performance metrics (like *PlayeRank*) with other critical game factors, such as fatigue and disciplinary risk, to assist in the specific tactical decision of player substitution. This work proposes to fill this gap by uniting performance evaluation approaches with a decision engine based on fuzzy logic, with the goal of developing a system aimed at supporting substitutions. .

3 Dataset

The database selected for this project is the *"Soccer match event dataset"*, a detailed public repository of soccer match events (Pappalardo, 2020). The choice of this dataset is based on its high granularity. The *dataset* details player-level actions and aggregated performance metrics, allowing for the modeling of *in-game* situations, which is an essential requirement for our system. We used the complete *dataset* available on Kaggle (Pappalardo, 2020), which comprises 27 interrelated CSV tables. They cover seven main competitions and can be grouped into five logical categories:

- **Match Data (`matches_*.csv`):** Contains information about the games, such as dates, lineups, substitutions, results, and tactical formations.
- **Event Data (`events_*.csv`):** The core of the *dataset*, recording millions of individual actions on the field.
- **Entity Data (`players.csv`, `teams.csv`, etc.):** Dimensional tables with demographic and static data.
- **Performance Metrics (`playerank.csv`):** A pre-processed file that provides the `playerankScore`, a multidimensional and role-aware performance evaluation metric.
- **Dictionaries (`tags2name.csv`, `eventid2name.csv`):** Metadata that translate event and *tag* IDs into readable descriptions (e.g., Tag `1702` = `'yellow_card'`).

The choice of `playerankScore` as our main "Performance" input is a central methodological decision. As proposed by PappalardoPappalardo et al. (2019), the PlayeRank *framework* was developed to solve the absence of a consolidated and universally accepted metric for evaluating player performance. The `playerankScore` is a metric derived from millions of game events that, according to the authors, surpasses other metrics when compared with assessments from professional scouts. Therefore, instead of trying to model performance from raw events, we adopted the `playerankScore` as an already validated and academically robust representation of a player's performance in a match.

Dataset available at:

https://www.kaggle.com/datasets/aleespinosa/soccer-match-event-dataset/data?select=playerank.csv

4 Data Integration and Pre-processing

The goal of pre-processing was to consolidate this ecosystem of 16 relevant files (7 from matches, 7 from events, `players.csv`, and `playerank.csv`) into a single structured *dataset*. Each row of the final *dataset* represents a single "player-match" instance. The integration process followed three main steps:

1. **Base and Main Variables:** We used `playerank.csv` as the main table, providing the inputs for `playerankScore` (Performance) and `minutesPlayed` (Fatigue), in addition to `roleCluster` (Player position).

2. **Creation of the Age Feature:** To create the `IdadeDoJogador` variable, we performed a data join. The date of birth (`birthDate`) was obtained from `players.csv` and the match date (`date`) was extracted from the concatenation of the seven `matches_*.csv` files. The age was then calculated in fractional years at the exact moment of the match.

3. **Creation of the Risk Feature:** For the `TemCartaoAmarelo` variable, we processed the concatenation of the seven `events_*.csv` files (totaling millions of events). We filtered all events that contained the *tag* `1702` (identified via `tags2name.csv` as `'yellow_card'`) and created a boolean indicator for each (`matchId`, `playerId`) pair that received a card.

Rows where crucial information (like `birthDate` or `date`) was missing were removed. This process resulted in a final clean and validated dataset of **46,897 instances**. Table 1 summarizes the final attributes used as inputs for the fuzzy system.

Table 1: Attributes of the Consolidated Dataset for the Fuzzy System

	Attribute	Description	Data type
Identifiers	`matchId`, `playerId`	Identification keys.	Integer
Context (Rules)	`roleCluster`	Player's tactical position.	Categorical
Fuzzy Input 1	`playerankScore`	Performance metric.	Real
Fuzzy Input 2	`minutesPlayed`	Total minutes played.	Integer
Fuzzy Input 3	`IdadeDoJogador`	Player's age on the match day.	Real
Fuzzy Input 4	`TemCartaoAmarelo`	(1) if received card, (0) otherwise.	Integer

5 Exploratory Data Analysis

The exploratory analysis was conducted on the final dataset of 46,897 instances. The objective was to define the **universes of discourse** (ranges of values) and to quantitatively substantiate the **design of the membership functions** for the control system.

5.1 Univariate Analysis: Defining the Universes

The analysis of the four input variables (Figure **??**) revealed their limits:

- **Input 1:** `playerankScore` **(Performance).** As justified in Section 3, this metric is a multidimensional performance evaluation Pappalardo et al. (2019). Although the theoretical *framework* might describe the score on a normalized scale (e.g., 0 to 1) for presentation purposes, the raw *dataset* used contains the *scores* standardized around zero. Our exploratory analysis confirmed this: the observed universe of discourse ranges from -0.119 (very low performance) to 0.173 (elite performance), with the mean (0.007) and median (0.004) very close to zero. This characteristic is ideal for fuzzy logic, as zero represents an "average performance," allowing for the creation of *Low* (negative), *Medium* (near zero), and *High* (positive) sets.

