Evaluating counter-pressing efficiency and intensity through hypothetical distributions of ghost defenders and closing-speed kinematics

Emaly A. Vatne[1,2], Kyle Pulvermacher[3], Matthew Vatne[3], Daniel Cencer[3]

[1] Medical and Performance Department, Denver Summit FC, Denver, CO, USA

[2] Human Performance Collaborative, The Ohio State University, Columbus, OH, USA

[3] Department of Athletics, The Ohio State University, Columbus, OH, USA

1 Introduction

Counter-pressing is typically defined as the immediate pressure that players apply to the new ball carrier following a loss of possession. This has become a widely discussed and adopted concept in modern football that was largely popularized by managers such as Jurgen Klopp, Jesse Marsh, and Ralf Rangnick. While counter-pressing is perceived to be highly effective, what constitutes a successful counter-press remains equivocal and there is no consensus on how counter-pressing success should be identified and measured. For example, a common definition of a counter-press is binary and its success or failure is based exclusively on a time constraint: a counter-press is successful if possession is regained within five seconds of the turnover. This operational definition and threshold has been an important starting point for measuring how often counter-pressing produces successful regains, but it is arbitrary and oversimplified in the transient game of football. For instance, a regain at 5.5 seconds would be classified as a failure and so would a pressure applied in a counter-press that results in the ball going out of play. Moreover, the binary success/failure dichotomization does not describe why a counter-press failed or whether it was the "right" moment or opportunity to commit numbers to a counter-press. Technically and importantly, every loss of possession in the run of play is a chance to counter-press; however, not every situation is equally well suited to successfully regain possession. Factors such as distances between players, the location on the pitch, and the number of opponents near the ball may influence whether a counter-press is likely to succeed. Although counter-pressing has become a commonly referenced component of game models, it remains unclear whether a failed counter-press reflects a poor decision or poor execution of the counter-press itself.

A counter-press can fail to regain possession for a number of different reasons and historical approaches to assess counter-pressing efficiency (e.g., proportion of opportunities when the regain takes place under a certain time threshold, passes per defensive action, etc.) have failed to describe which factors undermines counter-pressing performance. For instance, the distances between two teammates may be too large upon the loss of possession, leaving too much ground to cover in the reapplication of pressure to the new ball carrier. In another example, players may have

been well positioned but lacked the speed, acceleration, or capacity for repeated high-intensity efforts needed to close the space in time. These two failure potential outcomes have very different implications for coaching, player development, and future recruitment. The first is a tactical or decision-making problem that coaching may be able to address. The second is a physical constraint tied to the player's ability, conditioning, and identification. In order to understand which component drove the unsuccessful a counter-press, there is a need to quantify where the players could have been positioned upon the loss of possession and how quickly a particular player can close the distance to the ball.

Combining continuous tracking data with event data now makes it possible to observe every defender's position and movement through the few seconds following a turnover and to evaluate these moments without relying on arbitrary thresholds. With this, the present study aims to address two research questions. First, when a counter-press fails to regain possession, was it because a defender was poorly positioned, or positioned well but not quick enough to close the gap? Second, given the situation (e.g., player positionings) at the moment of the turnover, how likely was a regain in the first place?
In this paper, we introduce a multi-level framework designed to answer these questions with three primary components:

1. Ghost defender positioning: Informed by previous works and utilizing Random Forest Conditional Density Estimation, we construct a distribution of hypothetical "ghost" defenders that represent expected positioning given the game state at the moment of possession loss. Comparing each defender's actual position against this distribution measures the positional component of counter-pressing performance.
2. Individualized closing kinematics: Using a sprint-deceleration model fit with nonlinear mixed-effects regression, we estimate each player's maximal closing-speed capacity. Comparing each defender's actual counter-pressing trajectory against what that capacity allowed measures the athletic component. Bayesian mixed-effects models then pool these estimates across players and teams, producing stable player-level talent estimates even where observations are sparse.
3. Regain probability: Using XGBoost, we estimate the probability of regaining possession for each event from positioning, defender proximity, and kinematic features. This provides a benchmark for whether a given counter-press was likely to succeed regardless of how well it was executed.

Practically, this work has three applications. First, for coaches who value counter-pressing in their game model, it may be able to identify the situations (e.g., pitch location, player spacing, and contextual cues at the instance before, after, and during a loss of possession) in which a counter-press carries the highest probability of regaining possession and those in which reestablishing shape and defensive structure may be more advised because the regain probability is low. Second, by comparing actual regains against model-expected regains, it can evaluate teams' counter-pressing performance relative to the opportunities they faced and identify teams that counter-

press effectively, which may help inform a coaching staff who to learn from in terms of how training sessions, film review, and language are designed to instill this principle. Third, in recruitment, it can identify players whose positioning and physical attributes allow them to regain possession across a wider range of situations.

---

2 Related Works

This research questions and the described framework draws on three areas of research: the quantification of pressing and counter-pressing in football, ghosting and counterfactual evaluation of player positioning, and models of player movement and physical capacity.
Interest in measuring pressing has grown alongside its tactical popularity. Early tracking-data approaches described pressure geometrically. Andrienko et al. (2017) quantified the pressure on a player from the distances and angles between attackers and defenders. Others used team-level positional metrics, such as inter-player distances, dispersion, and line spacing, to separate high-pressing from deep-defending teams (Low et al., 2021). Forcher et al. (2022; 2024) identified spatial conditions and defensive principles associated with successful ball recovery. However, their analysis was restricted to possessions lasting at least five seconds and containing three or more passes, which excludes counter-pressing, since it typically occurs within seconds of a turnover. Bauer and Anzer (2021) proposed a precise definition of counter-pressing as a team or group behavior aiming to regain the ball immediately after a change of possession. They framed it as a risk-reward decision but were unable to compare it against alternatives such as reestablishing shape and compact structure. Robberechts (2019) introduced the Valuing Pressure decisions by Estimating Probabilities (VPEP) framework, and Merckx et al. (2021) extended it with rule-based detection of pressing situations and models of their effect on recoveries, pass selection, and scoring chances.

Most closely related to this work, Lee et al. (2025) proposed exPress, which uses XGBoost to estimate the probability that the defending team regains possession within five seconds of a pressing action. Importantly, the model model Hudl StatsBomb event and 360 freeze-frame data from 200 matches, which is not continuous tracking data. Thus, its best model reached a ROC AUC of 0.61. The authors attributed this modest performance partly to the 360 data capturing only players visible in the camera frame at the moment of each event. They also showed that predictive performance changed considerably depending on whether success was defined by a time threshold (3, 5, or 7 s), a number of actions, or a spatial radius. This supports our view that a fixed five-second, arbitrary definition of success is potentially oversimplistic. exPress also demonstrated counterfactual analysis by moving individual players and observing changes in regain probability. However, the alternative positions were chosen manually by the authors from video review, rather than drawn from a data-driven distribution of plausible positions. In addition, the analysis is anchored on recorded pressing actions, so it cannot assess turnovers in which a team chose not to press.

Ghosting evaluates a player by comparing their observed positioning or movement with that of a hypothetical baseline player, or "ghost." The underlying idea is that a player's observed position is only one of many positions they could plausibly have taken in that situation and ghosting judges it against those alternatives rather than in isolation. It was first described in basketball (Lowe, 2013) and later formalized through deep imitation learning, most notably by Le et al. (2017a; 2017b), who modeled where a soccer defender should have been based on league-average behavior. Subsequent work extended ghosting to basketball (Seidl et al., 2018) and American football (Schmid et al., 2021) and used deep generative models to simulate multi-agent movement in soccer (Gu & De Silva, 2023). However, a commonly cited limitation of these approaches is that they return a single point estimate of where a ghost should be without quantifying the uncertainty in that value.

Yurko et al. (2025) addressed this limitation with the first statistically grounded ghosting framework. They used random forests for conditional density estimation (RFCDE) to estimate a full distribution of plausible locations for the nearest defender at the moment an NFL receiver catches the ball, conditional on the play context but excluding information about that defender. Each hypothetical location was valued with an expected points model, and the defender was credited with the difference between the value of their actual position and the expected value across the ghost distribution. Because value is computed on the scale of expected points, the resulting metric is directly interpretable and can be aggregated across players and teams. The positional component of the present study adapts the framework described by Yurko et al. (2025) to counter-pressing in football but replacing expected points with the probability of regaining possession. However, Groom et al. (2026) argued that ghosts representing "average" behavior can lack tactical context. Subsequently, they introduced role-conditioned ghosts for defending corner kicks where a covariate-dependent hidden Markov model infers time-resolved man-marking and zonal assignments from tracking data without manual labels and each defender is compared with an average counterpart from the same role and team context. Related open-play approaches attribute defensive credit through graph attention (Everett et al., 2025) or marginal reductions in expected possession value (Kim et al., 2025). These evaluate defenders within the observed defensive structure rather than against plausible alternatives. In contrast, the ghosts in the present study are conditioned on game state rather than tactical role since that is more difficult to extrapolate in the broader context of the study. Therefore, they represent league-typical positioning which we complement with an "ideal" comparison against the highest-value plausible position.

Movement models have long been central to within-play evaluation from the possession value models of Cervone et al. (2016) in basketball to their adaptation for soccer (Fernández et al., 2021). Nguyen and Yurko (2026) recently proposed Bayesian multilevel step-and-turn models of ball-carrier movement in American football. Step length and turn angle are modeled with tracking-derived covariates and player random effects, and posterior predictive simulation generates a

distribution of hypothetical next steps against which observed movement is evaluated. The player-level random effects also serve as interpretable ratings of movement ability, such as variability in change of direction, acceleration, and deceleration. Separately, acceleration-speed profiling using measurements taken during matchplay estimates individual maximal acceleration and speed capacities and carry high ecological validity, but these capacities have rarely been linked to tactical outcomes.

