Contextual Expected Threat using Spatial Event Data

Paper Track

Greg Everett, Ryan Beal, Tim Matthews, Tim Norman, Gopal Ramchurn

1. Introduction

Measuring goal likelihood of possession sequences in football is a key analytical challenge that helps evaluate attacking threat beyond just goals and shots. Numerous studies have addressed this challenge using a number of probabilistic models to estimate goal likelihood from game situations [1, 2, 3]. In particular, the Expected Threat (xT) metric measures the probability of a goal by modelling a play sequence as a discrete-time Markov process which considers shot and move transition probabilities [1]. However, this Expected Threat model does not explicitly consider the player locations of teammates and opposition players, meaning that the attacking threat at a given state of play is considered the same irrespective of the defensive structure and attacking positioning.

Whilst the Expected Threat metric has proven to be a very useful metric for valuing attacking danger based on ball location, investigating how the positions of attackers and defenders influence the threat value will help coaches identify and better understand the strengths and weaknesses of a team's positional setup. Player positioning is a key part of team training, with off-ball movement playing a key role in football play outcomes. Whilst these play key parts in football training and game outcome, it is difficult to objectively quantify the effectiveness of a defensive unit or attacking setup. We therefore propose a spatial xT model which considers player positioning when modelling goal likelihood.

Specifically, we focus on introducing spatial context to the transition model that is at the core of Expected Threat. The emergence of widespread tracking and contextual event data (e.g. StatsBomb 360) in recent years has led to a surge of research into various spatial metrics. This research includes a pitch control model [4, 5] to measure ball control probabilities and an interception model [6] to quantify pass risk. In this paper, we adapt these pre-existing spatial metrics for use as input channels for a convolutional neural network (CNN) that uses them to predict refined ball transition probabilities. We validate this approach by comparing the predictive accuracy of our spatial xT to the original Expected Threat transition model, finding that our model outperforms this approach in calculating ball transition probabilities and goal likelihood.

Using this spatial xT model, we propose a novel way of valuing attacking and defensive units in football by computing a 'Threat Above Expected' (TAx) metric which measures the change in goal probability resulting specifically from spatial context by comparing spatial xT to the original xT value. Using this metric, we evaluate team defences and attacks over a set of games from the English Premier League 2021/22 season. We also demonstrate a concept for a positional defensive optimiser that can offer coachable insights by identifying improvements in defensive setups during post-match analysis to minimise the spatial xT value.

In summary, this paper presents the following contributions to research in football:

- A novel modification of the Expected Threat model, which uses a convolutional neural network to build a new spatial xT model with a 19.2% decrease in log loss score compared to the original model when predicting ball transitions.
- A combination of spatial maps designed to capture contextual information specifically from StatsBomb 360 data. Furthermore, we perform an ablation study to validate the impact of each spatial map.
- Based on our modified threat model, we propose a TAx metric to measure the success of attacking and defensive positioning in football.
- We apply our framework to ten teams in the English Premier League 2021/22 season, demonstrating several team evaluations and post-match analysis concepts that could be used by clubs using this model.
- We demonstrate a concept for a positional defensive optimiser that finds situations of high threat and calculates optimal defensive setups, presenting as a post-match analysis tool to improve team defences.

The rest of this paper is structured as follows: in Section 2, we give a background of related work in the field. Section 3 introduces the Expected Threat formulation and defines the problem space. Section 4 outlines the spatial maps we will use to add context to our Spatial xT model. Section 5 explains the CNN architecture used to predict ball transition probabilities. In Section 6 we validate our model by performing an ablation study on the spatial input channels and comparing its predictive accuracy to the original Expected Threat model. Section 7 is a case study of model applications. Finally, Section 8 concludes the paper.

2. Background

The idea of modelling play sequences in football as a discrete-time Markov process was first introduced in [7]. This idea was extended in [1], which introduced the Expected Threat (xT) metric, measuring goal likelihood from a play sequence. Other models also value actions and game states, such as the Valuing Actions by Estimating Probabilities (VAEP) framework introduced in [2]. This instead measures the probability of both scoring and conceding a goal at a given game state. A critical comparison is given between VAEP and xT in [8]. Van Roy et al. [9] use the xT model to analyse football decision-making, finding that teams could score more goals if they took more shots outside the penalty box. Van Roy et al. [10] build an extended xT model using StatsBomb 360 data which investigates the impact of teammate reachability on the success rate of a team reaching the attacking third during build-up play. We also modify the xT model but instead focus on refining transition probabilities using a number of spatial maps given to a CNN model. Furthermore, we specifically analyse both attacking and defensive positioning in the final attacking third and give an updated spatial xT value.

