A Fuzzy Logic External Reasoning Tool for High-Stakes Tactical Decision-Making: A Neuro-Symbolic Approach

Pedro Passos Farias,[1] **Patrick Barbosa Moratori**[2]

[1]Polytechnic School of the Federal University of Rio de Janeiro (Poli/UFRJ)
[2]Institute of Computing, Fluminense Federal University (UFF)
Av. Athos da Silveira Ramos, 149 - University City, Rio de Janeiro - RJ, 21941-909
pedropassos@poli.ufrj.br pmoratori@id.uff.br

Abstract

In the pursuit of reliable Artificial Intelligence, Large Language Models (LLMs) often struggle with logical consistency and precise state tracking in dynamic, high-stakes environments. While LLMs excel at semantic processing, they lack the grounded reasoning required for real-time prescriptive management. This paper introduces a specialized **Neuro-Symbolic Prescriptive AI Framework** designed to audit, rather than automate, human judgment. Unlike purely data-driven "black-box" models that encounter predictive ceilings, our architecture **synergizes machine learning-derived performance metrics (Role-Aware Cumulative Mean) with an interpretable fuzzy inference engine**. This hybrid approach processes uncertain inputs—such as performance decay, fatigue, and disciplinary risk—into actionable, rule-based priorities. We validate this approach through a complex case study in elite soccer (the 2018 FIFA World Cup), serving as a proxy for dynamic multi-agent environments. The system acts as a logical auditor, using the symbolic engine to interpret the statistical metrics and eliminate exposure bias. The results demonstrate that the neuro-symbolic engine correctly identified high-risk scenarios (e.g., the "Fagner Paradox") that were ignored by human intuition and statistical baselines. This work confirms that integrating fuzzy logic as a transparent reasoning layer is essential for future neuro-symbolic architectures, serving as a robust "tool" that grounds the generative capabilities of language models.

1 Introduction

Recent advancements in Large Language Models (LLMs) have revolutionized natural language understanding but have arguably reached a plateau regarding logical reasoning and real-time state management in complex environments. While capable of generating plausible narratives, these models frequently fail to maintain logical consistency when faced with imprecise, numerical, or rapidly changing data streams. To address this, the emerging field of **Neuro-Symbolic AI** advocates for the integration of connectionist models with external symbolic logic solvers—tools that can provide the "ground truth" that generative models lack.

This work formalizes a **state-of-the-art Neuro-Symbolic auditing architecture** designed precisely for this role: it functions as a domain-specific symbolic engine capable of grounding high-dimensional data streams into normative decision states. By integrating the robustness of statistical learning (via role-aware performance metrics) with the transparency of logic solvers, we bridge the gap between predictive accuracy and epistemic accountability. We apply this system to the domain of elite soccer management.

Current data-driven approaches in this domain encounter a "predictive ceiling". For instance, machine learning models attempting to predict tactical changes often achieve only 70% accuracy (Mohandas, Ahsan, and Haider 2023), as they fail to capture the qualitative reasoning behind expert decisions. Our contribution addresses this by shifting from prediction to **prescription**. We introduce a **Decision Support System (DSS)** that uses inputs like performance, fatigue, and disciplinary risk to generate a fuzzy output (*Substitution Priority*). By acting as an external logical solver, this system provides a quantitative, interpretable signal that can be used to ground the outputs of broader AI agents or human decision-makers.

2 Related Work

This work is grounded in prior research on performance evaluation, contextual interpretation, and symbolic reasoning under uncertainty, and addresses known limitations of purely predictive approaches in sports analytics.

2.1 Performance Evaluation and Exposure Bias

Quantifying individual contribution in multi-agent environments remains a central challenge. Several approaches aggregate technical and statistical indicators to evaluate player performance. The **PlayeRank** framework (Pappalardo et al. 2019) represents a major advance by introducing a multidimensional, role-aware metric validated against professional scout assessments. While PlayeRank explicitly normalizes action vectors to mitigate play-time bias, its match-level formulation relies on a weighted cumulative sum of contributions. When inspected temporally, this formulation produces strictly increasing trajectories, which can reintroduce exposure bias and obscure performance decay, momentum reversals, or tactical suppression.

