# Player-aware EPV: Separating Skill From Tactical Context in Possession Value

Minho Lee[1,2], Miru Hong[3], Hyeokje Cho[3], Songyun Jeong[3], Dario Andrea Prifti[1,4], Pascal Bauer[1,5], Sang-Ki Ko[3]

[1] *Saarland University, Germany*
[2] *VfB Stuttgart 1893 AG, Germany*
[3] *University of Seoul, South Korea*
[4] *SV 07 Elversberg, Germany*
[5] *Deutscher Fussball-Bund (DFB), Germany*

## 1. Introduction

Quantifying player contributions in football requires evaluating actions beyond goals and assists. While expected goals (xG) and expected assists (xA) characterize shot quality and chance creation [1, 2], descriptive measures such as progressive passes, carry distance, and pass length do not directly quantify how individual actions influence subsequent outcomes. Action valuation frameworks address this limitation by estimating the expected contribution of on-ball actions to future scoring and conceding. Expected Threat (xT) assigns value to ball progression between pitch locations [3], Valuing Actions by Estimating Probabilities (VAEP) measures changes in scoring and conceding probabilities [4], and On-Ball Value (OBV) evaluates actions through changes in expected goals for and against [5]. By expressing heterogeneous actions on a common value scale, these frameworks enable systematic evaluation and aggregation of player contributions.

Expected Possession Value (EPV) estimates the expected outcome of a possession from a given game state, incorporating the spatial configuration and movement of players and the ball to account for available options and defensive positioning. Fernández et al. [6] introduced a decomposed EPV framework that combines action-selection probabilities with conditional values for passes, drives, and shots, enabling the evaluation of both observed actions and available alternatives. However, their framework models a population-average action policy without explicitly accounting for individual decision-making preferences. Moreover, aggregated player values remain influenced by differences in tactical contexts and available opportunities, even when individual actions are evaluated conditionally on their observed states. For scouting and recruitment, this makes it difficult to determine whether differences in player value reflect individual decision-making ability or the tactical environments in which players operate. Although player-specific EPV modelling has been explored in basketball [7], passes have been evaluated against the choices of a typical player [8], and player-conditioned event modelling in football [9, 10], evaluating individual decision policies under comparable game situations remains challenging. Existing tracking-based EPV has also been evaluated on relatively large datasets, such as the 633 matches used by Fernández et al. [6], leaving its applicability to smaller datasets less well understood. These challenges motivate the evaluation of player-specific decision policies over shared role-specific game states using a common option-value function.

In this work, we develop a tracking-based EPV framework that combines a shared player-token encoder with task-specific prediction heads for action selection and value estimation, alongside a separate convolutional network for pass-related probabilities and values based on SoccerMap [11]. Unlike the original EPV framework, which uses scoring outcomes, we employ the signed expected goals (xG) of the first subsequent shot within 15 seconds as the value prediction target. Building on this framework, we investigate player-specific decision-making at two levels: action selection among passing, driving, and shooting, and pass-recipient selection among available teammates. We assess player-specific action-selection tendencies and their decision values relative to an average action policy, while modelling pass-recipient preferences using a hierarchically regularized conditional logit. To account for tactical context, we compare these decisions across tactical situations within positions and evaluate recipient-selection policies over shared role-specific game states using a common option-value function.

We evaluate the proposed framework using tracking and event data from 30 Serie A matches. Despite the limited data, the model outperforms constant baselines for drive success, shot value, and pass completion, and its EPV accuracy plateaus at approximately 20 training matches. Compared with the original EPV formulation reimplemented on our data, our model improves the correlation between predicted EPV and the realized 15-second xG outcome on non-shot states from 0.217 to 0.237 and reduces the action-selection negative log-likelihood by 1.12%. Using this model, we evaluate player-specific decision policies across tactical contexts. Players' action-selection tendencies show high between-player reliability, their recipient preferences carry over to held-out matches, and the change in pass share under pressure differs by position while being largely reproduced by the average-player policy. These results support the use of the framework for context-standardized player evaluation.

---

## 2. Methods

### 2.1 Data and Preprocessing

We use tracking data from Hudl and event data from StatsBomb for 30 Serie A matches in the 2025/26 season, comprising all home and away matches among six teams. The tracking data, recorded at 25 Hz, provide the positions and movements of players and the ball, while the event data provide information on on-ball actions, their outcomes, and shot-level expected goals (xG).

To construct the EPV training dataset, we synchronize event timestamps with tracking frames and normalize player and ball coordinates to a common attacking direction. We identify individual ball-possession intervals and extract decision states at one-second intervals. Each state is labelled as a pass or shot if the player releases the ball with that action within the next second, and as a drive otherwise, including when the player holds the ball. Observations with insufficient tracking or alignment quality are excluded, resulting in 51,483 states: 19,313 passes, 31,766 drives/holds, and

404 shots. Model inputs are restricted to information available at the observation time, while subsequent events are used to construct prediction targets.

### 2.2 Expected Possession Value

Expected Possession Value (EPV) represents the expected future outcome of a possession given the current game state. Following Fernández et al. [6], EPV is decomposed into the probabilities of selecting available actions and their corresponding expected values:

$$\mathrm{EPV}(s) = \mathrm{E}[G \mid s] = \sum_a \overline{\pi}(a \mid s) Q(s, a) \,. \#(1)$$

Here, $s$ denotes the game state, $G$ the future outcome, $\pi^{-}(a \mid s)$ the probability of selecting action $a$, and $Q(s, a)$ its expected value. The action space consists of passing, driving, and shooting. For passes, the expected value additionally integrates over potential destinations $z$, accounting for pass-success probability $p_s(s, z)$ and the expected values following successful and unsuccessful passes, $V_s(s, z)$ and $V_f(s, z)$, respectively:

$$Q(s, \mathrm{pass}) = \int p(z \mid s, \mathrm{pass}) \left[ p_s(s, z) V_s(s, z) + (1 - p_s(s, z)) V_f(s, z) \right] \mathrm{d}z \,. \#(2)$$

The original framework estimates pass-related probability and value surfaces using convolutional neural networks [11], while action selection, driving, and shooting are modelled using scalar prediction networks. Its value prediction target is based on signed scoring outcomes, assigning +1 when the team taking the action scores the next goal, –1 when the opponent scores, and 0 otherwise. In their implementation, only goals occurring within 15 seconds of an action contribute a nonzero value, corresponding to the average possession duration in their data.

### 2.3 Proposed EPV Framework

We retain the EPV decomposition of Fernández et al. [6] while modifying the value prediction target and model architecture. Given the limited number of matches, using actual scoring outcomes within a 15-second window would result in highly sparse nonzero training labels. To address this, we use the signed expected goals (xG) of the first subsequent shot by either team within 15 seconds as the value prediction target, providing a more informative learning signal than realized goals alone. The target is positive for shots taken by the team in possession, negative for opposing shots, and zero when no shot occurs within the window. For states in a possession spell that ends with a shot, that shot is always used as the first subsequent shot, so that small event-tracking timing differences cannot exclude it from the target.

Figure 1: Architecture of the proposed tracking-based EPV model.

Figure 1 illustrates the proposed architecture, which consists of two separately trained branches: a shared player-token encoder for scalar predictions and a convolutional network for pass-related spatial predictions. Each game state is represented by up to 22 player tokens, each containing 13 features describing position, velocity, relative spatial relationships, and team, ball-carrier, and goalkeeper indicators. A shared multilayer perceptron (MLP) embeds the tokens into 64-dimensional representations, followed by two residual processing blocks. Team-specific mean and max pooling, together with the ball-carrier representation, produce a 320-dimensional vector that is projected to 32 dimensions and concatenated with 20 additional spatial and contextual features (Appendix B).

