Chasing Shadows: Detecting and Valuing Decoy Runs

Francesco Esposito[1,2], Lorenzo Cascioli[3], Maria Angonese[1]

[1] *SportAnalytics, Milan, Italy*
[2] *Department of Medicine and Aging Sciences, University "G. d'Annunzio" of Chieti-Pescara, Chieti, Italy*
[3] *KU Leuven, Dept. of Computer Science; Leuven.AI, B-3000 Leuven, Belgium*

1. Introduction

Off-ball behaviour accounts for the vast majority of a football player's activity during a match, playing a central role in space creation, defensive organisation, and team coordination [1]. Each player possesses the ball for less than three minutes per match [2] and spends most of the game performing off-ball movements such as marking, pressing, or running into open space. However, many widely used performance metrics focus on actions performed with the ball, such as expected goals from shots [3] or the value of passes and other on-ball actions [4]. In practice, much of what happens on the pitch remains unaccounted for by these metrics. For example, consider a forward who drops short and pulls a centre-back out of the defensive line, allowing a running teammate to receive the ball at a location previously controlled by the defender. Even if the forward never touches the ball, their movement can be crucial to developing the play. Such movements, commonly referred to as *decoy runs*, are an important aspect of attacking play, yet their impact remains difficult to quantify.

Until recently, off-ball behaviours were difficult to study at scale, since widely available event stream data primarily describe on-ball actions and do not capture how off-ball players move during an event. Therefore, a player performing a decoy run without touching the ball leaves no direct trace in the event data. However, the increasing availability of tracking data has helped overcome this limitation. Tracking data record the positions of all players and the ball several times per second, providing the necessary context to evaluate off-ball positioning and movement.

Researchers have increasingly investigated off-ball defensive behaviours, such as the ability of defenders to anticipate opponents' play [2] or block passing lanes through cover shadows [5], as well as offensive behaviours, including the generation of space for teammates by attracting defenders [6], the evaluation of spatial control and space creation through player movement models [7], and the creation of scoring opportunities through alternative off-ball trajectories [8]. While these approaches demonstrate the value of off-ball movements, they do not explicitly identify decoy runs based on a defender's deviation from the coordinated movement of their defensive line or quantify the resulting loss of defensive control relative to a counterfactual

reference. Moreover, assessing the opportunities created by such movements requires understanding whether and how teammates could exploit the resulting space.

In this paper, we propose a data-driven approach to detect and quantify the effect of decoy runs in soccer. First, we identify attacking runs that engage a defender, causing them to follow the attacker along a trajectory that deviates from the coordinated movement of the defensive line.

Second, we measure the value generated by a decoy run through the space left by the engaged defender. We estimate a counterfactual trajectory describing how the defender might have moved without the run, and identify the areas they would have controlled in this scenario but no longer dominate in the observed play. We then quantify the potential offensive value of this space by combining its positional threat with the modelled likelihood that the ball is played towards it.

Finally, we assess whether and how the decoy runner's teammates could have better exploited the opportunities created by the decoy run. We formulate an optimization problem to identify alternative, physically feasible movements for individual teammates that maximize the attacking team's modelled spatial value. Importantly, these movements may generate greater value not only by exploiting the space directly opened by the decoy run, but also by gaining control of valuable areas elsewhere on the pitch.

In short, this work makes three key contributions:

- **(1):** We provide an analytical method to detect decoy runs from tracking data by identifying attacking movements that engage defenders and disrupt their defensive line.
- **(2):** We introduce a counterfactual-based metric to quantify the potential value of the space generated by a decoy run.
- **(3):** We propose an optimization approach to identify alternative teammate movements that better exploit the opportunities created by a decoy run, both within and beyond the space directly opened.

2. Data and Definitions

2.1 Dataset Overview

We use StatsBomb tracking data synchronised with StatsBomb event stream data. The tracking data record the position of all players and the ball 25 times per second. The event stream data is extracted from broadcast video and describes each on-the-ball action with standard metadata such as timestamp, event type, player, team, body part, outcome and location. Player and ball positions are expressed in meters on a standardised 105 x 68 pitch coordinate system. We exclude periods of play in which ball tracking data are unavailable.

Our dataset consists of 30 matches from the 2025/26 Serie A season. Specifically, we consider matches between two of the top six teams of the season, so six teams and ten matches per team. We divide these matches into two disjoint sets: a *development set* (20 matches) used to calibrate

the parameters for the detection and evaluation procedures, and a *held-out set* (10 matches) used to evaluate the proposed method and develop the use cases.

2.2 Defining an Attacking Run

An attacking run is any off-ball movement in which a player moves at a high enough intensity. Concretely, we define this requirement as reaching a speed of at least 3.98 m/s. This threshold corresponds to the lower bound of the "running" category according to Bradley et al. [9] and the lowest entry speed for high-speed running among the studies reviewed by Gualtieri et al. [10]. Walking and jogging movements are excluded from the candidate population. The start of an attacking run is defined as the estimated onset of acceleration preceding the speed-threshold crossing. The run ends when the velocity drops below 3.50 m/s.

2.3 Defining a Decoy Run

Traditionally, a decoy run has been described as an off-ball movement intended to distract defenders and create space for teammates without the intention of receiving the ball. However, FIFA also includes decoy runs among movements that constitute an offer to receive, highlighting their role in displacing opponents and creating space [11].

Therefore we consider attacking runs regardless of whether the runner subsequently receives the ball, and we adopt an operational definition based on the observed defensive behaviour in relation to the attacking run. Specifically, we define as decoy runs the attacking runs where at least one defender closely follows the runner and leaves the defensive structure in doing so. Such displacement can open up space for the runner's teammates [6], although the extent and tactical value of this space must be assessed separately. This perspective aligns with existing research highlighting that off-ball movements can be assessed through changes in individual defensive pressure [12].

Moreover, as we are interested in runs that attract a defender out of an organised defensive line, we only consider attacking runs occurring during organised in-possession phases and exclude transition phases. Specifically, we use a phase-of-play classification informed by FIFA's Enhanced Football Intelligence framework and focus on build-up, progression, final third and long ball [13]. In these phases, the defending team usually adopts an organised defensive structure and the attacking team uses off-ball movements to create space for ball progression.

3. Detecting a Decoy Run

The goal of the first research question is to identify decoy runs from the tracking data, which allows us to study their spatial and tactical impact in the subsequent sections. Our operational definition of a decoy run encompasses two conditions. An attacking run is a decoy run if (1) a defender exhibits movement consistent with following the attacker and (2) the defender's movement deviates from the coordinated movement of their teammates in the defensive line. We now examine these two conditions in detail.

3.1 Is a defender following the run?

We first analyse each defending player's movement individually to determine whether they follow a running attacker. We restrict our analysis to defenders belonging to the last defensive line at the start of the attacking run, as their displacement may create valuable attacking opportunities in advanced areas of the pitch. Extending the analysis to other defensive lines would require considering additional tactical benefits, such as opening passing lanes between defensive units. Defensive-line membership is established at run onset and remains fixed throughout the episode.

Given an attacking run, we consider a defender to be following the attacker if:

1. They are positioned near the attacker at the start of the attacking run.
2. They closely follow the runner while moving at running intensity for a sufficiently long period during the attacking run.

The first rule simply excludes defenders that are distant from the attacker at the start of the run. Even if they begin following the attacker later in the run, they may not be reacting to the original run, as player movement can be influenced by more complex team-level movement patterns. For example, consider a defender who is initially far from the runner and is already following another attacker moving towards the goal. Suppose that the runner later approaches the defender and continues alongside them. While it appears that the defender is now following the attacker, they simply kept following the runner's teammate and did not respond to this run, hence they should not be considered as engaged. In practice, we focus on defenders positioned within a 5m radius from the attacker at the start of the run. This spatial scale is informed by Narizuka and Yamazaki's [14] finding of stronger directional alignment below approximately 5 m and by Fernández and Bornn's [6] use of a 5 m closeness criterion in their space-generation framework.