Taken together, prior work has evaluated where defenders are positioned and how players move, but have not extended this work specifically to counter-pressing actions. Moreover, it has not evaluated whether a failed counter-press reflects poor positioning or insufficient physical capacity to close space in time, which may have major implications for coaching and recruitment decisions. Thus far, pressing models have been anchored on event-level freeze frames or on recorded pressing actions and existing ghosting frameworks for football have focused on set pieces or have not incorporated individual physical capacity. Therefore, our framework combines distribution-based ghost defenders (Yurko et al., 2025) with individualized, Bayesian multilevel models of acceleration and closing capacity in the spirit of Nguyen and Yurko (2026). It evaluates every open-play turnover using continuous tracking data and expresses both components on the common scale of regain probability.

---

3 Framework for Evaluating Counter-Pressing

This section details the methodological framework used to evaluate counter-pressing opportunities. The framework is underpinned by the idea that a counter-press can be evaluated against two hypothetical baselines: where a comparable player would typically have been positioned at the moment possession was lost, and how quickly that same player could have closed the space given their own physical capacity. Both baselines are expressed on a common scale, the probability of regaining possession, so that positional and athletic contributions can be compared directly and aggregated to players and teams.

We first describe the data and how the event and tracking feeds were integrated. We then define counter-pressing opportunities and their outcomes, describe the regain probability model that serves as the value function for the framework, and detail the two hypothetical evaluation components: ghost defender positioning and individualized closing kinematics. Finally, we describe the Bayesian multilevel models used to produce player- and team-level estimates, and how failed counter-presses were diagnosed.

This project uses data provided by Hudl covering the 2025/26 Italian Serie A season. Two related datasets were used: Hudl StatsBomb event data, which provides timestamped information on all on-ball actions, possession sequences, play patterns, and a counter-press flag for defensive actions performed shortly after a turnover; and Hudl broadcast X&Y tracking data, which provides

the positions and velocities of all players visible in the broadcast feed and the ball at 25 frames per second. Tracking coordinates are in meters with the origin at the center of the pitch and players are identified by jersey number and team.

3.1 Data Integration and Synchronization

Three integration steps were required before any modeling to prepare the separate event and tracking datasets. First, each event file was paired with its tracking file by match date and normalized home and away team names. Second, jersey numbers with team were mapped to StatsBomb player identifiers using the Starting XI and Tactical Shift lineups in the event data. Substitutes who did not appear in a Tactical Shift lineup were assigned to the unmapped jersey number that first appeared in the tracking data closest to, and after, the time of their substitution. Third, all tracking coordinates were transformed into a team-relative frame (origin at the team's own goal line, attacking toward positive x), matching the convention of the event data. The attacking direction of each team in each period was inferred from the mean x-coordinate of its goalkeeper during the first five seconds of the period. And finally, the orientation of the y-axis and the clock offset between feeds were then calibrated jointly for each match by grid search.

3.2 Counter-Pressing Opportunities and Outcomes

As operationally defined earlier, this study operates under the assumption that any loss of possession in the run of play is technically a chance to counter-press, which is importantly distinguished from organized team pressing opportunities that may originate, for example, from an opposition goal kick. Therefore, the counter-pressing opportunities were defined as any change of possession within the same period in which the new possession began in open play. The moment of possession loss, t_0, was defined as the timestamp of the first event performed by the gaining team in the new possession. The way in which possession was lost (e.g., incomplete pass, dispossession, miscontrol, failed dribble, lost duel, etc.) was also extracted from the losing team's final actions in the preceding possession.

An opportunity was classified as a successful regain if the losing team's next possession began within a horizon of $H = 8$ seconds of t_0 and started either in open play, from a restart caused by the opponent putting the ball out of play (throw-in, goal kick, or corner), or from a free kick won. A foul committed by the losing team before any regain ended the opportunity as a failure. The continuous time to regain was retained for every opportunity regardless of success or failure, and to evaluate how sensitive the conclusions are to an arbitrary time threshold, the context model was refit with horizons of 3, 5, 8, and 10 seconds and the stability of team rankings across horizons was assessed with Spearman correlations. We acknowledge that defining success with a time threshold partly repeats the limitation of previous counter-pressing work described earlier. We recognize that refitting the models across several horizons reduces this dependence but does not remove it. Future work may aim to replace the fixed time threshold by identifying the counter-press directly from tracking data such as when a team initiates it and then defining its end as when the team either regains possession or abandons the counter-press attempt and retreats into shape.

Separately from the outcome, each opportunity was labelled as committed if at least one engaged outfield player of the losing team who was more than 3 meters from the ball at t_0 came within 3 meters of the ball within 3 seconds, or if StatsBomb recorded at least one counter-press action by the losing team. This distinction allowed the framework to evaluate not only how well teams counter-pressed but also when they chose to do so.

3.3 Situation Representation at the Moment of Possession Loss

For each opportunity, a snapshot of the game state at t_0 was constructed in the frame of the pressing (possession-losing) team, with the origin at that team's own goal line. Goalkeepers were excluded from both teams. Following the anchoring approach of Yurko et al. (2020; 2025), all players were expressed relative to the ball carrier and ordered by their distance to the ball. The resulting feature vector contained five groups of information:

- Ball and carrier: ball location, distance to each goal, lateral position, and the carrier's speed and direction of travel relative to the pressing team's goal.
- Opponents: distance, relative position, and speed of the three nearest teammates of the ball carrier, the number of opponents within 10 and 20 m, the number of opponents closer to the pressing team's goal than the ball, and the lateral spread of the attacking team.
- Pressing team: distance, relative position, speed, and closing velocity (the component of velocity directed at the ball) of the four nearest pressing players; the number of pressers within 5, 10, and 15 m and the resulting numerical overload relative to opponents; the number of pressers goal-side of the ball; team depth, width, defensive line height, and centroid relative to the ball; and the number of pressers covering the passing lane between the carrier and the nearest opponent.
- Possession context: how possession was lost, and the duration and number of passes of the possession that was lost.
- Match context: minute and score differential from the pressing team's perspective.

Because broadcast tracking only captures players within the camera view, players outside the frame were treated as missing rather than imputed and the number of visible players on each team was included as a feature. This presents another opportunity for future, improved work where tracking data is obtained from cameras designed for optical tracking instead of extracted from broadcast footage.

3.4 Regain Probability Model

The value function at the center of the framework is the probability that the losing team regains possession, referred to here as expected regains (xRegain). Two gradient-boosted classifiers (XGBoost) were trained. The context model estimates $xR_{context} = P(\text{regain} \mid X_{context})$ from the ball, carrier, opponent, possession, and match context features only, and therefore contains no information about the pressing team. The full model estimates $xR_{full} = P(\text{regain} \mid X_{context}, X_{press})$ and adds the positions and velocities of the pressing team at t_0.

Each model serves a different purpose where the context model describes the difficulty of the opportunity a team faced and is the baseline for team-level regains over expected, which ensures that a team's own structure at the moment of the turnover is credited to the team rather than absorbed into its expectation and the full model is the value function used to evaluate hypothetical ghost positions (Section 3.5), analogous to the expected points model in Yurko et al. (2025). The difference between the two decomposes a team's performance into a structural component (how well positioned the team was at t_0) and an execution component (what happened after t_0).
Given the size of a single-season sample, both models were strongly regularized with 400 trees, a learning rate of 0.03, a maximum depth of 3, row subsampling of 0.8, column subsampling of 0.6, a minimum child weight of 20, an L2 penalty of 5, and a minimum split loss of 1. To avoid optimistic in-sample estimates, all predictions used in subsequent analyses were generated out of fold using five-fold cross-validation grouped by match, analogous to the leave-one-week-out approach of Yurko et al. (2025). Model performance was evaluated with log loss, Brier score, area under the ROC curve, and calibration by decile of predicted probability, and compared against a regularized logistic regression baseline.

3.5 Positional Component: Ghost Defenders

The positional component follows the ghosting framework of Yurko et al. (2025). For the k-th nearest pressing player at t_0 ($k = 1, 2, 3$), the quantity of interest is the distribution of where that player would typically be positioned relative to the ball, given the situation but without any information about the pressing team. This conditional density, $h(\ell \mid X_{context})$, where $\ell = (\Delta x, \Delta y)$ is the player's location relative to the ball, was estimated with random forests for conditional density estimation (RFCDE; Pospisil & Lee, 2018; Dalmasso et al., 2020).

In RFCDE, trees are grown to minimize a conditional density estimation loss rather than a regression or classification loss. Here, this was implemented by projecting each response onto a tensor-product cosine basis (six functions per dimension, 35 basis functions excluding the constant) and growing a multi-output regression forest on the basis coefficients, which is equivalent to splitting on the CDE loss. The forest contained 500 trees with a minimum of 20 observations per leaf. For a new opportunity, each training observation receives a weight proportional to how often it shares a terminal leaf with the new opportunity across trees, and the density is obtained by weighted Gaussian kernel density estimation over the 400 highest-weighted training observations, with a bandwidth set by a weighted Scott's rule and a minimum of 1 meter.

The density was evaluated on a 2-meter grid spanning ±40 meter around the ball, restricted to locations on the pitch. The highest-density cells covering 95% of the probability mass (up to 150 cells) were retained as ghost locations and renormalized. Because a ghost's value depends on its movement as well as its location, $B = 5$ trajectories were resampled for each ghost location from observed players of the same distance rank, with sampling weights inversely proportional to the difference between each observed player's distance to the ball and the ghost's distance to the ball, following equation 9 of Yurko et al. (2025). Trajectories were stored as speed and direction relative

to the bearing to the ball, so that a ghost moving toward the ball continues to move toward the ball from its hypothetical location.