Related work highlights the importance of temporal analysis in performance modeling. Schmidt, Lillo, and Bustos (2024) demonstrate that player influence fluctuates significantly over the course of a match due to tactical reconfig-

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

uration, opponent pressure, and physical factors, advocating temporal windowing to capture such dynamics. Complementary studies focusing on individual evaluation using fuzzy or hybrid models, such as Zeng and Li (2014), Bazmara et al. (2014), Sałabun et al. (2020), and Onwuachu et al. (2022), aggregate heterogeneous indicators into interpretable performance scores or rankings. However, these approaches typically operate at the match or season level and do not explicitly address temporal exposure effects within a match.

Building on these insights, the present work adopts short temporal slices and introduces a role-aware cumulative mean aggregation. This formulation preserves interpretability while enabling meaningful comparison between players with different playing times, eliminating monotonic exposure effects and allowing performance trajectories to reflect true temporal variability.

2.2 Contextual Interpretation and Symbolic Reasoning

Another relevant research stream focuses on contextual interpretation of football matches through spatial and relational indicators. Early work by Zhou et al. (2009) employs fuzzy contextual reasoning to map relative positioning and game situations into action-oriented decisions. In contrast, Vallejo, Vellido, and Roddick (2015) emphasize the semantic recognition of tactical actions from positional data, remaining primarily descriptive and retrospective.

Fuzzy logic has been widely adopted as a symbolic reasoning paradigm for handling imprecision and uncertainty. Studies such as LACCEI (2023) demonstrate how fuzzy systems transform quantitative descriptors into interpretable linguistic assessments (e.g., stamina, agility), while Marliere (2017) apply fuzzy inference to arbitrate between tactical states in uncertain environments. These works motivate the use of fuzzy logic as a symbolic grounding mechanism, enabling explicit reasoning over uncertain and heterogeneous inputs.

Distinct from prior applications focused on low-level control, robotics, or action recognition, the present work transposes fuzzy reasoning to high-level managerial decision making. Fuzzy inference is used to integrate temporally resolved performance indicators, physiological signals, and contextual risk factors into a single, interpretable *Substitution Priority* signal.

2.3 Substitution Analysis and Predictive Modeling

Player substitutions have also been studied from an analytical perspective. Several works analyze substitution timing, frequency, and impact using statistical or causal methodologies, including Goes et al. (2017), Power et al. (2017), Pamukkale et al. (2022), and Wu et al. (2024). These studies provide valuable empirical insights into substitution effects but are largely retrospective and do not directly support real-time decision making.

More recently, substitution decisions have been framed as supervised prediction problems. Mohandas, Ahsan, and Haider (2023) analyze over 50,000 substitutions using machine learning models such as Random Forests and Support Vector Machines, achieving a maximum predictive accuracy of approximately 70%. This observed ceiling reflects a fundamental limitation of predictive approaches: supervised models are trained to reproduce historical human behavior and thus inherit implicit biases and contextual heuristics, rather than explicitly reasoning about optimal actions under uncertainty.

Rather than treating substitutions as a prediction task, the present work adopts a prescriptive decision support perspective. The proposed symbolic solver audits temporally evolving performance against explicit, interpretable rules, serving as a bias-mitigating layer that supports—rather than replaces—human coaching judgment.

2.4 Research Gap

Taken together, existing studies address performance evaluation, contextual interpretation, symbolic reasoning, or substitution analysis in isolation. To the best of our knowledge, no prior work integrates a validated, temporally resolved performance metric with symbolic reasoning over contextual match factors for the explicit purpose of informing in-game substitution decisions. This work addresses this gap by combining role-aware temporal performance evaluation with fuzzy symbolic reasoning within a unified, interpretable decision support framework.

2.5 The Prescriptive AI Paradigm

Unlike general Prescriptive Analytics, which primarily focuses on optimization and outcome selection, we adopt a stricter notion of **Prescriptive AI** centered on epistemic auditing and normative justification. In this paradigm, the role of the AI system is not merely to recommend actions based on predicted outcomes, but to make explicit the normative assumptions, decision criteria, and constraints under which a given action is considered justifiable in a high-stakes context.

A complete formalization of this paradigm—including its ethical differentiation from predictive mimicry and post-hoc explainability—is presented in an extended theoretical framework in a complementary theoretical manuscript (Passos and Moratori 2025). The present work focuses on the operational instantiation of these principles, demonstrating how a prescriptive and auditable reasoning layer can be integrated into a real-time tactical decision-making system.