The second rule captures the concrete concept of engagement between defender and attacker, which we define as follows. Let R denote the runner and D_R the set of candidate defenders identified at run onset (i.e., close enough to the runner). Then, let t be any time instant between the start of the run and its end, $d_{R,D}(t)$ be the distance between runner R and defender D, and

$$\theta_{R,D}(t) = arccos\left(\frac{\mathbf{v}_R(t) \cdot \mathbf{v}_D(t)}{\left\|\mathbf{v}_R(t)\right\|\left\|\mathbf{v}_D(t)\right\|}\right)$$

the angle between their velocities. D is engaged with R at time t when the following three conditions hold:

1. $\left\|\mathbf{v}_R(t)\right\| \geq 3.98\ m/s$ and $\left\|\mathbf{v}_D(t)\right\| \geq 3.98\ m/s$: both runner and defender are moving at running intensity.
2. $d_{R,D}(t) \leq 3m$: D is close to R.
3. $\theta_{R,D}(t) \leq 30°$: D is moving in a similar direction to R.

We define an *engagement episode* as a maximal uninterrupted sequence of engaged frames lasting at least $T_{max} = 0.8s$. The episode ends at the last frame before any of the three given conditions fails, or at the end of the attacking run. We discuss the sensitivity of our detection algorithm to the value of T_{max} in Appendix 1.

Finally, a defender is *engaged* by an attacking run if we find at least one engagement episode.

Note that we use a 3 m proximity threshold to capture defenders following attackers without requiring tight man-marking. The selected spatial scale is consistent with the 3 m attraction parameter used by Fernández and Bornn [6] in their space-generation framework. This is different from the 5m threshold that determines candidate defenders as the latter only captures defenders close to the attacker at run onset. In order to be engaged, a candidate defender must clearly follow the runner which is established by the stricter proximity threshold. We further combine proximity with directional alignment because proximity alone does not establish that two players are moving together. This combination is consistent with approaches that identify marking interactions through both spatial proximity and directional alignment [15].

Therefore, at the end of this first step we have identified the subset of defenders whose movements were consistent with following the attacker during the run. We refer to these players as the *engaged defenders*.

3.2 Did the attacking run alter the defender's trajectory of motion?

An attacking run engaging a defender is not necessarily a decoy run. Consider a defender dropping deep to follow an attacking run while the rest of the defensive line retreats by a similar amount and in the same direction: the defender may have been engaged by a run, but has barely deviated from its expected movement in coordination with his teammates. In this case, the defender is not leaving the defensive structure, so the run is not a decoy run under our definition. More generally, it is well known that individual movements are often synchronized as a consequence of collective movements, especially along the length of the pitch [16].

Therefore the second step of our pipeline analyzes the movement of each engaged defender in relation to their teammates, to understand whether the defender deviates from their typical trajectory of motion as a consequence of the attacking run. We do so following a three-stage procedure.

I. Estimate Defensive Line Movement

First, we estimate the collective velocity of the defensive line at each frame $\mathbf{v}_{line}(t) = median_{d \in L \setminus \{D\}} \mathbf{v}_d(t)$, where L is the set of defenders belonging to the last defensive line, D is the defender being analyzed, and $\mathbf{v}_d(t)$ is a line-mate's velocity vector. Line membership remains fixed throughout the episode and is established at run onset. Using this local tactical unit as a reference is consistent with evidence that subgroup-level analysis can reveal movement patterns that whole-team measures obscure.

We calculate the median separately for the longitudinal and lateral velocity components. Excluding D prevents their own movement from influencing the reference against which they are assessed. We use the median because it is less sensitive than the mean to atypical movement by an individual line-mate.

II. Quantify Defender Displacement from the Defensive Line

Second, we quantify how far the engaged defender deviates from the defensive line's collective movement while following the attacking run. We do so by comparing the position of the defender at the end of the engagement episode $\mathbf{p}_{run}$ with an alternative position $\mathbf{p}_{line}$ that tracks the typical movement of a defender who simply follows the coordinated movement of the defensive line.

We define $\mathbf{p}_{line}$ as follows. Let t_R be the starting time of the attacking run and t_{end} the end time of the identified engagement episode. The reference trajectory is obtained by starting from the defender's position at t_R , which we call $\mathbf{p}_D(t_R)$, and accumulating the defensive line's velocity over each considered time instant:

$$\mathbf{p}_{line}(t_{end}) = \mathbf{p}_D(t_R) + \sum_{t=t_R}^{t_{end}} \mathbf{v}_{line}(t)\Delta t$$

where Δt=0.04 is the interval between two tracking frames. The reference trajectory estimates where a defender positioned at $\mathbf{p}_D(t_R)$ at the beginning of the run would have ended up if they had simply moved with the defensive line.

Then, we extract the defender displacement from the defensive line as:

$$meters\ out\ of\ line = \left\| \mathbf{p}_{run}(t_{end}) - \mathbf{p}_{line}(t_{end}) \right\|.$$

We illustrate the metric with an example in Figure 1, looking at an attacking run from Napoli's Matteo Politano against AC Milan. In the example, Politano comes short and Davide Bartesaghi, the engaged defender, steps out of his line. Bartesaghi ends up 10.2 meters from the position indicated by the line-movement reference.

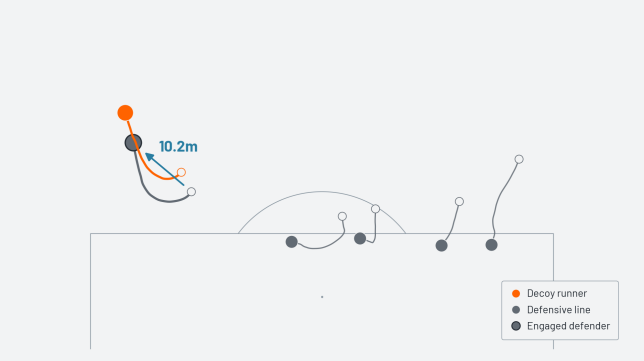


Figure 1. Measuring displacement from the defensive line. Politano's run is shown in orange and Bartesaghi's observed trajectory in dark grey. The dashed path shows where the collective movement of his line-mates would have taken him; the open circle marks his final line reference position. At the end of the measurement interval, Bartesaghi's observed position is 10.2 m from this reference

III. Identify Exceptional Displacements from the Defensive Line

The final step identifies engagements where the defender's displacement is exceptionally large and therefore indicates that the defender has left the defensive line to follow the runner. Our key assumption is that a defender following a decoy run departs from the line's collective movement more markedly than their non-engaged line-mates. We therefore use the displacements of non-engaged last-line defenders as our references to understand when the displacement of the engaged defender is normal and when it is not.

To build this reference set, we measure the displacement from the defensive line of all non-engaged defenders from each of the 645 identified engagement episodes in the 20 matches in our development set. We calculate displacement as described above for the engaged defender, considering non-engaged defenders' movements over the same interval as the engaged defender: from the attacker's movement onset to the end of the engagement.

Two factors clearly influence the raw values of these displacements: the duration of the interval over which displacement is measured and defensive line speed. Longer measurement intervals allow defenders more time to accumulate displacement, while faster-moving defensive lines may undergo greater internal reorganisation. Consequently, a single displacement threshold applied to all engagements could classify longer episodes as exceptional more frequently than shorter ones, regardless of how unusual the defender's movement is.