The resulting representation is used by five task-specific prediction heads: action selection, drive success, successful-drive value, failed-drive value, and shot value. Each head consists of a two-hidden-layer MLP, with cross-entropy losses for action and success probabilities and mean squared error for value predictions. The shared encoder and prediction heads are jointly optimized using AdamW with a learning rate of 0.001 and batch size of 256, with early stopping based on validation performance. Each head's loss is computed over the states for which it is defined and averaged over the full batch, so that heads with few eligible states, such as the failed-drive value, contribute in proportion to their frequency, and each head is taken from the epoch with its lowest validation loss.

The separate pass convolutional network, based on SoccerMap [11], estimates four spatial surfaces: pass-destination selection, pass-success probability, and the expected values following successful and unsuccessful passes. Together, the two branches provide nine prediction components. Their outputs are combined at inference according to the EPV decomposition in Equations (1) and (2) to estimate possession value for each game state.

### 2.4 Player-aware Decision Analysis

We analyse player decisions at three levels, each measured against the average-player policy: action selection among passing, driving, and shooting (Section 2.4.1), pass-recipient selection among available teammates (Sections 2.4.2–2.4.3), and the variation of both across positions and tactical contexts (Section 2.4.4). In each case, we separate the decision preference from its modelled value and assess whether player-level differences are reliable.

#### 2.4.1 Action Selection

We first examine individual preferences among passing, driving, and shooting relative to the population-average action policy. Player-specific action tendencies are measured by averaging the difference between observed actions and their predicted selection probabilities:

$$r_i(a) = \frac{1}{|S_i|} \sum_{s \in S_i} \left[ 1(a_s = a) - \bar{\pi}(a \mid s) \right]. \#(3)$$

Here, $S_i$ denotes the observed states of player $i$, and $a_s$ represents the selected action. These residuals quantify how frequently a player selects an action relative to its expected frequency under the average policy. We further evaluate action-selection value by comparing the expected value of the chosen action, $Q(s, a_s)$, with the average-policy value, $EPV(s)$. This analysis is performed across all eligible states and at the initial and terminal states of individual possession intervals.

#### 2.4.2 Player-conditioned Pass-recipient Selection

We further investigate individual decision-making through pass-recipient selection. For each pass-ending possession interval, we extract the state immediately before the pass and identify available teammates, excluding the ball carrier and goalkeeper. The primary analysis includes 11,596 completed passes with identified recipients.

For each pass, the passer chooses one recipient among the available teammates. We model recipient-selection probabilities using a hierarchically regularized conditional-logit model [12]. Each passing option is characterized by 11 features capturing pass-success probability, expected value, passing geometry, defensive pressure, and recipient positioning. The probability that player $i$ selects option $o$ in state $s$ is

$$p_i(o \mid s) = \frac{\exp\left(\beta_i^{\mathrm{T}} x_{s,o}\right)}{\sum_{o' \in O(s)} \exp\left(\beta_i^{\mathrm{T}} x_{s,o'}\right)}, \#(4)$$

$$\beta_i = \beta + \beta_r + \beta_{r,\mathrm{side}} + b_i, \#(5)$$

where $O(s)$ denotes the available passing options, $x_{s,o}$ the 11 option features, $\beta$ the shared coefficients, $\beta_r$ and $\beta_{r,\mathrm{side}}$ the adjustments for the passer's positional role and for the side of the pitch on which the passer usually plays within that role, respectively, and $b_i$ the player-specific adjustment. Positional roles comprise eight groups: centre-back, left and right full-back/wing-

back, defensive midfielder, central midfielder, attacking midfielder, winger, and centre-forward. The average-player policy $p_0$ for the passer's role and side uses the same expression without $b_i$. Each adjustment is drawn from a zero-mean normal distribution whose variance is estimated from the data, so players with few passes remain close to their role average.

### 2.4.3 Role-standardized Policy Valuation

To compare recipient-selection policies under comparable game situations, we estimate a common value function for available passing options. Pass-success probabilities are obtained from the pass CNN, while expected values following successful and unsuccessful passes are estimated using separate gradient-boosting models trained on calibrated EPV predictions at observed reception and possession-loss states.

We evaluate each player's recipient-selection policy over a common set of game states: for positional role $r$, $S_r$ contains the pre-pass states of all completed passes made by players of that role, pooled across all teams and matches. The standardized decision value is defined as:

$$D_i = \frac{1}{|S_r|} \sum_{s \in S_r} \sum_{o \in O(s)} \left[ p_i(o \mid s) - p_0(o \mid s) \right] Q(s, o), \#(6)$$

where $Q(s, o)$ is the expected value of option $o$ under the common value function. For each state, the inner sum is the difference between the expected option value under the player's policy and under the average-player policy for the player's role and side; $D_i$ averages this difference over $S_r$, so a positive $D_i$ indicates that, in the same situations, the player places more weight on higher-value options. Evaluating players on identical reference states with a common value function allows differences in modelled decision preferences to be assessed independently of their observed opportunity distributions. The resulting measure represents the expected value of recipient selection conditional on passing, rather than the overall value of action selection.

### 2.4.4 Contextual Decision Analysis

We further examine how decisions vary across positions and tactical contexts.
Each state is described from two perspectives at the decision time. From the ball carrier's perspective, we record the carrier's location by pitch third (defensive, middle, or final third) and by lateral channel (left, central, or right third of the pitch width, from the attacking team's perspective), and whether the carrier is under close pressure (nearest defender within 2 m, or able to reach the carrier within 1 s). From the defending team's perspective, we record whether the opponent is pressing high (at least eight opposing outfield players in the attacking team's own half). Players are grouped into six positions by merging the eight roles defined in Section 2.4.2 (left and right full-backs; defensive and central midfielders), so that each position group has sufficient states. For each position, we compare the pass rate, the pass direction, and the value difference between passing and driving, $Q(s, \text{pass}) - Q(s, \text{drive})$, across contexts. These comparisons measure differences between two conditions while holding the remaining context variables fixed. We then test whether these differences vary between positions and compare the observed changes with those predicted by the average-player policy, which receives no position information.

At the player level, we decompose between-player variation in action tendencies into position, team, and individual components, and measure how consistently each player's tendencies hold across contexts.

**2.5 Evaluation Protocol**

To evaluate the EPV model (Sections 2.2–2.3), we split the 30 matches into six folds of five matches; each fold serves once as the test set, with 20 of the remaining matches used for training and five for validation, and each model is trained with three random seeds. The same match-based folds are used for cross-fitting throughout the study: the recipient-selection policies and option-value models are also fitted without the matches on which they are evaluated. EPV prediction performance is assessed using the correlation between predicted EPV and the realized target (signed xG of the first shot within 15 s), reported separately for all states and for non-shot states, together with action-selection negative log-likelihood and task-specific prediction losses. Non-shot states serve as the primary measure because at shot states the target is largely the xG of the shot being taken. We compare our proposed player-token architecture with the original EPV formulation [6] reimplemented on our data and an attention-based encoder, while evaluating the pass-success model against provider estimates. Uncertainty is estimated using 2,000 match-cluster bootstrap resamples unless stated otherwise.

We first assess whether player-specific action-selection tendencies (Section 2.4.1) represent reliable between-player differences rather than match-level sampling noise, using a reliability measure that compares between-player variation with this noise. Reliability is also evaluated after team-by-role centring to examine whether individual differences persist beyond team and role averages. The value of these choices is assessed using the same reliability measure and its association with subsequent 15-second xG outcomes. For comparison, we apply the same reliability measure to the provider's on-ball value (OBV) and EPV-added, the change in EPV from the start of a player's possession to the start of the next possession. This comparison assesses the stability of different player-value metrics rather than model prediction accuracy.