For each ghost location and trajectory, the observed player was replaced by the ghost, the full feature vector was recomputed (including re-ordering players by distance), and xR_{full} was predicted using the model from the cross-validation fold that excluded that match. The positional contribution of the player was then defined as:

$$\delta_{pos} = xR_{full}(\text{actual}) - \Sigma_{\ell}\, h(\ell) \cdot (1/B)\, \Sigma_b\, xR_{full}(\text{ghost}_{\ell}, b)$$

A positive δ_{pos} indicates that the player's position at the moment of the turnover produced a higher regain probability than the distribution of positions typically occupied in comparable situations. Because the ghost distribution represents typical rather than optimal positioning, an additional ideal comparison, δ_{ideal}, was computed against the single ghost location with the highest xRegain within the 50% highest-density region. Finally, the highest-density region level of the observed position (the share of ghost probability mass more typical than the actual location) was recorded as a measure of how unusual the player's position was. Density estimates were evaluated out of sample with the CDE loss of Izbicki & Lee (2017).

3.6 Athletic Component: Individualized Closing Kinematics

The athletic component estimates how quickly each player could have closed the distance to the ball carrier and compares this with how quickly they actually did. This required an individualized model of each player's acceleration and deceleration capacity.

3.6.1 Acceleration-Speed Profiles

Following in-situ acceleration-speed profiling (Morin et al., 2021), each player's maximal acceleration was assumed to decline linearly with running speed, $a_{max}(v) = A_0(1 - v / V_0)$, where A_0 is the theoretical maximal acceleration and V_0 the theoretical maximal speed. Speed and acceleration were derived from every continuous visible segment of at least one second using a Savitzky-Golay filter (11 frames, second order), and implausible values (speed above 12.5 m/s or absolute acceleration above 10 m/s^2) were removed. Above 3 m/s, the highest acceleration and highest deceleration within each 0.2 m/s speed bin were extracted for each player in each match, and the median of these per-match extremes across matches formed the player's acceleration and deceleration envelopes. Using the median across matches, rather than the season maximum, prevents isolated noisy frames in broadcast tracking from defining a player's capacity.
Player profiles were estimated with a Bayesian nonlinear multilevel model, mirroring the player random-effect structure of Nguyen & Yurko (2026). Envelope accelerations were modeled with a Student-t likelihood ($\nu = 4$) with mean $A_{0j}(1 - v / V_{0j})$ for player j, and ($\log A_{0j}$, $\log V_{0j}$) followed a bivariate normal distribution with an LKJ prior ($\eta = 2$) on their correlation, allowing players with few observations to be pooled toward the league profile. Weakly informative priors centered on 6 m/s^2 and 9.5 m/s were used for the population means of A_0 and V_0. Maximal deceleration was modeled

separately as a multilevel log-normal model of envelope decelerations above 4 m/s, giving a player-specific maximal deceleration D.

3.6.2 Observed Closing

For every outfield player of the pressing team within 25 m of the ball at t_0, the distance to the ball was tracked frame by frame. The observed closing time, t_{reach}, was the first time the player came within 3 m of the ball. Observation was censored at the earliest of the 8 s horizon, the time of a regain, or a foul by the pressing team; players who did not reach the ball before censoring were retained as censored observations. Two benchmarks were computed for each run. The minimum closing time, t_{min}, is the time needed to close the gap present at t_0 (the distance to the ball minus 3 m) from the player's initial velocity toward the ball. Closing efficiency is the minimum time needed to cover the displacement the player actually travelled, divided by the observed time; it reflects effort relative to capacity and is unaffected by the ball carrier moving toward or away from the player. Peak speed during the run as a percentage of V_0 and reaction time (the first frame with a closing velocity above 2 m/s) were also recorded. Players were considered engaged in the counter-press if they were flagged in the event data or reacted within 1.5 s.

For engaged players, closing execution was modeled with a censored log-normal model of t_{reach}:

$$\log t_{reach} = \alpha + \beta_1 \log t_{min} + \beta_2 d_0 + \beta_3 v_0 + \beta_4 v_{carrier} + u_{player} + u_{team} + \varepsilon$$

where d_0 is the initial distance, v_0 the initial velocity toward the ball, $v_{carrier}$ the ball carrier's speed toward the player, and u_{player} and u_{team} are player and team random effects. Because t_{min} already accounts for each player's own physical capacity, a negative player effect indicates that a player closes faster than expected given their own profile, and the sign was reversed to produce a closing rating.

3.7 Player and Team Estimates

All player- and team-level estimates were obtained from Bayesian multilevel models so that estimates for players and teams with few opportunities are pooled toward league-level baselines and reported with uncertainty.

- Positional ratings: δ_{pos} from all three distance ranks was modeled with a Student-t likelihood ($\nu = 4$) with a fixed effect for rank and random effects for player and team, and the estimates were returned on the regain-probability scale.
- Team regains over expected (RoE): a logistic model with the logit of $xR_{context}$ as an offset, a global intercept, and a team random effect, so that the team effect represents how much more (or less) often a team regained possession than the opportunities it faced would predict. Refitting with xR_{full} as the offset isolates execution after t_0, and the difference between the two (the summed difference between xR_{full} and $xR_{context}$) represents the team's structure at the moment of the turnover. The raw sum of regains minus expected regains, and the same figure per 100 opportunities, were also reported.

- Decision quality: a logistic model of whether an opportunity was committed, with team-varying intercepts (commit rate) and team-varying slopes on the standardized logit of $xR_{context}$ xt. A larger slope indicates a team that is more selective, committing more often when a regain is likely and less often when it is not.

All models were fitted in PyMC with the No-U-Turn sampler (two chains of 1,000 draws after 1,000 tuning iterations, target acceptance of 0.9) using non-centered parameterizations for random effects. Convergence was assessed with $\hat{R}$ and effective sample sizes.

3.8 Evaluating Failed Counter-Presses

To answer the first research question, each failed opportunity was assigned a primary reason based on the nearest pressing player at t_0. The failure was labelled positional if that player's location fell outside the 80% highest-density region of the ghost distribution and had a negative δ_{pos}, and athletic if the player engaged in the counter-press and either reached the ball with a closing efficiency below 75%, or failed to reach the ball even though their minimum closing time was shorter than the time available and their peak speed remained below 70% of their theoretical maximal speed. Failures in which the nearest player did not engage were labelled separately. Failures meeting both criteria were labelled positional and athletic. Remaining failures were labelled low-probability situations when $xR_{context}$ was below 0.20, and otherwise attributed to execution or other factors (e.g., the actions of supporting players or the opponent).
Finally, to support the applied aim of identifying teams that counter-press effectively, the five teams with the highest RoE estimates were selected, and for each the opportunities with the largest regain over expectation and the failed opportunities with the highest $xR_{context}$ were extracted with their match, half, and game clock. These lists provide a structured starting point for video review and interviews with the coaching staff of those teams.

---

4 Results

This section presents the results of applying the framework to the available sample. We first describe the sample and the behavior of the outcome definition, then evaluate the regain probability model that serves as the value function, and finally report the team-level, positional, and athletic components and the diagnosis of failed counter-presses.

Sample and Outcome Description

Across 30 Serie A 2025/2026 matches, 2,640 open-play losses of possession were identified as counter-pressing opportunities, an average of 44 per team per match (Table 1). Possession was regained within 3, 5, 8, and 10 seconds in 7.0%, 14.7%, 24.8%, and 29.4% of opportunities, respectively.

Table 1: Sample details.

Measure	Value
Matches	30
Teams	6
Counter-pressing opportunities	2,640
Opportunities per team per match	44.0
Committed counter-presses, %	68.9%
Regained within 3 seconds, %	7.0%
Regained within 5 seconds, %	14.7%
Regained within 8 seconds, %	24.8%
Regained within 10 seconds, %	29.4%
Pressing-player closing runs	11,877

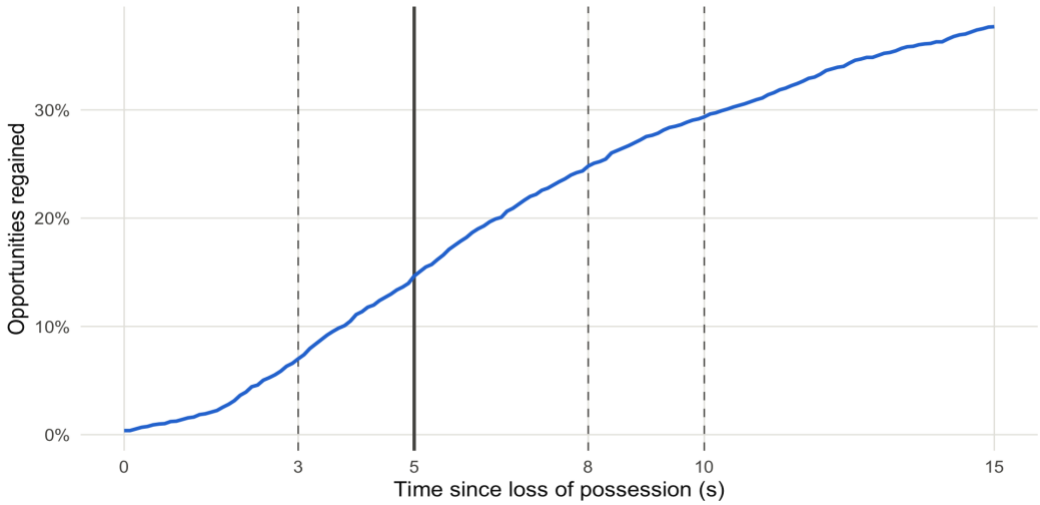


Figure 1: Cumulative proportion of opportunities in which possession was regained as a function of time since the turnover. Dashed lines mark the 3, 5, 8 and 10 second thresholds; a fixed 5 second rule (bolded line) treats every regain to its right as a failure.

3.2 Regain Probability Model

Both regain probability models showed modest discrimination when evaluated out of fold (Table 2). The context model, which contains no information about the pressing team, achieved an AUC of 0.577 (log loss 0.555, Brier score 0.185). Adding the positions and velocities of the pressing team at t_0 improved discrimination modestly (AUC 0.599, log loss 0.548, Brier score 0.183). The full

model also outperformed the regularized logistic regression baseline (AUC 0.570, log loss 0.564). This suggests that non-linear interactions between pressing-team positions carry some information. These values approach the best model reported by Lee et al. (2025) for pressing actions (AUC 0.61), despite the much smaller sample.

