Machine Learning (2026) 115:6
https://doi.org/10.1007/s10994-025-06935-6

Temporal Graph Network Framework for Quantifying Pass Reception Probabilities Against Defensive Structures

Pegah Rahimian[1] · **Jesse Davis**[2] · **Laszlo Toka**[3]

Received: 5 October 2023 / Revised: 3 May 2025 / Accepted: 3 November 2025

Abstract
Passing decisions in soccer are heavily influenced by the opposing team's defensive organization. Existing approaches often decompose the problem into pass selection and success probabilities, sometimes incorporating pressure-related features to account for defensive constraints. In this study, we propose a framework for evaluating passing decisions against defensive structures using temporal graph networks (TGNs). Rather than separately modeling selection and success, we estimate the probability of a pass being received by each teammate or intercepted by an opponent, leveraging temporal, spatial, and relational data to capture dynamic interactions. We focus on forward passes originating in the middle third of the pitch, where teams frequently encounter structured defensive shapes. Specifically, we analyze passes that (1) bypass defensive lines, (2) reach teammates inside defensive structure, or (3) penetrate, exit, and reach teammates outside defensive structure. Our evaluation compares pass reception prediction accuracy against baselines, examines receiver availability against defensive structure, and assesses the situational value of passing options. The results suggest that TGNs improve pass probability estimates while offering practical insights for decision-making against organized defenses.

Keywords Soccer analytics · Temporal graph network · Deep learning

---

Editor: Ulf Brefeld.

✉ Pegah Rahimian
pegah.rahimian@it.uu.se

Jesse Davis
jesse.davis@cs.kuleuven.be

Laszlo Toka
toka.laszlo@vik.bme.hu

[1] Department of Information Technology, Uppsala University, Uppsala, Sweden

[2] Department of Computer Science, KU Leuven, Leuven, Belgium

[3] HUN-REN-BME Cloud Applications Research Group, Budapest, Hungary

1 Introduction

In soccer, passing is fundamental for moving the ball and creating goal-scoring opportunities. Consequently, there has been substantial interest in analyzing passing behavior such as quantifying the value of a pass (Decroos et al., 2019; Karun: Expected threat, 2024; Gyarmati & Stanojevic, 2016; Fernandez et al., 2019; Liu et al., 2020; Bransen & Van Haren, 2019; Bransen et al., 2019), understanding the riskiness of a pass (Power et al., 2017; Rahimian et al., 2023; Spearman et al., 2017; Fernandez & Born, 2020; Szczepanski & McHale, 2016), or assessing the creativity of a pass (Robberechts et al., 2023; Memmert, 2015; Bénézet & Hasler, 2018). The key commonality among these approaches is that they evaluate a pass by considering richer contextual information about the game state in which the pass was undertaken.

One relevant way to contextual passing decisions is to evaluate them in relation to the opposing team's defensive structure. In particular, this work has focused on using geometric constructs to identify relevant defensive structures and defining passes types that disrupt this structure in some way. For example, *line-breaking passes* have been defined as passes that bypass an automatically identified line of defenders (https://bit.ly/39uQe6Q) or advance the ball closer to the opponent's goal while intersecting a pair of defenders (https://statsbomb.com/articles/soccer/statsbomb-launch-new-360-metrics-line-breaking-passes-and-ball-receipts-in-space-2/). Similarly, *penetrative passes* target teammates occupying space inside the opponent's defensive zone, represented as a polygon formed by nearby opponents (Rahimian et al., 2022a; Sotudeh, 2021). These approaches focus more on identifying pass types than quantifying aspects of them (e.g., success probability, impact, etc.). Moreover, they focus use static snapshots that consider the defensive structures in isolation and hence neglect how the preceding moments in time affect both passing opportunities and defensive responses.

In this paper, we build on this literature by filling the gap. Specifically, we focus on passes in the middle third of the pitch when the player possessing the ball faces a specific set defensive structure (i.e., a convex hull consisting of the opposing players between the ball and the goal). In contrast to existing work, we consider a taxonomy of different passes that may be possible in this situation. We evaluate these passing options by looking at the reception probability, which represents the likelihood of a pass being successfully received by any of the teammates and opponents available on the pitch (Rahimian et al., 2023; Stöckl et al., 2021). The available approaches estimate this probability based on positional data which records the location of all players and the ball are recorded multiple times per second. A key challenge in this regard is to how to exploit the rich spatiotemporal characteristics of the problem. To this end, we propose using Temporal Graph Networks (TGNs) (Rossi et al., 2020). By representing players as nodes and their interactions as edges, the TGN leverages both node features (e.g., player positions, velocities) and edge features (e.g., distances, relationships) to capture the temporal evolution of player positions and interactions. In contrast to prior approaches that use a decomposed approach that combines multiple models to predict reception probabilities (Robberechts et al., 2023; Fernandez et al., 2019), we train the TGN to directly predict the probability that a specific player will receive a pass. Our proposed approach jointly models selection intent and success likelihood under defensive pressure. We show that this results in better predictive performance for estimating reception probabilities.

This paper makes three core contributions to soccer analytics: First, we demonstrate that TGNs outperform the counterpart decomposed methods in predicting pass reception probabilities by (1) jointly modeling selection intent and success likelihood, and (2) explicitly memorizing event sequences from possession start through facing defensive structure. Second, we apply this framework to (i) quantify player positional availability for different pass types, (ii) identify high-value receivers through post-reception scoring probabilities, and (iii) assess how restricted passing options affect offensive threat. The TGN's ability to maintain possession-phase memory enables more realistic evaluation of passing opportunities under defensive pressure compared to snapshot-based approaches.

The paper is structured as follows: Sect. 2 reviews related work on soccer datasets and methods for identifying passes against defensive structures and quantifying passing performance. Section 3 introduces the dataset, explains the construction of the defensive structure model, and details the modeling of key pass types. Section 4 presents the TGN and its application in estimating reception probabilities. Section 5 outlines the experimental setup, evaluates the TGN's performance in pass outcome prediction, and explores further use cases. Finally, Sect. 6 concludes the paper.

2 Related Work

Analyzing passing in soccer requires event data and tracking data. Event data captures key match actions (e.g., passes, shots, tackles) along with their type, outcome, location, involved players, and timestamp. It is generated through manual annotation or computer vision but lacks information on players not directly involved, offering only a partial game context. Tracking data, in contrast, records the real-time positions of all players and the ball at high frequency (e.g., 10 frames per second) using multi-camera systems and computer vision. While event data provides discrete insights, tracking data offers a continuous view of player movement and positioning. Synchronizing both data types creates a comprehensive dataset, enabling deeper analysis of team dynamics, tactical patterns, and game strategies beyond what event data alone can provide. In this section we summarize related work on identifying key passes and analyzing passing behavior based on such datasets.

2.1 Passes Through Defensive Structures

Given the importance of passes that face the opponent's defensive zones, there is interest in developing data-driven approaches to recognize such passes. Typically, these approaches attempt to identify geometric properties of the defense's structure.

One line of work looks at line-breaking passes. A *defensive line* refers to a linear formation of defending players positioned to block forward progress, often identified through clustering algorithms based on player positions and movements. As illustrated in Fig. 1, such passes attempt to bypass a line of defending players. From a computational point of view, the challenge lies within identifying these lines despite continuous player movement. Researchers have proposed various clustering algorithms to tackle this problem. Fernandez et al. (2019) utilize spectral clustering based on the mean positions of opponent players to detect the dynamic lines of their formations. Rahimian et al. (2021) employ opponents' locations and velocities to detect pressure on the ball holder. [1], in his Stats Perform blog

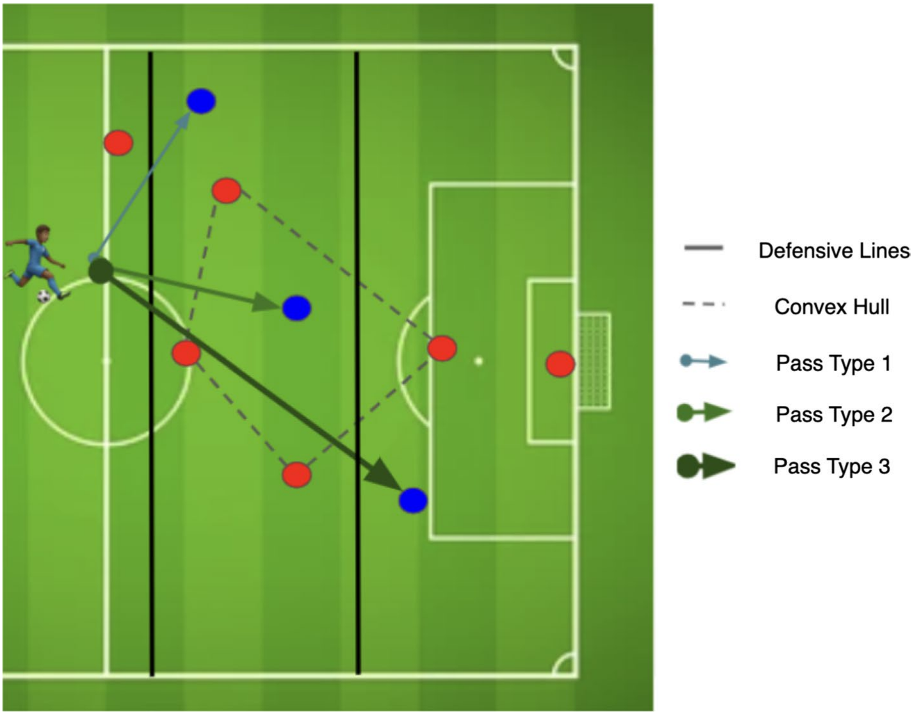

Fig. 1 Examples of three pass types in our analysis

post, suggests using Jenks natural breaks optimization with three clusters on outfield players. Another approach, proposed in studies by Power et al. (2017) and Bialkowski et al. (2021), adopts a role-based methodology to identify the formation structure and determine if a pass breaks the line of defenders.