We then examine whether decision values derived from player-conditioned recipient-selection policies (Sections 2.4.2–2.4.3) generalize to held-out matches. For each pass, empirical recipient-selection value is defined as the estimated value of the chosen recipient minus the expected value under the average recipient-selection policy for the passer's role and side, $p_0(o \mid s)$, in the same state. We test whether a player's standardized decision value $D_i$, estimated without a given fold, predicts their mean empirical selection value in that fold. We further assess whether this relationship persists after controlling for forward-pass frequency and team-by-role averages. We also examine whether passes with higher empirical selection values are followed by greater signed xG within 15 seconds.

Finally, the contextual analysis (Section 2.4.4) examines whether decisions change across tactical contexts in a position-specific manner and whether the average-player policy reproduces these changes. To ensure sufficient data coverage, we restrict the analysis to comparisons in which each

position–context group contains at least 200 states from at least 10 matches. We test whether the observed context effects differ between positions, applying Holm correction for multiple comparisons. We assess how well the average-player policy reproduces observed changes using sign agreement, correlation, and calibration slope across position-specific contrasts. Robustness is evaluated by leaving out each team in turn and repeating the analysis across 500 random splits of the 30 matches into two halves. At the player level, consistency of action-selection tendencies across contexts is assessed using split-half correlations corrected for measurement unreliability.

## 3. Experimental Results

Table 1: Comparison of EPV prediction architectures on 30 Serie A matches.

| Model | Input features | Non-shot EPV corr. [95% CI] | All-state EPV corr. | Action NLL |
|---|---|---|---|---|
| Tabular baseline [6] | 22 tabular | 0.2173 [0.1874, 0.2473] | 0.2980 | 0.5988 |
| Attention encoder | 22 + player tokens | 0.2335 [0.2021, 0.2638] | 0.3242 | 0.5943 |
| Token MLP | 22 + player tokens | 0.2333 [0.2021, 0.2649] | 0.3240 | 0.5941 |
| Attention encoder | 20 + player tokens | 0.2314 [0.1982, 0.2644] | 0.3245 | 0.5902 |
| **Proposed token MLP** | **20 + player tokens** | **0.2370 [0.2038, 0.2701]** | **0.3288** | **0.5921** |

*Note: EPV correlation (corr.) is the Pearson correlation between predicted EPV and the 15-second xG target. NLL denotes negative log-likelihood. Results are averaged over three random seeds, with 95% confidence intervals estimated using 2,000 match-cluster bootstrap resamples.*

### 3.1 EPV Model Performance

We evaluate the proposed EPV model against a tabular baseline following the feature-based design of Fernández et al. [6] and alternative player-token encoders (Table 1). The 22 tabular features include two pitch-control summaries [13]. Since player tokens directly represent player positions and velocities, we also evaluate models without these summaries (20 features) to assess whether they remain necessary.

Adding player tokens improves non-shot EPV correlation from 0.2173 to 0.2333 with the token MLP. The proposed model, using 20 tabular features, achieves the highest correlation of 0.2370, an improvement of 0.0197 (95% CI: 0.0084–0.0320) over the baseline, while reducing action-selection NLL by 1.12%. Compared with the 22-feature attention encoder, the proposed model satisfies the prespecified non-inferiority margins of −0.02 for correlation ($\Delta$ = +0.0035, 95% CI: −0.0061 to 0.0131)

and +0.0025 for action NLL (Δ = −0.0022, 95% CI: −0.0044 to −0.0001). These results support the use of a simpler token-wise architecture without explicit pitch-control summaries. Figure 2 further illustrates how predicted EPV changes across successive frames of an example possession, reflecting the evolving spatial configuration of players and the ball.

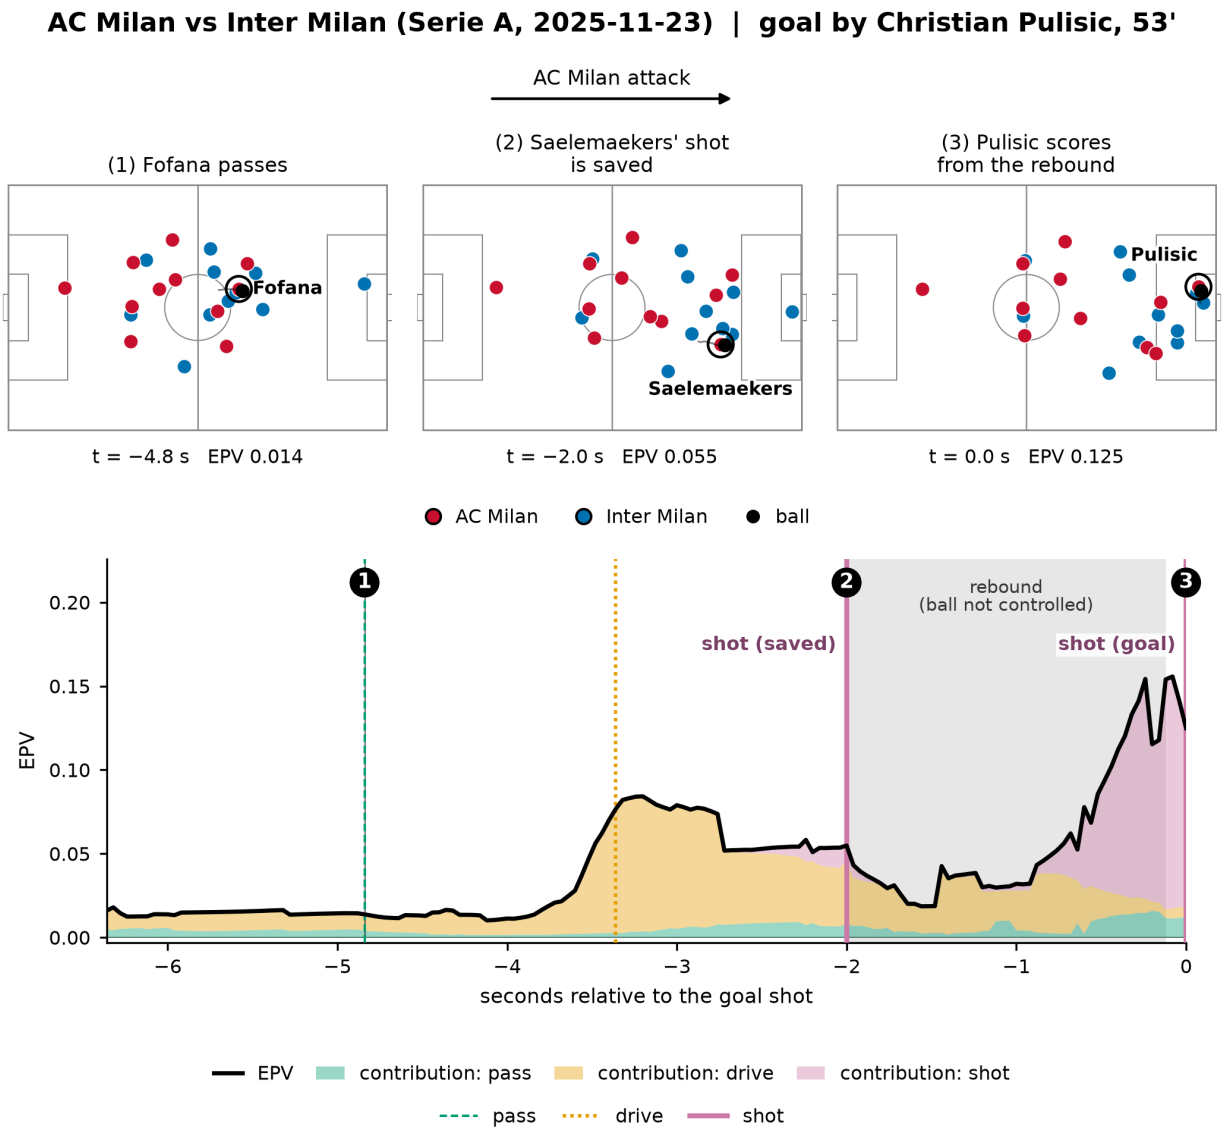



Figure 2: Evolution of predicted EPV during an example possession leading to a goal. The top row shows three key moments: (1) Youssouf Fofana's pass, (2) Alexis Saelemaekers' saved shot, and (3) Christian Pulisic's goal from the rebound. The lower panel shows the EPV trajectory relative to the goal shot, decomposed into pass, drive, and shot contributions.