Several approaches have been considered to value space and positioning in football. Spearman [4] uses a probabilistic physics-based model to objectively value off-ball scoring opportunities in football. Defensive positioning is focused on in [11], where the impact of defences on opposition pass availability is measured using a graph convolutional network with tracking data input. Fernandez et al. [3] calculate goal likelihood probability using the context of player positions and velocity with tracking data by using a series of deep neural networks to predict pitch surfaces. We similarly model goal likelihood using player positions, but instead use StatsBomb's contextual event data rather than tracking data. Furthermore, we build an adapted spatial model using the structure of the xT model, which doesn't consider player locations, so that the effect of player positioning on goal likelihood can be compared effectively against an average. Fernandez et al. [3] also calculate the likelihood of either team scoring at a given game instance, whereas the xT metric focuses only on the threat of the attacking team.

Spatiotemporal tracking data has furthered the development of physics-based models in football. Spearman et al. [5] estimate pass success probabilities by modelling ball and player trajectories. Link et al. [12] focus more on estimating imminent goal threats using factors such as player ball control and opponent pressure. These models use tracking data to model trajectories, however, StatsBomb 360 data contains only player locations. Burriel et al. [6] consider this by using StatsBomb 360 data to model pass risk and reward using interception probabilities which assume initial player velocities of 0. Our updated transition model uses the formulation introduced in previous ball control [5] and interception [6] models to model transition probabilities with StatsBomb 360 data. However, adjustments to these models are made to match the structure of the Expected Threat model.

There have been several studies that use CNNs to spatially interpret game situations in football. Fernandez et al. [13] use a CNN architecture to estimate pass probability surfaces in football. This example uses tracking data input, whereas in [14], player locations are used from StatsBomb 360 data to build a CNN which probabilistically classifies penetrative passes into the convex hull of a defence. Our model also spatially interprets game situations using a CNN but instead estimates transition probabilities within the Expected Threat model framework.

3. Standard Expected Threat

We aim to calculate an improved Expected Threat framework that considers the locations of attackers and defenders and identifies the most likely play transitions during a possession sequence. Expected Threat [1] is a probabilistic framework that iteratively models on-field actions to calculate the likelihood of a goal from a possession sequence. This is made computationally tractable by dividing the pitch into an $(M \times N)$ spatial grid of states, as well as adopting the Markov assumption such that state transitions are a function only of the current state. We will adopt the same assumptions throughout this paper. The Expected Threat equation is:

$$\mathrm{xT}_{x,y} = \left(s_{x,y} \times g_{x,y}\right) + (m_{x,y} \times \sum_{n=1}^{N} \sum_{m=1}^{M} T_{x,y \to n,m} \, \mathrm{xT}_{n,m})$$

Where $\mathrm{xT}_{x,y}$ is the likelihood of a goal occurring during the possession sequence starting at state (x, y), $s_{x,y}$ is the probability of a shot being taken, $g_{x,y}$ is the probability of a goal being scored given that a shot is taken, $m_{x,y}$ is the probability of the ball being moved (e.g. passed to a teammate), and $T_{x,y \to n,m}$ is a transition matrix which, for the ball-carriers pitch state, details the probability of successfully transitioning to each state on the pitch given that a move action is taken.

Evidently, the equation above has a recursive structure. For this paper, we iterate through the xT equation five times meaning that we model the likelihood of a goal in the next five actions for the attacking team. We also split the pitch into a 16x12 grid, meaning that the total number of pitch states is 192, as visualised in Figure 1.

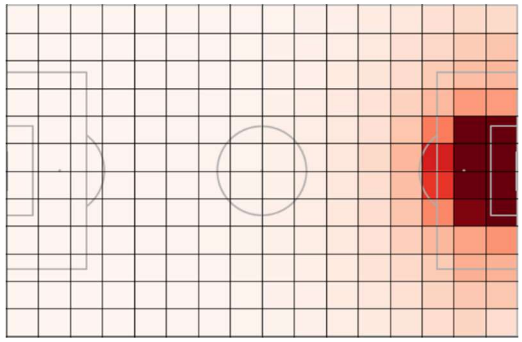

Figure 1 – Pitch grid showing all possible game states, with darker cells showing areas of higher xT.

In the Expected Threat framework, all probabilities, s, g, m and transition probability matrix T, are estimated from observed frequencies in past event data. The probability of possession loss is also factored into the model through failed shots and failed move transitions. As this learning process only uses event data, transition probabilities only consider ball location, and do not explicitly consider the location of attackers and defenders.