Table 2: Out-of-sample performance of the regain probability models (five-fold cross-validation grouped by match).

Model	Log loss	Brier	AUC	Observed rate	Mean prediction
xRegain – context (no pressing-team data)	0.555	0.185	0.577	0.248	0.247
xRegain – full (adds pressing-team positions)	0.548	0.183	0.599	0.248	0.247
Logistic regression baseline (full features)	0.564	0.189	0.570	0.248	0.249

Calibration of the full model was good across most of the probability range, with observed regain rates closely tracking predictions from the lowest decile (predicted 0.11, observed 0.11) to the middle of the range (Figure 2). In the highest decile, the full model slightly over-predicted, with a mean predicted regain probability of 0.39 against an observed rate of 0.34. The context model was less well calibrated at both extremes, under-predicting regains in the lowest decile (0.13 vs. 0.18) and over-predicting in the highest (0.37 vs. 0.30). Feature importance was distributed broadly rather than dominated by a few features.

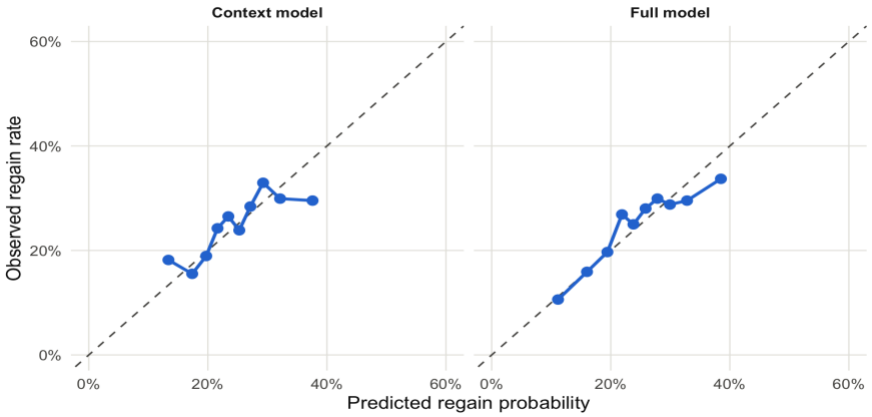


Figure 2: Calibration of the regain probability models by decile of predicted probability. Points on the dashed line indicate perfect calibration.

The outcome horizon had little effect on model discrimination (AUC 0.558-0.578 across horizons of 3-10 second) or on team-level conclusions (Table 3). The rank correlation between team regains over expected at 8 s and at every alternative horizon was ρ = 0.94. The ordering of teams was therefore robust to the choice of threshold, even though the classification of individual opportunities was not.

Table 3: Sensitivity of the outcome definition: context model refitted for each regain horizon.

Horizon (s)	Regain rate	AUC	Brier	Team RoE rank agreement with main horizon (Spearman rho)
3	7.0%	0.578	0.065	0.943
5	14.7%	0.558	0.126	0.943
8	24.8%	0.577	0.185	1.000

4.3 When and Where Counter-Pressing Succeeds

Regain rates by turnover location are shown in Figure 3 with attacking directions standardized from left to right.

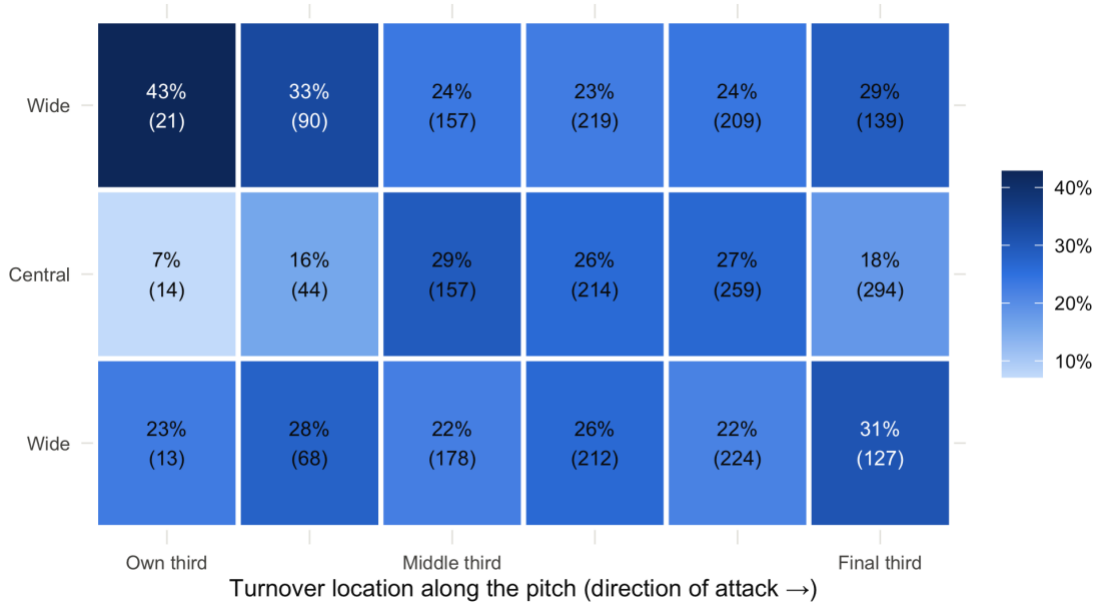


Figure 3: Regain probability by location of the turnover (pressing team attacking left to right). Cells show the observed regain rate and, in brackets, the number of opportunities; cells with fewer than 10 opportunities are left blank.

Committing to a counter-press was associated with a higher regain rate than not committing (27.4% vs. 19.1%; Table 4). The perceived cost of a failed counter-press was not borne out in the immediate possession. When a committed counter-press failed, the opponent produced a shot in

14.2% of possessions (0.017 xG per opportunity), almost identical to 13.9% (0.016 xG) when the team did not commit and did not regain possession. Successful regains almost entirely prevented an opponent shot in that possession (1.2% after a committed regain; 0.0% otherwise).

Table 4: Risk and reward of committing to a counter-press: opponent shots and expected goals (xG) in the possession following the turnover.

Decision	Outcome	*n*	Share of decision, %	Opponent shot in possession, %	Opponent xG per opportunity
Committed	Not regained	1,321	72.6%	14.2%	0.017
Committed	Regained	498	27.4%	1.2%	0.001
Did not commit	Not regained	664	80.9%	13.9%	0.016
Did not commit	Regained	157	19.1%	0.0%	0.000

4.4 Team Counter-Pressing Performance

Team regains over expected (RoE), relative to the opportunities each team faced, are shown in Table 5 and Figure 4.

- Juventus regained 19.6 more possessions than expected from 425 opportunities (123 vs. 103.4; +4.6 per 100; posterior probability of RoE > 0 = 0.85).
- Como followed (+9.3; +2.3 per 100; 0.70).
- Inter Milan (+1.2; +0.3 per 100; 0.52), AS Roma (–5.1; –1.1 per 100; 0.35) and Napoli (–8.3; –1.7 per 100; 0.28) performed close to expectation.
- AC Milan regained 13.5 fewer possessions than expected (–3.4 per 100; 0.18).

The 90% credible intervals for all six teams included zero. The evidence was strongest for Juventus (log-odds 0.11, 90% CrI –0.03 to 0.32) and AC Milan (–0.09, –0.28 to 0.05).

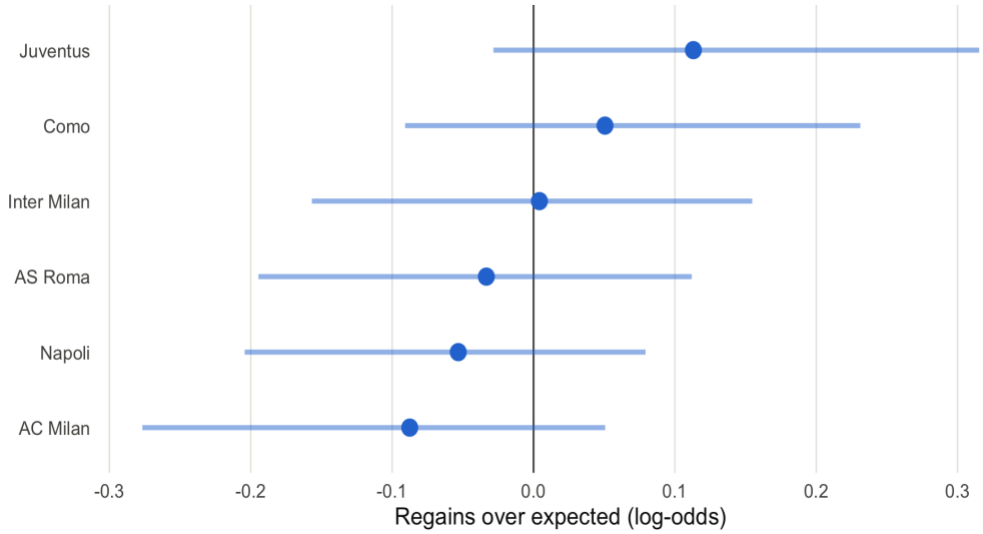


Figure 4: Team regains over expected (posterior mean and 90% credible interval, log-odds scale). Values above zero indicate more regains than the opportunities faced would predict.

Table 5: Team counter-pressing ratings. Regains over expected (RoE) compares actual regains with xRegain-context.

Team	n	Regains (expected)	RoE per 100	RoE, log-odds (90% CrI)	P(RoE > 0)	Structure	Execution
Juventus	425	123 (103.4)	4.61	0.11 (-0.03, 0.32)	0.85	1.26	18.33
Como	412	112 (102.7)	2.26	0.05 (-0.09, 0.23)	0.70	-0.25	9.56
Inter Milan	432	108 (106.8)	0.27	0.00 (-0.16, 0.15)	0.52	1.41	-0.23
AS Roma	474	110 (115.1)	-1.08	-0.03 (-0.19, 0.11)	0.35	0.72	-5.86
Napoli	494	114 (122.3)	-1.69	-0.05 (-0.20, 0.08)	0.28	0.16	-8.50
AC Milan	403	88 (101.5)	-3.36	-0.09 (-0.28, 0.05)	0.18	-2.58	-10.95

Structure = regains attributable to the pressing team's positioning at the turnover (xR full – xR context, summed); execution = regains beyond xR full. CrI = credible interval.