Individual prediction components are evaluated against fold-specific constant baselines (Appendix C). Drive-success and shot-value predictions reduce their respective losses by 30.18% and 40.24%, while failed-drive value prediction does not improve over the baseline. The pass CNN also improves destination-free pass-completion prediction (AUC 0.80; log-loss reduction 18.84%). At observed pass destinations, it achieves an AUC of 0.9396 compared with 0.9178 for the provider estimates

across 19,221 passes. At the action level, EPV-added is moderately correlated with the independently supplied OBV (Pearson r = 0.49, 95% CI: 0.43–0.55). Aggregated per player, the two rankings also agree moderately (Appendix C). A learning-curve analysis further shows that non-shot EPV correlation changes little beyond 20 training matches (Appendix A).

Overall, these results show that the player-token architecture improves EPV prediction over the tabular baseline while maintaining comparable performance without explicit pitch-control summaries. They also demonstrate that useful EPV predictions can be learned from a relatively small tracking dataset, providing a basis for the player-level decision analyses that follow.

### 3.2 Player-level Action Selection

Building on the EPV model, we examine whether players exhibit consistent differences in their action-selection tendencies and the expected values of their choices. State-conditioned action tendencies show substantial between-player reliability, reaching 0.719 for both passing and driving, compared with 0.401 for shooting (Table 2). After team-by-role centring, reliability decreases to 0.457 for passing and 0.464 for driving, suggesting that these differences are partly associated with team and role context.

In contrast, action-selection value shows lower between-player reliability, ranging from 0.000 across all states to 0.213 at terminal possession states.

Table 2: Between-player reliability of action-selection tendencies.

| Action | Reliability | Team-by-role centred |
|---|---|---|
| Pass | 0.719 | 0.457 |
| Drive | 0.719 | 0.464 |
| Shot | 0.401 | 0.000 |

### 3.3 Validation of Standardized Pass-recipient Decision Value

We next examine whether player-specific pass-recipient preferences provide reliable decision-value estimates that generalize to held-out matches. The standardized decision value ($D_i$) is positively associated with empirical recipient-selection value in held-out matches (slope 0.917, 95% CI: 0.455–1.387; Table 3). This association remains positive after controlling for forward-pass frequency, suggesting that the measure captures information beyond simple passing tendencies.

Empirical recipient-selection value is positively associated with subsequent signed xG (Table 3), with a between-player reliability of 0.371. Robustness checks retain a positive held-out association when failed passes are included (slope 0.849, 95% CI: 0.411–1.291), while synthetic policy tests support sensitivity to value-relevant choices. However, the association after team-by-role control remains inconclusive (Table 3). Overall, these results support the use of standardized decision values to evaluate player-specific pass-recipient policies.

Table 3: Validation of player-specific pass-recipient decision values.

| Evaluation | Estimate | 95% CI |
|---|---|---|
| Held-out recipient-selection value predicted by $D_i$ | 0.917 | [0.455, 1.387] |
| After controlling for forward-pass frequency | 1.170 | [0.579, 1.578] |
| Subsequent signed xG associated with empirical selection value | 0.433 | [0.252, 0.600] |
| Held-out prediction after team-by-role control | 0.562 | [−0.137, 1.282] |

## 4. Player-level Applications

Using the average-player policy as a common reference, we examine how action selection, pass-recipient preferences, and modelled decision values vary across tactical contexts, positions, and individual players. Tactical contexts are defined by pitch third (defensive, middle, or final), close pressure on the ball carrier, and the opponent's high-press structure (Table 4). Of the 12 possible contexts, 10 are observed in the dataset, with no high-press situations occurring in the final third.

Table 4: Distribution of decision states across tactical contexts.

| Ball location | Opponent structure | Without pressure | Pressure |
|---|---|---|---|
| Defensive third | High press | 5,779 | 1,400 |
| Defensive third | No high press | 7,599 | 494 |
| Middle third | High press | 642 | 453 |
| Middle third | No high press | 19,412 | 3,330 |
| Final third | High press | 0 | 0 |
| Final third | No high press | 8,994 | 3,380 |
| **Total** | | **42,426** | **9,057** |

### 4.1 Action Selection Across Tactical Contexts

We first examine how action selection varies across tactical contexts and positions. In the middle third, pressure is associated with an increase in pass rate of 20.7 percentage points (pp) for centre-backs, 11.2 pp for defensive/central midfielders, and 10.5 pp for full-backs, compared with only 1.2 pp for centre-forwards. These differences between positions are statistically significant (Holm-adjusted p = 0.002). The average-player policy correctly predicts the direction of pass-rate changes in 21 of 24 position-specific comparisons. However, it does not fully capture how responses to pressure differ between positions.

Within positions, players also exhibit different responses to pressure (Figure 3). Alessandro Bastoni shifts towards passing by 17.2 pp more than the average-player policy predicts, whereas Gianluca Mancini shifts 12.9 pp less than predicted. These responses differ from overall action tendencies: Bastoni is near the average policy across all non-shot states (–1.2 pp), while Mancini is more pass-leaning (+7.5 pp). Among full-backs, Federico Dimarco (+21.7 pp) and Leonardo Spinazzola (–5.7 pp) also show contrasting overall passing tendencies.

