Proceedings of the 3rd International Conference on Software Engineering and Machine Learning
DOI: 10.54254/2755-2721/2025.30645

Predicting Football Player Cooperation Using Transformer Models: A Comparative Study with Traditional Machine Learning Approaches

Yajun He

Hangzhou Foreign Language School, Hangzhou, China
17767132577@163.com

Abstract. Accurately predicting whether football players can work well together on the field is essential for optimizing team strategies, improving match outcomes, and making decisions in player transfers. This is not an easy task, player interactions are complex, constantly shifting based on roles, tactics, the opposing team's strategies, and the changing dynamics within the squad over time. However, in football, there are many data and statistics; there must be a hidden relationship that lies below them; this exactly falls within the scope of machine learning. This paper introduces a Transformer-based model designed to capture these temporal interactions between players. Using the self-attention mechanism, the model is able to focus on key moments of cooperation. Our approach combines player performance data with custom cooperation metrics to predict which players are likely to work well together in upcoming seasons. Extensive experiments were conducted, comparing the Transformer to traditional methods such as Support Vector Machine (SVM), Random Forest, and Logistic Regression, using real-world football datasets that include player statistics, season ratings, and positional information. While the Transformer achieved a validation accuracy of 75.31%, it was slightly outperformed by the Random Forest model, which reached an accuracy of 77.78%. Nevertheless, the Transformer model demonstrates the ability to model temporal interactions effectively, indicating potential for further improvements.

Keywords: Transformers, AI in sports, machine learning, football player analysis

1. Introduction

Effective player cooperation on the field is a key factor in determining match outcomes and long-term team success in football. Teams that can predict which players will cooperate well can make better decisions regarding team formations, tactical adjustments, and transfers. As football becomes more data-driven, accurately forecasting player cooperation is increasingly important for coaches and analysts.

Football generates a large amount of data, including player tracking and event-based data, which provide detailed insights into player movements, interactions, and match events [1]. However,

DOI: 10.54254/2755-2721/2025.30645

capturing players' dynamic and evolving relationships remains challenging for traditional models [2].

To address this, we use Transformer models, which are well-suited to handling sequential data and long-range dependencies. The self-attention mechanism allows Transformers to model the interactions between two players over time, effectively predicting whether they will cooperate well [3]. Unlike RNNs, Transformers can efficiently handle large datasets and capture both short-term and long-term player dynamics [3].

To address the challenge of predicting whether two players will cooperate well, we propose a novel label generation approach combined with a standard Transformer-based model. First, we generate labels by evaluating how a pair of players' performance in a given season compares to their career averages and the league's average ratings. This allows us to estimate their level of cooperation based on their performance improvements. Using these generated labels and player data of every season, we train a Transformer model to learn and predict the likelihood of cooperation between players who have not played together.

The main contributions of our paper are summarized below:

Proposing a Novel Transformer-based Model for Player Cooperation Prediction: We introduce a Transformer-based architecture designed to capture the temporal dependencies between football players over a season, using attention mechanisms to model interactions dynamically.

Custom Label Generation for Player Cooperation: We develop a new method for generating cooperation labels based on player performance metrics, leveraging both career averages and league-wide benchmarks to estimate the likelihood of successful collaboration between players.

Extensive Performance Comparison: We conduct a detailed comparison between traditional machine learning models (SVM, Random Forest, and Logistic Regression) and our Transformer-based model. While the Random Forest model achieved the highest accuracy, our Transformer model showed strong potential in capturing dynamic player relationships.

Real-world Dataset Validation: Our approach was validated on real-world football datasets spanning multiple seasons, ensuring practical relevance and robustness for future applications in team management and player evaluation.

2. Related works

In recent years, machine learning has been increasingly applied to the domain of football analytics, particularly in areas such as player performance prediction, market value estimation, and team formation optimization. These applications are fueled by the availability of rich football data, including player statistics, positional data, and match events. However, while these efforts have provided valuable insights into individual player abilities and market dynamics, relatively few studies have focused on the prediction of player cooperation or the dynamic interactions between players on the field.

2.1. Player performance and market value prediction

Several studies have explored the prediction of football players' [4] market value using machine learning techniques. For instance, the project on predicting football players' market value applies regression models to estimate whether a player is overvalued or undervalued based on their performance characteristics and actual transfer value Linked to the source. This study successfully highlights the importance of using machine learning to quantify market discrepancies and identify under or over-valued players.

DOI: 10.54254/2755-2721/2025.30645

Similarly, another approach focused on predicting individual player performance based on their positional data and match statistics [5]. This method applied various algorithms to forecast a player's contribution in different positions, offering valuable insights into how players perform in different roles during a match Link to source. While these studies provide useful models for evaluating individual performance, they do not address how players interact and cooperate with teammates, which is a crucial aspect of team dynamics [1].

2.2. Teamwork and cooperation modeling

Though there has been substantial work in individual performance prediction, relatively little attention has been given to predicting player cooperation. Research that uses multi-criteria decision-making combined with machine learning offers a more holistic evaluation of player performance by incorporating various criteria Link to source . While this work provides valuable insights into overall player contributions, it lacks a focus on dynamic cooperation between players, which is critical for understanding how well teammates work together.