Decomposing RoE into structure and execution showed that differences between teams arose almost entirely after the turnover (Figure 5):

- Juventus and Como: their positive RoE came from execution (+18.3 and +9.6 regains), with structure contributing little (+1.3 and -0.3).
- AC Milan: its shortfall was also mainly execution (-11.0), alongside the largest structural deficit of any team (-2.6).

- Napoli and AS Roma: both lost regains through execution (-8.5 and -5.9) despite neutral or slightly positive structure (+0.2 and +0.7).
- Inter Milan: it had the most favorable structure (+1.4) but no gain from execution (-0.2).

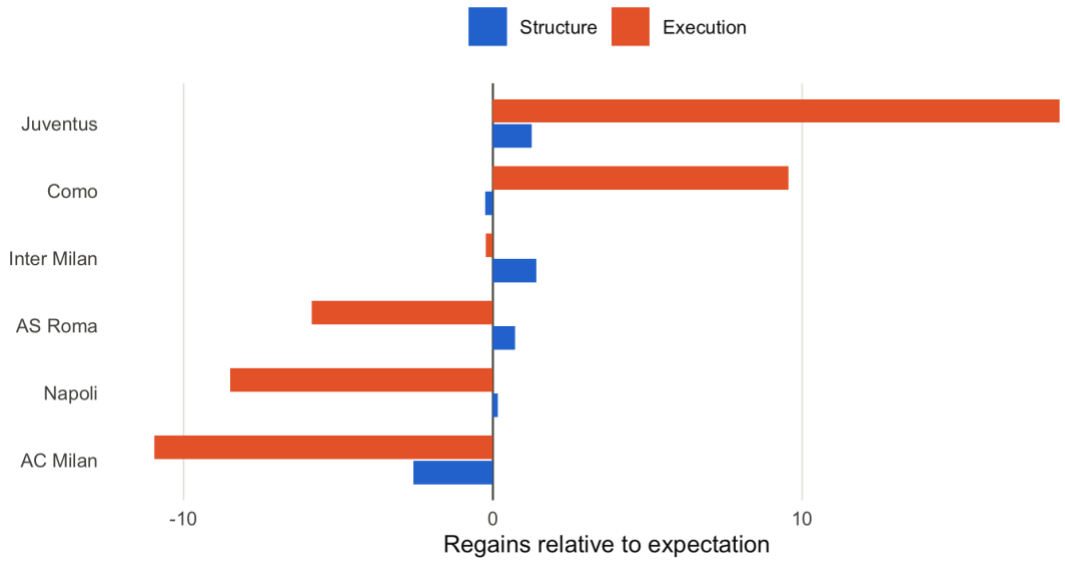


Figure 5: Decomposition of each team's regains over expected into structure (positioning of the pressing team at the moment of the turnover) and execution (what happened after the turnover), in regains over the season sample.

All six teams were selective in their decisions to commit (Table 6, Figure 6). Commit rates ranged from 63.6% (AC Milan) to 74.6% (Como). The team-specific slope of commit probability on xRegain was positive for every team (0.44–0.49), with all credible intervals excluding zero. Differences between teams in selectivity were small relative to their uncertainty.

Table 6: Evaluation of team decision-making.

Team	Commit rate	Selectivity (90% CrI)
AS Roma	70.4%	0.49 (0.38, 0.62)
AC Milan	63.6%	0.46 (0.35, 0.58)
Napoli	68.1%	0.45 (0.35, 0.56)
Inter Milan	69.3%	0.45 (0.32, 0.56)
Juventus	72.2%	0.44 (0.32, 0.55)
Como	74.6%	0.44 (0.32, 0.54)

Model-estimated commit rate and selectivity (the team-specific slope of commit probability on standardised xRegain-context; higher = the team commits more selectively when a regain is likely).

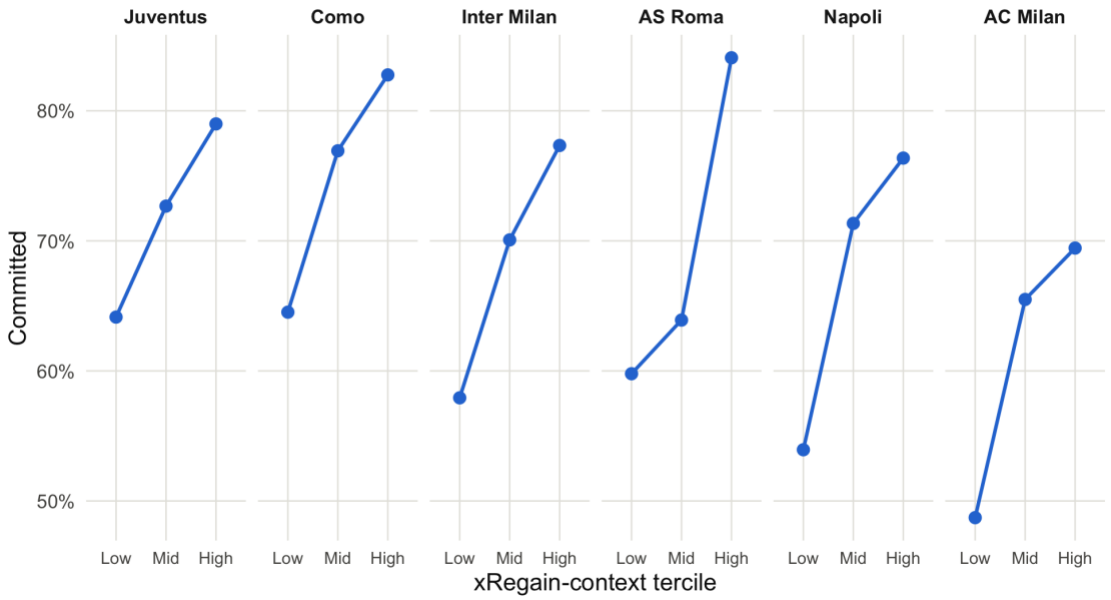


Figure 6: Decision-making: proportion of opportunities in which each team committed to a counter-press, by tercile of xRegain-context. Rising lines indicate teams that commit more when a regain is more likely.

4.5 Positional Component: Ghost Defenders

Ghost evaluations were computed for the three nearest pressers in all 2,640 opportunities. On average, the nearest presser was 5.3 m from the ball at t_0, the second-nearest 10.4 m, and the third-nearest 14.9 m. Observed positions were close to those modeled as typical: the average highest-density-region level was 0.43-0.46, so roughly half of the ghost density was more typical than the actual position. The mean positional value was close to zero for the nearest presser (δ_{pos} = 0.0004) and slightly positive for the second- (0.0049) and third-nearest pressers (0.0030). On average, supporting pressers therefore stood in slightly more favorable positions than their ghost distributions.

The comparison with the ideal ghost demonstrated that the best plausible ghost location (within the 50% highest-density region) carried a higher regain probability than the actual position by an average of 1.5, 2.2, and 3.5 percentage points for the first, second, and third nearest pressers, respectively. The largest available gains were therefore associated with the positioning of supporting pressers rather than the player closest to the ball.

After partial pooling, player-level positional ratings were small:

- Highest: Rasmus Højlund (+0.0060; P(> 0) = 1.00; 203 evaluations), Christopher Nkunku (+0.0054; 0.997), Donyell Malen (+0.0049; 0.998), Francesco Pio Esposito (+0.0048; 0.992) and Christian Pulisic (+0.0044; 0.98).

- Lowest: Luka Modrić (–0.0050; P(> 0) = 0.002), Alessandro Bastoni (–0.0045), Stanislav Lobotka (–0.0043; 0.003), André-Frank Zambo Anguissa (–0.0042) and Bryan Cristante (–0.0037).

The highest-rated players were mostly forwards and wide attackers, who are typically the first to lose and then contest possession high up the pitch. The lowest-rated were mostly central midfielders and a center-back, whose positions at the turnover tend to be set by rest-defense responsibilities rather than by immediate pressure on the ball.

Table 7: Players with the highest (top 5) and lowest (bottom 5) positional value at the moment of the turnover, with their closing ratings (minimum 20 observations).

Player	Positional value (Δ xRegain)	P(> 0)	n	Closing rating	Closing efficiency	V0 (m/s)
Rasmus Winther Højlund	0.0060	1.0000	203	-0.1915	0.7362	10.4902
Christopher Alan Nkunku	0.0054	0.9967	81	-0.0949	0.7492	10.1758
Donyell Malen	0.0049	0.9983	113	-0.1628	0.7175	10.1589
Francesco Pio Esposito	0.0048	0.9917	82	-0.1473	0.7544	10.5150
Christian Pulisic	0.0044	0.9800	90	-0.0896	0.6974	10.0498
Bryan Cristante	-0.0037	0.0217	147	0.0790	0.7254	10.0669
André-Frank Zambo Anguissa	-0.0042	0.0083	86	0.0580	0.5631	9.6928
Stanislav Lobotka	-0.0043	0.0033	108	0.1020	0.6907	9.8786
Alessandro Bastoni	-0.0045	0.0217	39	0.0652	0.6439	9.7282
Luka Modrić	-0.0050	0.0017	91	0.1494	0.6700	9.6364

Positional value is the posterior mean difference in regain probability between the player's actual position and the ghost distribution; closing rating > 0 = closes faster than expected given their own acceleration-speed profile.

4.6 Athletic Component: Closing Speed Kinematics

Acceleration-speed profiles were estimated for 141 players (Figure 8). The median theoretical maximal speed was 9.9 m/s (interquartile range 9.4 – 10.2 m/s), the median maximal deceleration 3.41 m/s^2, and the median theoretical maximal acceleration 12.3 m/s^2.