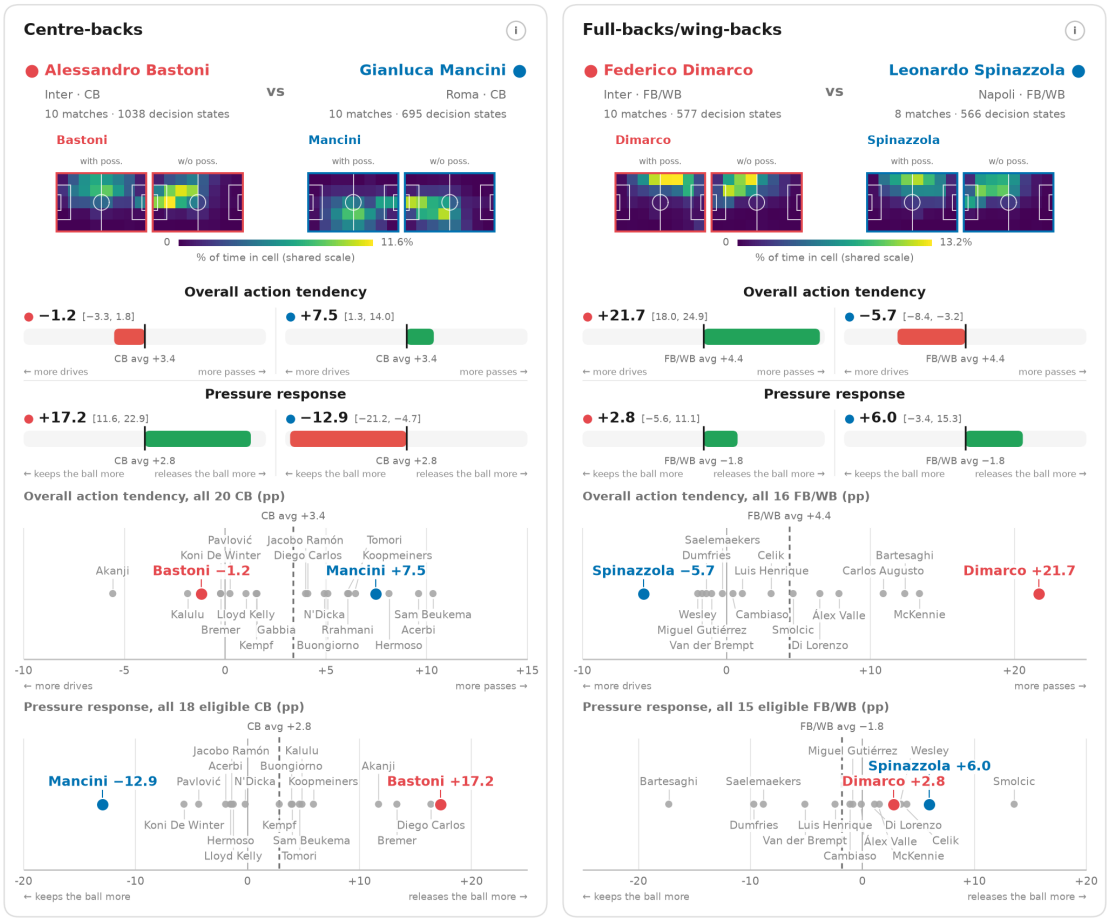



Figure 3: Player-level action tendencies and pressure responses for selected centre-backs and full-backs/wing-backs. Heatmaps show player positioning, while distributions compare players within each position. Values are expressed in pp, with 95% match-cluster bootstrap confidence intervals.

### 4.2 Pass-recipient Preferences Across Tactical Contexts

We next examine how pass-recipient selection varies across tactical contexts, focusing on forward and backward passes. In the middle third, the proportion of passes directed backward is higher under pressure than without pressure for centre-backs (+13.0 pp), defensive/central midfielders (+12.9 pp), and full-backs (+12.2 pp). Unlike the pass-rate changes in Section 4.1, differences in passing-direction responses between positions are not statistically significant (Holm-adjusted p = 0.80 for forward passes). This suggests that positions differ more in how often players pass under

pressure than in how they adjust passing direction. In the defensive third, centre-backs also show a higher forward-pass share under a high press (+13.3 pp).

Player-conditioned recipient policies also reveal differences between players in similar situations (Figure 4). In two illustrative final-third situations, Bryan Cristante's policy assigns a higher probability to forward recipients than Hakan Çalhanoğlu's policy. For Cristante's situation, the predicted forward-pass probabilities are 44% and 32%, respectively, compared with 36% under the average-player policy. The same pattern appears in Çalhanoğlu's situation (36% and 26%, compared with 34% under the average-player policy). These comparisons illustrate differences in player-conditioned recipient preferences under similar tactical conditions.

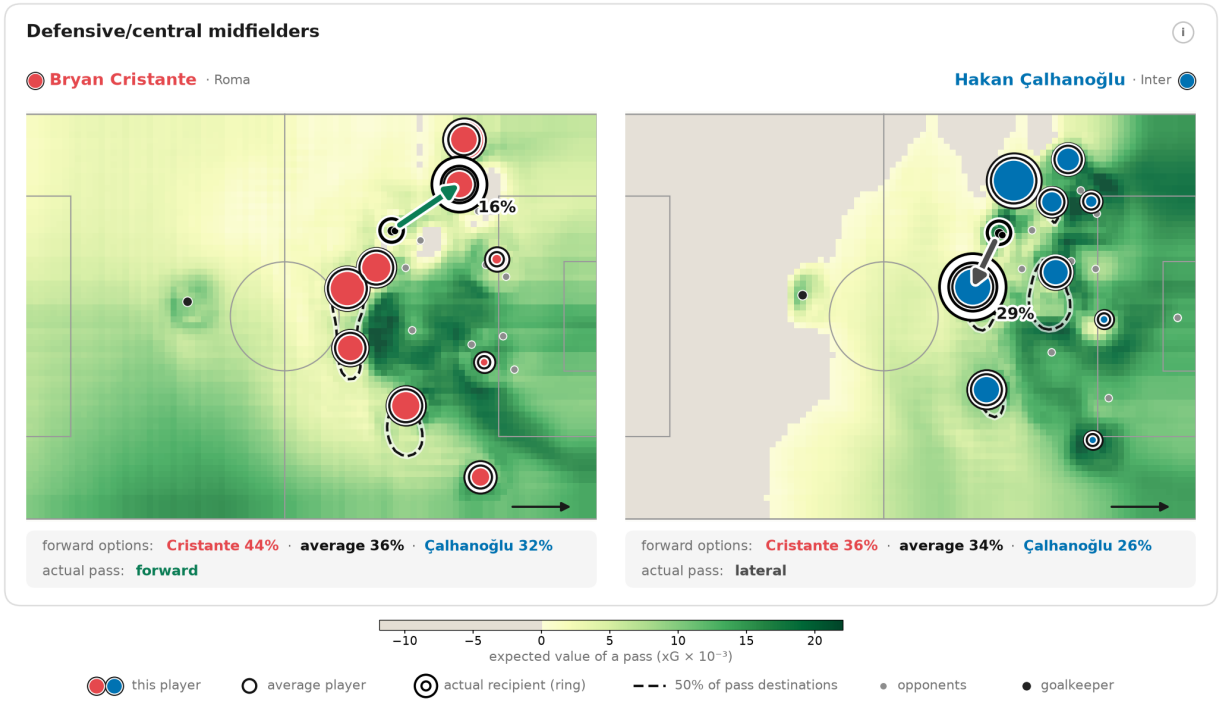



Figure 4: Illustrative comparison of player-conditioned pass-recipient policies in two matched tactical situations. Circles indicate model-estimated recipient probabilities, arrows show observed passes, and the background represents expected pass value ($\times 10^{-3}$ xG). The two player policies are evaluated in both situations.

### 4.3 Integrated Player Profiles

Overall, the preceding analyses highlight different aspects of player decision-making, including action tendencies, responses to pressure, and pass-recipient preferences. To examine how these measures complement spatial positioning, we compare player similarities based on positioning heatmaps and decision profiles. Within positions, the overlap between the three nearest neighbours identified by the two representations is 0.305 (95% CI: 0.244–0.371), close to the chance expectation of 0.290. Their pairwise distance matrices also show a low correlation (Spearman $\rho = 0.091$, permutation $p = 0.060$).

Figure 5 compares Manuel Locatelli and Çalhanoğlu, two defensive/central midfielders with similar positioning patterns, pass shares (42.1% and 42.8%), and OBV per 90 (0.30 and 0.34). However, their decision profiles differ. On common reference states, Locatelli shows a higher preference for forward recipients (+5.1 pp relative to the average-player policy), whereas Çalhanoğlu shows a lower preference (−8.3 pp). Their responses to pressure also differ (+5.5 and +16.2 pp, respectively), as do their EPV-added values (0.074 and 0.020 per 90). These comparisons illustrate how decision-specific measures complement conventional player profiles.