Other researchers have considered *penetrative passes* Rahimian et al. (2022a) and Sotudeh (2021), which arise when a player is in a passing situation and there is at least one teammate positioned inside the opponent's defensive zone, represented by a polygon created by the opponent players in front of the passer (c.f. Fig. 1). This work adopts a similar definition for such passes. Sotudeh's work (Sotudeh, 2021) focuses on identifying such passes to measure players' and teams' abilities to complete the pass, whereas in this work we additionally value the impact of such passes. This current paper goes beyond (Rahimian et al., 2022a)'s approach by modeling the passer's passing options and differentiating among a broader set of pass categories. The work of Goes et al. (2019) assumes that an important aim of passing is to disrupt the organization of the defense, thus creating space. Yet, movement alone does not necessarily result in space creation. For that reason, they developed separate measures: first, the total individual movement of the defensive players on the field, second, the disruption of the defensive organization. One advantage of their approach is that it is based on a continuous performance measure (defensive disruptiveness), not depending on the occurrence of rare events like goals or goal-scoring opportunities.

2.2 Quantifying Passing Performance

There is a large body of work on quantifying aspects of passing performance such as their value (Gyarmati & Stanojevic, 2016; Decroos et al., 2019; Fernandez et al., 2019), effectiveness (Rein et al., 2017), risk-reward tradeoff (Power et al., 2017), and creativity (Robberechts et al., 2023). Some approaches are designed to work with event data (e.g. Bransen and Van Harren (2019); Bransen et al. (2019); Gyarmati and Stanojevic (2016); Decroos et al. (2019)), however, deeper analysis is attainable when using synchronized tracking and event data. Knowing the locations of other players is needed for being able to model which passing options are available and to get accurate estimates of completion probabilities (Peralta Alguacil et al., 2020). A fundamental challenge in developing models to quantify passing performance using tracking data lies in effectively representing the game context. This complexity arises from the continuous movement of players and their dynamic role and position changes throughout a match. Consequently, significant efforts have been dedicated to structuring this data for meaningful analysis. Researchers typically address this challenge via two approaches: leveraging domain knowledge to encode tactical insights or employing data-driven methods to learn representations directly from the data. Examples for the prior are Power et al. (2017) and Robberechts et al. (2023), in which features of the game state are determined by domain knowledge about the sport, while an example for the latter is Spearman et al.'s work (2017), in which the authors use a physics-based approach to model the probability of successfully passing the ball to each location on the pitch.

The data-driven approaches exploit various deep neural architectures (Fernandez & Born, 2020; Kim et al., 2023; Mehrasa et al., 2018; Peralta Alguacil et al., 2020; Rahimian et al., 2021, 2022a, b, 2023). Of particular promise in this regard are approaches that model interaction between players and the ball use graph-based (Battaglia et al., 2018) or transformer-based (Vaswani et al., 2017) neural networks. These models are well-suited for handling tracking data in sports because they can effectively model the complex, dynamic interactions between players and the ball. Graph neural networks capture spatial and relational dependencies between players, while transformers excel at capturing temporal sequences, making both approaches ideal for understanding movement patterns and decision-making. For example, Dick and Brefeld (2022) represent players and the ball as a fully connected graph and utilize graph recurrent neural networks (GRNN). However, this approach fails to adequately account for the continuous time differences between actions and player substitutions in different matches. More recently, Anzer et al. (2022) and Bauer et al. (2023) have constructed graph neural networks (GNN) to model the game state in order to detect overlaps (i.e., an overlapping run, when a player runs on the outside of a team mate who is in possession of the ball, a frequently applied drill in wide areas) and to divide a match into multiple phases of play, respectively. To address the limitations of previous works, this paper builds upon the application of Rossi et al. (2020) to effectively account for continuous time differences between actions and player substitutions across different matches.

3 Defining Considered Pass Types

We analyze forward passes attempted in the middle third of the pitch against a set defensive structure. This zone is particularly insightful for evaluating team and player decision-making against defensive structure because, while the defense tends to be more organized than in the defensive third, the space is not as restricted or congested as in the final third. This balance allows for meaningful forward progression without the extreme spatial limitations typically encountered closer to goal. Moreover, successful forward passes from the middle third are often pivotal in transitioning from buildup to chance creation, offering a valuable window into how teams attempt to break structured defense. In this section, we describe the data, the selection criteria for considered passes, and formally define the three key pass types that we analyze.

3.1 Data

The dataset utilized in this study comprises a full season (380 matches involving 20 teams) of an undisclosed league, provided by Stats Perform (https://www.statsperform.com/). This dataset combines detailed event data (describing on-ball actions such as passes, shots, and dribbles) and player tracking data (capturing player positions at 25 frames per second), enabling a fine-grained analysis of match situations. According to Stats Perform's definition, a pass is an attempt by one player to play the ball to a teammate (including set pieces). On average, over 850 passes occur per match, with a range from 721 to 1,003. After excluding backward passes, cutbacks, crosses, goalkeeper throws, throw-ins and offside passes, the dataset includes approximately 185,000 passes, of which 151,885 were successful, yielding a pass completion rate of 82.1%.

3.2 Constructing the Defensive Structure

In our dataset, 56% of all passes are started from middle third and out of all passes starting from middle third, 61% are forward. For each of these forward passes that originated from the middle third of the pitch, we assess their eligibility for our analysis. Specifically, we select the tracking frame associated with moment in time that the pass was given and apply the following checks. First, we filter out offside situations by checking that the last defending player (excluding the goalkeeper) is positioned closer to their goal line than the attacking pass recipient at the moment the pass is played. Second, a convex hull is generated using the positions of defenders in front of the ball, excluding the goalkeeper. Third, we check that at least one potential receiver is positioned inside the hull.

To ensure the defensive hull represents an active defensive block, only defenders within a preset percentage of the distance between the ball and the goal line are considered. Various percentage thresholds were tested in 5% increments to balance retaining enough passing frames with preserving meaningful defensive structures. When using a threshold of 100% (i.e., all opponent players located between ball and goal-line), a convex hull can be constructed in 93% of cases with the remaining 7% being excluded due to the absence of a receiver inside the hull or an insufficient number of defenders to form a meaningful shape. As shown in Table 1, stricter thresholds reduce the number of retained frames, with an 85% threshold preserving approximately 80% of passing frames. This setting effectively captures

Table 1 Proportion of passing frames (i.e., starting frame of forward passes starteing from middle third of the pitch) retained at different filtering thresholds for constructing defensive zones. The thresholds represent the percentage of the distance between the ball and the goal line used to filter defending players

Fraction of distance between ball and goal-line (%)	Proportion of retained frames (%)
100	93
95	89
90	86
85	**80**
80	69
75	55
70	42
65	38
60	35
55	28
50	21

realistic defensive formations while minimizing the inclusion of overly dispersed or loosely structured defensive zones.

Throughout this study, the term “Defensive Hull” refers to the convex hull formed by the positions of opposing players located in front of the ball within 85% of the distance to the goal line. This terminology is used consistently to describe the defensive context surrounding passing decisions.

3.3 Analyzed Pass Types

We classify passes into three distinct types based on their interaction with defensive structures. These pass types are defined as follows.

Pass Type 1 refers to passes that bypass at least one of the opponent’s defensive lines, advancing the ball toward the opponent’s goal without entering the boundaries of the defensive hull. These passes traverse defensive lines but avoid penetrating the core defensive structure. We employ a clustering approach to identify the opponent’s defensive lines using KMeans clustering: we group the x-coordinates (along the touchline) of the opponent players’ positions to identify vertical defensive lines. In this study, we set $K = 3$ corresponding to typical formations like 4-4-2 or 4-3-3. The resulting clusters represent defensive lines such as forwards, midfielders, and defenders. Then the centroids of these clusters are calculated to determine the x-coordinates (relative to the touchline) of each defensive line.

Pass Type 2 involves passes that successfully penetrate the defensive hull and reach a teammate positioned within this hull. This pass type is designed to capture moments when the ball is delivered into a tightly organized defensive structure. A pass is marked as Pass Type 2 if the receiving teammate is positioned inside the hull at the moment the pass is played.

Pass Type 3 refers to passes that traverse the defensive hull and exit it, targeting a teammate positioned outside the defensive hull. The key distinction of Pass Type 3 is that its endpoint lies outside the defensive hull after passing through it. Unlike Pass Type 1, which avoids the defensive hull, and Pass Type 2, which ends within the hull, Pass Type 3 fully penetrates and exits the defensive hull. As such, passes of this type do not necessarily bypass defensive lines (as in Pass Type 1).