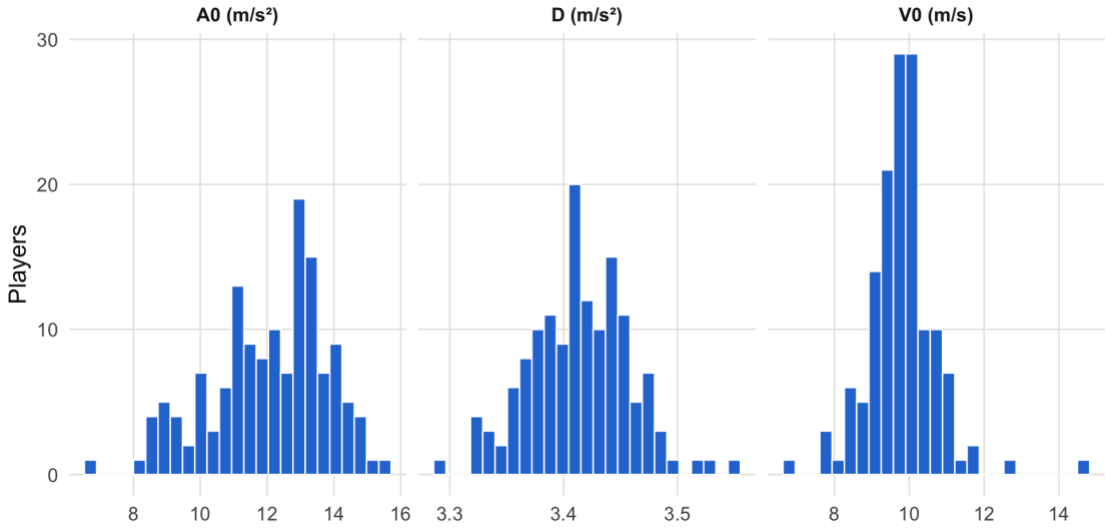


Figure 7: Distribution of individual acceleration-speed profile parameters estimated from broadcast tracking: theoretical maximal acceleration (A0), theoretical maximal speed (V0) and maximal deceleration (D).

A total of 11,877 closing runs by pressing players within 25 meters of the ball were analyzed, of which 51.5% were classified as engaged in the counter-press. Engaged players reached the ball carrier in 38.8% of runs, and those who did so ran at a median closing efficiency of 0.72. On average, players reached 40% of their theoretical maximal speed during a closing run. Player closing ratings, which reflect how quickly players closed relative to their own physical capacity, were highly uncertain. More specifically, the 90% CrI included zero for 132 of 133 players.

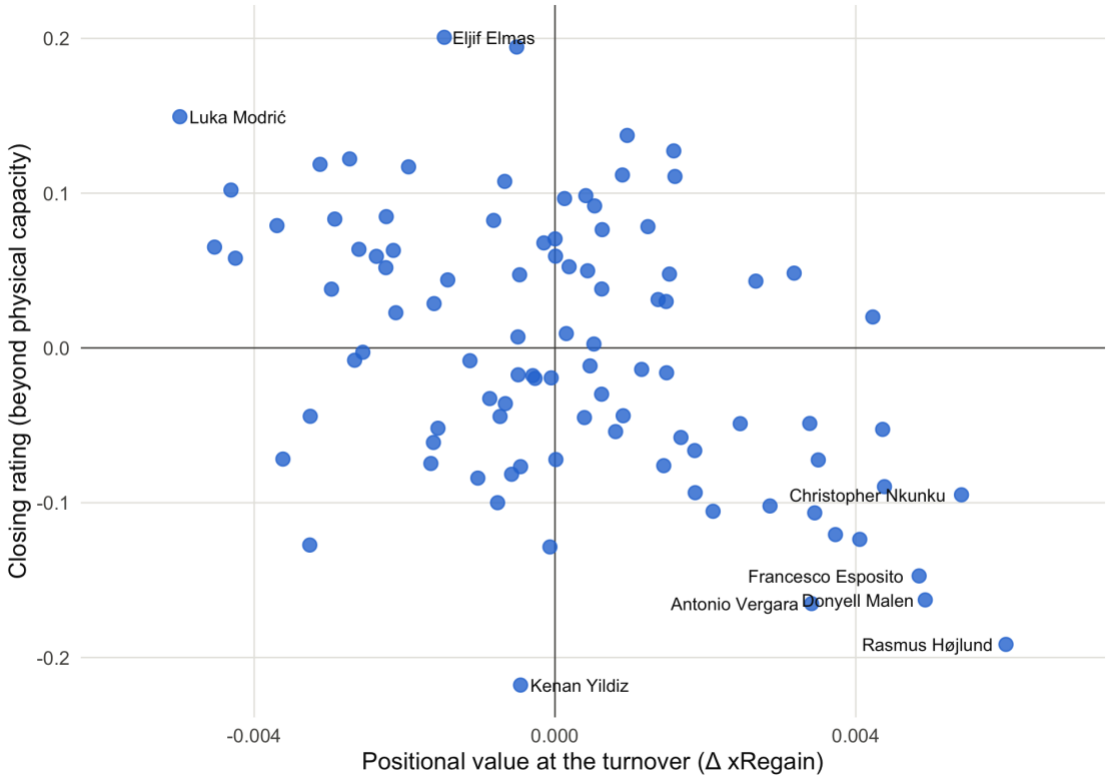


Figure 8: Positional value versus closing rating for players with sufficient observations. Upper-right players are well positioned and close quickly relative to their physical capacity; the most extreme players are labelled.

4.7 Why the Counter-Press Failed

Of the 1,985 failed opportunities, failures were attributed as follows (Table 8, Figure 9):

- Athletic factors were the most common single attribution: 705 (35.5%), closing at less than 75% of the presser's own capacity or never reaching the carrier despite being able to.
- Nearest player did not engage: 515 (25.9%).
- Execution/other: 305 (15.4%) could not be attributed to the nearest presser and likely reflect supporting players or the opponent.
- Low-probability situation: 271 (13.7%) had an xRegain-context below 0.20, so failure would be expected regardless of execution.
- Positioning of the nearest presser: 118 (5.9%).
- Both positioning and athletic: 71 (3.6%).

Positional failures were most common for AC Milan (10% of its failures), which was also the team whose nearest presser most often did not engage (30%). Napoli (40%) and Juventus (38%) had the highest shares of athletic failures, AS Roma the highest share of low-probability situations (19%), and Como the highest share of failures attributed to execution or other factors (20%).

Table 8: Primary reason for failed counter-pressing opportunities, based on the nearest pressing player at the turnover.

Team	Positional	Athletic	Positional + athletic	Nearest player did not engage	Low-probability situation	Execution /Other
AC Milan	32 (10%)	101 (32%)	14 (4%)	96 (30%)	34 (11%)	38 (12%)
AS Roma	20 (5%)	114 (31%)	9 (2%)	96 (26%)	70 (19%)	55 (15%)
Como	22 (7%)	104 (35%)	11 (4%)	64 (21%)	39 (13%)	60 (20%)
Inter Milan	15 (5%)	118 (36%)	10 (3%)	81 (25%)	41 (13%)	59 (18%)
Juventus	11 (4%)	115 (38%)	14 (5%)	73 (24%)	43 (14%)	46 (15%)
Napoli	18 (5%)	153 (40%)	13 (3%)	105 (28%)	44 (12%)	47 (12%)

Data are presented as n (% of the team's failures). Positional = outside the 80% region of the ghost distribution with a negative positional value; athletic = engaged but closed at <75% of own physical capacity, or never reached the carrier despite being physically able to and without approaching top speed.

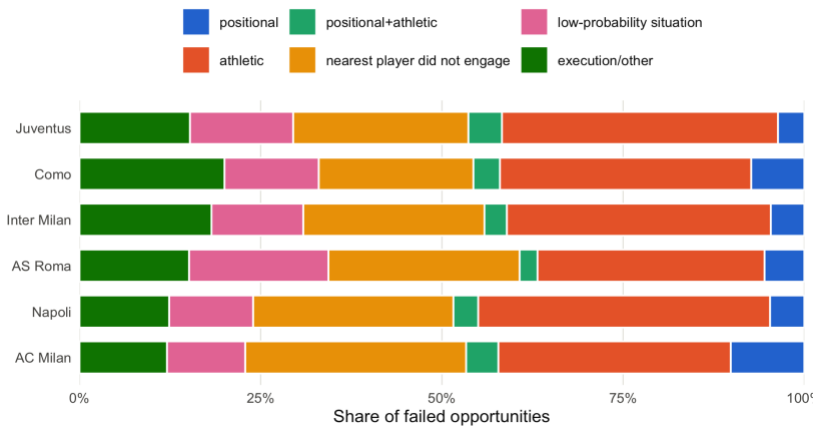
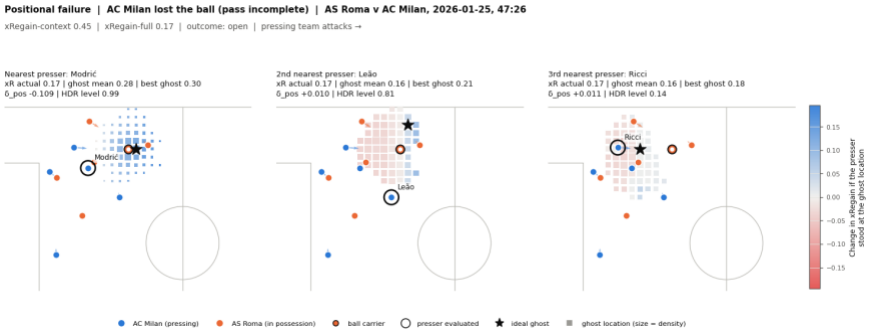


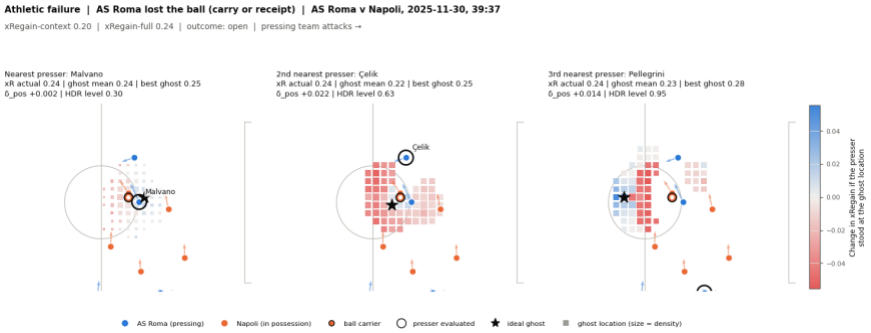
Figure 9: Composition of failed counter-pressing opportunities by team (see Table 7 for counts).