In order to introduce spatial context to the xT framework, we choose to focus on the transition matrix T as we believe this is the parameter most heavily affected by player location, and this also allows us to evaluate model accuracy more rigorously by treating the transition matrix as a distinct probabilistic model, as goals are substantially rarer than transitions in football. We will denote our contextual transition matrix as T^*, and in turn use it to define a spatially-aware xT model capable of more accurate threat predictions One challenge for such a model is that Expected Threat is calculated recursively across multiple iterations, and hence forecasts of future player locations are required for a complete spatially-aware xT. This is beyond the scope of this paper, and so we use T^* only in the first xT iteration computed from known player locations from StatsBomb 360 data, reverting to the conventional xT formulation for subsequent iterations. Refinement of shot, goal and move probabilities, s, g and m, are left for future work, and are estimated in the same way as for the original model for our new model. In the following sections, we discuss different spatial contexts and how they can be used to learn the improved transition matrix T^*.

4. Contextual Expected Threat Model

Using Statsbomb 360 data, we can now use added game context to build a spatially-aware transition model T^* which factors in defensive positioning and attacking space to learn a distribution of probabilities over state transitions or possible possession loss. To do this, we produce several spatial maps with dimensions equal to that of our $(M \times N)$ pitch grid (see Figure 1). Each of these spatial maps is somewhat heuristical in design, intended to provide a signal to the CNN model we introduce later in the paper by capturing different notions of what might motivate the ball carrier to make a particular pass. The spatial maps we will focus on in this paper are as follows:

1. **Teammate Distribution –** The distribution of attacking teammates in each zone. This considers the distance of the zone from the ball carrier and the closeness of attacking teammates to this zone.

2. **Interception Probability –** The probability of the ball being intercepted by the defensive team on the way to each zone.
3. **Control Probability –** The probability that the attacking team will successfully control the ball given that it arrives at each zone.
4. **Original xT Transitions -** A spatial map of $T_{i,j\rightarrow x,y}$ for each grid cell (X, Y) where (I, J) is the location of the ball carrier.
5. **Original xT outputs –** A spatial map of $\mathrm{xT_{x,y}}$ calculated for each grid cell (X, Y). Hence this map is independent of the current ball location.

These spatial maps give information on the likelihood of pass completion for each state. Teammate distribution is used to show whether there are teammates in or near that state to be passed to, whilst the control and intercept probabilities measure the likelihood of a pass to that state being successful. Finally, the xT spatial maps give information about the threat of a potential transition and allow T^* to possibly be refined by knowledge already encapsulated in T and xT. We validate the predictive value of each spatial map in an ablation study in Section 6.1. It is worth noting that the first three spatial maps only calculate values within the visible area given in StatsBomb 360 data, as player locations outside of this area are unknown. Therefore, for the first three spatial maps, pitch states outside the visible area are given values of 0. In the following subsections, we will go into further detail on how the first three of these spatial input channels are modelled.

4.1 Teammate Distribution

The teammate distribution model is a 2D Gaussian mixture model with components specified by ball location and teammate locations. These represent areas of the pitch close to the ball and teammates on the pitch to model likely ball transitions. For each teammate and the ball, we therefore generate distributions within the mixture model formulated as:

Ball Distribution: $N(\mu_B, \sigma_B)$
Teammate Distributions: $N(\mu_m, \sigma_m)$ for each m in M

Where μ_B is the ball location, σ_B is the standard deviation of ball motion, μ_m is the location of teammate $m \epsilon M$ where M is the set of all visible teammates, and σ_m is the standard deviation of teammate motion. We define σ_B as 23.9 metres, which is a maximum a posteriori estimate from [4], which uses a normal distribution around ball location as the basis for a physics-based pass probability model. For the teammate distributions, we set σ_m as the square root of the Euclidean distance between the ball and the teammate m, representing that the distribution around nearby teammates is smaller as a pass to them is likely to be more accurate. We then discretise the mixture model for the spatial grid and normalise to unity. An example heatmap is given in Figure 2.

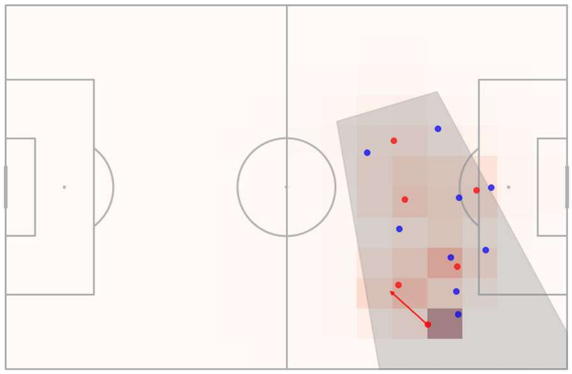

Figure 2 – Pitch grid showing a heatmap of the teammate distribution. The darkest cells are those with the highest discretised mixture probability.