Figure 1 illustrates the difference between there three pass types in our analysis. It is important to note that Type 2 passes (reaching a teammate within the defensive hull) are always an option in our dataset, as we filter for situations where a well-organized defensive hull exists and at least one teammate is positioned inside the hull. In contrast, Type 1 (bypassing the defensive line) and Type 3 (penetrating and exiting the hull) passes depend on the availability of teammates beyond the defensive line and outside the hull, respectively.

3.4 Excluded Pass Types

For clarity and precision, we summarize pass types that do not meet are selection criteria and hence are excluded from our analysis.

- *Passes not performed in the middle third* Passes in defensive and attacking thirds are excluded to ensure a reasonable formation of defensive hull of defenders.
- *Backward passes* Backward passes are excluded as they usually do not involve significant risk, nor reward.
- *Cutbacks* Passes which involve passing the ball backward or sideways from advanced wide areas are excluded from our analysis as they are infrequent in the middle third of the pitch and do not typically interact with defensive structures in the same way as forward passes.
- *Crosses* Crosses are excluded as they are typically high-risk, high-gain, aerial passes aimed at crowded areas in the final third of the pitch. Early crosses from the middle third are rare and distinct in nature from the pass types analyzed in this study, which focus on passes that interact with defensive structures in the middle third.
- *Failed deliveries* Passes that go out of bounds, offside or fail to reach a receiver are excluded.

4 Temporal Graph Networks for Receiver Prediction

The problem that we tackle can be formalized as follows: given contextual information (e.g., event history, player locations and velocities, etc.) from the beginning of a ball possession leading up to the moment a pass is played, predict the probability of each remaining player gaining control of the ball after the pass. This includes both teammates (in the case of a successful pass) and opponents (in the case of an interception or unsuccessful pass).

We address this problem using an extended version of the temporal graph network (TGN) (Rossi et al., 2020), which allows us to model both the spatial structure and temporal evolution of play. At each timestep, the match state is represented as a fully connected dynamic graph where nodes correspond to players, and edges encode pairwise interactions (e.g., successful passes or interceptions). The temporal dynamics of the game are captured by updating node representations over time through a memory module and message-passing mechanism. Specifically, TGN maintains a memory for each node that stores historical information, which is updated at each timestep using both past messages (capturing interactions) and positional/contextual features (capturing the evolving state of the match). This architecture allows the model to learn not only who is nearby at the moment of the pass, but also how those players have been moving and interacting over time.

At prediction time, TGN outputs a probability distribution over all 21 other players (i.e., 10 teammates and 11 opponents). In the case of a successful pass, the predicted receiver is expected to be a teammate; in the case of an interception, it is one of the opponents.

4.1 Model Design: Unified Reception Probability Estimation

The original TGN, introduced by Rossi et al. (2020), is designed for link prediction in dynamic graphs, such as social media, where it estimates the likelihood of interactions between users. We adapt the original TGN to estimate *reception probabilities* in football: the probability that any other player on the field (teammate or opponent) will control the ball after a pass is attempted. Figure 2 illustrates these probabilities visually represented by the width of the connecting lines between the ball carrier and potential receivers. Our unified model integrates both intent and success likelihood into a single probabilistic prediction of each player on the pitch being in possession of the ball immediately after the pass. This section details our approach.

4.1.1 Overall Setup

Each possession is treated as a sequence of timestamped on-ball actions, represented as events $\{g(t_1), g(t_2), \dots\}$. These events include passes, interceptions, and other relevant interactions, occurring relative to the possession start time ($t = 0$). For each pass event at time t_k, the goal is to predict the likelihood that any of the 21 other players (10 teammates, 11 opponents) will receive or intercept the ball.

Formally, each event is expressed as:

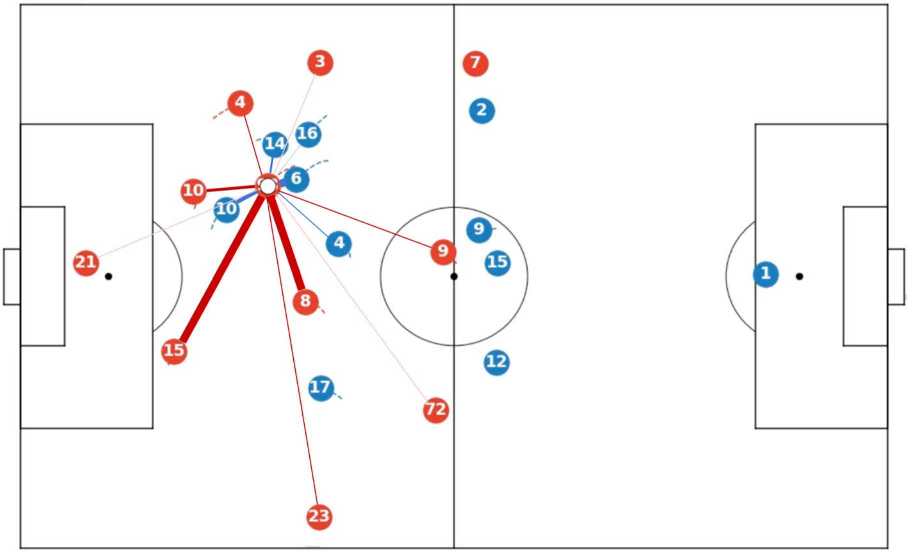

Fig. 2 An example snapshot of a match, where the width of the line between the ball carrier and each player represents the estimated probability of that player (whether a teammate or an opponent) receiving the pass from the ball carrier. The sum of all probabilities for all potential teammates and opponents adds up to 1

$$g(t_k) = \begin{cases} (v_b, v_i, \phi_{\text{pass}}) & \text{if ball carrier } b \text{ passes to teammate } i \\ (v_b, v_j, \phi_{\text{intercept}}) & \text{if opponent } j \text{ intercepts the pass from } b \end{cases} \tag{1}$$

Here, *v* represents the nodes in the underlying graph structure: v_b corresponds to the ball-carrying player (the passer), v_i denotes a potential receiving teammate, and v_j denotes a potential intercepting opponent. The ϕ terms represent the edges in the graph, encoding the interaction types between players-either a successful pass (ϕ_{pass}) or an interception ($\phi_{\text{intercept}}$).

TGN maintains memory across these sequential events within a possession window and resets it at the beginning of each new possession.

4.1.2 Node and Edge Features

Each player *p* (teammate or opponent) is represented as a node with the following features:

Node Features

- **Spatial features**
  - Absolute position: Cartesian coordinates (x_p, y_p)
  - Distance to ball carrier: $d_p = \|(x_p, y_p) - (x_b, y_b)\|_2$
  - Angle to ball carrier: $\theta_p = \arctan\left(\frac{y_p - y_b}{x_p - x_b}\right)$
- **Kinematic features**
  - Velocity vector: (u_p^x, u_p^y) (in m/s)
  - Speed: $\|(u_p^x, u_p^y)\|_2$
- **Role features**
  - Binary indicator $\mathbb{I}[p = b]$ if player p is the ball carrier**Edge Features** For every player pair (p_1, p_2):

- Euclidean distance between p_1 and p_2
- Binary team relationship: 1 for teammates, 0 for opponents

4.1.3 Model Architecture

The TGN constructs a fully connected dynamic graph of players at each timestep (c.g., Fig. 3). Messages are passed between player nodes based on edge features. Each player node maintains a memory state m_i updated over time:

$$\mathrm{m}_i(t_k) = f_{\text{mem}}\left(\mathrm{m}_i(t_{k-1}), \{e_{ij}(t_k)\}_{j \in \mathcal{N}_i}\right) \tag{2}$$

where f_{mem} is a learnable function (e.g., an LSTM or GRU), and e_{ij} denotes incoming edge messages from neighbors $\mathcal{N}_i$. These temporal updates allow the model to integrate historical context into current predictions.

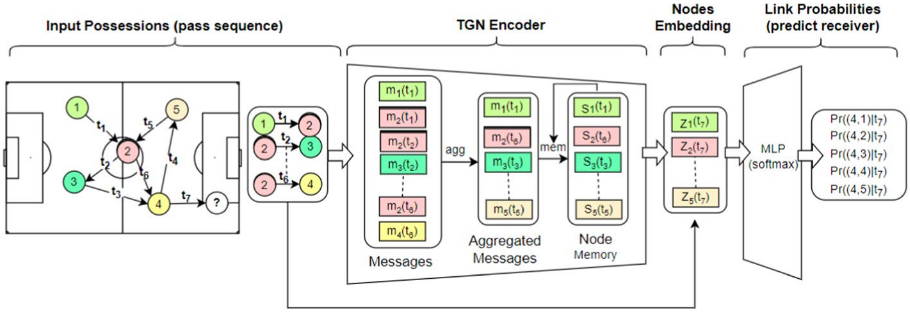

Fig. 3 TGN architecture for pass receiver prediction. Nodes represent players; edges encode spatial and team-based relations. The model outputs a probability distribution over all potential receivers 0

After processing the full history up to the moment of a pass, the TGN produces a logit for each potential receiver *k*:

$$\operatorname{logit}(k) = \begin{cases} \log P(S_i) + \log P(R_i \mid S_i) & \text{for teammate } i \\ \log P(\text{Intercept by } j) & \text{for opponent } j \end{cases} \tag{3}$$

Here, $P(S_i)$ denotes the probability that teammate i is the intended target of the pass (the selection), and $P(R_i \mid S_i)$ represents the probability that the pass is successfully received by teammate i, given that they were selected. The model also estimates the probability of an interception by any opponent j, treating it as a competing outcome.