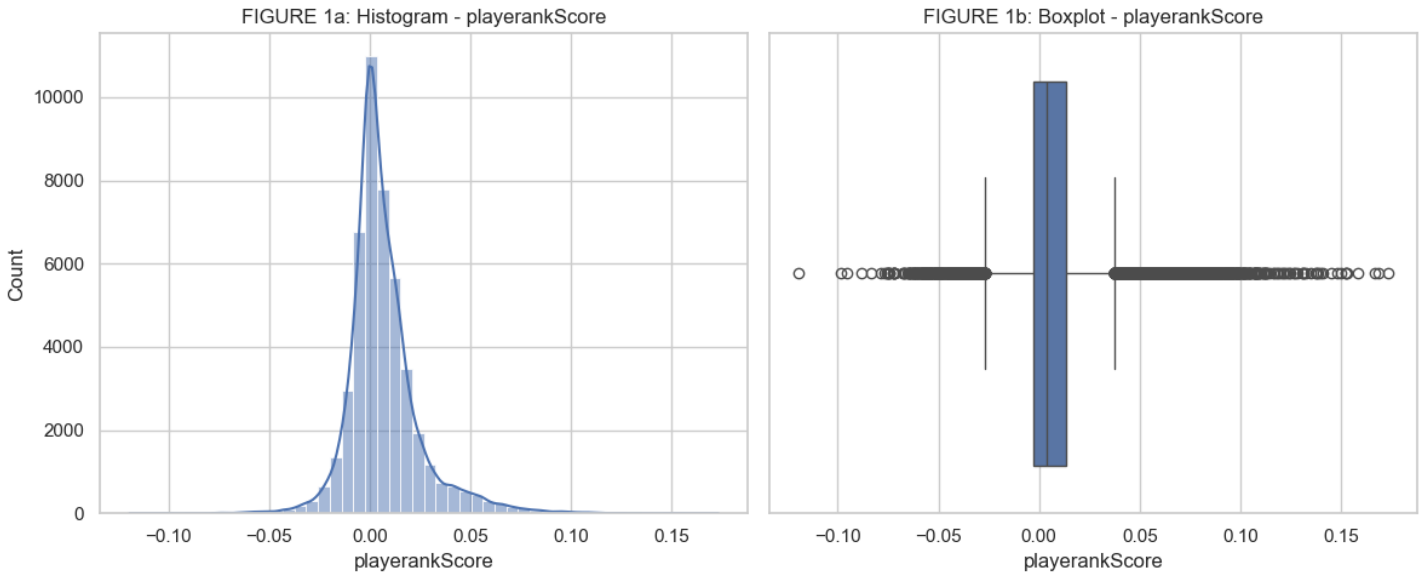

- **Input 2: `minutesPlayed` (Fatigue).** The fatigue universe ranges from 1 to 120 minutes (extra time). The distribution shows strong negative skewness (*skewness* of -1.40), with the median at 90 minutes. In fact, 59.3% of the instances (n=27,828) represent players who played the full 90 minutes.

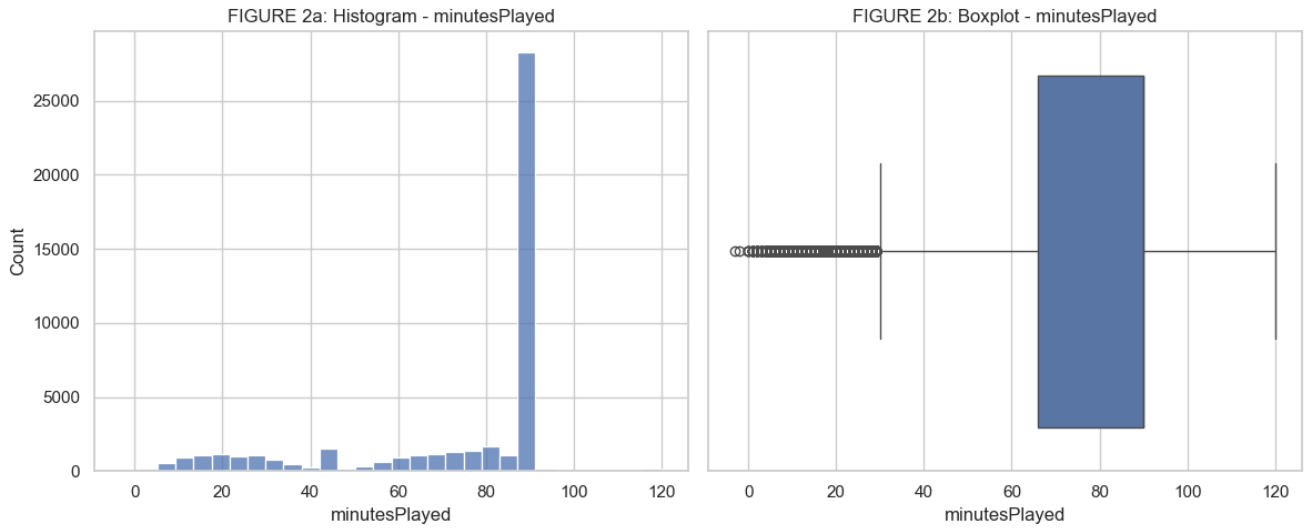

- **Input 3: `IdadeDoJogador` (Age).** Age presents a normal distribution, with the mean (26.68 years) and median (26.51 years) being almost identical. The universe is defined between 15.89 and 41.37 years (oldest age recorded in the dataset). This symmetrical distribution supports the creation of three balanced fuzzy sets: *Young*,

Peak (centered on the mean), and *Veteran*.