We therefore organize the reference set into three duration groups (short, intermediate, and long) and three defensive line speed groups (slow, intermediate, and fast). Their combinations form nine subsets. This allows us to assess each engaged defender's displacement relative to non-engaged defenders observed under comparable conditions, rather than applying a single threshold across all situations. We evaluate the benefit of this grouping against a global-threshold baseline in Section 3.3.

Specifically, we check how rarely displacements of that magnitude occur among non-engaged defenders in similar situations. We set the threshold at the 95th percentile within each group. An engaged defender's displacement is classified as exceptionally large when it exceeds this threshold.

We calculate the speed and duration group limits from the 20 non-held-out matches, counting each attacking run interval once regardless of how many defenders it contains (Table 1). These limits divide each variable into three intervals and are used to assign engagements to the nine comparison groups in both the development and held-out matches.

Third / Group Type	Displacement duration	Defensive line speed
Lower third	≤ 2.52 s	≤ 2.56 m/s
Middle third	> 2.52 and ≤ 3.88 s	> 2.56 and ≤ 3.634 m/s
Upper third	> 3.88 s	> 3.634 m/s

Table 1. Duration and defensive-line speed group limits, calculated from the 20 non-held-out matches.

In the Politano–Bartesaghi example, the line moves at a speed of 1.36 m/s and the measurement lasts 2.96 s, placing the interaction in the slow-line, intermediate-duration group. The group-specific threshold for exceptional displacements is 8.92 m. Bartesaghi's displacement of 10.2 m exceeds this value and is therefore classified as exceptionally large, meaning that this episode is detected by our algorithm as a decoy run.

3.3 Evaluating the detection algorithm

As ground-truth annotations of decoy runs are unavailable, we evaluate our detection procedure indirectly. First, we assess whether group-specific displacement thresholds provide a more consistent definition of exceptional displacement than a single global threshold. We then examine whether exceptional displacements are associated with defender engagement and whether selected defenders stand out within their defensive lines. Together, these analyses provide indirect evidence supporting our detection criterion and help identify the conditions under which it is most reliable.

Does grouping by speed and duration improve the comparison? We first ask whether dividing the reference into speed–duration groups justifies the added complexity, or whether a single threshold calculated from all non-engaged defenders would be sufficient.

We compare one 95th-percentile threshold across all non-engaged displacements with nine thresholds calculated separately for the speed–duration groups. For this comparison, both alternatives are estimated from the 20 non-held-out matches and applied unchanged to the 10 held-out matches. The single threshold is 11.99 m.

The 95th-percentile threshold defines exceptionally large displacement as the upper 5% of the non-engaged reference. We therefore assess whether each approach produces exceedance rates close to 5% within the held-out speed–duration groups. Similar rates across groups would indicate that "exceptionally large" represents a comparable degree of rarity across different line speeds and measurement durations. Table 2 reports the exceedance rates obtained with both approaches

Speed Line / Episode Duration	Short duration	Intermediate duration	Long duration
Slow speed	0.0 → 6.3%	1.0 → 4.7%	10.0 → 3.0%
Intermediate speed	0.0 → 13.6%	5.2 → 6.0%	13.4 → 6.3%
Fast speed	0.0 → 4.1%	4.3 → 3.2%	14.6 → 5.7%

Table 2. Percentage of non-engaged defenders–interval observations exceeding the threshold in the held-out matches. Each cell reports a single global threshold → group-specific threshold. Values closer to 5% indicate closer agreement with the reference rarity level.

The single threshold classifies no short-interval displacement as exceptionally large, compared with 10.0–14.6% of long-interval displacements. This illustrates the risk of confusing the time available to accumulate displacement with how exceptional the movement is. A prolonged adjustment to cover a teammate, for example, has more time to produce displacement relative to the line than a brief step out of position.

Group-specific thresholds reduce this duration imbalance, bringing exceedance rates close to the 5% reference level in most groups. They therefore help assess each movement against a more comparable reference, rather than allowing measurement duration to dominate its classification. However, the intermediate-speed, short-duration group remains an exception: 13.6% of its non-engaged displacements exceed the threshold. Grouping improves comparability, but does not fully resolve differences between groups.

Are exceptionally large displacements associated with engagement? Our second evaluation asks whether exceptionally large displacements are associated with engagement. Specifically, does knowing which defender is engaged help us identify an exceptionally large displacement, that is, one above the threshold? We assess this by comparing the engaged defender with a random choice within the same defensive line.

This comparison serves a different purpose from the reference collected across matches: that reference sets the threshold, whereas the within-line comparison checks whether the engaged defender exceeds it more often than a defender chosen at random from the same line.

For each engagement event, we compare choosing the engaged defender with choosing randomly among that defender and their non-engaged line-mates. If one of five defenders exceeds the threshold, a random choice has a one-in-five chance, or 20%, of identifying an exceptionally large displacement.

Across all engagements, these chances give the number of above-threshold cases expected from random selection. We compare this with how often the engaged defender *actually* exceeds the threshold.

Table 3 reports the observed count divided by the random expectation. A value of one means that choosing the engaged defender performs like random selection; a value of three means it identifies an exceptionally large displacement three times as often.

Defensive-line speed/ Dataset	Observed/Expected: non-held-out	Observed/Expected: held-out
Slow	3.32	3.28
Intermediate	1.38	1.10
Fast	1.14	0.84

Table 3. Association between defender engagement and exceptional displacement across defensive-line speed groups. Values represent the ratio of observed above-threshold engagements to the number expected from random defender selection within the same defensive line. Values above 1 indicate a stronger association than random selection.

The association is strongest in slowly moving lines. Engaged defenders exceed the threshold in 77 engagements, compared with 23.2 expected from random selection in the non-held-out matches, and in 42 engagements, compared with 12.8 expected in held-out, approximately 3.3 times the random expectation in both sets. In faster lines, the ratios are much closer to one, or below it. Collective reorganization may make an individual defender's departure harder to distinguish from movements occurring throughout the line. A relatively stable line provides a clearer reference against which that departure can stand out. This also fits the tactical situation we focus on: a decoy run creating space by disrupting an organized defensive structure when collective movement is limited. These results support our focus on the slowest defensive-line group, with a median speed of at most 2.56 m/s. For the subsequent spatial analysis in section 4, we retain engagements in this group whose displacement exceeds the relevant duration-specific threshold, including all three duration groups.

Do selected defenders stand out within their defensive line? Exceeding the threshold means that a displacement is exceptionally large compared with non-engaged movements in the same speed–duration group. We now ask whether the selected defender also shows the largest displacement among their non-engaged line-mates over the same interval. This provides an additional check that the selected defender stands out within the local defensive structure.

The selected defender shows the largest displacement among their non-engaged line-mates in 72 of 77 non-held-out interactions and 40 of 42 held-out interactions. This comparison describes the selected cases; it does not determine selection. Being first is automatic when the engaged defender is the only player above the threshold. We therefore also examine interactions in which at

least one non-engaged line-mate exceeds it. In this subset, the engaged defender still has the largest displacement in 25 of 32 interactions (78%). The selected defender therefore usually stands out within their line, even when another defender also shows an exceptionally large displacement.

Together, these checks show that exceptionally large displacement is associated with engagement most clearly in slowly moving defensive lines. We next examine the spatial consequences of the selected interactions, asking how the defender's departure changes the opportunities available to the runner and their teammates.

4. Valuing a Decoy Run

A decoy run may benefit the attacking team by drawing a defender away from areas they would otherwise cover. An intuitive way to assess its value is therefore to identify the space the defender leaves behind and estimate how valuable it is to the attack.