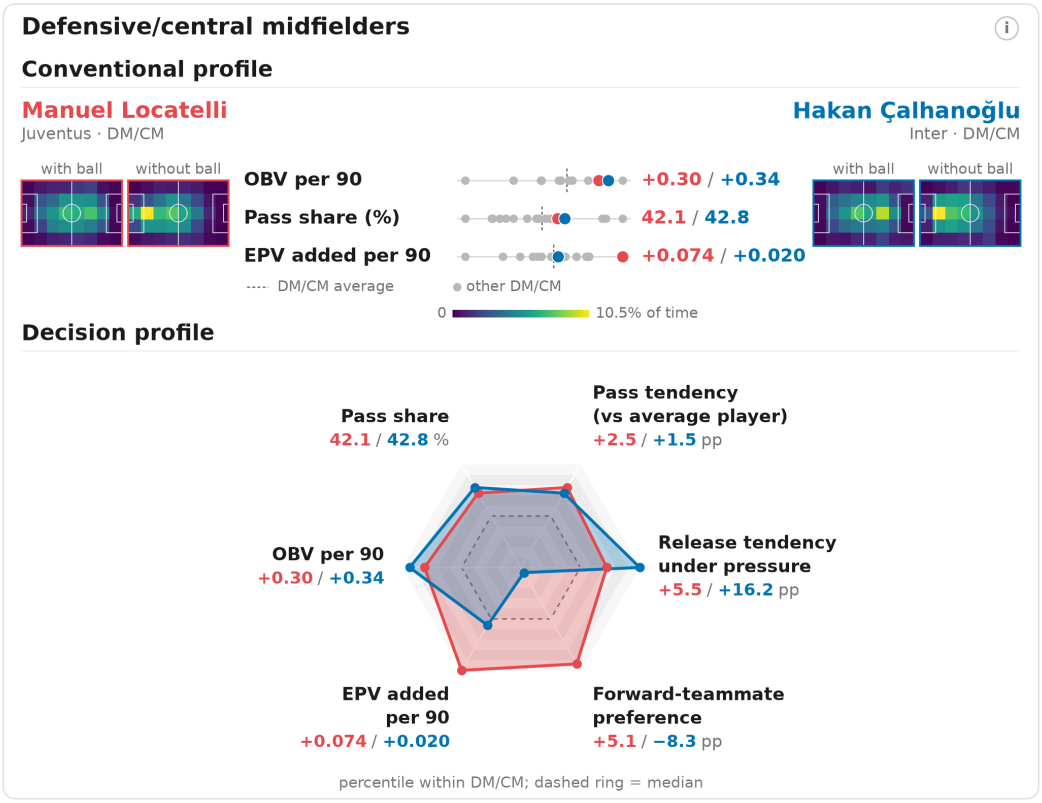



Figure 5: Comparison of conventional and decision-based profiles for two defensive/central midfielders. The upper panel shows positioning heatmaps and conventional performance measures. The lower panel compares six measures using within-position percentiles, including pass share, OBV per 90, EPV-added per 90, pass tendency, pressure response, and forward-recipient preference.

## 5. Conclusion

This study presents a player-aware EPV framework that combines tracking-based possession valuation with player-conditioned decision policies. Using 30 Serie A matches, the proposed player-token architecture improves EPV prediction over a tabular baseline while maintaining comparable performance without explicit pitch-control summaries. Standardized pass-recipient policy values are positively associated with held-out recipient-selection values, while empirical recipient-selection values are associated with subsequent signed xG. These findings support the use of the framework to examine both players' decision preferences and their estimated values.

The tactical context analyses reveal distinct patterns across positions and individual players. Pressure-related changes in pass rates differ substantially between positions, whereas changes in passing direction are more similar. The average-player policy captures much of the observed variation across contexts but does not fully explain position-specific responses. Player-conditioned recipient policies further illustrate how players can prefer different passing options under similar game situations. Together, these analyses provide a more detailed description of player decision-making than action frequencies or spatial positioning alone.

Several limitations remain. The dataset is relatively small, and shared reference states cannot fully account for tactical instructions, teammate quality, or unobserved contextual factors. The association between standardized decision value and held-out recipient-selection value remains inconclusive after team-by-role control, limiting claims about individual skill independent of tactical context. Moreover, the 15-second first-shot-xG target does not capture all longer-term possession benefits, and uncertainty in the fitted models is not fully propagated to player-level estimates. The proposed measures should therefore be interpreted as context-standardized descriptions of decision policies rather than causal estimates of individual ability. Further validation across more matches, teams, and tactical environments is needed to assess their stability and generalizability.

## References

[1] Anzer, G., & Bauer, P. (2021). A goal scoring probability model for shots based on synchronized positional and event data in football (soccer). Frontiers in Sports and Active Living, 3, 624475. https://doi.org/10.3389/fspor.2021.624475

[2] Stats Perform. (n.d.). Opta event definitions. https://www.statsperform.com/opta-event-definitions/

[3] Singh, K. (2019). Introducing expected threat (xT). https://karun.in/blog/expected-threat.html

[4] Decroos, T., Bransen, L., Van Haaren, J., & Davis, J. (2019). Actions speak louder than goals: Valuing player actions in soccer. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (pp. 1851–1861). https://doi.org/10.1145/3292500.3330758

[5] Hudl StatsBomb. (2021). Introducing on-ball value (OBV). https://statsbomb.com/news/introducing-on-ball-value-obv/

[6] Fernández, J., Bornn, L., & Cervone, D. (2021). A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. Machine Learning, 110, 1389–1427. https://doi.org/10.1007/s10994-021-05989-6

[7] Cervone, D., D'Amour, A., Bornn, L., & Goldsberry, K. (2016). A multiresolution stochastic process model for predicting basketball possession outcomes. Journal of the American Statistical Association, 111(514), 585–599. https://doi.org/10.1080/01621459.2016.1141685

[8] Robberechts, P., Van Roy, M., & Davis, J. (2023). un-xPass: Measuring soccer player's creativity. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 4768–4777). ACM. https://doi.org/10.1145/3580305.3599924
[9] Hong, M., Lee, M., Jo, G., So, J.-H., Bauer, P., & Ko, S.-K. (2026). EventGPT: Capturing player impact from team action sequences using GPT-based framework. arXiv preprint, arXiv:2512.17266v3. https://arxiv.org/abs/2512.17266v3
[10] Hong, M., Lee, M., Jo, G., Cho, H., Kim, H., Bauer, P., & Ko, S.-K. (2026). Modeling matches as language: A generative transformer approach for counterfactual player valuation in football (ScoutGPT). arXiv preprint, arXiv:2603.15212v2. https://arxiv.org/abs/2603.15212v2
[11] Fernández, J., & Bornn, L. (2021). SoccerMap: A deep learning architecture for visually-interpretable analysis in soccer. In Machine Learning and Knowledge Discovery in Databases: Applied Data Science and Demo Track (ECML PKDD 2020), Lecture Notes in Computer Science (pp. 491–506). Springer. https://doi.org/10.1007/978-3-030-67670-4_30
[12] Gelman, A., & Hill, J. (2007). Data Analysis Using Regression and Multilevel/Hierarchical Models. Cambridge University Press. https://doi.org/10.1017/CBO9780511790942
[13] Fernández, J., & Bornn, L. (2018). Wide open spaces: A statistical technique for measuring space creation in professional soccer. Proceedings of the MIT Sloan Sports Analytics Conference. https://barcainnovationhub.fcbarcelona.com/wp-content/uploads/2024/02/Wide-Open-Spaces-A-statistical-technique-for-measuring.pdf

# Appendix

### Appendix A. Learning Curve Analysis

To assess the effect of training data size, we retrain the scalar branch of the proposed model using 5, 10, 15, 20, and 25 matches while keeping the pass CNN outputs fixed. The analysis uses the same six match-based test folds, nested training subsets, and three random seeds. For the 25-match setting, the validation matches are added to training, using the stopping epoch selected with 20 training matches.