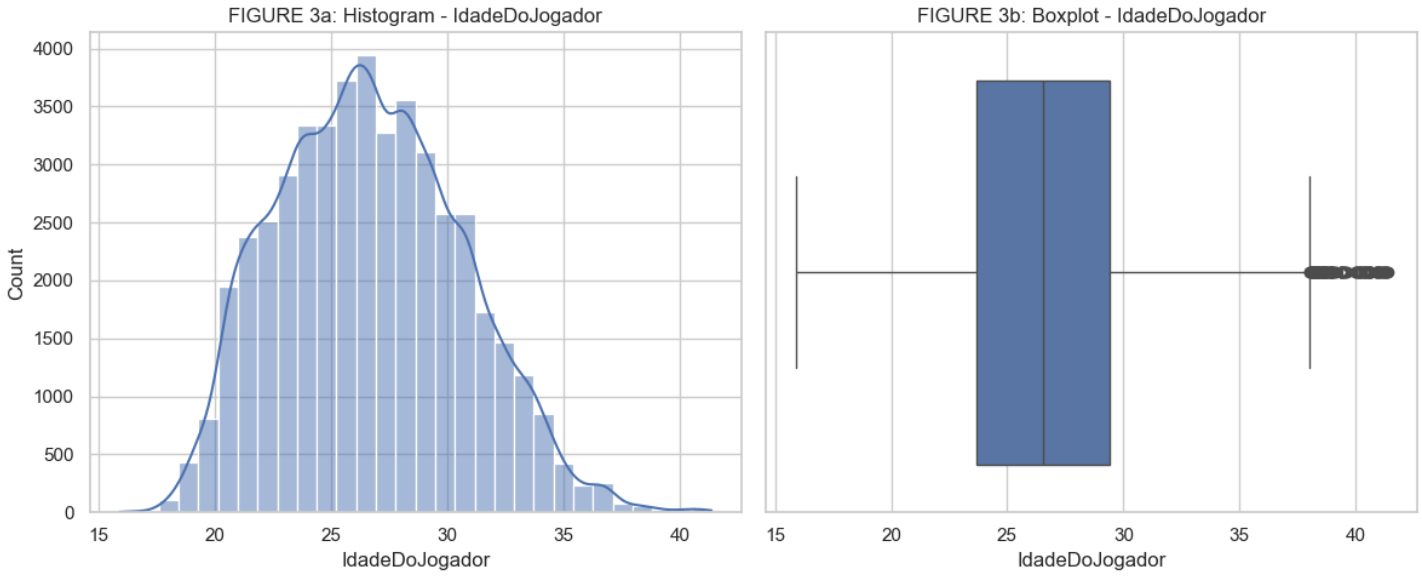

- **Input 4:** `TemCartaoAmarelo` **(Risk).** This is a binary input. The analysis revealed that **15.34%** of all player-match instances (n=7,194) involve a yellow card, confirming its importance as a risk factor.

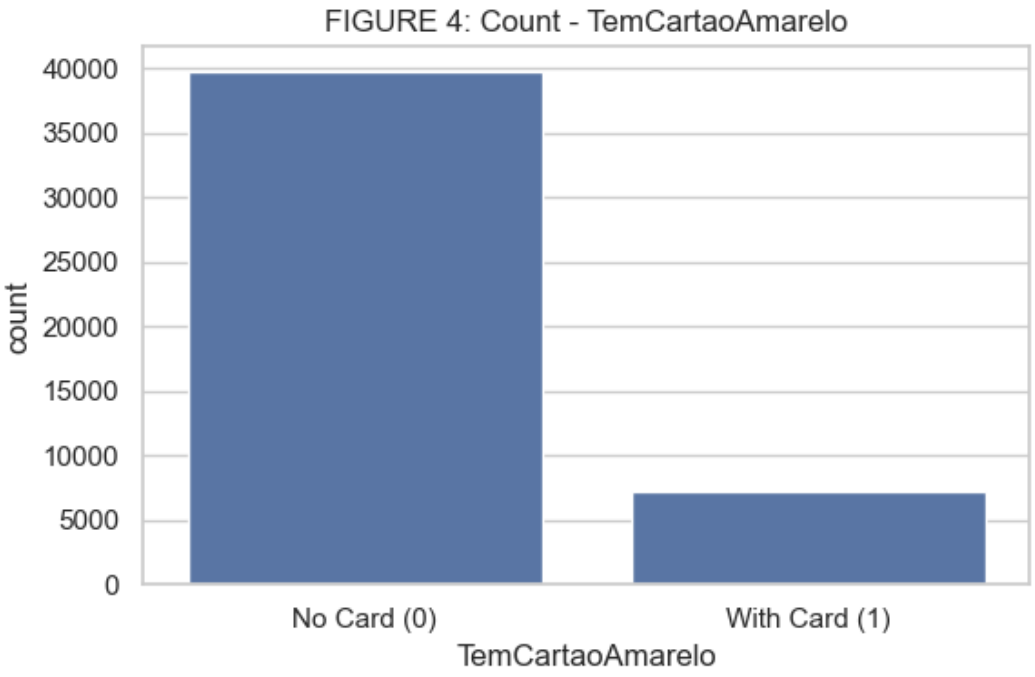

5.2 Bivariate and Multivariate Analysis: Validating the Rules

This analysis validates the common-sense soccer assumptions that we will use to build the system's rules.

5.2.1 Validation of the Context Rule: Risk vs. Position

The central premise is that the risk of a yellow card is position-dependent . The Figure below validates this hypothesis. Defensive positions (`right CB` 19.1%; `left CB` 18.5%) and central midfielders (`central MF` 17.3%) show a statistically higher risk. In contrast, attacking positions (`central FW` 11.1%; `right FW` 10.3%) show a significantly lower risk. This finding justifies the creation of **contextual rules**, where the `roleCluster` variable will modulate the importance of the `TemCartaoAmarelo` input.