We value each candidate decoy run identified in the previous section through three steps:

1. Estimate the defender's counterfactual trajectory: we estimate how the engaged defender might have moved without the decoy run, combining the defensive line's collective movement with the defender's initial velocity.
2. Identify the hole: we compare the defender's spatial dominance under the counterfactual and observed trajectories using pitch control, identifying areas they would dominate in the counterfactual but no longer dominate in the observed game state.
3. Value the hole: we weight these cells by their expected possession value and a distribution of possible next ball destinations, estimating the potential attacking value of the space freed up by the decoy run.

4.1 Estimating the defender's counterfactual trajectory

We first estimate how the engaged defender might have moved in the absence of the decoy run, assuming that their movement would have largely followed the coordinated motion of the defensive line. Concretely, this requires simulating a counterfactual trajectory for the defender's movement. Counterfactual trajectories have been used to assess off-ball attacking contribution [8] and defensive positioning at corners [17]. We construct our counterfactual trajectory by starting from the defender's observed position at the beginning of the attacking run and letting the defender move using the component-wise median velocity of the defensive line (identified as described in Section 3.2). We also account for how the defender was already moving relative to the defensive line at the onset of the attacking run t_R. We model this initial relative velocity as an additional component that decays over time, so that the defender's counterfactual velocity gradually approaches the line's velocity. Specifically, we express the counterfactual residual velocity at time t as:

$$\hat{\mathbf{v}}_D^{\,rel}(t) = [\mathbf{v}_D(t_R) - \mathbf{v}_{line}(t_R)]e^{-(t-t_R)/\tau}$$

where $\mathbf{v}_D(t_R)$ is the defender's observed velocity at the attacking-run onset, and $\mathbf{v}_{line}(t_R)$ is the defensive-line velocity defined in Section 3.2. The residual term preserves the movement the defender already had relative to the line when the attacking run began. We assume that, without a response to the attacking run, this initial relative motion would gradually fade. The parameter τ controls how quickly this initial relative velocity decays. We selected τ = 2.75 s by minimising the median endpoint prediction error relative to the defensive line for non-engaged defenders in the 20 development matches, evaluating values from 0.25 to 8s in 0.25s increments. Compared with a baseline model using only the defensive line's collective velocity, including the decaying residual reduced this error by 17.6% on the ten held-out matches (95% CI: 15.5–21.9%), with improvement in every match, supporting its inclusion in the counterfactual model.

The defender's counterfactual velocity combines the line's velocity with this decaying residual component:

$$\hat{\mathbf{v}}_D(t) = \mathbf{v}_{line}(t) + \hat{\mathbf{v}}_D^{\,rel}(t)$$

At t_R, the counterfactual velocity equals the defender's observed velocity. As the residual component decays, the defender increasingly follows the line's collective movement. Starting from the defender's observed position at the trajectory's initial time t_R, we obtain the counterfactual trajectory by accumulating these velocities:

$$\hat{\mathbf{p}}_D(t) = \mathbf{p}_D(t_R) + \sum_{s:t_R \le s < t} \hat{\mathbf{v}}_D(s)\Delta t$$

where s indexes tracking timestamps and $\Delta t = 0.04s$. The resulting trajectory estimates where the defender would be at each instant if they followed the line's collective movement while retaining a gradually decaying component of their initial relative motion. We use this position and the corresponding counterfactual velocity in the pitch-control comparison described in Section 4.2.

4.2 Defining the space freed up by the decoy run

The potential value generated by a decoy run comes from the space left behind by the engaged defender as they move out of position. Therefore the second step identifies which areas the engaged defender no longer dominates compared with the counterfactual defined in Section 4.1. To do so, we use a discretized cell grid to represent the pitch, describing each cell with its centre coordinates (x, y). We assume that the ball carrier can attempt a pass towards each of these cells, and use pitch control to identify the dominant player in each cell.

Pitch control is a well known physics-based model that estimates each player's chance of controlling the ball if the latter is played towards a given location, given the current game state [18]. We compute pitch control using a vectorised implementation based on LaurieOnTracking [19] and assign dominance of each cell to the player with the highest individual pitch-control contribution.

At each sampled instant, we compute these contributions for two configurations. In the first, the engaged defender has the counterfactual position and velocity defined in Section 4.1; in the second, we use their observed position and velocity. Their line-mates follow the same line-based

reference trajectories in both configurations, while the ball and all other players remain as observed. Thus, only the engaged defender's position and velocity differ between the two configurations.

We define the *hole* as the set of pitch cells, represented by their centre coordinates (x, y), dominated by the engaged defender in the first configuration but no longer dominated by that defender in the second:

$$H(t) = \{(x, y) : arg\,max_j\, PC_j^{(1)}(x\,y, t) = D \;\; and \;\; arg\,max_j\, PC_j^{(2)}(x, y, t) \neq D\}$$

where D is the engaged defender, and $PC_j^{(k)}(x, y, t)$ denotes player j's individual pitch-control contribution at the cell (x, y) and time t in configuration k. The maximisation considers all players and identifies the dominant player, with ties resolved consistently across configurations.

The hole therefore represents a loss of individual dominance by the engaged defender. Its cells may remain dominated by another defender, so the hole alone does not establish an attacking advantage. We assess who dominates these cells separately using all players' observed positions and velocities. We recalculate the hole at each sampled instant because its boundary can change as play unfolds.

Figure 2 illustrates this comparison for a representative frame, showing the engaged defender's observed and counterfactual positions, with the resulting hole.

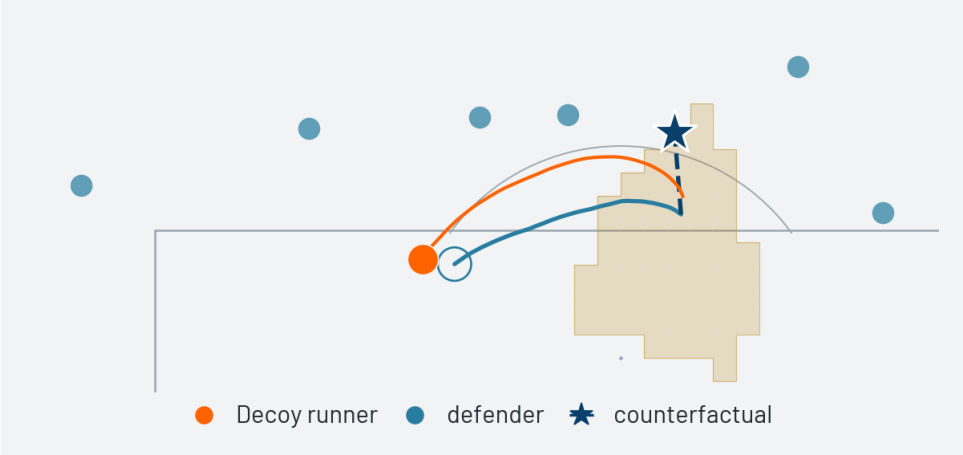


Figure 2. Observed trajectories of the decoy runner (orange) and defender (blue), with the defender's counterfactual position shown by the star. The shaded area marks the resulting hole.

4.3 Valuing the hole

After identifying the hole, we estimate how valuable this space is to the attacking team.
We define the value of the hole as the sum of its cell-level contributions. At each considered (x,y) location, we quantify the offensive threat as the combination of two components: (1) the positional value of an attacker possessing the ball at that location, $threat(x, y)$, and (2) a spatial weighting

function representing the assumed distribution of possible next ball destinations $p(x, y)$. Therefore:

$$V_H(t) = \sum_{(x,y)\in H_t} threat(x, y) * p(x, y).$$

Concretely, we estimate the two terms as:

- $threat(x, y) = EPV(x, y)$, i.e., we simply define positional threat using the expected possession value (EPV) grid from Teranishi et al. [8];
- $p(x, y) = (1/Z) * exp\{-[(x - b_x)^2 + (y - b_y)^2]/2\sigma^2\}$ i.e., the relative likelihood of the ball being played towards location (x, y) is modeled as a Gaussian distribution centred on the ball's position (b_x, b_y) at the evaluation instant. The parameter σ controls its spatial spread, while Z normalises the weights to sum to one across all pitch cells.