4.2 Interception Probability

The interception probability models the likelihood of a move action from a start zone s to an end zone s' being intercepted. We calculate this for every possible end zone to generate a spatial grid of intercept probabilities. The following model adapts and extends previous work done by Burriel et al. [6], which models the likelihood of a successful defensive interception as:

$$I(s', s) = 1 - \prod_{d \in D} (1 - P^{*}_{intercept}(d, s', s))$$

Where $I(s', s)$ is the interception probability between s and s', D is the set of visible defensive players, and $P^{*}_{intercept}(d, s', s)$ is our interception probability for defender d, which is calculated by extending the interception model in [6] with two added modifications. We define this as:

$$P^{*}_{intercept}(d, s', s) = P_{intercept}(d, s', s) \times \left(1 - P_{contested}(d, s', s)\right) \times (1 - P_{lobbed}(d, s', s))$$

Where $P_{intercept}(d, s', s)$ is the intercept model in [6] using extracted ball and player velocity parameters from [5], with the assumption that initial player velocity is 0m/s. $P_{contested}(d, s', s)$ is the likelihood of an attacker at the destination s' contesting the interception by defender d, and $P_{lobbed}(d, s', s)$ is the probability that the ball is successfully lobbed over the defender d. Essentially we modify the base intercept model in [6] by making successful interceptions conditional on i) the ball being below the head height of the defender, and ii) being uncontested by an attacker at the destination.

Specifically, $P_{contested}(d, s', s)$ models the likelihood of an attacker at s' reaching the trajectory point where the defender d is looking to intercept before the ball arrives there, which is calculated using the same intercept model as in [6]. To calculate the successful lob probability, $P_{lobbed}(d, s', s)$, we simulate ball trajectories for every integer pass height angle between $\theta_{start} = 20°$ and $\theta_{end} = 60°$ and find the proportion of angles where the pass would successfully go over the defender's head(defined as 2 metres). This is denoted as:

$$P_{lobbed}(d, s', s) = \frac{\sum_{\theta=20}^{60} \mathrm{H}(h_c(\theta) - 2)}{\theta_{end} - \theta_{start} + 1}$$

Where H is the Heaviside step function, c is the point on the ground beneath the trajectory where the defender will try and intercept (e.g. the closest point on the trajectory to the defender), and h_c is the height of the ball at point c given pass height angle θ. We calculate h_c using standard equations for projectile motion:

$$h_c(\theta) = \Delta_{s,c} \tan(\theta) - \frac{{\Delta_{s,c}}^2 g}{2{b_v}^2 \cos(\theta)^2}$$

Where $\Delta_{s,c}$ is the distance between the interception point and the ball's starting location, b_v is the initial ball velocity and g is the acceleration due to gravity (9.8m/s^2). Ball velocity for these calculations is taken to be: $b_v = \sqrt{g \times \Delta_{s,s'}}$ where $\Delta_{s,s'}$ is the distance between the ball location and the intended end location. This represents the velocity needed for the ball to be kicked at 45-degree trajectory and land exactly at the destination.

We can use the intercept equation to form a spatial grid of interception probabilities across the pitch. An example heatmap is given in Figure 3.

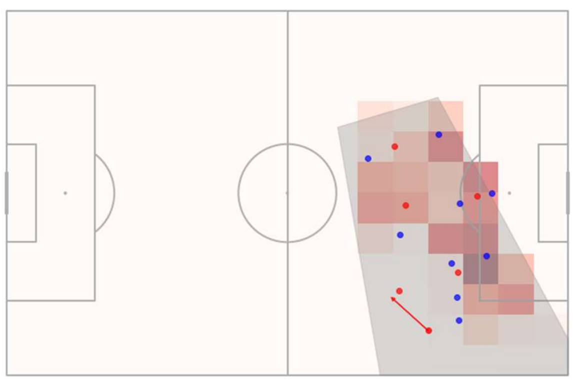

Figure 3 – Pitch grid showing a heatmap of the intercept probability relative to the current ball state. The highest value states are those blocked by the defending opposition players.

4.3 Control Probability

The probability of the attacking team successfully controlling the ball given that it arrives at a certain state is now considered. To compute this, we use the pitch control models developed in [5, 4]. These models consider a player's likelihood of controlling the ball as a Poisson point process which increases in likelihood when a player has more time in proximity to the ball without opposition interference. Parameters used for this model are detailed in [4], with the initial velocity instead set to 0m/s as these data points are not available. We discretise this as a spatial pitch control grid by finding control probabilities for the midpoint of each grid zone. An example is shown in Figure 4.