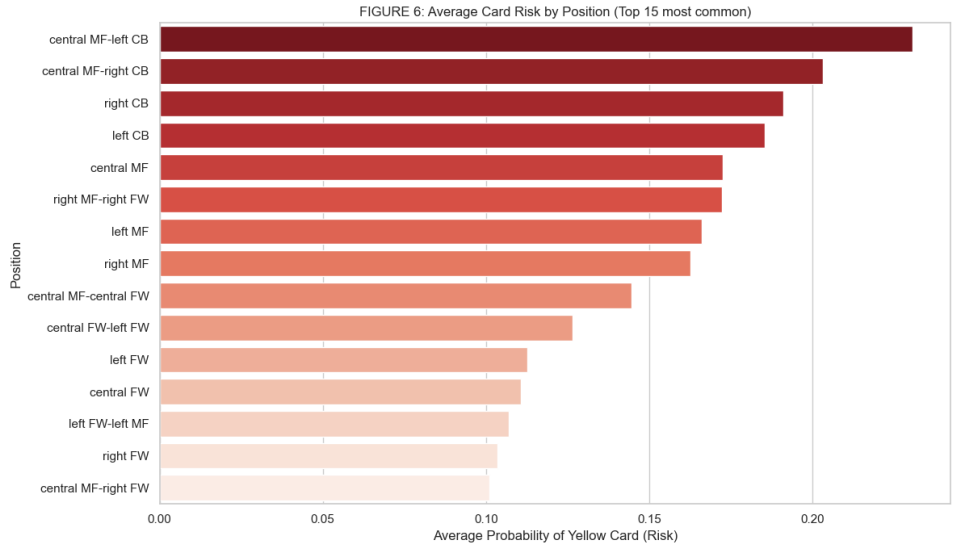

5.2.2 Independence and Relationships of Inputs

The Spearman correlation matrix (Figure 5.2.2) reveals that the four input variables are, for the most part, independent. All correlations are weak (below |0.2|). The strongest negative correlation is between `TemCartaoAmarelo` and `playerankScore` (-0.179). This suggests that receiving a card is associated with a drop in performance. Table 2 reinforces this: for the most common positions, the average `playerankScore` is consistently lower for players with a card.

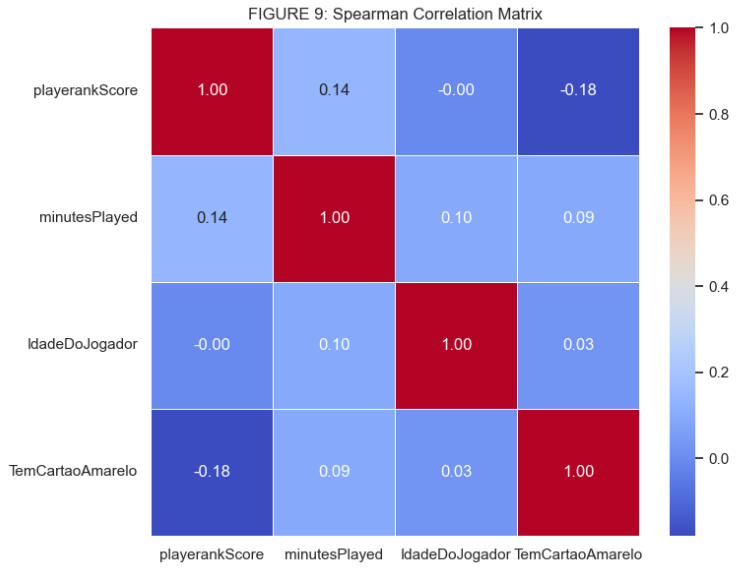

Table 2: Average Performance (`playerankScore`) by Position and Risk.

Position (roleCluster)	Score (No Card)	Score (With Card)
`left CB`	0.0112	0.0055
`right CB`	0.0108	0.0046
`central MF`	0.0070	-0.0007
`right MF`	0.0030	-0.0039
`left MF`	0.0018	-0.0050
`central FW`	0.0119	0.0063

5.3 Analysis Conclusion

The exploratory data analysis allowed for a detailed characterization of the final dataset of 46,897 instances, establishing a solid foundation for understanding the variables before modeling. The univariate analysis (Section 5.1) defined the practical limits (universes of discourse) for the numerical inputs: `playerankScore` [-0.119, 0.173], `minutesPlayed` [1, 120], and `IdadeDoJogador` [15.9, 41.37]. It also quantified their distributional characteristics, notably the normality of age (mean 26.7), the strong concentration of minutes played at 90 (median 90), and the peaked, right-skewed distribution of performance (median 0.004). The prevalence of the binary variable `TemCartaoAmarelo` was established at 15.34% . Subsequently, the bivariate and multivariate analysis (Section 5.2) confirmed two central hypotheses: the statistical dependence of disciplinary risk on tactical position (`roleCluster`), with defenders showing a higher propensity for cards, and the low mutual correlation among the four input variables (all $< |0.18|$ in the correlation matrix),

indicating their informational independence. With these characteristics quantified and relationships validated, the data exploration is complete.

6 Fuzzy Control System Design

Based on the theoretical foundation (Section 2) and the quantitative data analysis (Section 5), we designed the Fuzzy Control System (FCS) to evaluate *Substitution Priority*. The system was implemented using the `scikit-fuzzy` library in Python, following the standard Mamdani architecture for fuzzy inference and the Centroid method for defuzzification.

6.1 Definition of Fuzzy Variables

The system comprises four input variables (antecedents) and one output variable (consequent). The universes of discourse were defined from the minimum and maximum values observed in the Exploratory Analysis (Section 5.1), with a 5% margin to ensure the model can accommodate values more extraordinary than those recorded in the dataset.