These logits are transformed into a probability distribution across all players using softmax (Goodfellow et al., 2016):

$$P_{\mathrm{TGN}}(k) = \frac{e^{\operatorname{logit}(k)}}{\sum_{n \in \mathcal{V}_t} e^{\operatorname{logit}(n)}} \tag{4}$$

Here, $\mathcal{V}_t$ denotes the set of all 21 non-passing players (i.e., all teammates and opponents excluding the ball carrier) available at time *t* as potential pass recipients or interceptors.

This formulation ensures that the output probabilities across the 21 non-passing players sum to 1. A high $P_{\mathrm{TGN}}(i)$ implies that teammate *i* is both a likely target and has a high probability of receiving the ball successfully.

4.1.4 Training Objective

The model is trained using cross-entropy loss over the 21-class softmax output. For each training sample (i.e., a pass attempt), the ground truth label corresponds to the actual receiver:

$$\mathcal{L}_{\mathrm{CE}} = -\log P_{\mathrm{TGN}}(k^*) \tag{5}$$

where k^* is the index of the actual receiver (either a teammate in case of successful pass, or an opponent in case of interception).

During training, we use batches of passes sampled across games and possessions, allowing the TGN to generalize across different team formations, match contexts, and player roles. Further implementation details and ablation studies are provided in "Appendices A and B". Additionally, we provide the code for the general architecture of TGN in our GitHub repository [19].

4.1.5 Comparison to Decomposed Models

State-of-the-art decomposed approaches treat receiver prediction as a two-step problem: first estimating the selection probability $P(S_i)$ [e.g., using CNNs or gradient-boosted trees (Fernandez & Born, 2020; Robberechts et al., 2023)], and then independently estimating the success probability $P(R_i \mid S_i)$ [often via binary classification (Power et al., 2017)]. While intuitive, these decoupled pipelines require careful coordination between modules and separate supervision signals, which can introduce inconsistencies and make end-to-end optimization difficult.

In contrast, our TGN-based framework jointly models selection and success probabilities within a unified representation. This enables the model to learn implicit dependencies between offensive decision-making and spatial-temporal conditions without transition logic between stages. Moreover, the single softmax layer over all possible receivers eliminates the need for candidate selection heuristics, simplifying deployment and enhancing training stability.

4.2 TGN Limitations

The TGN's predictions are based on the game state at the moment the pass is executed, not when it is received. While this approach offers insights into the decision-making process, it may not fully capture the anticipatory nature of proposed passes, where players often aim for an anticipated arrival location rather than the receiver's current position. Additionally, this study focuses exclusively on passes that are either successfully received by a teammate or intercepted by an opponent. Passes that are deflected or go out of bounds are considered out of scope, as they do not result in a clear outcome for either team.

5 Experimental Results

In this section, we address the following research questions:

- **RQ1:** How does the predictive performance of TGN in estimating probabilities of reception compare to the existing decomposed models?
- **RQ2:** Which player positional groups (e.g., strikers, wingers, midfielders) tend to be more available for each considered pass type?
- **RQ3:** Which player positional groups are most effective at creating scoring opportunities upon reception of each of the considered pass type?
- **RQ4:** Are there differences in situations where only certain of our considered pass types are available?

5.1 TGN Performance Evaluation

In this section, we evaluate the performance of TGN in predicting the reception probability, which represents the likelihood that a given player (teammate or opponent) will receive the ball from a pass attempt. As mentioned in Sect. 4, the TGN directly predicts this probability by leveraging temporal and spatial relationships between players, encoded as node and edge features in a graph structure. In contrast, *decomposed models* (Robberechts et al., 2023; Fernandez et al., 2019) separately predict the probability of selecting a target location and the probability of success in delivering the pass using different methods and then combine them.

In the evaluation, we compare the TGN's performance against seven decomposed baselines. Among these baselines, SoccerMap (Fernandez & Born, 2020) is a deep learning model that predicts pass selection probabilities for every location on the field, rather than for specific players. SoccerMap uses a convolutional neural network (CNN) architecture to process spatial input channels, such as player positions, distances to the ball and goal, and angles between locations (details in Table 3). To adapt SoccerMap's location-based predictions to a receiver-based framework, we map the most likely pass destination to the closest teammate or opponent, as proposed by Robberechts et al. (2023). Specifically, for each pass, we identify the player nearest to the predicted pass destination and assign the selection probability to that player. This allows us to compare SoccerMap's performance with other receiver-based models, such as XGBoost and the TGN.

To evaluate the performance of the TGN and decomposed models, we consider the following combinations of models for predicting pass receivers:

1. *event-based XGBoost for Selection + SoccerMap for Success* Uses an XGBoost model trained on event data to predict selection and a SoccerMap model to predict success.
2. *contextual XGBoost for Selection + SoccerMap for Success* Uses an XGBoost model trained on event and tracking data to predict selection and a SoccerMap model to predict success.
3. *SoccerMap for Selection + event-based XGBoost for Success* Uses a SoccerMap model to predict selection and an XGBoost model trained on event data to predict success.
4. *SoccerMap for Selection + contextual XGBoost for Success* Uses a SoccerMap model to predict selection and an XGBoost model trained on event and tracking data to predict success.
5. *event-based XGBoost for Both* Uses XGBoost models trained on event data for both selection and success.
6. *contextual XGBoost for Both* Uses XGBoost models trained on event and tracking data for both selection and success.
7. *SoccerMap for Both* Uses SoccerMap models for both selection and success.

In order to estimate pass selection and success probabilities using the SoccerMap and XGBoost baselines, we rely on the publicly available implementation by Robberechts et al. (2023), provided in the GitHub repository: https://github.com/ML-KULeuven/un-xPass. The SoccerMap models are trained using the PyTorch Lightning framework with the adaptive moment estimation optimizer. We perform a grid search over the learning rate and batch size (16, 32, 64). Early stopping is used with a patience of 10 epochs and a delta of 1×10^{-3} for the pass success probability model, and 1×10^{-5} for the pass selection model.

For the XGBoost models, we apply the Tree-structured Parzen Estimator (TPE) algorithm to optimize hyperparameters including the maximum tree depth (ranging from 1 to 9), learning rate (between 0.01 and 0.25), L1 and L2 regularization terms, and the minimum loss reduction required to make a further partition on a leaf node (ranging from 1×10^{-8} to 1). Early stopping is applied with a patience of 100 boosting rounds. To produce a dense probability surface that covers the entire field, we simulate passes to every cell in a coarse 26×17 field grid and compute the associated feature representations. Predictions for each simulated pass are then mapped to their corresponding grid cell and later upscaled via bilinear interpolation for visualization.

It is important to note that our feature set differs substantially from the one used by Robberechts et al. (2023). Whereas their approach is based on StatsBomb 360 freeze-frame data,[1] which does not include visibility of all players, our models leverage complete tracking data, ensuring full-field visibility of all 22 players at every moment. In their XGBoost implementation, Robberechts et al. (2023) use a feature set composed of the ball height (categorized as ground level, shoulder height, or above shoulder level), the speed during the two preceding actions in the possession sequence, the time the passer held the ball before passing, the distance to the nearest defender from both the passer and receiver, the distance of the nearest defender to the line between the passer and receiver, and the number of opposing players located within a triangular corridor connecting the pass origin and destination, with a base of one meter at the receiver's end.

In contrast, we experiment with two XGBoost models trained with different feature sets. The first, an event-based model, uses only event-level data including normalized start and end locations of the pass, pass length, pass angle, pass direction, a binary indicator for whether the pass enters the penalty box, angle and distance to goal, and pass type. The second, a contextual model, enriches the input with tracking-derived features such as the passer's velocity, the velocities and distances of the nearest defenders relative to the passer, the angles of the nearest defenders to the passing line, and the time elapsed since possession was regained. The feature sets used for training the XGBoost and SoccerMap models are detailed in Tables 2 and 3, respectively. These features capture various aspects of player positioning, movement, and game context, enabling the models to make informed

Table 2 Feature sets of different models

Model	Selected features
Event-based XGBoost (trained with event data)	Normalized locations of the passes, pass length, pass angle, pass direction, a flag indicating whether the pass goes inside the penalty box, angle and distance to goal, pass type
Contextual XGBoost (trained with event and tracking data)	Velocity of the passer, velocities of the nearest defenders toward the passer, distances from the passer to the nearest defenders, the nearest defenders' angles to the passing line
TGN	(1) Node features: player's location, velocity, distance and angle from the ball carrier, and a flag indicating whether the player is the ball carrier for each time-step. (2) Edge features: distance and relationship (i.e., teammates or opponents) between the two interacting nodes

---

[1] https://statsbomb.com/what-we-do/soccer-data/360-2/

Table 3 Input channels for SoccerMap

Input channels	Description
`Home X,Y Locations Matrix`	Matrix of x,y locations of home team players at each time step
`Away X,Y Locations Matrix`	Matrix of x,y locations of away team players at each time step
`Home X Velocity Matrix`	Matrix of x-axis velocity of home team players at each time step
`Home Y Velocity Matrix`	Matrix of y-axis velocity of home team players at each time step
`Away X Velocity Matrix`	Matrix of x-axis velocity of away team players at each time step
`Away Y Velocity Matrix`	Matrix of y-axis velocity of away team players at each time step
`Distance to Goal Matrix`	Matrix of Euclidean distances of each pitch zone to the goal at each time step
`Distance to Ball Matrix`	Matrix of Euclidean distances of each pitch zone to the ball at each time step
`Angle-based Matrices`	Sine, cosine, and radian angles between zones, ball, and goal position
`Angle-based to Teammate Matrices`	Sine and cosine of angles between ball velocity vector and teammates
`Label`	Outcome of the pass in the next time step

predictions. For instance, the event-based XGBoost model uses pass-related features such as start and end locations, pass length, and pass angle, while the contextual XGBoost model incorporates additional tracking data, such as player velocities and defender positions. The TGN, on the other hand, leverages node and edge features that encode temporal and spatial relationships between players.