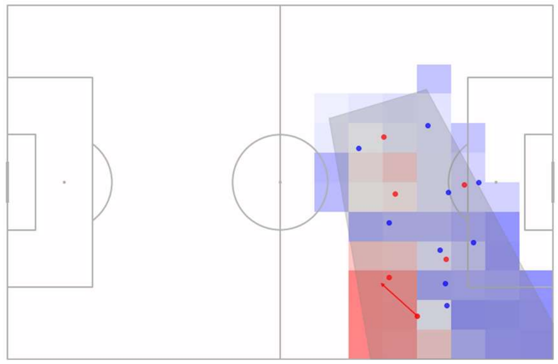

Figure 4 – Pitch control grid, red zones highlighting high probability of the attacking team retaining the ball, blue zones highlighting low probability of the attacking team retaining the ball.

5. Contextual Model Architecture

After calculating the spatial input maps, we wish to use them to compute more accurate ball transition probabilities. This will produce a new contextualised transition model T^* within Expected Threat, and consequently a new xT value. We use a convolutional neural network (CNN) architecture, as these are an established architecture for learning from spatial data in a flexible way. At every event, we pass in a $M \times N \times C$ matrix, where $M \times N$ is the 2D pitch grid and C is the number of input channels. The CNN architecture we use has similarities to the architecture used in [13], which builds pass success probability surfaces in football. For example, the use of symmetric padding is used in convolutional layers to maintain pitch dimensionality, and a max pooling layer is used to learn wider spatial context. However, there are also many differences between our models, and we therefore describe our model architecture below. Figure 5 details the structure of the CNN, with reasoning explained further for each section of the architecture.

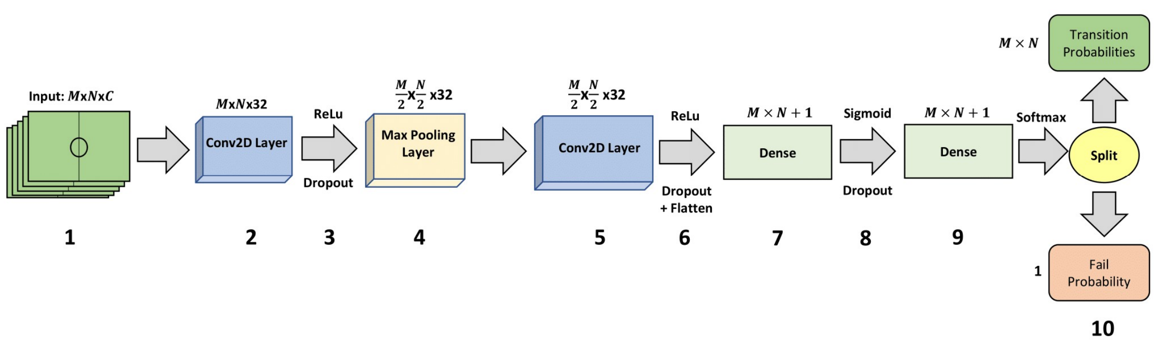

Figure 5 – CNN Model Architecture for $M \times N \times C$ input

1 – Input

We use an input of $\mathrm{M} \times \mathrm{N} \times \mathrm{C}$. Spatial grids are inputted as input channels, and we therefore include our previously defined spatial grids in the final model architecture.

2&3 – Conv2D Layer with ReLu activation

The inputs are passed into a 2D convolutional layer with 32 filters, kernel size of 2×2 and stride of 1. Symmetric padding is used to maintain pitch dimensionality when the kernel is applied. A ReLu activation function is applied, and a dropout layer with a dropout rate of 0.4 is used to limit overfitting.

4 - Max Pooling
A 2D max pooling layer with a pool size of 2×2 is used to represent the field at half its length and width. This reduces overfitting and computation costs, whilst helping capture wider spatial features.

5&6 – Conv2D Layer with ReLu activation
This layer uses the same parameters, activation function and dropout as layer 2&3. Reduced dimensionality from the pooling layer will help identify wider spatial features. The output is flattened to be passed into a dense layer.

7&8 – Dense Layer with Sigmoid activation
This dense layer uses outputs from the convolutional layer to generate probabilities with size $M \times N + 1$, which is equivalent to the number of possible end states, including a possession loss. A sigmoid activation function and dropout layer is used with a dropout rate of 0.2.

9&10 – Dense Layer with Splitting
We pass our model outputs into a final output layer of $M \times N + 1$ neurons. To appropriately receive a distribution of probabilities for the transition outcome, we use a Softmax activation function. After this, we split the output into the $M \times N$ output states corresponding to pitch zones, and the final output relating to the probability of a possession loss.

6. Model Evaluation

Now that we have set up a spatial xT model which uses a CNN architecture to predict transition matrix T^*, we will evaluate the accuracy of its predictions. Unless otherwise stated, the CNN is trained using all of the spatial maps outlined in Section 3, where the original xT maps are computed using training data. The way we evaluate our model is three-fold. Firstly, we will perform an ablation study on our model input where each spatial map given to the CNN will be independently removed to evaluate its impact on model performance. Secondly, we evaluate the predictive accuracy of the model for state transitions compared to the original xT model. Finally, we evaluate the accuracy of the final xT value in predicting goal likelihood using the spatial xT model in comparison to the original xT model.