6.1.1 Antecedents (Inputs)

1. **Performance (`desempenho`):** Based on `playerankScore`.
    - Universe of Discourse: [-0.134, 0.188].
    - Fuzzy Sets: *Low*, *Medium*, *High*.
2. **Fatigue (`fadiga`):** Based on `minutesPlayed`.
    - Universe of Discourse: [0, 126] minutes.
    - Fuzzy Sets: *Low*, *Medium*, *High*.
3. **Age (`idade`):** Based on `IdadeDoJogador`.
    - Universe of Discourse: [14.6, 42.7] years.
    - Fuzzy Sets: *Young*, *Peak*, *Veteran*.

4. **Risk (`risco`):** Based on `TemCartaoAmarelo`.

    - Universe of Discourse: [0, 1].
    - Fuzzy Sets: *NoRisk*, *WithRisk*.

6.1.2 Consequent (Output)

- **Substitution Priority (`prioridade`):** Indicates the calculated urgency .
    - Universe of Discourse: [0, 10].
    - Fuzzy Sets: *Low*, *Medium*, *High*.

6.2 Membership Functions

The membership functions were designed using triangular (`trimf`) and trapezoidal (`trapmf`) functions for the continuous variables, and narrow Gaussian functions (`gaussmf`) for the binary variable `Risk`, aiming for interpretability and reflecting the distributions observed in the EDA (Section 5.1). The exact parameters used in the Python implementation are:

- **Performance:** *Low* = `trapmf`([-0.134, -0.134, -0.03, 0.004]); *Medium* = `trimf`([-0.03, 0.004, 0.04]); *High* = `trapmf`([0.004, 0.04, 0.188, 0.188]). The narrow *Medium* set reflects the high kurtosis of the distribution.
- **Fatigue:** *Low* = `trapmf`([0, 0, 40, 70]); *Medium* = `trimf`([40, 70, 95]); *High* = `trapmf`([70, 95, 126, 126]). The *High* set is concentrated after 70 minutes, reflecting the median at 90.
- **Age:** *Young* = `trapmf`([14.6, 14.6, 22, 27]); *Peak* = `trimf`([22, 27, 32]); *Veteran* = `trapmf`([27, 32, 42.7, 42.7]). The sets are approximately symmetrical around the mean of 26.7 years.
- **Risk:** *NoRisk* = `gaussmf`(mean=0, sigma=0.1); *WithRisk* = `gaussmf`(mean=1, sigma=0.1). Narrow Gaussians provide a smooth transition and robustness if the input is not exactly 0 or 1.

- **Priority (Output):** *Low* = `trimf`([0, 0, 5]); *Medium* = `trimf`([0, 5, 10]); *High* = `trimf`([5, 10, 10]). Overlapping triangular functions to allow interpolation in the output.

6.3 Rule Base and Contextual Logic

The knowledge base was built by combining domain knowledge inspired mainly by the statements of coach Josep Guardiola (ESPN Brasil, s.d.; Terra Esportes, s.d.; CNN Brasil Esportes, s.d.), with the validations from the EDA. A set of 17 base rules was defined (see Table 3 for examples), using the `AND` (minimum) operator to combine antecedents.

Table 3: Complete Set of Base Fuzzy Rules.

ID	Rule
A1	IF (*Performance* is Low AND *Fatigue* is High) THEN (*Priority* is High)
A2	IF (*Performance* is Low AND *Age* is Veteran AND *Fatigue* is Medium) THEN (*Priority* is High)
A3	IF (*Fatigue* is High AND *Age* is Veteran) THEN (*Priority* is High)
A4	IF (*Risk* is WithRisk AND *Fatigue* is High) THEN (*Priority* is High)
A5	IF (*Risk* is WithRisk AND *Performance* is Low) THEN (*Priority* is High)
B1	IF (*Performance* is High AND *Fatigue* is Low AND *Risk* is NoRisk) THEN (*Priority* is Low)
B2	IF (*Performance* is High AND *Age* is Young AND *Risk* is NoRisk) THEN (*Priority* is Low)
B3	IF (*Performance* is Medium AND *Fatigue* is Medium AND *Age* is Peak AND *Risk* is NoRisk) THEN (*Priority* is Low)
B4	IF (*Performance* is High AND *Fatigue* is Medium AND *Risk* is NoRisk) THEN (*Priority* is Low)
M1	IF (*Performance* is Low AND *Fatigue* is Medium) THEN (*Priority* is Medium)
M2	IF (*Performance* is Medium AND *Fatigue* is High AND *Risk* is NoRisk) THEN (*Priority* is Medium)
M3	IF (*Age* is Veteran AND *Fatigue* is Medium AND *Risk* is NoRisk) THEN (*Priority* is Medium)
M4	IF (*Performance* is High AND *Fatigue* is High AND *Risk* is NoRisk) THEN (*Priority* is Medium)
R1	IF (*Risk* is WithRisk AND (*Fatigue* is Medium OR *Fatigue* is High)) THEN (*Priority* is High)
R2	IF (*Risk* is WithRisk AND *Fatigue* is Low) THEN (*Priority* is Medium)
R3	IF (*Risk* is WithRisk AND *Age* is Veteran) THEN (*Priority* is Medium)
R4	IF (*Risk* is WithRisk AND *Performance* is Medium) THEN (*Priority* is Medium)

However, the EDA (Section 5.2) demonstrated the crucial importance of the categorical variable `roleCluster` (Position) in modulating `Risk`. As this variable is not fuzzy, a **post-inference contextual logic** was implemented. The system first calculates a `base_priority` using the 17 fuzzy rules. Then, an external function ('calcular_prioridade_contextual' in the Python code) evaluates the player's position and risk level:

- If the player is a **Defender** ('CB') and the risk is high, the final priority is adjusted to a high minimum value (e.g., 8.0).
- If the player is a **Midfielder** ('MF'), has high risk, and is tired (Medium or High Fatigue), the final priority is adjusted to a medium-high minimum value (e.g., 7.5) .
- If the player is a **Forward** ('FW') with high risk, the base priority is generally maintained (considered Medium), but may be limited if the performance is high.