Figure 10 a-e illustrates one failed opportunity from each category. Each panel shows one of the three nearest pressers (ringed) at the moment of possession loss. Squares are ghost locations, sized by RFCDE density and colored by the change in regain probability had the presser stood there (blue = higher, red = lower). The star marks the best plausible (ideal) ghost.

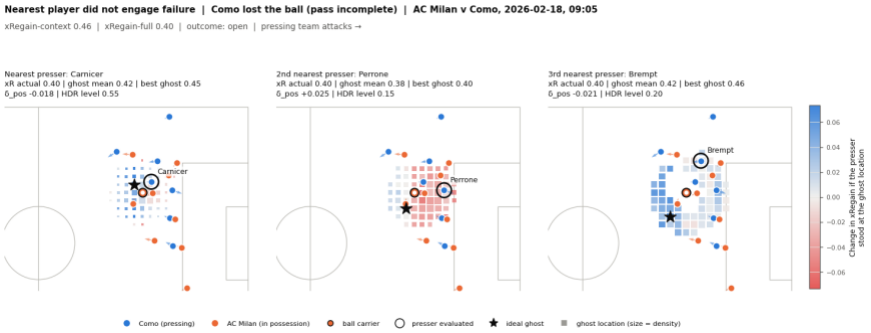
- (a) Positional. Napoli lost the ball against Como (2025-11-01, 45:47) with an xRegain-context of 0.34. All three nearest pressers stood in highly atypical positions (highest-density-region level 1.00), with positional values of –0.047 to –0.036: the ghost means were 0.33–0.34 against an actual xRegain of 0.29.



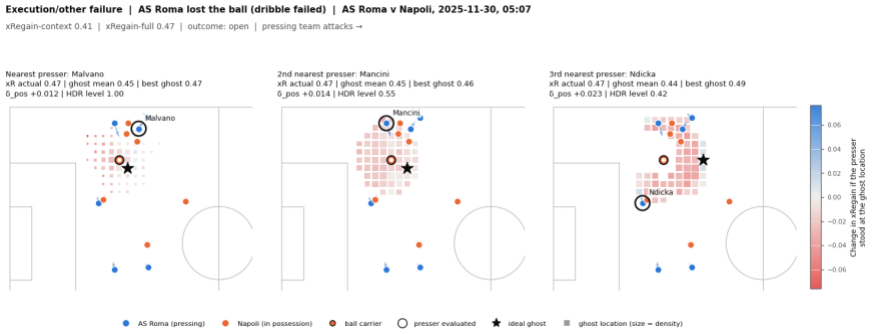
- (b) Athletic. AC Milan lost the ball against Juventus (2025-10-05, 33:36) with an xRegain-full of 0.34. The pressers were close to typical positions (positional values –0.015 to +0.007), but the nearest presser did not close the carrier down efficiently relative to his capacity.



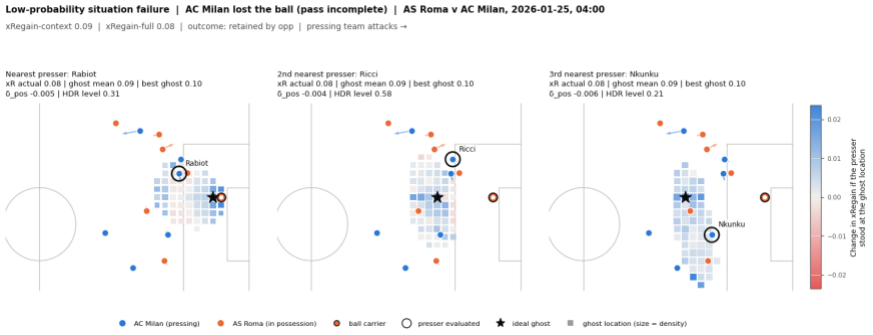
- (c) Nearest player did not engage. Como lost the ball against AC Milan (2026-02-18, 09:05) in one of the highest-probability opportunities in the sample (xRegain-context 0.46). The nearest presser did not engage, and possession was not regained within the horizon.



- (d) Execution/other. AC Milan lost the ball against Napoli (2025-09-28, 63:35) with an xRegain-full of 0.46. Two of the three pressers were better positioned than their ghost distributions (+0.038 and +0.033), so the failure cannot be attributed to the nearest pressers.



- (e) Low-probability situation. AC Milan lost the ball against Inter Milan (2026-03-08, 85:07). The xRegain-context was 0.09, and no ghost location changed the regain probability by more than about 0.01.



5 Discussion

This study set out to determine whether failed counter-presses reflect poor positioning or insufficient physical capacity to close space, and how likely a regain was given the situation at the moment of possession loss. The framework produced a coherent set of player- and team-level estimates on the common scale of regain probability. Applied to 2,640 opportunities, it showed that most failures were attributable not to where players stood when possession was lost but to what happened afterwards. Players either did not engage or closed the ball carrier down well below their own physical capacity. However, the strength of the conclusions is limited by the predictive signal available in the value function and by the size of a single-season, six-team sample. We discuss each component in turn, then the lessons learned and directions for future work.

The most important finding for interpreting the remainder of the results is that the regain probability model had modest discrimination (AUC 0.58–0.60). Information about the pressing team's positioning at t_0 improved discrimination (from 0.577 to 0.599), and the full XGBoost model outperformed a regularized logistic baseline. This indicates that the spatial configuration of the pressing team carries real, if limited, information, largely through non-linear interactions between players. Nevertheless, both the ghost evaluation and the structure-execution decomposition are expressed as differences in predicted regain probability, so they can only be as sensitive as this value function. A model that changes little when a presser is moved will necessarily assign small positional values such that average positional values were close to zero, and the largest player-level ratings amounted to less than one percentage point of regain probability per opportunity.

This may have been observed because the outcome is noisy and depends on what happens after t_0, because broadcast tracking misses distant players, or because the sample is small for flexible models. First, regaining possession is a multi-step outcome shaped by events after t_0, including the ball carrier's decision and technique, which a snapshot at the moment of possession loss cannot capture. Second, broadcast tracking captures only players within the camera frame, so the positions of players far from the ball were frequently missing, and these are often the players covering passing options. Lee et al. (2025) reported the same limitation with freeze-frame data, and the similarity of our model's performance to theirs (AUC 0.61) suggests that it reflects the difficulty of the task rather than a particular modelling choice. Finally, with 2,640 opportunities across six teams, flexible models remain prone to overfitting rare configurations, as indicated by the over-optimistic predictions in the highest decile of predicted probability.

Three primary findings address this research question. First, failures were rarely attributable to the positioning of the nearest presser alone: positional failures were 6% of failures, or 10% including those that were both positional and athletic. More often, the nearest presser either did not engage (26%) or engaged but closed down the ball carrier at less than 75% of their own capacity (36%). Because the athletic category is defined relative to each player's individual acceleration-speed profile, these failures do not indicate that players lacked the physical capacity

to arrive in time. Instead, they indicate that the capacity was not used, as players reached on average only [X]% of their theoretical maximal speed during a closing run. Taken together, the dominant explanation for failed counter-presses in this sample was engagement and intensity after the turnover, not starting position or athletic ceiling. Secondly, the ideal-ghost comparison showed that the potential gain from better positioning increased from the nearest presser (1.5 percentage points) to the second (2.2) and third (3.5). This is consistent with the view that counter-pressing is a collective behavior, in which the positioning of supporting players who restrict passing options matters as much as the pressure applied by the nearest player (Bauer & Anzer, 2021; Lee et al., 2025). The nearest presser is usually the player who has just lost the ball, so their position is largely determined by the preceding action. The positions of supporting players are where coaching may have the greatest influence. And finally, the positional and athletic components separated players by role. The highest positional ratings belonged mainly to forwards and wide attackers (Højlund, Nkunku, Malen, Esposito, Pulisic, Leão), who typically lose possession high up the pitch and are already close to the new ball carrier. However, these players closed down the ball carrier more slowly than their physical capacity allowed. Conversely, central midfielders such as Modrić, Lobotka, Cristante and Zambo Anguissa had the lowest positional ratings but the highest closing ratings. Their positions at the turnover are constrained by rest-defense responsibilities, which leaves them further from the ideal pressing position but closing with greater intensity relative to their capacity. This role-dependent trade-off has a practical implication: a single "counter-pressing ability" rating would conflate two distinct attributes. It also motivates the role-conditioned ghosts proposed for future work.

At the team level, the decomposition of regains over expected into structure and execution produced the most interpretable distinctions. Differences between teams arose almost entirely after the turnover:

- Juventus and Como owed their positive RoE to execution (+18.3 and +9.6 regains).
- AC Milan's shortfall was mainly execution (-11.0), compounded by the largest structural deficit of any team (-2.6).
- AC Milan's failure profile points the same way: it had the highest shares of both positional failures and failures in which the nearest player did not engage.
- Inter Milan illustrates the opposite case, with the most favorable structure of any team (+1.4) but no gain from execution.