When training and evaluating our model, we use StatsBomb 360 data from games for 10 teams in the English Premier League seasons 2020/21 and 2021/22. We only use events in the final attacking third to build a model focussed on transitions in dangerous attacks. Based on this, we use 24259 events and 24215 events in the 2020/21 and 2021/22 seasons respectively. The way this data is split is varied for each evaluation and is therefore explained in each subsection. It is noted that the model is trained over 100 epochs with a batch size of 32. Early stopping is also implemented which monitors the validation loss with a patience of 2 to limit overfitting. As the model is probabilistically classifying outcome states of events, we use categorical cross-entropy loss as the loss function. Finally, adaptive moment estimation (Adam) [15] is used as the optimization algorithm during model training with a learning rate of 0.001.

6.1 Evaluation of Model Channels

We first perform an ablation study by evaluating each spatial map used by the model to review its influence on predictive performance. To do this, we perform 10-fold cross-validation on 2020/21 season move events with an 80:10:10 split for training, validation and testing at each fold. We also include the original xT model transition probabilities as a comparative baseline. This evaluates predictive accuracy for move transition probabilities outputted by the contextual transition model, T^*, meaning that there are 193 classification outcomes for our 16x12 grid with one outcome labelled as a failed action. For each event, the target variable will be a one-hot encoding with a 1 in the end state for the event i.e. the end zone of the move action. Figure 6 details the predictive accuracy of each model.

Figure 6: Comparisons of the predictive accuracy of the transition models.

Model	Log Likelihood
Empirical transition model, T	-5.065
T^* without Teammate Distribution channel	-4.176
T^* without control channel	-4.199
T^* without intercept channel	-4.176
T^* with intercept model from [6]	-4.167
T^* without xT channels[1]	-4.192
Contextual transition model, T^*	-4.159

It is shown here that the contextual transition model, T^*, with all spatial channels passed to the CNN has the highest predictive accuracy, justifying the use of each channel. As shown above, we also show that the added parameters to the interception model in [6] have led to an improved prediction accuracy of state transitions.

6.2 Evaluating the Transition Model

We now evaluate the predictive accuracy of our contextual transition matrix T^* over a season of testing data. Similarly to the evaluation in Section 6.1, the predictive accuracy of the transition probabilities is evaluated when considering the actual end zone for each event. In this evaluation, we perform a 90:10 split on the 2020/21 season move events into training and validation, and test the model on move events from the whole 2021/22 season. Now that each spatial map in our model is justified, we compare this model directly to the empirical transition model in xT to see if the contextual transition model, T^*, has improved predictive accuracy for ball transition probabilities. Results can be seen in Figure 7.

[1] Refers to both the spatial xT grid and the original xT transition grid.

Figure 7: Comparisons of the predictive accuracy of the transition models.

Model	Log Likelihood
Empirical transition matrix, T	-5.137
Contextual transition matrix, T^*	-4.151

It is shown here that the contextual transition model has improved the predictive performance of ball transitions compared to the empirical transition matrix T.

6.3 Evaluating Expected Threat Values

Finally, we evaluate the predictive accuracy of the model for the spatial xT value. For this evaluation, the model is trained and tested over the same datasets as those in Section 7.2. However, shot events are also included in the test dataset, leading to an additional 2232 test events. We compare our model accuracy to the original xT model. To calculate predictive accuracy, as the xT value models goal likelihood in the next 5 actions, we use the xT value as a probabilistic prediction of a goal within 5 actions and use a binary target variable which is set to 1 for events where there is a goal within the next 5 actions in the data. Results can be seen in Figure 8.

Figure 8: Comparisons of the predictive accuracy for the final xT value.

Model	Log Likelihood
Original xT Model	-0.08796
Spatial xT	-0.08759

As shown, the new model shows slight improvement over the original xT model. As our model currently modifies the transition matrix, an interesting step for future work would be to also model the shot probability model within xT and evaluate its influence on these results.

7. Model Application: TAx Case Study

In this section, we present an example of an application of our spatial xT model. We recall that the parameters of the original xT model are estimated by counting the frequencies and grid location of events from past data, without incorporating any further spatial context. In some sense then, this xT definition represents the average threat, independent of further spatial context. Because our spatial xT model explicitly includes this context, we can compare the two xT values to better understand how the spatial context has increased or reduced the threat of a matchplay situation versus the average expected for the same pitch location. For example, we could understand whether the observed defensive structure has reduced the threat of a goal versus the xT baseline - or even versus other hypothetical defensive structures that could have been employed.