This hybrid approach allows combining the flexibility of fuzzy inference with "crisp" contextual rules essential for the domain.

6.4 Inference and Defuzzification

The system uses the Mamdani inference method, where the result of each rule's activation is aggregated (using the 'OR' = maximum operator) to form a resulting membership function for the output variable `Priority`. The conversion of this aggregated fuzzy function into a final numerical (crisp) value is performed by the **Centroid** method, which calculates the center of mass of the area under the curve of the resulting membership function. This value, on a scale of 0 to 10, represents the *Substitution Priority* calculated by the system.

Codes available at `https://github.com/Pedro-Passos77/Sistema-de-Apoio-a-Decisao-Baseado-em-Logica-Nebulosa`

7 Results

In this section, we evaluate the behavior of the Fuzzy Control System (FCS). The goal of the validation is not to measure predictive accuracy against actual substitutions (as we do not have a *ground truth* for "ideal priority"), but rather to verify if the system produces outputs that are logically consistent with the inputs, the fuzzy rules, and the implemented contextual logic. To do this, we applied the FCS to four representative scenarios, covering different combinations of performance, fatigue, age, risk, and tactical position. The system's output, *Substitution Priority*, is a numerical value on a [0, 10] scale, interpreted as Low (approx. 0-3.5), Medium (approx. 3.5-6.5), and High (approx. 6.5-10).

7.1 Analysis of Typical Scenarios

The following four scenarios were simulated using the `calcular_prioridade_contextual` function, which integrates the base fuzzy inference with contextual position adjustments.

7.1.1 Scenario 1: Solid Performance, Low Risk

In this scenario, we evaluate a young player (21 years old, *Age* = Young/Peak), playing as a midfielder ('central MF'), with good performance (`playerankScore` = 0.05, *Performance* = High), low time on field (30 min, *Fatigue* = Low), and no yellow card (`TemCartaoAmarelo` = 0, *Risk* = NoRisk).

Result: The system calculated a **Priority = 1.67**. **Discussion:** This value, clearly in the *Low* range, is aligned with expectations. The low priority rules (like B1 and B2) were activated by the combination of favorable performance, low fatigue, and absence of risk. The system correctly identifies that there is no urgency to substitute a player in these conditions, reflecting the philosophy of keeping players who are contributing positively and without restrictions.

7.1.2 Scenario 2: Veteran Defender Under Risk

We evaluate a defender ('right CB'), veteran (34 years old, *Age* = Veteran), with considerable time on field (85 min, *Fatigue* = High), average performance (`playerankScore` = 0.001, *Performance* = Medium), and with a yellow card (`TemCartaoAmarelo` = 1, *Risk* = WithRisk).

Result: The system calculated a **Priority = 8.00**. **Discussion:** This *High Priority* result demonstrates the effectiveness of the contextual logic. Although the base rules (like A3, A4, R1, R3) were already activated by the combination of high fatigue, advanced age, and risk, the contextual function, by identifying that the player is a defender ('CB') with risk, applied a minimum adjustment, raising the final priority to 8.00. This correctly reflects the high tactical risk associated with a yellow-carded and potentially tired defender at the end of the game, a scenario where substitution is often considered a priority.

7.1.3 Scenario 3: Yellow-carded and Exhausted Midfielder

We consider a midfielder ('left MF') at peak age (29 years old, *Age* = Peak), with average performance (`playerankScore` = 0.0, *Performance* = Medium), but with very high playing time (95 min, *Fatigue* = High) and with a yellow card (`TemCartaoAmarelo` = 1, *Risk* = WithRisk).

Result: The system calculated a **Priority = 7.50**. **Discussion:** Again, a *High Priority* value. The base rules A4 and R1 (Risk + High/Medium Fatigue) were strongly activated. The contextual logic for midfielders ('MF') with risk and fatigue also contributed, ensuring a minimum priority of 7.50. Although slightly lower than the defender scenario (perhaps reflecting a marginally lower tactical risk), the system correctly identifies this as a high-urgency situation for substitution, due to the combination of physical exhaustion and disciplinary risk in a central field position.

7.1.4 Scenario 4: Yellow-carded, Tired, but Decisive Forward

This scenario explores the trade-off between risk/fatigue and high performance. We evaluate a forward ('central FW') at peak age (25 years old, *Age* = Peak), with excellent

performance (`playerankScore` = 0.08, *Performance* = High), high playing time (80 min, *Fatigue* = Medium/High), and with a yellow card (`TemCartaoAmarelo` = 1, *Risk* = With-Risk).

Result: The system calculated a **Priority = 6.00**. **Discussion:** The result (in the *Medium-High* range) captures the complexity of the decision. The rules related to risk (R1, R3, R4) and fatigue push the priority up. However, rules B4 (High Performance + Medium Fatigue) and M4 (High Performance + High Fatigue) exert an opposing force, reflecting the principle of keeping a decisive player on the field. The contextual logic for forwards ('FW') with risk and good performance limited the final priority to 6.00, suggesting that substitution should be considered, but is not as critical as in scenarios 2 and 3, given the player's positive impact on the game. The system demonstrates the ability to balance conflicting factors.