3 Methodology: The Neuro-Symbolic Tool-Use Architecture

We propose a conceptual integration framework based on the **Tool-Use paradigm** (Function Calling). In this architecture, the Fuzzy Logic Control System acts as a deterministic "Oracle" that the Large Language Model (LLM) queries to ground its reasoning. The interaction follows a strict three-step protocol: Formalization, Execution, and Synthesis.

3.1 Step 1: Intent Recognition and Formalization

The workflow begins with the human agent (e.g., the coach) inputting a query containing the processed variables of a player (e.g., *"Analyze Fagner: Performance 0.35, Defender, Yellow Card..."*). The LLM acts as a semantic parser. Instead of answering immediately, it recognizes the need for computation and converges the natural language input into a structured **JSON Payload** compatible with the Solver's API, ensuring all eight critical variables are correctly mapped:

```
{
  "tool": "FuzzySolver",
  "parameters": {
    "player_name": "Fagner",
    "p_cum": 0.35,       // Cumulative Performance
    "momentum": -0.5,    // Performance Trend
    "minutes_played": 45,// time on the field
    "age": 29,
    "role": "Defender",  // Tactical Context
    "yellow_card": true, // Disciplinary Risk
    "goals": 0,          // Offensive Output
    "assists": 0         // Offensive Output
  }
}
```

3.2 Step 2: Symbolic Execution (The Solver)

The Fuzzy System receives the structured payload. It creates no hallucinations; it executes the predefined mathematical rules (detailed in the Appendix) upon the input variables. The Solver returns a deterministic **JSON Output** containing the calculated priority score and the tokens of the activated rules:

```
{
  "status": "success",
  "data": {
    "substitution_priority": 100.0,
    "risk_level": "VeryHigh",
    "activated_rules": ["R03_DefensiveRisk",
    "R04_Collapse"]
  }
}
```

3.3 Step 3: Semantic Synthesis and Visualization

Upon receiving the Solver's output, the LLM resumes generation. It combines the quantitative "ground truth" with its semantic capabilities to produce a two-part response:

1. **Narrative Explanation:** The LLM translates the token `R03_DefensiveRisk` into natural language: *"Fagner requires immediate substitution (Priority 100). The Solver indicates a critical defensive liability due to the combination of low performance and a yellow card."*
2. **Data Visualization:** Optionally, if the urgency is high or explicitly requested, the LLM generates code (e.g., Python/Plotly) to plot the specific membership function that triggered the alert, visually proving to the coach why the decision is necessary.

This architecture ensures that the LLM functions solely as an **Interface Layer**, while the decision logic remains strictly grounded in the auditable Fuzzy Solver.

4 Data Science and Hypotheses

For the "Tool-Use" architecture to be effective, the input variables (fed into the JSON) must be scientifically robust. We formulate specific hypotheses to validate the quality of these inputs.

4.1 Problem Formulation and Hypotheses

Tactical decision-making suffers from the *Metric Exposure Bias* (where cumulative metrics favor players with more time on field) and the *Predictive Ceiling* (where ML models merely mimic human errors). To address this, we test:

- **H1 (Metric Hypothesis):** Reformulating performance evaluation from a cumulative sum to a *Cumulative Mean with Role-Aware Normalization* effectively eliminates exposure bias, allowing the system to detect negative momentum.
- **H2 (Prescriptive Hypothesis):** A Fuzzy Logic-based Solver can identify high-risk tactical scenarios (e.g., "invisible" defensive liabilities) earlier and more accurately than human intuition.
- **H3 (Contextual Hypothesis):** The modulation of disciplinary risk by tactical role (e.g., weighting a yellow card more heavily for a defender) significantly alters the priority score.

4.2 Metric Engineering (Pre-processing)

To validate H1 and generate the `p_cum` variable used in the JSON payload, we implemented a pipeline that transforms raw event logs into a unified temporal representation.

First, the match timeline is discretized into 5-minute intervals. Within each interval, two complementary components are calculated: the **Technical Score** (aggregation of passes, shots, duels) and the **Network Score** (Eigenvector Centrality). These are linearly combined to form the raw performance indicator for each slice:

$$P_{\text{raw}}^{(i)} = (1-\alpha) \cdot S_{\text{tech}}^{(i)} + \alpha \cdot S_{\text{net}}^{(i)} \tag{1}$$

Where $\alpha = 0.2$ controls the relative contribution of the network-based component, balancing structural influence with technical efficiency as reported in sensitivity ranges by Pappalardo et al. (2019).