We evaluate model performance using three complementary metrics: *top-1 accuracy* measures the proportion of passes where the true receiver matches the highest-probability prediction; *ROC-AUC* evaluates the ranking quality of all potential receivers by their predicted probabilities; and the *Brier score* quantifies calibration error between predicted probabilities and empirical outcomes, ensuring probabilistic predictions are trustworthy for downstream applications.

Table 4 reports the evaluation metrics across all compared approaches. While the TGN-based model shows improved performance across the reported metrics, we are cautious in interpreting these results. The unified graph-based formulation may offer advantages in modeling the temporal and spatial dependencies between players, especially by jointly considering pass selection and outcome in a single framework. This integrated formulation avoids the need for separate modules and may help reduce inconsistencies introduced by decomposing the problem into multiple stages. Additionally, representing all players-including opponents-explicitly in the graph allows the model to condition its predictions on a fuller view of the match context. However, we refrain from making strong claims about the nature of what is being captured, and acknowledge that further work is needed to understand the specific types of decision-making the model learns.

Table 4 Evaluation metrics for TGN and decomposed models

Model	Top-1 Accuracy	ROC-AUC	Brier score
event-based XGBoost (Selection) + SoccerMap (Success)	0.61	0.65	0.022
contextual XGBoost (Selection) + SoccerMap (Success)	0.73	0.76	0.018
SoccerMap (Selection) + event-based XGBoost (Success)	0.58	0.66	0.028
SoccerMap (Selection) + contextual XGBoost (Success)	0.73	0.75	0.018
event-based XGBoost (Both)	0.69	0.71	0.019
contextual XGBoost (Both)	0.79	0.81	0.015
SoccerMap (Both)	0.69	0.71	0.012
TGN	**0.81**	**0.85**	**0.010**

5.2 Analyzing Player Positional Availability for Each Pass Type

We analyze player positional availability by estimating two key probabilities for each labeled positional group (defender, midfielder, winger, striker) and pass type (Type 1–3). This builds on the concept of availability as defined by Dick et al. (2022)-the probability that a pass reaches its target player successfully-but extends it to account for both positional roles and defensive interactions. To do so, we measure:

$$\text{Actual Reception Rate} = \frac{\#\ \text{passes received by position } p}{\text{Total passes of type } t} \in [0,1] \tag{6}$$

$$\text{Reception Potential} = \frac{1}{N_t}\sum_{n=1}^{N_t} P_{\text{TGN}}(i_n)\cdot \mathbb{I}[i_n \in p] \in [0,1] \tag{7}$$

where N_t is the number of type-t pass attempts, i_n denotes the n-th receiver, and $\mathbb{I}[\cdot]$ is an indicator function for position p.

In football terms, the actual reception rate shows how often each positional group actually gets the ball for different pass types in real game situations—a high percentage means they are frequently used in that phase of play, while low numbers suggest they are less involved. The Reception Potential measures how well-positioned players are to receive passes, where values closer to 1 indicate perfect availability behind defensive lines, and lower values show either poor positioning or strong marking by opponents. Together they reveal whether teams are making the most of their players' positioning - if a position has high Reception Potential but low Actual Reception, it suggests untapped passing opportunities, while the opposite may indicate forced passes into tight spaces. Our analysis assumes players adhere to their labeled positional groups (e.g., defender, winger) as defined in the dataset. However, dynamic scenarios-such as defenders advancing into attacking spaces-may temporarily alter

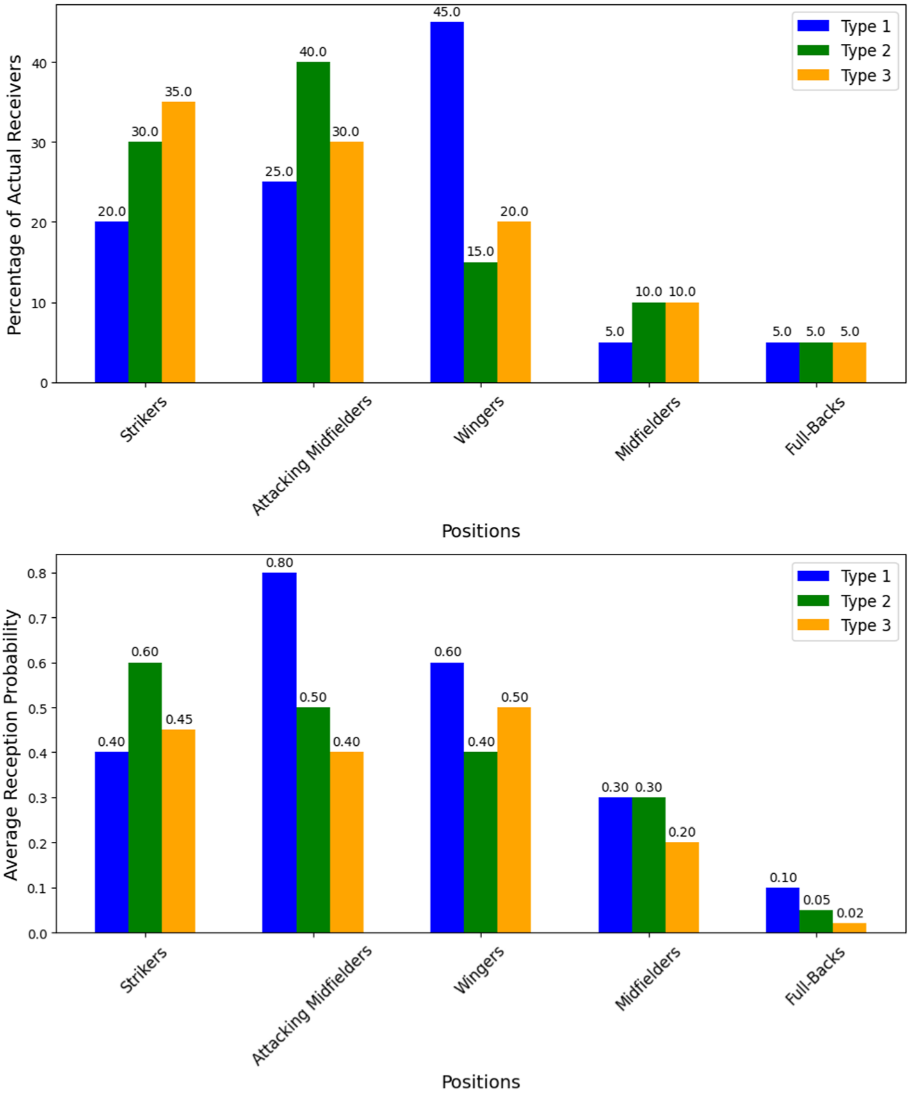

Fig. 4 Analysis of positional availability and actual receivers for Type 1, Type 2, and Type 3 passes. The top plot shows the percentage of actual receivers by position, normalized to sum to 100% for each pass type. The bottom plot shows the receiving potential, which is the sum of pass receiving probabilities of all players with a particular position, averaged over all frames for each pass type

their effective position. While this simplification allows systematic evaluation, it does not account for transient positional overlaps or position fluidity during play.

Figure 4 shows the distribution of these metrics across positions. The top panel displays actual reception rates, while the bottom panel shows average Reception Potential values. Key insights emerge when comparing these metrics:

Type 1: Bypassing defensive lines

1. *Actual reception* Wingers are the most targeted position (45% of passes), followed by attacking midfielders (25%) and strikers (20%). Full-backs receive the fewest passes (5%), consistent with their limited role in bypassing plays.
2. *Reception potential* Attacking midfielders demonstrate the highest availability (0.8), significantly exceeding their actual reception rate (+55 percentage points). Wingers show moderate availability (0.6) despite being the most targeted position, while strikers' potential (0.4) doubles their actual reception frequency.

Type 2: Defensive zone play

1 *Actual reception* Attacking midfielders receive the most passes (40%), followed by strikers (30%). Full-backs again show minimal involvement (5%).

2 *Reception potential* Strikers exhibit the highest availability (0.6), suggesting they're frequently open but underutilized (30% actual reception). Attacking midfielders show slightly lower potential (0.5) despite being the primary targets.