7.2 Aggregated Analysis of Results

To evaluate the overall behavior of the system, the SCN was applied to all 46,897 instances of the *dataset*, generating a *benchmark dataset* with the `CalculatedPriority`. The distribution of this variable provides an initial view of the model's recommendation profile. The histogram shows that most results fall within the range of 4 to 6, indicating a predominance of priorities classified as *Medium*. This pattern is partially influenced by the high incidence of players who played 90 minutes or more (59.3%, see Section 5.1), frequently activating rules associated with *High Fatigue* (e.g., A1, A3, A4, M2, M4, and R1), which naturally increase the priority score. However, the distribution does not concentrate in the *High* range, suggesting that other factors moderate the priority in most instances, such as median performance (`playerankScore` median of 0.004), peak-age players (median age 26.5 years), and low disciplinary risk (only 15.34% with a card).

A more detailed analysis compares the distribution of `CalculatedPriority` between players who were substituted ('FoiSubstituido = 1', n=10,977) and those who remained on the field ('FoiSubstituido = 0', n=35,920). Descriptive statistics are presented in Table 4. Contrary to initial expectations, both the mean (4.48 vs 4.68) and median (4.74 vs 5.00)

are slightly lower for substituted players. This result reflects the influence of fatigue, the contextual nature of the metric, and the complexity of substitution decisions, which are often driven by tactical factors (strategy changes, opponent response, score management) not captured by the SCN's individual-condition variables.

It is important to note that this analysis considers the players' *final state*, rather than an *in-game state*. The 'Not Substituted' group mainly consists of players who completed 90 minutes or more, maximizing the `minutesPlayed` input and thus receiving high fatigue scores. In contrast, the 'Substituted' group often left the game earlier (e.g., at 70 minutes). Therefore, a direct comparison of final scores may be distorted by differences in minutes played. An *in-game* analysis, performed at the exact moment of substitution, would be necessary to evaluate predictive accuracy, but this is outside the scope of the current validation.

Table 4: Comparative Statistics of Calculated Priority by Substitution Status.

Statistic	**Not Substituted (0)**	**Substituted (1)**
Count (n)	35,920	10,977
Mean	4.6805	4.4765
Standard Deviation (std)	2.1235	1.4758
Minimum	0.0	0.0
25th Percentile	4.7174	3.8105
Median (50th Percentile)	5.0035	4.7426
75th Percentile	5.7324	5.1170
Maximum	8.2778	8.3333

Analyzing data dispersion reveals an important insight. The standard deviation of `CalculatedPriority` is lower for substituted players (1.48) compared to non-substituted players (2.12). This lower variability suggests that when substitutions are driven by individual conditions, the SCN identifies these situations consistently, producing scores concentrated around the medium range. In contrast, the higher dispersion among non-substituted players, including some with relatively high priority scores who remained on the field, indicates that other contextual and tactical factors may have influenced the coach's final decision, even when the model signaled an individual condition justifying a substitution.

In summary, the lower variability in substituted players' scores and the higher disper-

sion in non-substituted players reinforce the interpretation that coaches prioritize tactical schemes and game-context factors over isolated individual metrics. This pattern highlights the importance of analyzing SCN results within the game context, rather than solely as absolute priority values.

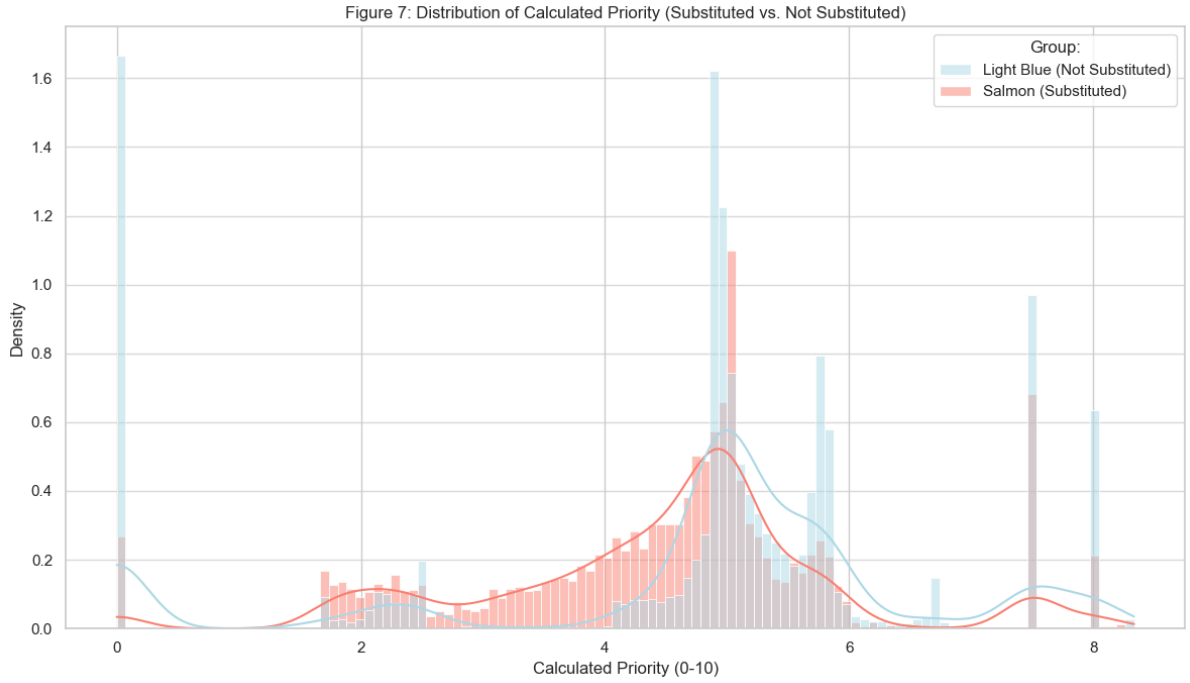

Figure 1: Distribution of Calculated Substitution Priority for all 46,897 instances in the dataset.