These profiles suggest different coaching emphases. For teams such as AC Milan, the priority is the immediate reaction to possession loss, followed by positioning in rest defense. For teams like Inter Milan, the priority is converting favorable structure into regains.

Team rankings were robust to the outcome horizon ($\rho = 0.94$ between 8 seconds and every alternative from 3 to 10 seconds), even though the classification of individual opportunities as successes or failures depended strongly on the threshold. A fixed threshold therefore misclassifies many individual moments, but with a full season of opportunities it does not

substantially change which teams counter-press effectively. Nevertheless, with six teams and roughly 400-500 opportunities per team, no team estimate was conclusive. Juventus had the strongest evidence of outperforming expectation (posterior probability 0.85), and AC Milan the strongest evidence of underperforming it (0.18).

The absence of greater opponent threat after failed committed counter-presses also merits attention. When a committed counter-press failed, the opponent produced a shot in 14.2% of possessions (0.017 xG), almost identical to the 13.9% (0.016 xG) when a team did not commit and did not regain possession. The common concern that a failed counter-press exposes a team to counter-attacks was therefore not evident in the possession immediately following the turnover. This comparison is observational. All six teams were selective, committing more often when a regain was likely, so the situations in which teams chose not to commit may already have been more dangerous. Nonetheless, the result suggests that the risk side of the risk-reward trade-off is smaller than is often assumed, while committing was associated with a higher regain rate (27.4% vs. 19.1%).

Several directions could strengthen future applications of this framework. First, the value function should be improved before the ghost and decomposition components are interpreted at the level of individual players. This could be achieved by training on multiple seasons or leagues, modelling the outcome as time to regain with a survival model rather than a binary horizon, and including features from the frames immediately after t_0, such as the first pass option chosen by the ball carrier. Second, the ghosts in this study were conditioned on game state and therefore represent league-typical positioning, so a forward and a holding midfielder are compared against the same baseline. Conditioning ghosts on tactical role (Groom et al., 2026) or on team, and modelling the joint positioning of several pressers rather than one at a time, would better reflect the coordinated nature of counter-pressing. Third, the athletic component could be extended from single closing runs to frame-by-frame movement models with posterior predictive simulation (Nguyen & Yurko, 2026). This would allow the full trajectory of each presser to be evaluated, including whether and when they abandon the press. Finally, optical tracking with full pitch coverage would remove the missing-data limitations of broadcast tracking, particularly for the supporting players whose positioning appeared to matter most.

---

6 Conclusions

This paper introduced a framework for evaluating counter-pressing that combines regain probability modelling, ghost defenders estimated with random forests for conditional density estimation, and individualized acceleration-speed profiles within Bayesian multilevel models. Applied to broadcast tracking and event data from Serie A, the framework showed that a fixed five-second definition of success misclassifies a substantial share of regains, that regain probability is difficult to predict from the situation at the moment of possession loss, and that the

largest improvements in positioning were available for supporting pressers rather than the player nearest the ball. Players rarely approached their maximal physical capacity when closing down the ball carrier, suggesting that failed counter-presses were more often a matter of positioning and engagement than of athletic ceiling in this specific dataset.

At the team level, regains over expected and its decomposition into structure and execution distinguished teams whose counter-pressing performance came from post-turnover execution from those whose shortfall was rooted in positioning at the moment of possession loss, although no team estimate was conclusive in this relatively small sample. The present study therefore contributes to the broader field by formally introducing a framework that decomposes counter-pressing into positional, athletic, and situational components, and to document the findings and limitations of this specific application.

---

7 Practical Applications

For coaching staff who value counter-pressing in their game model, there are three takeaways from the present study. First, regain probability by situation can inform the cues that should trigger a counter-press, such as the manner and location of possession loss, and the situations in which re-establishing defensive shape is preferable. The results also suggest that the immediate risk of committing to a counter-press may be lower than commonly assumed. In other words, when a counter-press failed, the opponent was no more likely to produce a shot in that specific possession sequence than when the team did not commit to the counter-press (14.2% vs. 13.9%). Second, regains over expected identifies teams that regain possession more often than their opportunities would predict, and the structure–execution decomposition indicates whether this advantage arises from positioning before the turnover or from behavior after it. Combined with the list of their most over-expected regains and missed high-probability opportunities, this provides a structured basis for video study and for interviews with the coaching staff of those teams about how counter-pressing is trained. And finally, in player development and recruitment, the positional and closing components describe whether a player tends to be well positioned at the moment of possession loss and whether they close down the ball carrier efficiently relative to their own physical capacity. Because these estimates reported herein are uncertain at the sample sizes available in a single season, they should be used alongside video and coaching assessment rather than as stand-alone rankings; nevertheless, this could be replicated with more comprehensive datasets to return more narrow credible intervals.

In summary, this framework moves counter-pressing evaluation beyond whether a team won the ball back within an arbitrary time window towards understanding why it did or did not, giving coaches, analysts, and recruiters support for deciding when to press, whom to learn from, and which players are best suited to execute it.

References

Andrienko, G., Andrienko, N., Budziak, G., Dykes, J., Fuchs, G., von Landesberger, T., & Weber, H. (2017). Visual analysis of pressure in football | Data Mining and Knowledge Discovery | Springer Nature Link. Data Minining and Knowledge Discovery. https://link.springer.com/article/10.1007/s10618-017-0513-2

Bauer, P., & Anzer, G. (2021). Data-driven detection of counterpressing in professional football. Data Mining and Knowledge Discovery, 35(5), 2009–2049. https://doi.org/10.1007/s10618-021-00763-7

Bauer, P., Anzer, G., & Shaw, L. (n.d.). Putting team formations in association football into context. Retrieved October 23, 2026, from https://journals.sagepub.com/doi/10.3233/JSA-220620

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

Everett, G., Beal, R. J., Matthews, T., Norman, T. J., & Ramchurn, S. D. (2025). Evaluating Defensive Influence in Multi-Agent Systems Using Graph Attention Networks. 2025 IEEE 12th International Conference on Data Science and Advanced Analytics (DSAA), 1–10. https://doi.org/10.1109/DSAA65442.2025.11248003

Forcher, L., Forcher, L., Altmann, S., Jekauc, D., & Kempe, M. (2024). Is a compact organization important for defensive success in elite soccer? – Analysis based on player tracking data. International Journal of Sports Science & Coaching, 19(2), 757–768. https://doi.org/10.1177/17479541231172695

Forcher, L., Preine, L., Forcher, L., Wäsche, H., Jekauc, D., Woll, A., Gross, T., & Altmann, S. (2023). Shedding some light on in-game formation changes in the German Bundesliga: Frequency, contextual factors, and differences between offensive and defensive formations. International Journal of Sports Science & Coaching, 18(6), 2051–2060. https://doi.org/10.1177/17479541221130054

Groom, S., Wang, S., Belo, F., Rice, A., & Anderson, L. (2026). A Machine Learning Framework for Off Ball Defensive Role and Performance Evaluation in Football (arXiv:2601.00748). arXiv. https://doi.org/10.48550/arXiv.2601.00748

Gu, C., & Silva, V. D. (2023). Deep Generative Multi-Agent Imitation Model as a Computational Benchmark for Evaluating Human Performance in Complex Interactive Tasks: A Case Study in Football (arXiv:2303.13323). arXiv. https://doi.org/10.48550/arXiv.2303.13323

Kim, H., Seo, S., Choi, H., Boomstra, T., Yoon, J., & Park, C. (2025). Better Prevent than Tackle: Valuing Defense in Soccer Based on Graph Neural Networks (arXiv:2512.10355). arXiv. https://doi.org/10.48550/arXiv.2512.10355

Le, H. M., Carr, P., Yue, Y., & Lucey, P. (2017). Data-Driven Ghosting using Deep Imitation Learning.

Le, H. M., Yue, Y., Carr, P., & Lucey, P. (2017). Coordinated Multi-Agent Imitation Learning. Proceedings of the 34th International Conference on Machine Learning, 1995–2003. https://proceedings.mlr.press/v70/le17a.html

Lee, M., Jo, G., Hong, M., Bauer, P., & Ko, S.-K. (2025). exPress: Contextual Valuation of Individual Players Within Pressing Situations in Soccer.
Low, B., Rein, R., Raabe, D., Schwab, S., & Memmert, D. (2021). The porous high-press? An experimental approach investigating tactical behaviours from two pressing strategies in football. Journal of Sports Sciences, 39(19), 2199–2210. https://doi.org/10.1080/02640414.2021.1925424
Lowe, Z. (90392). » Lights, Cameras, Revolution. https://grantland.com/features/the-toronto-raptors-sportvu-cameras-nba-analytical-revolution/
Merckx, S., Robberechts, P., Euvrard, Y., & Davis, J. (n.d.). Measuring the Effectiveness of Pressing in Soccer.
Nguyen, Q., Jiang, R., Ellingwood, M., & Yurko, R. (2025). Fractional tackles: Leveraging player tracking data for within-play tackling evaluation in American football. Scientific Reports, 15(1), 2148. https://doi.org/10.1038/s41598-025-85993-1
Nguyen, Q., & Yurko, R. (2026). Bayesian multilevel step-and-turn models for evaluating player movement in American football (arXiv:2603.17866). arXiv. https://doi.org/10.48550/arXiv.2603.17866
Robberechts, P. (n.d.). Valuing the Art of Pressing.
Schmid, M., Blauberger, P., & Lames, M. (2021). Simulating Defensive Trajectories in American Football for Predicting League Average Defensive Movements. Frontiers in Sports and Active Living, 3, 669845. https://doi.org/10.3389/fspor.2021.669845
Seidl, T., † A., † A., Carr, P., & Lucey, P. (2018). Bhostgusters: Realtime Interactive Play Sketching with Synthesized NBA Defenses.
Yurko, R., Nguyen, Q., & Pelechrinis, K. (2025). NFL Ghosts: A framework for evaluating defender positioning with conditional density estimation (arXiv:2406.17220). arXiv. https://doi.org/10.48550/arXiv.2406.17220