Type 3: Zone exits

1. *Actual reception* Strikers are the primary targets (35%), with attacking midfielders close behind (30%). Full-backs remain minimally involved (5%).
2. *Reception potential* Wingers show the highest availability (0.5) but receive only 20% of passes, while strikers' reception rate (35%) nearly matches their potential (0.45).

Three key tactical insights emerge:

- *Attacking midfielders* are consistently available (0.8 potential for Type 1) but may be overutilized in defensive zones (40% actual vs 0.5 potential)
- *Wingers* show strong availability for Types 1 (0.6) and 3 (0.5).
- *Full-backs* demonstrate minimal involvement across all pass types ($< 5\%$ actual, < 0.1 potential), suggesting systemic positional constraints

5.3 Identifying High-Value Receivers Through Threat Generation

We identify high-value receivers by estimating the scoring probability generated after each pass reception, using the action-based expected threat (xT) framework (https://soccermatics.readthedocs.io/en/latest/lesson4/xTAction.html). As detailed in Appendix C, our logistic regression model calculates the probability that a possession starting with the receiver will lead to a goal, considering factors like field position, defensive pressure, and game context (Table 8). This approach differs from traditional receiver analysis by quantifying the direct offensive value created through passing choices rather than just completion rates.

We quantify each player's offensive value by calculating the expected scoring probability for every pass reception using our logistic regression model (Appendix C), which estimates the likelihood of a possession ending in a goal from that point onward. These probabilities are then averaged across all receptions for individual players and positional groups (strikers, midfielders, etc.), enabling us to rank receivers by their threat generation capability.

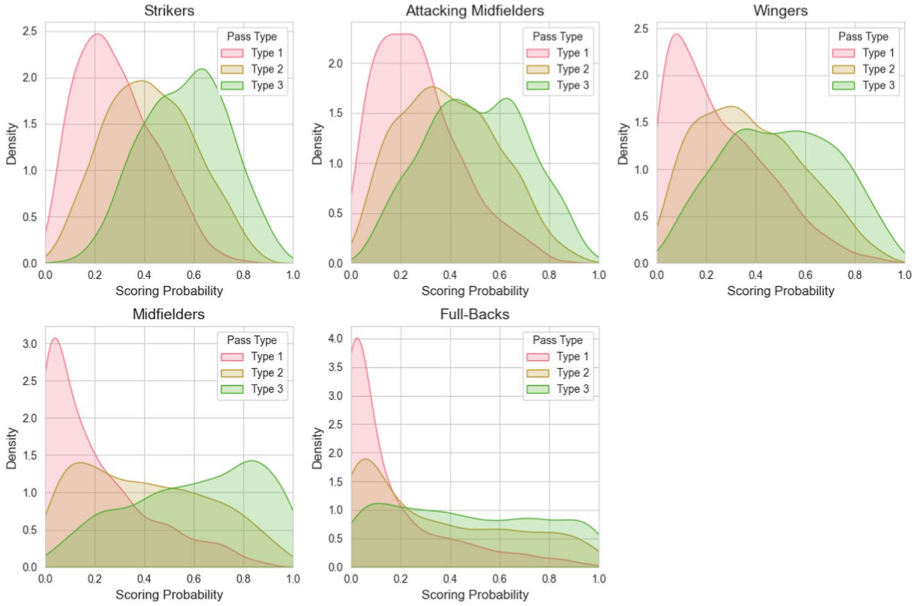

Fig. 5 KDE plots of scoring probabilities by player position and pass type. Each plot shows the distribution of scoring probabilities for a specific position and pass type, with darker areas indicating higher density

The resulting distributions, visualized in Fig. 5, reveal how scoring potential varies across positions and pass types. The analysis confirms that strikers and attacking midfielders are the most effective at converting receptions into goal-scoring opportunities, with strikers having the highest impact across all pass types. Attacking midfielders demonstrate versatility, creating chances from different types of passes. Wingers also contribute, particularly in transitional play. Midfielders primarily facilitate ball progression rather than direct scoring, while full-backs play a more limited position, occasionally contributing in attacking transitions. As expected, Type 3 passes, which penetrate and exit defensive structures, are the most likely to result in goal-scoring opportunities, followed by Type 2 passes that reach teammates within the defensive hull, while Type 1 passes, which bypass defensive lines without penetrating, are the least likely to lead to goals. Interestingly, the scoring distribution of Type 3 passes seems bimodal for attacking midfielders, suggesting that through passes towards this player position have two distinct categories in terms of effectiveness. These findings highlight the importance of targeting key attacking players with high-value passes while recognizing the varying effectiveness of different pass types in creating scoring opportunities.

5.4 Evaluating Situation Value Based on Pass Type Availability

In soccer, the value of a passing situation depends on both pass quality and the feasible pass types within the defensive structure. This analysis assesses situation value based on the availability of three pass types, regardless of the player's actual choice. We consider three scenarios:

- *Scenario A* All three pass types (Type 1, Type 2, and Type 3) are available, meaning that we have teammates available as the potential receivers over the defensive lines (Type 1), inside the defensive hull (Type 2) and beyond the defensive hull (Type 3).
- *Scenario B* Only Type 1 and Type 2 passes are available, meaning that we have teammates available as the potential receivers only over the defensive lines and inside the defensive hull. But there is no potential receiver beyond the hull.
- *Scenario C* Only Type 2 and Type 3 passes are available, meaning that we have teammates available as the potential receivers only inside and beyond the defensive hull, but no teammate over the defensive lines.

This analysis quantifies how the availability of different pass types influences potential offensive threat before execution. By examining these pre-pass situations, we identify systematic patterns in how passing options correlate with scoring probability under various defensive configurations. These findings offer coaches evidence-based insights about which passing lanes typically yield the highest value in specific game states, supporting more informed tactical planning.

To compute the *Situational Value V* for each passing situations, we calculate the maximum expected threat across all available receivers in front of the ball in that situation:

$$V = \max_{i \in \mathcal{A}} \left(P_{\mathrm{TGN}}(i) \cdot SP(x_i) \right) \tag{8}$$

where:

- $\mathcal{A}$ is the set of available receivers in passing situation
- $P_{\mathrm{TGN}}(i)$ is the reception probability for player *i* (Sect. 4)
- $SP(x_i)$ is the probability of scoring the goal at the end of possession from Appendix C if *i* receives the pass

This captures the optimal offensive potential for each passing situation. We then calculate the average Situation Value *V* for each scenario to measure the value of the scenarios where only certain pass types are available.

Figure 6 illustrates the average situation values and their standard deviations in our dataset. These results provide several key insights:

- *Scenario A* has the highest average situation value (0.08), indicating that the availability of all three pass types maximizes the offensive threat. The small standard deviation (0.0009) suggests that the situation values are highly consistent across situations, reflecting a stable and reliable offensive threat.
- *Scenario B* has a lower average situation value (0.04) compared to Scenario A, as the absence of Type 3 passes reduces the overall offensive threat. The slightly higher standard deviation (0.01) indicates moderate variability in situation values, reflecting the dependence on the effectiveness of Type 1 and Type 2 passes.
- *Scenario C* has an intermediate average situation value (0.05), higher than Scenario B but lower than Scenario A. The higher standard deviation (0.02) suggests greater variability in situation values, reflecting the risks and rewards associated with Type 3 passes.

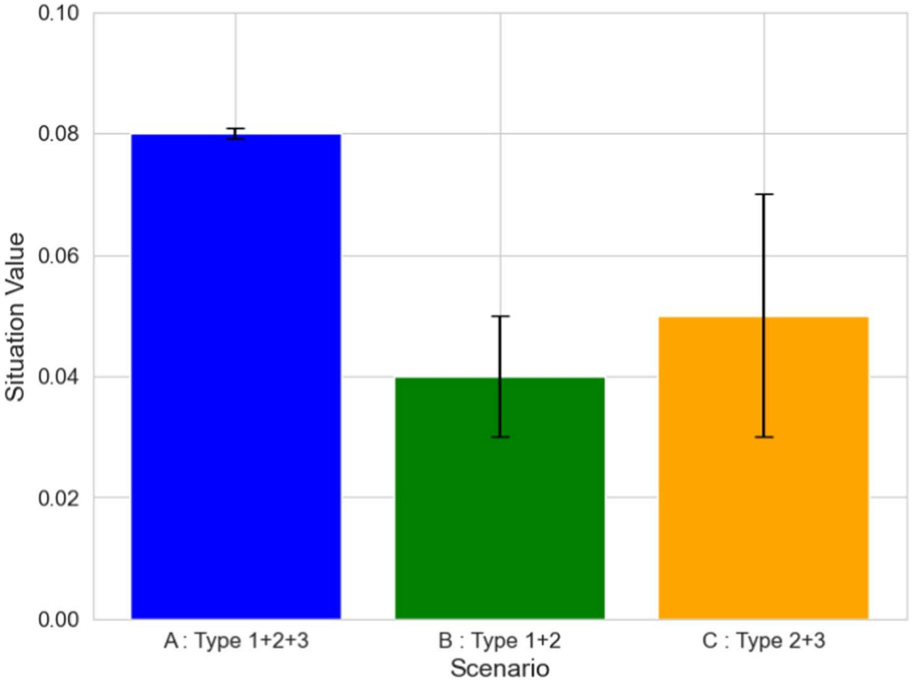

Fig. 6 Comparison of situation values for three scenarios based on pass type availability

6 Conclusion

This paper presents a framework for evaluating passing decisions against a defensive structure in soccer. By employing temporal graph networks (TGN), we directly estimate the probability of a pass being successfully received as an alternative solution to those decomposed models that rely on separate predictions of destination location selection and success probabilities. The TGN captures the dynamic, spatial, and temporal aspects of player interactions, providing a more accurate and unified approach to pass receiver estimation.