8 Conclusion

This work demonstrated the feasibility of applying fuzzy logic to create a Decision Support System (DSS) focused on the substitution priority of soccer players. The designed Fuzzy Control System (FCS) successfully integrated academically validated performance metrics (`playerankScore` Pappalardo et al. (2019)) with indicators of physical condition (`minutesPlayed`, `IdadeDoJogador`) and disciplinary risk (`TemCartaoAmarelo`), using a rule base inspired by expert knowledge ("Pep" Guardiola) and validated by Exploratory Data Analysis. The validation results indicate that the system behaves in a logically consistent manner. The aggregate analysis and case studies demonstrated the system's ability to differentiate urgency levels, identify and reflect a conservative tactical philosophy of

coaches, and, crucially, implement the contextual logic that adjusts the impact of risk based on the player's position. The FCS thus provides a quantitative and interpretable index of the player's individual condition, offering a valuable starting point to assist coaching staff in the complex task of managing fatigue and risk during the match. The scope of this project focused on evaluating the player's condition using aggregated data common in event datasets. **Promising future extensions** could refine and expand this base. The integration of more granular data, such as Global Positioning System metrics for a more direct assessment of fatigue instead of the `minutesPlayed` *proxy*, could increase the accuracy of the physical assessment. Similarly, incorporating game context variables (score, remaining time, current strategy) would allow the system to provide recommendations even more aligned with immediate tactical needs. The calibration of the membership functions and, especially, the rule base through iterative feedback with coaches and analysts represents a fundamental step to **maximize practical applicability** and trust in this tool in the professional environment. The complete source code of the FCS implementation in Python with `scikit-fuzzy`, including the pre-processing, exploratory analysis, and validation scripts, is publicly available in the following repository: `https://github.com/Pedro-Passos77/Sistema-de-Apoio-a-Decisao-Baseado-em-Logica-Nebulosa` . The generated *benchmark dataset* can also be found in the same location, promoting the reproducibility of the research. We thank Luca Pappalardo and collaborators for making the *Soccer match event dataset* publicly available on Figshare (Pappalardo, 2020), which was essential for conducting this study. I highlight that generative artificial intelligence tools were used as an assistant in the writing and revision of parts of this article, always under full human supervision to ensure the accuracy and appropriateness of the final content.

‘‘‘

References

CNN Brasil Esportes. Guardiola é o técnico que menos faz substituições na premier league; veja ranking [guardiola is the coach who makes the fewest substitutions in the premier league; see ranking]. `https://www.cnnbrasil.com.br/esportes/futebol/futebol-`

internacional/guardiola-e-o-tecnico-que-menos-faz-substituicoes-na-premier-league-veja-ranking/, s.d. Accessed on: 24 Oct. 2025.

ESPN Brasil. Por que guardiola não fez nenhuma substituição no city contra o rb leipzig? ele mesmo responde [why didn't guardiola make any substitutions for city against rb leipzig? he answers]. https://www.espn.com.br/futebol/manchester-city/artigo/_/id/11661194/por-que-guardiola-nao-fez-nenhuma-substituicao-no-city-contra-o-rb-leipzig-ele-mesmo-responde, s.d. Accessed on: 24 Oct. 2025.

ESPN.com.br. Quanto o fluminense ganhou em premiação no mundial de clubes após eliminação para o chelsea na semifinal. https://www.espn.com.br/futebol/fluminense/artigo/_/id/15405687/quanto-fluminense-ganhou-premiacao-mundial-de-clubes-apos-eliminacao-para-o-chelsea-semifinal, s.d. Acesso em: 26 out. 2025.

ge. Premiação da premier league: veja quanto cada clube recebeu em 2023/24. https://ge.globo.com/futebol/futebol-internacional/noticia/2025/02/07/premiacao-da-premier-league-veja-quanto-cada-clube-recebeu-em-202324.ghtml, 2025. Acesso em: 26 out. 2025.

LACCEI. A system for the control of the performance of high level soccer players applying fuzzy logic. In *2023 LACCEI International Multi-Conference for Engineering, Education, and Technology*, 2023.

F. T. Marliere. Sistema de apoio à decisão baseado na lógica fuzzy e aplicado ao futebol de robôs [decision support system based on fuzzy logic and applied to robot soccer]. https://www2.ufjf.br/eletrica_automacao/wp-content/uploads/sites/647/2017/02/TCC_Frederick-Tavares-Marliere.pdf, 2017.

L. Pappalardo. Soccer match event dataset. https://doi.org/10.6084/m9.figshare.c.4415000.v5, 2020. [Data set].

L. Pappalardo, P. Cintia, A. Rossi, E. Massucco, P. Ferragina, D. Pedreschi, and F. Giannotti. Playerank: Data-driven performance evaluation and player ranking in soccer via

a machine learning approach. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(5):1–24, 2019.

Terra Esportes. Guardiola explica não ter feito substituições no duelo contra o real madrid [guardiola explains not making substitutions in the match against real madrid]. https://www.terra.com.br/esportes/futebol/internacional/guardiola-explica-nao-ter-feito-substituicoes-no-duelo-contra-o-real-madrid,e3fd37046389eb17727c7e0d1ef1230ahuy6brqn.html, s.d. Accessed on: 24 Oct. 2025.