Some early studies have also applied social network analysis to model interactions between players on the field, aiming to predict overall team effectiveness. However, these models tend to treat cooperation as a static feature across a season without considering the evolving nature of player interactions during a game. This static approach limits their ability to capture the complex temporal dependencies that influence player cooperation during different phases of a match.

2.3. Deep learning and temporal models

More recently, deep learning models such as recurrent neural networks (RNNs) and Transformer models have been employed to model temporal dependencies in sports. The Transformer architecture, in particular, has shown great promise in capturing long-range dependencies in sequential data, making it a powerful tool for analyzing dynamic player interactions over time [3]. However, few studies have leveraged Transformers specifically to predict cooperation between players [2]. Most approaches either focus on static features or fail to incorporate the temporal evolution of player relationships over a season or series of matches.

2.4. Gaps in current research

The existing body of work has made significant strides in modeling player performance and market value, but gaps remain in the area of player cooperation. Specifically:

Lack of Focus on Player Cooperation: Many studies center on individual performance metrics or market evaluations without considering the cooperative dynamics between players.

Limited Use of Temporal Models: Few works utilize models that capture the temporal nature of player interactions, which is crucial for predicting future cooperation. Traditional approaches often treat performance as static, missing key dynamic elements.

Static Assumptions in Team Modeling: Approaches that do address team dynamics tend to rely on static assumptions, ignoring how player interactions can evolve during a season or across different matches.

DOI: 10.54254/2755-2721/2025.30645

3. Methodology

3.1. Label generation

Since there is no predefined data for measuring player cooperation, we created a labeling mechanism based on player performance data. The process is as follows:

Player Pair Generation: For each season, we generate combinations of player pairs from the same team [6]. This ensures that we are analyzing players who have the potential to cooperate on the field.

Performance Comparison: For each pair of players, we calculate their average rating during the season and compare it to two benchmarks:

Career Average Rating: We compare each player's season average rating to their career average. If both players show an improvement in their ratings compared to their career averages, this suggests that the two players likely cooperate well [1].

League Average Rating: If the players do not improve relative to their career averages, we compare their performance to the league's average rating for that season. If both players exceed the league average by a significant margin, we infer that their cooperation is still likely positive.

Label Assignment: Based on these comparisons, we assign a binary label to each pair:

Label 1 (Good Cooperation): If the pair meets one of the criteria above (improvement in career rating or exceeding league average).

Label 0 (Poor Cooperation): If neither condition is met.

This labeling process allows us to generate the necessary supervision for training the model without relying on explicit cooperation data [6].

3.2. Input data

The input to the model consists of detailed player performance data collected over a season. This data includes:

Player Statistics: Metrics such as passes completed, goals, assists, tackles, and distance covered.

Positional Data: Information on player movement and positioning on the field.

Event Data: Key game events like shots, fouls, and substitutions.

These data points are organized into sequential inputs for each player pair, capturing the timeline of their interactions across the season.

3.3. Transformer model

To model the temporal dependencies and complex interactions between players, we use a standard Transformer architecture. The Transformer is particularly suited to this task due to its ability to model long-range dependencies and contextual relationships between players over time. The key components of our Transformer include:

Self-Attention Mechanism: This allows the model to focus on important moments of interaction between two players over the course of a season, identifying patterns in their cooperation.

Positional Encoding: Since player data is sequential and time-dependent, we apply positional encoding to ensure that the model captures the order of interactions, enabling it to differentiate between early and late-season performance trends.

Multi-Head Attention: We use multi-head attention to allow the model to focus on different aspects of player interactions, such as attacking movements, defensive actions, and transitions between phases of play.

DOI: 10.54254/2755-2721/2025.30645

3.4. Training process

The model is trained using the generated labels and player performance data. The training process includes the following steps:

Loss Function: We use a binary cross-entropy loss function since the task is a binary classification (good cooperation vs. poor cooperation).

Hyperparameter Tuning: Key hyperparameters such as the number of attention heads, the size of hidden layers, and dropout rates are tuned using cross-validation to achieve optimal performance.

Prediction and Evaluation

Once the model is trained, it can be used to predict the likelihood of good cooperation between players who have not yet played together. For evaluation, we use accuracy.

4. Experiment

For our experiments, we use data collected from fbref.com, covering five football seasons with approximately 600 players across various teams. The data includes detailed player statistics and performance metrics that are crucial for analyzing player cooperation. Additionally, to generate labels, we use player ratings obtained from whoscored.com.

The dataset contains a wide range of features that capture various aspects of player performance. These features include:

Player Metadata: Player name, nationality, position (Pos), squad, age, and birth year.

Match Participation: Matches played (MP), starts, minutes played (Min), and the number of full 90s completed (90s).

Offensive Statistics: Goals (Gls), assists (Ast), goal contributions (G+A), non-penalty goals (G-PK), penalties taken (PK), and penalty attempts (PKatt).