Our analysis focuses on forward passes initiated in the middle third of the pitch, where defensive structures are most organized, and classifies passes into three types based on their relation to the defending team: bypassing defensive lines (Type 1), reaching teammates within the defensive hull (Type 2), and penetrating and exiting the hull (Type 3). Through four key analyses, we demonstrate the practical applications of the TGN: (1) evaluating its performance in receiver prediction compared to baseline models, (2) analyzing the availability of players in different positions to receive each pass type, (3) identifying high-value receivers based on expected scoring probabilities, and (4) assessing the situational value of passing options based on the availability of pass types.

The results highlight the effectiveness of the TGN in capturing the dynamics of soccer through its availability of spatiotemporal modules and provide potential insights for coaches and analysts. By quantifying the availability and value of different passing options, this work offers a data-driven approach to potentially optimizing passing strategies. Future work could extend this framework to other areas of the pitch, considering players' locations after pass attempt, incorporate additional contextual factors, and explore its application in real-time decision-making systems.

TGN Architecture

The proposed TGN consists of an encoder-decoder pair, where an encoder is a function that maps players' interactions to node embeddings and a decoder takes the node embeddings as input and performs link prediction of the future time-steps. Figure 3 depicts the architecture of our network, which consists of the following modules:

- *Memory* A memory of the model $\{\mathrm{s}_i(t)\}_{i \in P}$ at time *t* is a representation of the node's history that the model has seen until *t*. It consists of a state vector $\mathrm{s}_i(t)$ for each player $i \in P$ with the set *P* of all the players in the game at *t*, which is updated after an event *x*(*t*). Note that *x*(*t*) can be either node-wise or interactive. When a substitute is sent onto the pitch and a new node is created, the network initializes a zero vector for it, and then updates the memory after each event the player is involved in.
- *Message function* For each event *x*(*t*) involving player *i*, the model computes a message $\mathrm{m}_i(t)$ to update *i*'s memory. When a node-wise event $\mathrm{v}_i(t)$ happens, a single message for *i* is computed as:

$$\mathrm{m}_i(t) = \mathrm{msg_n}(\mathrm{s}_i(t^-), t, \mathrm{v}_i(t)). \tag{9}$$

Likewise, an interaction event $\mathrm{e}_{ij}(t)$ induces the computation of messages for the passer *i* and the receiver *j* as follows:

$$\mathrm{m}_i(t) = \mathrm{msg_s}(\mathrm{s}_i(t^-), \mathrm{s}_j(t^-), \Delta t, \mathrm{e}_{ij}(t)) \tag{10}$$

$$\mathrm{m}_j(t) = \mathrm{msg_d}(\mathrm{s}_j(t^-), \mathrm{s}_i(t^-), \Delta t, \mathrm{e}_{ij}(t)) \tag{11}$$

where $\mathrm{s}_i(t^-)$ is the memory of *i* at the time of the last event before *t* in which the player is involved and $\mathrm{msg_n}, \mathrm{msg_s}, \mathrm{msg_d}$ are learnable message functions such as Multilayer Perceptron (MLP).

- *Message aggregator* Since each player *i* can be involved in multiple events until time *t*, we aggregate all the memories $\mathrm{m}_i(t_1), \ldots, \mathrm{m}_i(t_b)$ of *i* generated before *t* by averaging them, i.e.,

$$\bar{\mathrm{m}}_i(t) = \mathrm{mean}(\mathrm{m}_i(t_1), \ldots, \mathrm{m}_i(t_b)). \tag{12}$$

- *Memory updater* For each event *x*(*t*), a learnable memory update function (mem) updates the memory s_i of each player *i* involved in the event:

$$\mathrm{s}_i(t) = \mathrm{mem}(\bar{\mathrm{m}}_i(t), \mathrm{s}_i(t^-)). \tag{13}$$

In this work, we employ the structure of Long Short-Term Memory (LSTM) (Hochreiter & Schmidhuber, 1997) for this memory update function.

- *Embedding* Even if s_i is not updated at time t because player i is not involved in event $x(t)$, the context around i can change by interactions of other players. To reflect this, we also deploy the embedding module to generate the temporal embedding $\mathrm{z}_i(t)$ of i at any time t by

$$\mathrm{z}_i(t) = \sum_{j \in P_{-i}} h(\mathrm{s}_i(t), \mathrm{s}_j(t), \mathrm{e}_{ij}, \mathrm{v}_i(t), \mathrm{v}_j(t)) \tag{14}$$

where P_{-i} is the set of all 21 players other than i and h is Temporal Graph Attention proposed in Rossi et al. (2020).

TGN Training

Our TGN is trained to learn the future edge probabilities to predict the most likely receiver. Since the fundamental task is link prediction, we evaluate its performance through both transductive and inductive settings. In the transductive setting, we predict future links of the nodes observed during training, whereas in the inductive setting we predict future links of nodes never observed before. For all tasks, we used a split of 80%-10%-10% for the train-validation-test sets respecting the chronological order of the games to make sure actions from the same matches do not end up in both train and test sets and to avoid temporal information loss. Furthermore, we perform an ablation study to design the best combination of different components of TGN modules as described in Table 5. *msg* stands for the message function that we set as either a Multi-layer Perceptron (MLP) or identity (id) that is the concatenation of the inputs. *agg* is the aggregation function; in this work, we experiment with setting it as either to the most recent message (i.e., keep only most recent message for a given node) denoted as (last), and to the mean message (i.e., average all messages for a given node) denoted as (mean). *mem* is a learnable memory update function; in this work, we experiment with setting *mem* as either a Long Short-Term Memory (LSTM) or a Gated Recurrent Unit (GRU). *emb* is again a learnable function; in this work, we experiment with learning as one the following cases:

- *Identity (id)* $emb(i, t) = s_i(t)$, which uses the memory directly as the node embedding.
- *Temporal graph attention (attn)* A series of L graph attention layers compute i's embedding by aggregating information from its $L - hop$ temporal neighborhood. The temporal graph attention is able to select the neighbours that are the most important according to their features and timing information.The final results of Average Precision and AUC for different versions of TGN are illustrated in Table 6. According to this table, TGN-att outperforms the counterparts (by achieving 90% of AUC in the transductive setting) and we used its prediction results for all of our experiments for our deployed system.

Table 5 Multiple variants of TGN used in our ablation studies. msg, agg, mem, and emb stand for message function, aggregation function, memory updater, and embedding, respectively

	msg	agg	mem	emb
TGN-id	id	last	GRU	id
TGN-mean	MLP	mean	GRU	attn
TGN-last	MLP	last	GRU	attn
TGN-att	MLP	last	LSTM	attn

Table 6 Area under curve (AUC) and average precision (AP) for future edge prediction tasks in transductive and inductive settings of different versions of TGN. All the results are averaged over 10 runs

	Transductive		Inductive	
	AUC	AP	AUC	AP
TGN-id	0.76	0.73	0.63	0.64
TGN-mean	0.78	0.74	0.66	0.68
TGN-last	0.80	0.78	0.68	0.68
TGN-att	**0.90**	0.87	0.85	0.79

Scoring Probability Estimation

To calculate the reward metric introduced in Section 5, we need to estimate the scoring probability at the end of a possession which starts by the receiver of the pass. To do so, we employ a typical machine learning-based approach to estimate the probability of scoring a goal by the end of the possession from that moment.

We construct an action-based model using the following machine learning methods: Logistic regression, random forest, and XGBoost, which allows us to estimate the probability that an action will lead to a goal by the end of the possession. This is treated as a binary classification problem, where the label is 1 if the possession results in a goal, and 0 otherwise. This is not our contribution, this model is also known as action-based expected threat (xT) in soccer analytics.[2]

To develop this model, we consider the following features for each action within a possession:

- *Starting location* The location on the pitch where the ball is received.
- *Time to end* The remaining time until the match half ends.
- *Goal difference* The difference in goals between the teams at the moment of the action.
- *Zone on the pitch* The specific zone where the action takes place.
- *Distance and angle to goal* The action's proximity and orientation relative to the goal.
- *Number of bypassed opponents and teammates* The count of opponents and teammates bypassed by the action.
- *Pass type* Whether the pass is penetrative, long, short, or any other pass type.
- *Set piece or open play*: Whether the action occurs during a set piece or open play.Using these features, we train the machine learning models on our dataset to predict the probability of scoring from each point. The training process is conducted using an 80%-10%-10% split for train-validation-test sets. Based on the performance metrics shown in Table 7, Logistic Regression was selected as the best-fitting model for estimating xT across all actions, achieving a suitable balance between interpretability and predictive accuracy.

---

[2] https://soccermatics.readthedocs.io/en/latest/lesson4/xTAction.html.