Subsequently, to ensure metric interpretability independent of match intensity, this raw score is normalized against the historical distribution of the agent's specific tactical role. This transforms the unbounded scalar into a **Role-Aware Percentile** ($P_{\text{pct}}^{(i)}$).

Finally, to capture performance evolution through time without exposure bias (H1), we define the **Cumulative Mean**:

$$\texttt{p_cum} = \frac{1}{t} \sum_{i=1}^{t} P_{\text{pct}}^{(i)} \tag{2}$$

This formulation departs from the additive model of the original PlayeRank. By averaging the role-normalized percentiles up to time t, the metric represents the player's efficiency rather than volume, maintaining logical comparability among players with different durations on the pitch.

5 The Symbolic Reasoning Engine

To function as a reliable external tool, the Fuzzy Logic Control System (FLCS) must be transparent. Unlike the opaque weights of a neural network, every decision boundary here is explicitly defined.

5.1 Fuzzy Universes and Membership Functions

The inference engine processes the JSON inputs by mapping the eight discrete variables into linguistic sets. The membership functions (MFs) are calibrated to capture the non-linear dynamics of soccer:

- **Performance (`p_cum`):** Universe $[0, 1]$. Sets: *Very Low, Low, Medium, High, Very High*. The *Very Low* set is defined by a trapezoidal MF $[0.0, 0.0, 0.10, 0.35]$, explicitly anchoring the logic to detect "disastrous" performances immediately.
- **Momentum (`momentum`):** Universe $[-1, 1]$. Sets: *Falling, Stable, Rising*. The *Falling* set is strictly defined in the negative quadrant $[-1.0, -1.0, -0.03, -0.01]$ to distinguish genuine decline from statistical noise.
- **Fatigue (`minutes_played`):** Universe $[0, 100]$. Sets: *Low* (0-45'), *Medium* (40-80'), *High* (70-100'). The overlapping MFs model the gradual and contextual onset of physical exhaustion.
- **Contextual Age (`age`):** Universe $[15, 45]$. Sets: *Young* (< 22), *Peak* ($24 - 30$), *Veteran* (> 32). This variable modulates the fatigue logic (e.g., veterans trigger fatigue rules earlier).
- **Offensive Output (`goals`, `assists`):** Integer Universe $[0, 10]$. Mapped to fuzzy sets *None* (0), *Some* (1), and *Many* (≥ 2). High values in these variables trigger "Immunity Rules," protecting the player even if performance metrics are low.
- **Logical Switches (`role`, `yellow_card`):**
  - `role`: Categorical (*GK, Defender, Midfielder, Forward*) used to select the active rule subset.
  - `yellow_card`: Binary $[0, 1]$ acting as a risk amplifier, specifically for defensive roles (Rule R03).

5.2 Symbolic Rule Base

The core of the "Reasoning Tool" is the rule base (detailed in Appendix Table 1). For example, Rule **R08 (Inefficacy)** states: *IF (Forward) AND (P_{cum} is Low) AND (Goals is None) THEN (Priority is High).* This provides a causal trace for every decision.

5.3 Inference Output Calculation

The system employs Mamdani inference with Centroid defuzzification to compute a modifier, which adjusts the baseline statistical performance:

$$P_{final} = \text{clip}(\text{Baseline} + (\text{Modifier} \times \alpha), 0, 100) \quad (3)$$

Where $\alpha = 0.25$ is the weighting factor. This ensures that the "logic" acts as a tuner for the "data".

6 Results: Validating Symbolic Reasoning

To assess the validity of the FCS as a reasoning tool (H2), we applied it to historical data from the 2018 FIFA World Cup (Brazil vs. Belgium) (FIFA 2018).

6.1 Adherence and Divergence

The system successfully identified the same high-priority substitution candidates as the human coaches (e.g., Gabriel Jesus, Priority 99.1). However, the system's true value was demonstrated in the divergences.

The Fagner Paradox (Logic vs. Bias) The system flagged the right-back, **Fagner**, as **Maximum Priority (100.0)** from minute 45 onwards due to "Critical Risk" (Low Performance + Defensive Role + High Exposure). The human decision-makers did not substitute him. The result was a defensive breakdown and a critical yellow card at the 90th minute. The symbolic engine correctly diagnosed the state that human intuition ignored.