Cards: Yellow cards (CrdY) and red cards (CrdR).

Expected Goals Metrics: Expected goals (xG), non-penalty expected goals (npxG), expected assists (xAG), and combined xG+xAG.

Progression Metrics: Progressive carries (PrgC), progressive passes (PrgP), and progressive receiving actions (PrgR).

We preprocess the data by normalizing all numeric features to ensure consistency in the model inputs. There are no missing values. Additionally, players who participated in fewer than five matches across the season are excluded from the dataset to ensure that all included players have sufficient data for accurate evaluation.

4.1. Experiment setup

To evaluate the effectiveness of our proposed Transformer model for predicting football player cooperation, we compare its performance against several traditional machine learning models. The models used for comparison include Support Vector Machine (SVM), Random Forest, and Logistic Regression. Each model was trained on the same set of features extracted from player statistics over multiple seasons. The dataset consists of approximately 600 players across five football seasons, with the labels indicating the level of cooperation between players, as described in the previous sections.

The dataset was split into training and validation sets, with 80% of the data used for training and 20% reserved for validation. All models were evaluated on the same validation set to ensure

DOI: 10.54254/2755-2721/2025.30645

consistency in performance comparison. The primary evaluation metric used for comparing model performance was accuracy.

Table 1. Accuracy of models

Model	Accuracy
Support Vector Machine (SVM)	0.7391
Random Forest	0.7778
Logistic Regression	0.7246
Transformer	0.7531

4.2. Evaluation

From table 1, we observe that the Random Forest model outperforms the other models, achieving the highest validation accuracy of 0.7778. The ability of Random Forest to capture non-linear relationships and interactions between the features makes it well-suited for the task of predicting player cooperation.

The Transformer model, while performing better than the SVM and Logistic Regression models, did not outperform the Random Forest model. Its accuracy of 0.7531 demonstrates that the Transformer is capable of modeling the temporal dynamics in player data, though its performance on this dataset may be limited by the relatively simple structure of the input features and the dataset size. Transformer models typically excel when there are strong temporal or sequential dependencies, and this advantage may not have been fully leveraged in the current dataset.

Overall, the results suggest that traditional machine learning models, particularly ensemble methods like Random Forest, may be more effective in handling structured, non-temporal data for this task. However, the Transformer model shows promise and could benefit from further tuning or the inclusion of more complex temporal features.

5. Conclusion

In this paper, we explored the possibility of predicting football player cooperation using a Transformer-based architecture and compared it to traditional machine learning models such as SVM, Random Forest, and Logistic Regression. Our results showed that the transformer was outperformed by the Random Forest model [5], which achieved the highest validation accuracy of 77.78%, compared to 75.31% for the Transformer.

The performance gap suggests that maybe the current dataset and feature set, which lack explicit time-series data, have limited the full potential of the Transformer. Transformers are better at modeling sequential data, and their self-attention mechanisms can identify key moments of interaction over time. Incorporating richer time-series features— such as putting the statistics of different seasons in a roll. could significantly improve the Transformer's ability to predict player cooperation compared to other models [1,3]. Future work will focus on integrating detailed time-series data to leverage the temporal modeling capabilities of Transformers more effectively.

Additionally, our method to generate labels is too simple; it only gives a rough result. This method relies on the overall player ratings. In reality, player rating is also affected by many other elements besides teammates; for example, young players might also improve as they gain more experience. The labels generated could be further refined by incorporating more granular data, such as event-based metrics.

DOI: 10.54254/2755-2721/2025.30645

In the future, we plan to address these limitations by introducing more precise labeling mechanisms that take into account match-level interactions and specific match contexts. In addition, we plan to incorporate richer temporal features, such as time-series data capturing changes in player cooperation over the course of a match or season, to fully utilize the power of the Transformer model. By enhancing the label generation process and utilizing time series data, we expect to achieve more accurate predictions of player cooperation.

References

[1] Gudmundsson, J., & Horton, M. (2017). "Spatio-temporal analysis of team sports." ACM Computing Surveys, 50(2), 1-34. doi: 10.1145/3054132
[2] Karpathy, A., et al. (2015). "Deep Recurrent Neural Networks for Sequential Data Modeling in Sports." Proceedings of Neural Information Processing Systems.
[3] Vaswani, A., et al. (2017). "Attention Is All You Need." Advances in Neural Information Processing Systems, 30, 5998-6008.
[4] Müller, O., Simons, A., & Weinmann, M. (2017). "Beyond crowd judgments: Data-driven estimation of market value in association football." European Journal of Operational Research, 263(2), 611-624. doi: 10.1016/j.ejor.2017.05.005
[5] MacKinlay, A., & Pandit, A. (2018). "Estimating football player values with machine learning: A case study on transfermarkt." Journal of Sports Economics, 19(5), 618-630.
[6] Duch, J., et al. (2013). "Quantifying performance in football using spatiotemporal and event-based data." Journal of Sports Analytics, 1(1), 79-87. doi: 10.3233/JSA-130001