Non-shot EPV correlation increases from 0.216 with five training matches to 0.238 with 20 matches, with only a small further increase to 0.240 with 25 matches (Table A.1; Figure A.1). Action-selection NLL decreases from 0.608 to 0.578 across the full range.

Table A.1: Learning curve of the proposed model.

| Training matches | Non-shot EPV correlation [95% CI] | Action NLL [95% CI] |
|---|---|---|
| 5 | 0.216 [0.188, 0.241] | 0.608 [0.602, 0.614] |
| 10 | 0.224 [0.196, 0.250] | 0.595 [0.589, 0.601] |
| 15 | 0.233 [0.201, 0.263] | 0.588 [0.582, 0.594] |
| 20 | 0.237 [0.204, 0.270] | 0.592 [0.585, 0.599] |
| 25 | 0.240 [0.208, 0.270] | 0.578 [0.572, 0.584] |

*Note: Results are computed from seed-averaged predictions over six match-based folds, with 95% confidence intervals estimated using 2,000 match-cluster bootstrap resamples. The pass CNN is held fixed.*

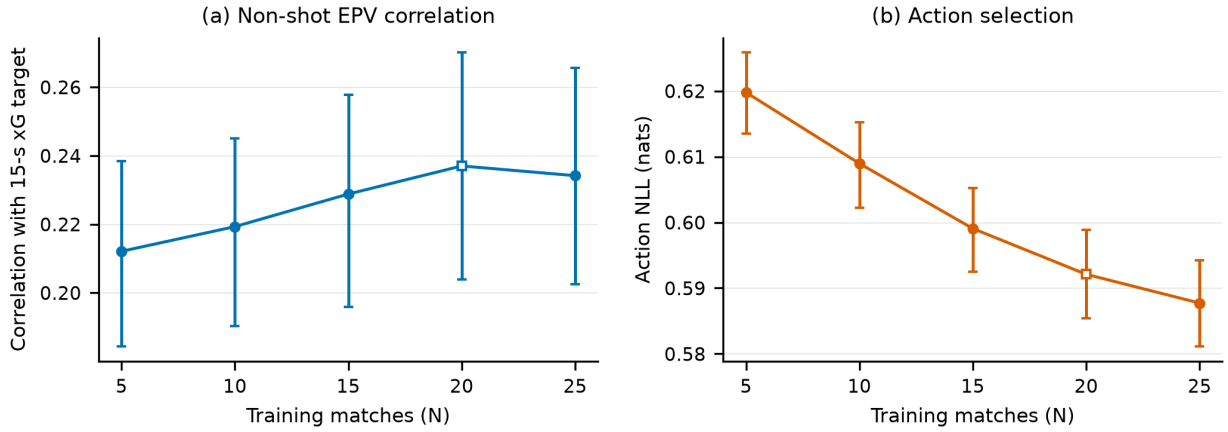



Figure A.1: Learning curves of the proposed model across training-set sizes: (a) non-shot EPV correlation and (b) action-selection NLL. Error bars indicate 95% match-cluster bootstrap confidence intervals.

## Appendix B. Model Input Features

The proposed model uses 22 player tokens and 20 tabular state features extracted from the decision frame. Player and ball coordinates are expressed in metres in the attacking direction on a 105 × 68 m pitch, with velocities estimated from the preceding second. Input features do not include the match score or time.

**Player-token features.** Each player token contains 13 features describing the player's position, velocity, spatial relationships to the ball and attacking goal, and team-related indicators (Table B.1). The possessing team's tokens are ordered first, with the ball carrier in the first position. Tokens do not encode player identity or positional role, and missing players are masked.

Table B.1: Player-token input features.

| Feature group | Dimensions |
|---|---|
| Player position | 2 |
| Player velocity | 2 |
| Vector to attacking goal | 2 |
| Distance to attacking goal | 1 |
| Vector to ball | 2 |
| Distance to ball | 1 |
| Team, ball-carrier, and goalkeeper indicators | 3 |
| **Total** | **13** |

**Tabular state features.** The 20 tabular features describe ball-carrier geometry, local pressure, team positioning, and player and ball movement (Table B.2). Features are standardized using training-fold statistics. Two last-defensive-line features are constant in the present dataset.

Table B.2: Tabular state features.

| No. | Feature | Definition |
|---|---|---|
| 1 | Distance to goal | Distance from the ball carrier to the centre of the goal line being attacked, divided by 105 m. |
| 2 | Angle to goal (sine) | Sine of the angle between the carrier-to-goal-centre direction and the long axis of the pitch (lateral offset of the goal centre divided by the distance); signed by pitch side. |
| 3 | Distance to ball | Distance between the ball carrier and the ball, divided by 105 m (near 0 when the carrier is in control). |
| 4 | Teammates within 3 m | Number of players of the possessing team within 3 m of the ball carrier; the carrier itself is counted. |
| 5 | Opponents within 3 m | Number of opposing players (goalkeeper included) within 3 m of the ball carrier. |
| 6 | Opponents within 5 m | Number of opposing players (goalkeeper included) within 5 m of the ball carrier. |
| 7 | Outplayed-opponent balance | Number of possessing-team players ahead of the ball carrier (larger x) minus the number of opposing players ahead of the carrier (goalkeeper included); negative when more opponents than teammates are ahead. |
| 8 | Last-defensive-line position | Column reserved for the position of the opposing last defensive line (scaled by 105 m). |
| 9 | Last-defensive-line height | Column reserved for the height of the opposing last defensive line (scaled by 105 m). |
| 10 | Carrier speed | Speed of the ball carrier (displacement over the preceding 1 s), clipped to [0, 12] m/s and divided by 12 m/s. |
| 11 | Ball speed | Speed of the ball (displacement over the preceding 1 s), clipped to [0, 30] m/s and divided by 30 m/s. |
| 12 | Opponents within 5 m of the ball | Number of opposing players (goalkeeper included) within 5 m of the ball. |
| 13 | Possessing-team line index | The x positions of the possessing team are grouped into three lines (complete-linkage clustering in one dimension; lines ordered from own goal to attacked goal); index 0, 1 or 2 of the line whose mean x is closest to the ball. |
| 14 | Opposing-team line index | The same index for the three lines of the opposing team. |
| 15 | Teammates between ball and goal | Number of possessing-team players whose x lies between the ball and the goal line and whose y lies within 5 m of the band spanned by the ball and the goal centre. |
| 16 | Opponents between ball and goal | The same count for opposing players (goalkeeper included). |
| 17 | Teammates ahead of the ball | Number of possessing-team players with x larger than the ball x (any y). |
| 18 | Opponents ahead of the ball | Number of opposing players (goalkeeper included) with x larger than the ball x. |
| 19 | Opponents in the shooting triangle | Number of opposing players inside the triangle formed by the ball and the two goal posts (posts at y = -3.66 m and +3.66 m on the goal line). |

| No. | Feature | Definition |
|---|---|---|
| 20 | Opponents within 3 m of the ball | Number of opposing players (goalkeeper included) closer than 3 m to the ball. |
| 21* | Pitch control of the possessing team at the ball | Probability that the possessing team controls the ball location, from a physics-based pitch-control mode; 0.5 if the ball position is missing. Only in 22-feature variants. |
| 22* | Pitch control of the opposing team at the ball | Probability that the opposing team controls the ball location, same model; 0.5 if the ball position is missing. Only in 22-feature variants. |

*Note: Features 1-20 are the inputs of the proposed model; features marked * (21, 22) are added only in the 22-feature variants. Distances in metres are divided by the pitch length (105 m) and speeds by their clipping bound; counts are in natural units. All columns are standardized with the training-fold mean and standard deviation before entering the model.*

The 22-feature variants in Table 1 additionally include two pitch-control summaries evaluated at the ball location. These summaries are excluded from the proposed 20-feature model.