The Lukaku Paradox (Handling Anomalies) **Romelu Lukaku** provided an assist. However, the fuzzy logic detected that his general participation was near zero (14 touches in 87 minutes). The "Ineffective Forward" logic correctly overpowered the "Assist Protection" rule, identifying him as the #1 candidate for substitution 50 minutes before the human coach reacted. This validates H1: the Cumulative Mean metric prevented the "assist bias" from masking the performance decay.

7 Comparative Analysis

As shown in Table 2 (Appendix), standard ML models hit a predictive ceiling ($\approx 70\%$) because they try to mimic human errors (Mohandas, Ahsan, and Haider 2023). Our symbolic approach provides a prescriptive audit, making it a superior component for reasoning tasks where explainability is paramount.

8 Contributions

This work advances the field of Neuro-Symbolic AI and sports analytics through five distinct contributions:

- **Methodological Innovation (Cumulative Mean):** We propose a mathematical reformulation of the PlayeRank framework. By replacing the cumulative sum with a **Role-Aware Cumulative Mean** (H1), we eliminate the "exposure bias" inherent in standard metrics, enabling the first objective detection of intra-match performance decay.
- **Domain Adaptation (Control Theory to Management):** While fuzzy logic is traditionally applied to robotic actuation, we transpose this control theory to **tactical-managerial reasoning**. We demonstrate that intrinsic agent variables can model temporal risk, moving beyond reactive navigation to complex state management.
- **Neuro-Symbolic Tool-Use Architecture:** We introduce a formal integration protocol where the LLM acts as a

semantic interface and the Fuzzy System as a deterministic "Oracle." This architecture validates the **"perception-reasoning" loop**, where ML metrics provide the rigorous perception and Fuzzy Logic provides the normative reasoning.

- **Codification of Tactical Heuristics:** We formalize subjective expert intuitions into computable variables. Specifically, the introduction of **"Momentum"** (rate of change) and **Role-Modulated Risk** allows the system to detect subtle "invisible" risks (like the Fagner Paradox) that purely statistical models fail to capture.
- **Paradigm Shift (Prescriptive Auditing):** We demonstrate that Supervised Machine Learning hits a "predictive ceiling" ($\approx 70\%$) by mimicking human biases. Our approach shifts the paradigm to **prescriptive auditing**, prioritizing objective logical utility over the replication of historical patterns.

9 Conclusion

This study presents a Fuzzy Logic Control System that serves as a robust symbolic reasoning tool for high-stakes environments. By codifying expert heuristics into a computable framework, we verified that the system allows for real-time logical auditing of agent performance. The reformulation of metrics into Cumulative Means (H1) and the use of role-aware logic (H3) enabled the system to outperform human intuition in identifying complex risks (H2). In the context of Neuro-symbolic AI, the proposed "Tool-Use" architecture provides the necessary bridge: the LLM handles the semantics and visualization, while the Fuzzy Solver guarantees the logical consistency of the decision.

10 Limitations and Future Work

Although the experimental evaluation is conducted on a single match, the proposed system operates at a fine-grained temporal and individual level. Each match is decomposed into 18 five-minute windows, evaluated independently for each of the 10 outfield players on both teams. As a result, a single match yields 360 distinct fuzzy inference executions (2 teams × 10 players × 18 windows), each representing an independent logical reasoning pass under different contextual and performance conditions. This structure enables a dense temporal validation of the system's reasoning behavior, despite the limited number of matches.

While the proposed FCS demonstrates significant utility, certain limitations exist. First, the reliance on minutes played as a proxy for fatigue assumes a linear accumulation of load. Second, the static rule base requires manual calibration for different tactical styles. Future work will address these gaps through:

1. **Biometric Integration:** Incorporating GPS data to refine the "Fatigue" token in the JSON payload.
2. **Neuro-Fuzzy Adaptation (ANFIS):** Developing an Adaptive Neuro-Fuzzy Inference System to learn membership functions from expert data automatically.

References

Bazmara, A.; et al. 2014. A Novel Fuzzy Approach for Determining the Best Position of Soccer Players. *International Journal of Intelligent Systems and Applications*.

FIFA. 2018. 2018 FIFA World Cup Russia™ - Match Report: Brazil - Belgium. Official Match Report.

Goes, F. R.; et al. 2017. Data-Driven Analysis of Performance Indicators in Elite Soccer. *Journal of Sports Sciences*.

LACCEI. 2023. 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*.

Marliere, F. T. 2017. 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.