At each evaluation instant, we compare the observed and counterfactual configurations at that same moment. We then integrate the resulting hole value over time to account for both the value of the space and how long it persists. An opening that maintains the same value for longer contributes more than one that disappears quickly. The episode-level measure therefore captures the cumulative potential opportunity associated with the hole.

Alternative formulations of the two components are possible. Similar approaches in the literature combine a spatial value surface with models of ball control, pass success, or scoring likelihood [6], [20].

5. Optimizing teammate movement in response to a decoy run

Drawing a defender away from his line does not necessarily make the hole available to the attack. Other defenders may still be better placed to control a ball played there. Across the 119 episodes examined here, the defence dominates 69.3% of the total hole value, with the displaced defender's line-mates alone accounting for 36.4% of the total. The runner dominates 9.0%, while his other teammates dominate only 21.7%. The runner may have drawn his opponent away, yet his teammates do not necessarily gain control of the hole.

Could the runner's teammates have made better use of the hole, or was there more value in moving elsewhere? To address this question, we use an optimization approach to identify alternative player behaviors, an idea previously explored by Stein et al. [21] and Peralta Alguacil et al. [22]. We optimize one teammate's movement at a time, searching for a physically admissible trajectory that maximizes the value added to the attacking team compared with his observed route. We then examine whether this movement would have taken him into the hole or towards other areas of the pitch, helping us assess when the hole created by the decoy was worth attacking.

5.1 The optimization objective

A supporting run can help the attack control a valuable area, while leaving another area less well controlled. Whether the move is worthwhile depends on where these changes occur and how likely the ball is to be played there. The optimization identifies the feasible movement that offers the greatest overall gain in team value compared with the teammate's observed route.

Let π_i denote an alternative trajectory for teammate i, F_i the set of trajectories allowed by the movement constraints, and $G_i(\pi_i)$ the gain in team value relative to his observed movement. The constraints are described in Section 5.2. We seek:

$$max_{\pi_i \in F_i} G_i(\pi_i).$$

The teammate follows his observed trajectory for a response delay of 0.7 s after the decoy run begins. This delay is introduced to account for the time required to perceive and respond to the movement, drawing on the reaction-time assumption used in the pitch-control model [20]. We then evaluate his alternative movement during the decoy episode, from the first evaluated instant after the delay t_0 to the last instant used to follow the hole t_1.

To calculate G_i, we compare the attacking team's control under the alternative and observed movements. Let $PC_{att}(x, y, t \mid \pi_i)$ denote the team's modelled probability of controlling a ball played to cell (x, y) when teammate i follows π_i and let ${\pi_i}^{obs}$ be his observed trajectory. Then:

$$G_i(\pi_i) = \int_{t0}^{t1} \sum_{(x,y)\in\Omega} [PC_{att}(x, y, t \mid \pi_i) - PC_{att}(x, y, t \mid {\pi_i}^{obs})]\, threat(x, y) p_t(x, y) dt\,.$$

Where Ω is the whole-pitch grid. We reuse the positional value $threat(x, y)$ and Gaussian ball-destination model $p_t(x, y)$ from Section 4.3; the subscript in p_t indicates that its centre follows the observed ball position. These factors are the same for both trajectories, so G_i measures the change in team value due to the change in control.

Summing over the whole pitch accounts for both the control gained by moving into another area and the control lost by leaving the original position, while integrating over time accounts for how long those changes last. The optimization therefore seeks the largest overall gain across the pitch and throughout the evaluated part of the decoy episode, without requiring the teammate to enter the hole.

5.2 Eligible teammates and movement constraints

We consider the runner's outfield teammates, excluding the ball carrier from the final comparison. A permissive reachability check uses each teammate's position and speed at the start of the run to retain those who might reach the hole during the decoy episode. The optimization then applies the full movement constraints and may find a more valuable route elsewhere.

The alternative starts from the teammate's observed position and velocity after the response delay. From there, he can speed up, slow down or turn, while respecting limits on acceleration and speed and remaining within the pitch. These limits prevent instantaneous changes in speed or direction.

Each trajectory must satisfy:

$$\|\mathbf{a}_k\| \leq a_{max},$$
$$\|\mathbf{v}(t)\| \leq v_{max},$$
$$s_{k+1} \leq min\{s_k + g\Delta t_k,\ v_0 - (v_0 - s_k)e^{-\Delta t_k/\lambda}\},$$
$$\mathbf{r}(t) \in P\ .$$

Here, $\mathbf{r}(t)$ and $\mathbf{v}(t)$ are position and velocity, $\mathbf{a}_k$ is the acceleration over an interval of duration Δt_k, s_k is the speed at its start, and P is the pitch. We set $a_{max} = 6\ m/s^2$ by rounding up the 99.5th percentile of accelerations observed in the development matches. The third constraint, also derived from tracking data, reduces how quickly the player can gain speed as he runs faster. The fitted parameters and motion equations are given in Appendix 2.

We repeat the optimization under three maximum-speed limits, representing running (3.98 m/s), high-speed running (5.48 m/s) and sprint (7.0 m/s). A higher limit may allow a teammate to reach valuable areas that were inaccessible at a lower speed. Because teammates start from different positions and velocities, they benefit differently from this allowance, so both the best route and the teammate selected may change.

Comparing these scenarios gives practitioners a way to assess which teammate could best support the attack at each intensity, how much value his movement could add, and whether it would take him into the hole or towards a more valuable area elsewhere.

5.3 Finding and selecting a trajectory

For each eligible teammate and applicable speed limit, we use Sequential Least Squares Programming, as implemented in SciPy [23], to search for a trajectory that maximizes G_i. We try up to five starting routes to reduce dependence on a single starting point, retaining the highest-scoring route among those for which the search converges and a separate check confirms that the movement constraints are satisfied. We then recalculate its gain on the full pitch grid with the exact offside rule. At each speed limit, we select the teammate with the largest recalculated team gain. This is the best admissible solution found by the search, without a guarantee of a global optimum.

5.4 Can the optimization find a more valuable run?

We checked whether the algorithm could find a movement that respects the physical constraints and improves modelled team value compared with the teammate's observed route during the decoy episode. Across 119 episodes, the selected alternatives met both conditions at each speed ceiling.

Median gains were 6.0%, 7.9% and 8.9% at speed ceilings of 3.98, 5.48 and 7.0 m/s, respectively (Table 4), measured relative to observed team value during the same part of each decoy episode.

Speed ceiling (m/s)	Median gain	Interquartile range
3.98	+6.0%	3.7-9.8%

5.48	+7.9%	4.8-14.9%
7.00	+8.9%	5.5-17.8%

Table 4. Relative improvement in modelled attacking team value. Median and interquartile range of episode-level percentage gains relative to observed team value (119 episodes). The interquartile range contains the middle 50% of episodes.

The algorithm identified physically admissible movements with higher modelled team value than the observed routes across all 119 episodes. Higher speed ceilings can make more valuable spaces reachable during the decoy episode, although using that extra speed is more demanding athletically. Positive gains at jogging pace also suggest that useful support can be provided at lower intensity.

5.5 Does the selected run always target the hole?

Our aim is to find a better supporting movement for a teammate of the decoy runner. This means improving on the observed route by reaching more valuable space, given the game context and the player's physical reach during the episode. The hole opened by the decoy is a natural opportunity, but another reachable area may offer greater team value. The two examples below illustrate both possibilities at a speed ceiling of 5.48 m/s: one selected run attacks the hole, while the other favours space elsewhere.