As the spatial xT model modifies the Expected Threat model by adding player positions, we can consider a new metric that evaluates how the player locations have added or reduced the threat of an attacking situation. This could be used as a way to value the positioning of a team's attack or defence by reviewing how their positioning increased or decreased the threat of the situation in comparison to the original xT value, which is essentially the average threat for a given state. We therefore present a new 'Threat Above Expected' (TAx) metric, this is expanded on in the following subsections.

7.1 Threat Above Expected Metric

As discussed previously, the original xT model will output an average attacking threat as it is using past event data. Therefore, we can compare the attacking threat output from the spatial xT model to the original xT model to review how player positioning has affected goal likelihood. To properly evaluate this in an interpretable way, we propose a new TAx metric formulated as:

$$\mathrm{TAx} = 100 \times \frac{\left(\mathrm{xT}_{spatial} - \mathrm{xT}_{original}\right)}{\mathrm{xT}_{original}}$$

Where TAx is the TAx metric, $\mathrm{xT}_{original}$ is the goal probability in the next 5 actions using the original xT model and $\mathrm{xT}_{spatial}$ is the goal probability in the next 5 actions using the spatial xT model. The TAx metric is essentially the percentage increase or reduction in attacking threat. Therefore, a negative TAx value suggests that the observed spatial structure is favourable to the defending side compared to what might be expected on average at the same pitch location, and vice versa for positive values. We can use this metric to perform numerous evaluations of teams in the English Premier League, as shown in the upcoming sections. Assessing the contribution of the attack and the defence to the TAx metric is an interesting research task as both are interconnected with both attackers and defenders positioning themselves based on the opposition. During our evaluations, we examine TAx over many samples to identify trends for a team's defence and attack.

7.2 Team head-to-heads

Now that we have introduced the TAx metric, we can evaluate team performances to see how well they limit threat when defending and increase threat when attacking. In particular, we will calculate the mean TAx for the defensive attacking team in head-to-heads between each team. It is important to note that the TAx metric measures how player positioning has contributed to an added or reduced threat in comparison to the original Expected Threat model. Therefore, a team may have a higher mean spatial xT than another team, but a lower mean TAx value. In this case, the team is more susceptible to dangerous attacks, but their defensive positioning and the location of the opposition's attackers have led to less dangerous attacking space and passing options for the opposition on average. As the TAx value is calculated from the spatial xT model using spatial features, it is used as a metric to evaluate the effect of player positioning, independent of player skill. This evaluation uses StatsBomb's 360 data from the 2021-22 season.

Figure 9 presents a matrix of head-to-heads, showing a team's attack and defence on the x-axis and y-axis respectively, with matrix values indicating the mean TAx in each head-to-head. A mean column is also given, to show the mean TAx of a team's attack and defence over all head-to-heads.

Defending Team \ Attacking Team	Arsenal	Brighton & Hove Albion	Burnley	Chelsea	Leeds United	Liverpool	Manchester City	Manchester United	Southampton	Tottenham Hotspur	mean
Arsenal	0.0	1.19	1.63	3.45	5.97	-3.38	-0.01	-6.77	-1.27	4.0	**0.54**
Brighton & Hove Albion	5.34	0.0	2.45	3.9	0.46	-0.4	-2.62	4.35	-0.14	7.73	**2.34**
Burnley	0.48	0.04	0.0	4.89	10.06	-0.65	-2.74	-3.94	5.8	7.57	**2.39**
Chelsea	2.64	6.14	-0.62	0.0	-7.32	1.8	0.36	-3.48	-0.14	8.32	**0.85**
Leeds United	1.24	0.55	-1.18	-1.86	0.0	2.89	-0.3	-4.73	-0.93	-0.28	**-0.51**
Liverpool	3.19	1.5	-13.98	1.19	13.05	0.0	0.2	0.34	-4.32	11.47	**1.4**
Manchester City	-0.95	-2.47	-4.62	-1.28	2.68	-1.02	0.0	0.2	4.2	0.92	**-0.26**
Manchester United	-2.43	8.35	1.44	4.46	2.79	0.53	1.55	0.0	-1.86	6.32	**2.35**
Southampton	-1.3	-0.43	-1.46	-0.07	1.48	3.1	7.82	-0.02	0.0	15.82	**2.77**
Tottenham Hotspur	1.71	-3.09	-2.66	7.66	0.15	3.52	3.91	-1.48	3.19	0.0	**1.44**
mean	**1.1**	**1.31**	**-2.11**	**2.48**	**3.26**	**0.71**	**0.91**	**-1.72**	**0.5**	**6.88**	

Figure 9 – Mean TAx Values for head-to-heads over the 2021/22 English Premier League season.