Mohandas, A.; Ahsan, M.; and Haider, J. 2023. Tactically Maximize Game Advantage by Predicting Football Substitutions Using Machine Learning. *Big Data and Cognitive Computing*, 7(2): 117.

Onwuachu, W. C.; et al. 2022. A Neuro-Fuzzy Logic Model Application for Predicting the Result of a Football Match. *European Journal of Electrical Engineering and Computer Science*.

Pamukkale, A.; et al. 2022. An Analysis of Substitution Timing and Its Impact on Match Performance in Soccer. *International Journal of Performance Analysis in Sport*.

Pappalardo, L.; Cintia, P.; Rossi, A.; Massucco, E.; Ferragina, P.; Pedreschi, D.; and Giannotti, F. 2019. 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.

Passos, P.; and Moratori, P. 2025. Auditing Human Decision-Making in High-Stakes Environments via Prescriptive AI: A Stress-Test on Real-Time Tactical Management. *arXiv preprint arXiv:2512.04480*. Extended version.

Power, P.; Ruiz, H.; Wei, X.; and Lucey, P. 2017. Not All Substitutions Are Created Equal. *Journal of Sports Sciences*.

Sałabun, W.; et al. 2020. A Fuzzy Inference System for Players Evaluation in Multi-Player Sports: The Football Study Case. *Symmetry*.

Schmidt, L.; Lillo, C.; and Bustos, J. 2024. Revisiting PlayeRank. arXiv:arXiv:2410.20038v1.

Vallejo, D.; Vellido, A.; and Roddick, J. 2015. SIDANE: Towards the Automatic Analysis of Football Tactics and Actions. In *Proceedings of the International Conference on Knowledge Discovery and Information Retrieval (KDIR)*.

Wu, Y.; et al. 2024. Impact of Substitutions on Elite Soccer Team Performance Based on Player Evaluation Systems. *Sports, Health and Research*.

Zeng, D.; and Li, X. 2014. Fuzzy Logic and Its Application in Football Team Ranking. *Mathematical Problems in Engineering*.

Zhou, X.; et al. 2009. Fuzzy Decision Making for Robotic Soccer Based on Contextual Interpretation. In *Proceedings of the RoboCup International Symposium*.

Appendix: Tables

Table 1: Complete Symbolic Rule Base for Tactical Inference

Rule ID	Logical Antecedents (IF)	Symbolic Rationale
R01	IF (P_{cum} High/Very High) THEN (Mod Negative)	Star Protection: High performers shielded.
R02a	IF (P_{cum} Low) AND (Min High) THEN (Mod Large Pos)	Late-Game Drag: Unrecoverable fatigue.
R03	IF (Defender) AND (Yellow-Card Yes) THEN (Mod Med Pos)	Risk Management: Defensive liability.
R04	IF (P_{cum} Low) AND (Momentum Falling) THEN (Mod Large Pos)	Collapse Detection: Poor play + decline.
R07	IF (P_{cum} High) AND (Momentum Rising) THEN (Mod Med Neg)	Hot Hand: Reinforces peak state.
R08	IF (Forward) AND (P_{cum} Low) AND (Goals None) THEN (Mod Large Pos)	Inefficacy: No utility without goals.
R09	IF (Forward) AND (P_{cum} Low/Med) AND (Mom Falling) THEN (Mod V. Large Pos)	Striker Pressure: Decline is fatal for role.
R10	IF (Midfielder) AND (P_{cum} Low) AND (Assists None) THEN (Mod Med-Large Pos)	Invisible Playmaker: Failed facilitation.
R11	IF (Assists Some) OR (Goals Some) THEN (Mod Negative)	Creator Bonus: Output buys immunity.
R12	IF (Age Veteran) AND (Min High) THEN (Mod Med Pos)	Geriatric Risk: Steeper physical decay.
R13	IF (Age Young) AND (Output High) THEN (Mod V. Large Neg)	Youth Development: Encourages talent.
R15	IF (P_{cum} Medium) AND (Momentum Stable) THEN (Mod Zero)	Neutral State: No action required.

Table 2: Structural Comparison: Supervised ML vs. Proposed Fuzzy Solver

Dimension	Supervised ML (e.g., Mohandas et al.)	Proposed Fuzzy Solver
Objective	Minimize Pred. Error	Maximize Utility
Interpretability	Low (Black Box)	High (Explicit Rules)
Outliers	Treats as Noise	Handles via Logic
Role in AI	Standalone Predictor	External Reasoning Tool