Thuram's run draws Gatti out of the defensive line, leaving a wide gap for a teammate to attack. The optimizer favours a supporting run by Carlos Augusto through this opening and in behind the line (Figure 3a). This alternative increases modelled team value by 8.4%, with 58% of the net gain coming from inside the hole.

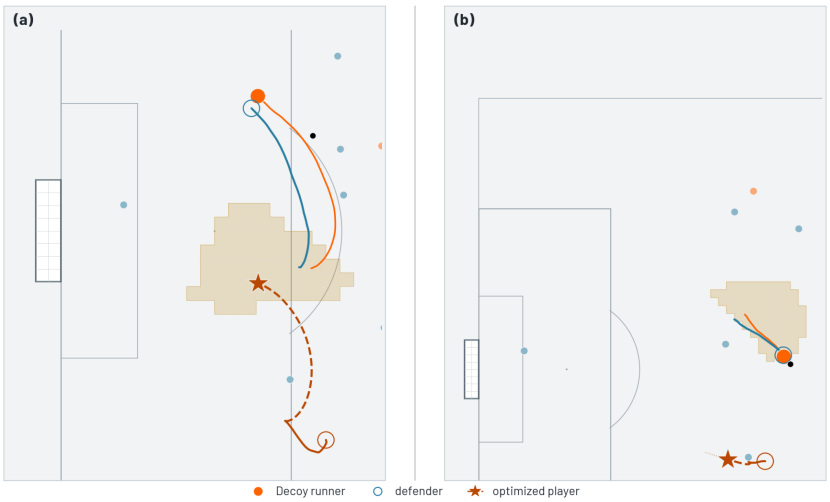


Figure 3. (a) Attacking the hole. Carlos Augusto's selected route ends inside the hole. Final measured snapshot; blue denotes attackers, purple defenders. (b) Preferring value elsewhere. Baturina's selected route favours space outside the hole. The dashed 3 m continuation and its endpoint are illustrative and excluded from the gain.

Diao, shown as a blue dot in the upper part of the Figure 3b, looks a natural choice to exploit the space opened by Nico Paz's run against Hermoso. The optimizer instead favours Baturina's run in behind on the opposite side, whose timing keeps him onside at every evaluated instant. The selected route adds 6.6% to modelled team value, entirely outside the hole.

Together, these results confirm the expected behaviour of the optimizer within the model: finding reachable supporting movements that improve team value over the observed play. The hole provides an opportunity, but does not predetermine the best response. Useful support may instead come from space elsewhere, depending on the game context, the available time and the player's movement limits.

6. Application and Use Cases

We now apply our methods for detecting decoy runs, valuing the openings they create and optimizing teammate movement to study players' contributions to an attack. Beyond quantifying individual off-ball actions, these applications aim to provide a tactical interpretation of how space creation interacts with teammates' positioning and movement. In particular, we distinguish between generating a potentially valuable opening and having teammates in a position to benefit from it.

Using our 2025/26 Serie A dataset, we examine:

- which players create the greatest cumulative potential value through their decoy runs;
- which teammates could add more value through alternative positioning, and whether this improvement arises inside the hole or elsewhere on the pitch.

Together, these analyses provide two complementary perspectives for performance analysis: identifying players whose movements disrupt defensive structures, and examining whether the attacking team's subsequent spatial organisation could make better use of the resulting opportunities.

6.1 Which players create the most valuable openings?

We aggregate the hole value generated by each decoy runner and compare it with the share controlled by eligible off-ball teammates in observed play (Figure 4)

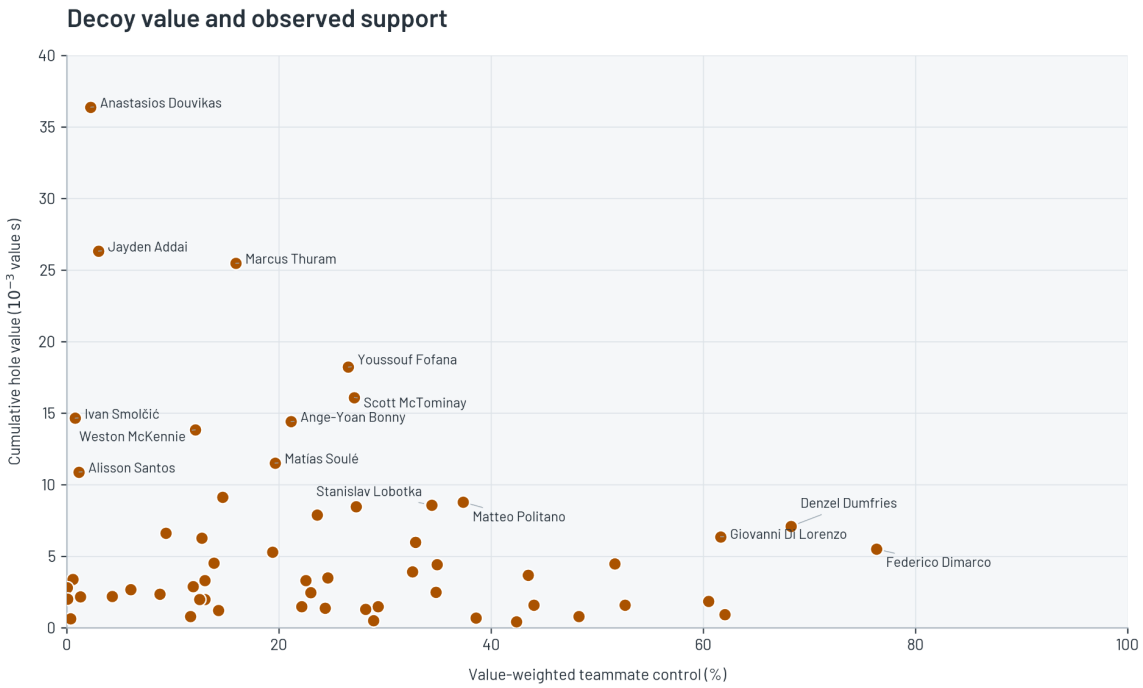


Figure 4. Decoy-runner profiles in the selected Serie A matches. Each point shows a runner's cumulative potential hole value (Y) and the value-weighted share controlled by eligible teammates in observed play (X). Selected names are shown for readability.

Douvikas, Addai and Thuram generate the largest cumulative hole values, but their teammates control relatively little of the space created. Douvikas leads this ranking, yet his teammates control only 2.3% of the associated hole value, compared with 15.9% for Thuram.

Douvikas' profile highlights a potential mismatch between individual space creation and collective attacking support. His runs generate substantial cumulative hole value, but teammates rarely gain control of the resulting openings. For analysts, this raises a concrete tactical question: are supporting players recognising and attacking the spaces his movements create, or are these opportunities being left unexploited?

Dumfries and Dimarco generate smaller cumulative hole values, potentially reflecting their less advanced positioning, but their teammates control a greater share of these openings.

These contrasting profiles suggest that evaluating decoy runs requires looking beyond the runner. The framework can help analysts distinguish players who generate valuable openings from attacking situations in which the team is better positioned to exploit them, providing a starting point for reviewing off-ball coordination.

6.2 Which decoy runner's teammates could improve their positioning?

We examine the potential for improving teammate positioning across the 18 player profiles analysed for Como, distinguishing gains generated inside the hole from those arising elsewhere on the pitch (Figure 5). All alternative movements are evaluated under a maximum speed of 5.48 m/s.