Table 7 AUC of different classification models for estimating goal probability of the team at the end of the possession

Model	Team goal	Opponent goal
Random forest	0.75	0.71
XGBoost	0.79	0.75
Logistic regression	0.84	0.81

Table 8 Estimated coefficients for logistic regression models

Feature	Team goal model	Opponent goal model
Starting location (x, y)	0.18	−0.29
Time to end	0.05	0.15
Goal difference	−0.1	−0.05
Zone on the pitch	−0.2	−0.1
Distance to goal	−0.4	0.4
Angle to goal	0.3	−0.2
Bypassed opponents	0.5	−0.3
Bypassed teammates	−0.1	0.4
Pass type	0.4	−0.2
Set piece or open play	0.25	0.3

The resulting models from logistic regression was applied throughout the paper to calculate the reward metric, providing a measure of how each pass contributes to the team's overall threat level. Table 8 illustrates the coefficients of the features in our trained models.

Acknowledgements Project no. 2021−1.2.4-TÉT-2021-00053 has been implemented with the support provided by the Ministry of Culture and Innovation of Hungary from the National Research, Development and Innovation Fund, financed under the 2021−1.2.4-TÉT funding scheme.

Author Contributions All authors have contributed.

Funding Open access funding provided by Uppsala University. Not applicable.

Data Availability Not applicable.

Declarations

Funding Not applicable.

Conflict of interest Not applicable.

Ethics Approval Not applicable.

Consent to Participate Not applicable.

Consent for Publication Not applicable.

Availability of Data and Material Not applicable.

Code Availability Not applicable.

References

Anzer, G., Bauer, P., Brefeld, U., & Fassmeyer, D. (2022) Detection of tactical patterns using semi-supervised graph neural networks. In *16th MIT Sloan sports analytics conference*

Battaglia, P.W., Hamrick, J.B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V.F., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gülçehre, Ç., Song, H.F., Ballard, A.J., Gilmer, J., Dahl, G.E., Vaswani, A., Allen, K.R., Nash, C., Langston, V., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M.M., Vinyals, O., Li, Y. & Pascanu, R. (2018). Relational inductive biases, deep learning, and graph networks. *CoRR***abs/1806.01261**. http://arxiv.org/abs/1806.01261

Bauer, P., Anzer, G., & Shaw, L. (2023). Putting team formations in association football into context. *Journal of Sports Analytics, 9*(6), 39–59.

Bénézet, J.M. & Hasler, H. (2018) Youth football training manual. https://digitalhub.fifa.com/m/1b3da6976c9290aa/original/mxpozhvr2gjshmxrilpf-pdf.pdf. Accessed 30-03-2025

Bialkowski, A., Lucey, P., Carr, P., Yue, Y., Sridharan, S. & Matthews, I. (2021) Large-scale analysis of soccer matches using spatiotemporal tracking data. In *IEEE International Conference on Data Mining (ICDM)*

Bransen, L. & Van Harren, J. (2019). Measuring football players' on-the-ball contributions from passes during games. In *In ECML PKDD 2019 Workshop on Machine Learning and Data Mining for Sports Analytics* (MLSA 2019)

Bransen, L., Van Harren, J., & van de Velden, M. (2019). Measuring soccer players' contributions to chance creation by valuing their passes. *Journal of Quantitative Analysis in Sports, 15*(2), 97–116.

Decroos, T., Bransen, L., Van Haaren, J. & Davis, J. (2019) Actions speak louder than goals: Valuing player actions in soccer. In: In ACM KDD

Dick, U. & Brefeld, U. (2022) Action rate models for predicting actions in soccer. *AStA Advances in Statistical Analysis*

Dick, U., Link, D., & Brefeld, U. (2022). Who can receive the pass? A computational model for quantifying availability in soccer. *Data Mining and Knowledge Discovery, 36*, 987–1014. https://doi.org/10.1007/s10618-022-00827-2

Fernandez, J. & Born, L. (2020). SoccerMap: A deep learning architecture for visually-interpretable analysis in soccer. In *ECML PKDD*

Fernandez, J., Bornn, L. & Cervone, D. (2019) Decomposing the immeasurable sport: A deep learning expected possession value framework for soccer. In *13th MIT Sloan sports analytics conference*

Goes, F. R., Kempe, M., Meerhoff, L. A., & Lemmink, K. A. (2019). Not every pass can be an assist: A data-driven model to measure pass effectiveness in professional soccer matches. *Big Data, 7*(1), 57–70. https://doi.org/10.1089/big.2018.0067. pMID: 30321059.

Goodfellow, I., Bengio, Y. & Courville, A. (2016) *Deep Learning*. MIT Press. https://www.deeplearningbook.org/ (2016)

Gyarmati, L. & Stanojevic, R. (2016). A merit-based evaluation of soccer passes. In *ACM KDD workshop on large-scale sports analytics*

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780.

How impactful are line-breaking passes? https://bit.ly/39uQe6Q

HSNLab: TGN code for soccer (nd). https://github.com/hsnlab/sportsanalitica.

Karun: Expected threat (xt) (2024). https://karun.in/blog/expected-threat.html. Accessed: 11-12-2024

Kim, H., Choi, H.J., Kim, C.J., Yoon, J. & Ko, S.K. (2023) Ball trajectory inference from multi-agent sports contexts using set transformer and hierarchical bi-LSTM. In *Proceedings of the 29th ACM SIGKDD international conference on knowledge discovery and data mining*

Line-breaking passes and ball receipts in space. https://statsbomb.com/articles/soccer/statsbomb-launch-new-360-metrics-line-breaking-passes-and-ball-receipts-in-space-2/

Liu, G., Luo, Y., Schulte, O. & Kharra, T. (2020) Deep soccer analytics: Learning an action-value function for evaluating soccer players. *Data Mining and Knowledge Discovery* **34**(2)

Mehrasa, N., Zhong, Y., Tung, F., Bornn, L. & Mori, G. (2018). Deep learning of player trajectory representations for team activity analysis. In: *12th MIT Sloan sports analytics conference*

Memmert, D. (2015). *Teaching tactical creativity in sport: Research and practice*. Routledge. https://doi.org/10.4324/9781315797618.

Peralta Alguacil, F., Fernandez, J., Piñones Arce, P. & Sumpter, D. (2020) Seeing in to the future: using self-propelled particle models to aid player decision-making in soccer. In *In Proceedings of the 14th MIT Sloan sports analytics conference*

Power, P., Ruiz, H., Wei, X. & Lucey, P. (2017) Not all passes are created equal: Objectively measuring the risk and reward of passes in soccer. In *ACM KDD*

Rahimian, P., da Silva Guerra Gomes, D.G., Berkovics, F. & Toka, L. (2022) Let's penetrate the defense: A machine learning model for prediction and valuation of penetrative passes. In: *Machine learning and data mining for sports analytics*

Rahimian, P., Kim, H., Schmid, M. & Toka, L. (2023) Pass receiver and outcome prediction in soccer using temporal graph networks. In *In ECML PKDD 2023 workshop on machine learning and data mining for sports analytics (MLSA 2023)*

Rahimian, P., Oroojlooy, A. & Toka, L. (2021) Towards optimized actions in critical situations of soccer games with deep reinforcement learning. In *IEEE DSAA*

Rahimian, P., Van Haaren, J. & Toka, L. (2023) Towards maximizing expected possession outcome in soccer. *International Journal of Sports Science and Coaching*

Rahimian, P., Van Haaren, J., Abzhanova, T. & Toka, L. (2022) Beyond action valuation: A deep reinforcement learning framework for optimizing player decisions in soccer. In *16th MIT Sloan sports analytics conference*

Rein, R., Raabeb, D., & Memmert, D. (2017). Which pass is better? Novel approaches to assess passing effectiveness in elite soccer. *Human Movement Science, 55*, 172–181.

Robberechts, P., Roy, M.V., & Davis, J. (2023) un-xPass: Measuring soccer player's creativity. In *Proceedings of the 29th ACM SIGKDD international conference on knowledge discovery and data mining*

Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F. & Bronstein, M. (2020). Temporal graph networks for deep learning on dynamic graphs. In *ICML 2020 workshop on graph representation learning*

Soccermatics: Soccermatics official website (nd). https://soccermatics.readthedocs.io/en/latest/lesson4/xTAction.html

Sotudeh, H. (2021) Potential penetrative pass (P3). In *STASBOMB conference*

Spearman, W.R., Basye, A.T., Dick, G.J., Hotovy, R. & Hudl, P.P. (2017). Physics-based modeling of pass probabilities in soccer. In: *MIT Sloan sports analytics conference*

Stats perform: Stats perform official website (nd). https://www.statsperform.com/. Accessed: 01 02 2025

Stöckl, M., Seidl, T., Marley, D. & Power, P. (2021) Making offensive play predictable: Using a graph convolutional network to understand defensive performance in soccer. In: *MIT Sloan sports analytics conference. Stats perform*. https://cdn.prod.website-files.com/5f1af76ed86d6771ad48324b/607a44a3c3d021c9cb376186_PaulPower-OffensivePlaySoccer-RPpaper.pdf

Szczepanski, L., & McHale, I. (2016). Beyond completion rate: evaluating the passing ability of footballers. *Journal of the Royal Statistical Society. Series A (Statistics in Society), 179*, 513–533.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L. & Polosukhin, I. (2017). Attention is all you need. In: *Advances in Neural Information Processing Systems*