**Pass-recipient choice features.** The conditional-logit model uses 11 features for each available recipient (Table B.3). Three represent pass-completion and value estimates from the pass CNN, while eight describe passing geometry and defensive positioning.

Table B.3: Pass-recipient choice features.

| Feature group | Features |
|---|---|
| Pass CNN estimates | Completion log-odds; value after completion; difference between successful- and unsuccessful-pass values |
| Passing geometry | Forward progress; passing distance; passing angle; lateral displacement towards the passer's own side |
| Defensive positioning | Nearest-opponent distance; opponents beyond the option; opponents between passer and option; position inside the opposing block |

## Appendix C. Additional EPV Results and OBV Comparison

### C.1 Component-level Performance

We evaluate individual EPV prediction components against simple baselines using out-of-fold predictions (Table C.1). The proposed model improves action selection, drive success, successful-drive value, shot value, pass-destination selection, and the successful-pass components. In contrast, failed-drive and failed-pass value predictions do not improve over constant baselines.

Pass completion without destination information achieves a log-loss reduction of 18.9%. These results support the use of the learned components for EPV estimation while showing weaker performance for infrequent failure outcomes.

Table C.1: Performance of individual EPV prediction components.

| Component | n | Metric | Baseline loss | Model loss [95% CI] | Reduction (%) |
|---|---|---|---|---|---|
| Action selection | 51,483 | NLL | 0.7040 | 0.5921 [0.5854, 0.5989] | 15.9 |
| Drive success | 31,765 | Log-loss | 0.1185 | 0.0828 [0.0758, 0.0909] | 30.2 |
| Successful-drive value | 30,958 | MSE | 0.001352 | 0.001264 [0.000952, 0.001598] | 6.5 |
| Failed-drive value | 807 | MSE | 0.000683 | 0.000740 [0.000304, 0.001312] | –8.3 |
| Shot value | 404 | MSE | 0.013006 | 0.007773 [0.005176, 0.010694] | 40.2 |
| Pass destination | 19,313 | Cell NLL | 7.931 | 6.495 [6.427, 6.566] | 18.1 |
| Pass success (actual destination) | 19,313 | Log-loss | 0.3844 | 0.2005 [0.1878, 0.2132] | 47.8 |
| Value after successful pass | 16,825 | MSE | 0.001391 | 0.001294 [0.000955, 0.001692] | 7.0 |
| Value after failed pass | 2,488 | MSE | 0.001165 | 0.001234 [0.000675, 0.001982] | –5.9 |
| Pass completion (no destination) | 19,313 | Log-loss | 0.3844 | 0.3119 [0.2981, 0.3261] | 18.9 |

*Note: Results use six-fold match-level cross-fitting and three random seeds. Baselines are fitted on the training data of each fold; the destination-selection baseline uses a radial density, while other baselines use class frequencies, completion rates, or mean targets as appropriate. Confidence intervals are estimated using 2,000 match-cluster bootstrap resamples. Loss reductions are comparable only within the same row.*

### C.2 Comparison with Provider OBV

We compare EPV-added with the independently supplied On-Ball Value (OBV) to examine agreement between the two action-valuation measures. EPV-added is defined as the change in EPV between successive possession spells, reversing the sign when the opponent gains possession. The comparison excludes spells ending in shots and limits the gap between spells to 15 seconds.

At the action level, EPV-added and OBV show moderate agreement (Pearson r = 0.49, 95% CI: 0.43–0.55; Table C.2). At the player level, the two measures also show moderate rank correlation across 63 outfield players with at least 450 minutes (Spearman $\rho = 0.42$), decreasing to 0.36 after removing position means.

Table C.2: Agreement between EPV-added and OBV.

| Evaluation level | n | Pearson r [95% CI] | Spearman ρ [95% CI] |
|---|---|---|---|
| All credited actions | 17,412 | 0.487 [0.425, 0.547] | 0.452 [0.435, 0.470] |
| Passes only | 17,228 | 0.480 [0.417, 0.542] | 0.446 [0.428, 0.463] |
| Players, per 90 | 63 | — | 0.417 [0.129, 0.459] |
| Players, per 90, position-demeaned | 63 | — | 0.358 [0.063, 0.468] |

OBV and EPV-added include different sets of actions, so their player-level comparisons are based on rank agreement rather than direct equality of values. Tables C.3 and C.4 report the ten highest-ranked players under each measure.

Table C.3: Top ten players ranked by EPV-added per 90 minutes.

| Rank | Player | Team | Position | Minutes | EPV-added/90 | OBV/90 | OBV/90 rank |
|---|---|---|---|---|---|---|---|
| 1 | Manuel Locatelli | Juventus | DM/CM | 805 | +0.0743 | +0.299 | 8 |
| 2 | Alessandro Bastoni | Inter Milan | CB | 786 | +0.0608 | +0.120 | 33 |
| 3 | Nico Paz | Como | AM | 813 | +0.0518 | +0.362 | 3 |
| 4 | Nicolò Barella | Inter Milan | DM/CM | 772 | +0.0457 | +0.391 | 2 |
| 5 | Manu Koné | AS Roma | DM/CM | 727 | +0.0444 | +0.155 | 24 |
| 6 | Piotr Zieliński | Inter Milan | DM/CM | 584 | +0.0433 | +0.287 | 13 |
| 7 | Kenan Yıldız | Juventus | AM | 780 | +0.0426 | +0.117 | 34 |
| 8 | Manuel Akanji | Inter Milan | CB | 952 | +0.0382 | +0.058 | 46 |
| 9 | Gianluca Mancini | AS Roma | CB | 902 | +0.0351 | +0.135 | 29 |
| 10 | Luka Modrić | AC Milan | DM/CM | 927 | +0.0351 | +0.296 | 9 |

*Note: Top 10 of 63 players by EPV-added per 90 minutes (descending). OBV/90 is the sum of the provider OBV of all of the player's events, per 90 minutes. EPV-added/90 and OBV/90 are on different scales and count different events, so only ranks are compared.*

Table C.4: Top ten players ranked by OBV per 90 minutes.

| Rank (OBV/90) | Player | Team | Position | Minutes | OBV/90 | EPV-added/90 | EPV-added rank |
|---|---|---|---|---|---|---|---|
| 1 | Wesley | AS Roma | FB/WB | 788 | +0.407 | −0.0077 | 49 |
| 2 | Nicolò Barella | Inter Milan | DM/CM | 772 | +0.391 | +0.0457 | 4 |
| 3 | Nico Paz | Como | AM | 813 | +0.362 | +0.0518 | 3 |
| 4 | Hakan Çalhanoğlu | Inter Milan | DM/CM | 674 | +0.336 | +0.0196 | 21 |
| 5 | Leonardo Spinazzola | Napoli | FB/WB | 505 | +0.323 | +0.0137 | 26 |
| 6 | Martin Baturina | Como | W | 562 | +0.319 | +0.0104 | 32 |
| 7 | Lucas Da Cunha | Como | DM/CM | 683 | +0.301 | +0.0147 | 25 |
| 8 | Manuel Locatelli | Juventus | DM/CM | 805 | +0.299 | +0.0743 | 1 |
| 9 | Luka Modrić | AC Milan | DM/CM | 927 | +0.296 | +0.0351 | 10 |
| 10 | Bremer | Juventus | CB | 530 | +0.295 | +0.0221 | 17 |

*Note: Top 10 of the same 63 players by OBV per 90 minutes (sum over all of the player's events), with their rank by EPV-added per 90 minutes.*