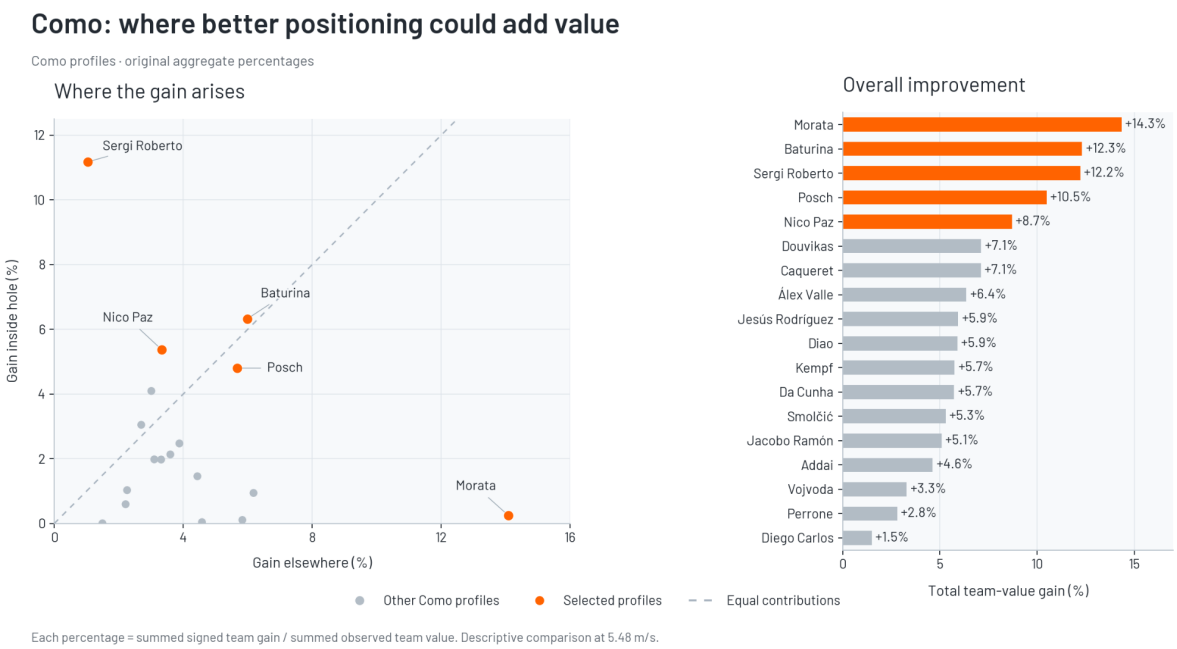


Figure 5. Potential gains from alternative supporting movements by Como players. The left panel shows additional team value inside the hole versus elsewhere; the right shows the combined gain relative to observed team value. Percentages aggregate each player's evaluated episodes; selected profiles are highlighted in orange.

Similar overall improvements can arise from different spatial opportunities. Sergi Roberto and Baturina illustrate this contrast particularly well. Both show potential gains of approximately 12.3% in modelled attacking team value. However, more than 90% of Sergi Roberto's gain originates inside the hole, whereas Baturina's improvement is distributed almost equally between the hole and other areas of the pitch. Posch (+10.5%) also exhibits relatively balanced contributions, slightly favouring space elsewhere, while Nico Paz (+8.7%) shows a stronger contribution from inside the hole (approximately 62%).

Morata represents an extreme case, with the largest relative improvement (+14.3%) and approximately 98% of the additional value generated outside the hole. However, his cumulative hole value is relatively low, and his profile is based on a limited number of episodes. This may partly reflect differences in playing time and exposure to decoy-run situations. His result should therefore be interpreted cautiously, particularly when comparing aggregated player profiles.

Overall, these findings suggest that similar potential improvements may require different supporting movements, either exploiting the hole directly or targeting more valuable areas elsewhere. Interestingly, some improvements are associated with players nominally classified as defenders, such as Posch. This is tactically plausible given Como's fluid positional rotations, which can bring traditionally defensive players into advanced attacking positions.

For performance analysts, these profiles offer a starting point for examining how supporting movements could be better coordinated with decoy runs made on the last defender line.

7. Related Work

Tracking data have enabled the development of methods to quantify the value of off-ball positioning and movement in soccer. Fernández and Bornn [6] distinguished between *space occupation*, where players gain control of valuable areas, and *space generation*, where their movements attract defenders away from teammates. Spearman [20] introduced a probabilistic framework combining pitch control, ball movement and scoring probability, while Martens et al. [7] developed data-driven measures of space control and generation. However, attributing spatial changes to a specific attacking run requires distinguishing the defender's response from broader defensive movements.

Related research has investigated attacker–defender interactions and defensive coordination. Herold et al. [12] evaluated off-ball actions through changes in defensive pressure, while Calero-Sanz et al. [15] examined marking relationships using spatial proximity and directional alignment. Although these approaches characterise defensive interactions, they do not explicitly establish whether a defender following an attacker departs from the coordinated movement of their defensive unit, a key distinction for identifying decoy runs.

Counterfactual modelling provides a means of evaluating individual movements against alternative behaviours. Teranishi et al. [8] assessed off-ball attacking contributions by comparing observed and predicted trajectories, while Groom et al. [17] applied defensive counterfactuals to evaluate positioning during corner kicks. Building on this principle, we estimate how an engaged defender would have moved without following the runner and quantify the value of the resulting spatial displacement.

Finally, previous studies have explored how alternative movements can improve attacking opportunities. Stein et al. [21] introduced region-based what-if analysis, while Peralta Alguacil et al. [22] optimised attacking movements considering passing probability and spatial control. Extending these approaches, we optimize teammates' movements in response to identified decoy runs, evaluating whether greater attacking value can be generated by exploiting the space opened by the defender or by moving into other valuable areas.

8. Conclusions

We introduced a novel framework to analyse decoy runs in soccer. First, we proposed a rule-based algorithm to detect decoy runs from tracking data. Specifically, we identify an attacking run as a decoy run if a defender closely follows the attacker and leaves the defensive structure in doing so. Second, we introduced a counterfactual-based approach to quantify the potential value of the space left by the engaged defender, identifying areas they would otherwise have controlled and estimating their positional threat. Third, we proposed an algorithm to optimize teammates' movements to better exploit the opportunities created by a decoy run. Across 119 episodes, the

optimised movements increased modelled attacking value by a median of 7.9% under a 5.48 m/s speed limit, with gains arising both inside and outside the space opened by the decoy run. Overall, this work provides tools to detect decoy runs, quantify their potential spatial value, and identify alternative teammate movements that could better exploit the opportunities they create.

Our framework has several limitations that should be addressed in future work. The detection procedure relies on handcrafted thresholds and lacks ground-truth annotations for direct validation. Our indirect evaluation showed the strongest association between engagement and exceptional displacement in slowly moving defensive lines, leading us to restrict the subsequent analysis to these conditions. Detecting decoy runs in more dynamic defensive situations remains a challenge for future research. Moreover, the components employed by our decoy run valuation metric could be improved by using more powerful models. For example, future work could use trajectory prediction models to better simulate the counterfactual defender trajectory, and more sophisticated pass selection and success models (e.g. [24]) to estimate possible ball destinations, alongside improved spatial-value models to assess the threat of the space opened by the decoy run. The optimization only changes one teammate's trajectory at a time while keeping all other players fixed, and therefore does not capture how defenders might respond to an alternative attacking movement. Finally, our dataset only covers 30 matches between top-six Serie A teams, which limits conclusions at the player and team level.

Acknowledgements

This research is supported by the Research Foundation Flanders (FWO, LC: 11I8125N).

References

[1] F. Esposito, M. Bertollo, D. Pompa, M. Angonese, and M. Beato, 'Off-the-ball behaviors in elite football using spatiotemporal tracking data: a scoping review', *Sci. Med. Footb.*, pp. 1–25, Aug. 2026, doi: 10.1080/24733938.2026.2713521.

[2] Y. Wu and T. Swartz, 'Evaluation of off-the-ball actions in soccer', *Stat. Appl. - Ital. J. Appl. Stat.*, no. 2, May 2023, doi: 10.26398/IJAS.0035-008.

[3] P. Lucey, A. Bialkowski, M. Monfort, P. Carr, and I. Matthews, 'Quality vs Quantity: Improved Shot Prediction in Soccer using Strategic Features from Spatiotemporal Data', in Proceedings of the MIT Sloan Sports Analytics Conference, 2015.

[4] T. Decroos, L. Bransen, J. Van Haaren, and J. Davis, 'Actions Speak Louder than Goals: Valuing Player Actions in Soccer', in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, Anchorage AK USA: ACM, Jul. 2019, pp. 1851–1861. doi: 10.1145/3292500.3330758.

[5] L. Cascioli, A. Wang, L. Stradiotti, M. Van Roy, P. Robberechts, M. Wouters, A. Jaspers, and J. Davis, 'Quantifying Off-Ball Defensive Impact through Cover Shadows', Hudl Performance Insights Conference, 2025.

[6] J. Fernández and L. Bornn, 'Wide Open Spaces: A statistical technique for measuring space creation in professional soccer', in Proceedings of the MIT Sloan Sports Analytics Conference, 2018.

[7] F. Martens, U. Dick, and U. Brefeld, 'Space and Control in Soccer', *Front. Sports Act. Living*, vol. 3, p. 676179, Jul. 2021, doi: 10.3389/fspor.2021.676179.
[8] M. Teranishi, K. Tsutsui, K. Takeda, and K. Fujii, 'Evaluation of creating scoring opportunities for teammates in soccer via trajectory prediction', Jul. 27, 2022, *arXiv*: arXiv:2206.01899. doi: 10.48550/arXiv.2206.01899.
[9] P. S. Bradley, W. Sheldon, B. Wooster, P. Olsen, P. Boanas, and P. Krustrup, 'High-intensity running in English FA Premier League soccer matches', *J. Sports Sci.*, vol. 27, no. 2, pp. 159–168, Jan. 2009, doi: 10.1080/02640410802512775.
[10] A. Gualtieri, E. Rampinini, A. Dello Iacono, and M. Beato, 'High-speed running and sprinting in professional adult soccer: Current thresholds definition, match demands and training strategies. A systematic review', *Front. Sports Act. Living*, vol. 5, Feb. 2023, doi: 10.3389/fspor.2023.1116293.
[11] 'Fifa Coaching Manual | PDF | Association Football | Sports'. Accessed: Nov. 01, 2026. [Online]. Available: https://it.scribd.com/document/390665265/Fifa-Coaching-Manual
[12] M. Herold *et al.*, 'Off-ball behavior in association football: A data-driven model to measure changes in individual defensive pressure', *J. Sports Sci.*, vol. 40, no. 12, pp. 1412–1425, Jun. 2022, doi: 10.1080/02640414.2022.2081405.
[13] 'EFI Explanation document (PDF) - FIFA Training Centre'. Accessed: Nov. 01, 2026. [Online]. Available: https://www.fifatrainingcentre.com/en/game/tournaments/fifa-womens-world-cup/2023/efi-metrics/explanation-document.php
[14] T. Narizuka and Y. Yamazaki, 'Statistical properties for directional alignment and chasing of players in football games', EPL (Europhysics Letters), vol. 116, no. 6, p. 68001, 2016, doi: 10.1209/0295-5075/116/68001.
[15] J. Calero-Sanz, R. L. Del Campo, R. Resta, and J. M. Buldú, 'Beyond marking networks in soccer: Coordination, similarity and entropy', *Chaos Solitons Fractals*, vol. 211, p. 118819, Oct. 2026, doi: 10.1016/j.chaos.2026.118819.
[16] R. Duarte, D. Araújo, V. Correia, K. Davids, P. Marques, and M. J. Richardson, 'Competing together: Assessing the dynamics of *team–team* and *player–team* synchrony in professional association football', *Hum. Mov. Sci.*, vol. 32, no. 4, pp. 555–566, Aug. 2013, doi: 10.1016/j.humov.2013.01.011.
[17] S. Groom, F. Belo, A. Rice, L. Anderson, and S. Wang, 'A machine learning framework for off-ball defensive evaluation applied to corner kicks', *Npj Artif. Intell.*, Jul. 2026, doi: 10.1038/s44387-026-00138-7.
[18] W. Spearman, A. Basye, G. Dick, R. Hotovy, and P. Pop, 'Physics-Based Modeling of Pass Probabilities in Soccer', 2017.
[19] L. Shaw, LaurieOnTracking, GitHub repository, 2020. [Online]. Available: https://github.com/Friends-of-Tracking-Data-FoTD/LaurieOnTracking
[20] W. Spearman, 'Beyond Expected Goals', in Proceedings of the MIT Sloan Sports Analytics Conference, 2018. [Online]. Available: https://static.hudl.com/craft/downloads/SSAC2018_Beyond_Expected_Goals.pdf
[21] M. Stein *et al.*, 'Where to go: Computational and visual what-if analyses in soccer', *J. Sports Sci.*, vol. 37, no. 24, pp. 2774–2782, Dec. 2019, doi: 10.1080/02640414.2019.1652541.
[22] F. Peralta Alguacil, J. Fernández, P. Piñones Arce, and D. Sumpter, 'Seeing in to the future: using self-propelled particle models to aid player decision-making in soccer', in Proceedings of the MIT Sloan Sports Analytics Conference, 2020.
[23] P. Virtanen *et al.*, 'SciPy 1.0: fundamental algorithms for scientific computing in Python', *Nat. Methods*, vol. 17, no. 3, pp. 261–272, Mar. 2020, doi: 10.1038/s41592-019-0686-2.
[24] P. Robberechts, M. Van Roy, and J. Davis, 'un-xPass: Measuring Soccer Player's Creativity', in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and*

Data Mining, Long Beach CA USA: ACM, Aug. 2023, pp. 4768–4777. doi: 10.1145/3580305.3599924.

Appendix

Appendix 1

An earlier analysis on the 17-match development set compared minimum engagement durations of 0.6, 0.8 and 1.0 s, yielding 150, 125 and 103 detected events, respectively. We use 0.8 s as a compromise between excluding brief co-movement and retaining sustained but relatively short engagements. A shorter threshold admits more transient interactions, whereas a longer threshold discards additional following episodes. The selected value is therefore an operational persistence requirement, rather than an empirically identified optimum.

Appendix 2

Alternative trajectories use constant vector acceleration within each interval of duration Δt_k:

$$\mathbf{r}_{k+1} = \mathbf{r}_k + \mathbf{v}_k \Delta t_k + (1/2)\mathbf{a}_k \Delta t^2{}_k$$

$$\mathbf{v}_{k+1} = \mathbf{v}_k + \mathbf{a}_k \Delta t_k$$

The initial position and velocity are the teammate's observed state at the response frame. Speed and pitch-boundary constraints are checked throughout each interval, including intermediate position extrema.

Movement limits were calibrated on the 20 non-held-out matches. The 99.5th percentile of observed acceleration magnitudes was $5.63\ m/s^2$, rounded upwards to $a_{max} = 6\ m/s^2$.

To constrain speed gains, we calculated the 99.5th percentile of observed speed increases over 0.2 s within initial-speed bins of 0.25 m/s, retaining bins with at least 1,000 observations. Fitting the speed-gain envelope by least squares yielded $g = 4.10\ m/s^2$, $v_0 = 12.23\ m/s$, and $\lambda = 1.40\ s$. This profile limits speeding up more strongly at higher speeds; braking and turning remain subject to the vector acceleration bound. The fitted v_0 shapes the envelope and is not the permitted running speed, which remains bounded by the scenario-specific v_{max}.