This highlights interesting insights into which teams can maintain the best defensive unit against particular opponents, and would be useful as a pre-match analytical tool to identify which team setup limited attacking threat most against a particular opposition, and traits of that system could therefore be practised in training. This also could be used as an attacking tool, by identifying which teams struggle to position defensively against different attacking styles.

For example, Tottenham Hotspur are an interesting case study, as they seem to struggle defensively against the teams nearer the top of the table, suggesting they may find difficulty in defending against sustained build-ups. Whereas they have lower TAx values for lower table teams, suggesting that they position well defensively against counterattacks. It is also clear that most teams struggle to maintain strong defensive units against Tottenham Hotspur as their mean TAx when attacking is highest. Teams may therefore use this model to study particular examples of poor defending and identify trends where Tottenham Hotspur are causing defensive vulnerabilities. This is likely due to their counter-attacking style, meaning that they were finding many situations with open spaces to attack. It may also be due to the impact of their attacking fullbacks creating overloads. Southampton also present as another team who struggle most defensively against teams at the top of the table. This may suggest that they are well prepared for counterattacks but struggle with player overloads.

Overall, the teams with the lowest mean TAx value in defence are Leeds and Manchester City. Leeds stand out here as their mean spatial xT is the highest of all teams but they have the lowest TAx value. This may be due to their high-intensity style presenting high interception likelihood in the spatial xT model, however, as the spatial xT only considers player positions myopically, space left when the press is bypassed cannot be accounted for. Southampton and Manchester United have the highest mean TAx values when defending, suggesting that they struggle to work as a defensive unit. Manchester United stand out as an interesting example here, and the statistics support this, showing that they conceded the 13th most goals in the league despite finishing 6th in the table.

7.3 Lowest TAx Examples

To demonstrate the model in action, we generate some heatmaps of transition probabilities to see how the transition probabilities in the T^* model compared to the original transition model T, giving further insight into the calculation of the TAx values. In these examples, the heatmaps detail the increase and decrease in transition probability for each zone in T^* compared to T. Therefore, if the red team is on the ball, and an area is shaded red, this is suggesting that the spatial transition model, T^*, model predicts the ball is more likely to go to this region than average. These heatmaps are interpolated for easier viewing. We will first present some examples of good defensive structure and limited attacking threat based on the best TAx values in the dataset, shown in Figure 10.

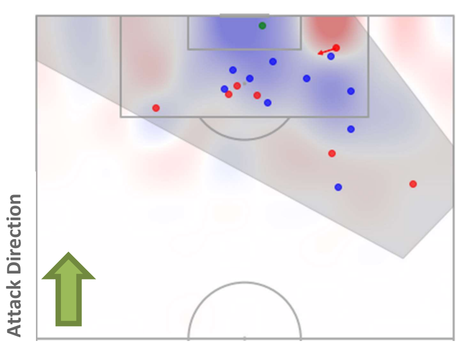

TAx: -78.5%

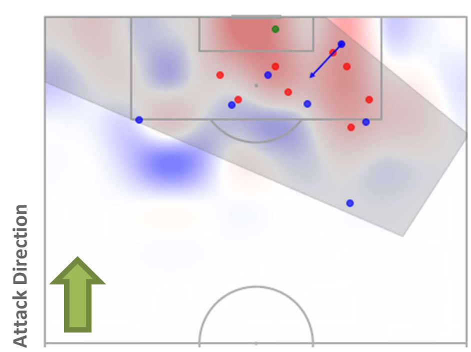

TAx: -75.9%

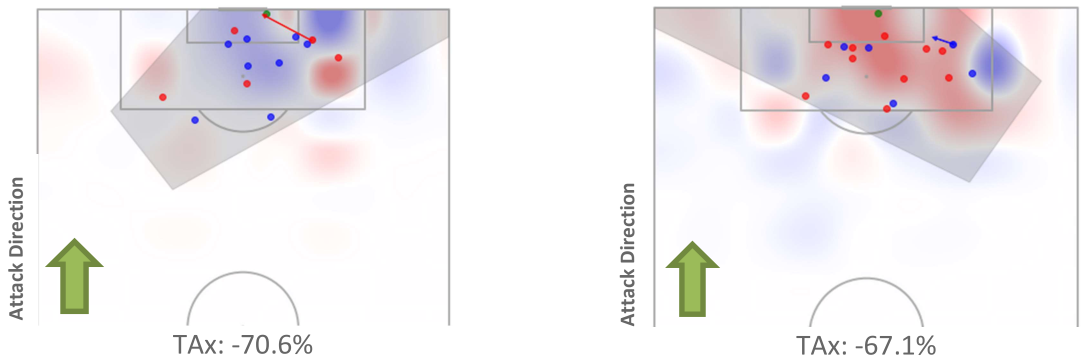

Figure 10 – Examples of lowest TAx values e.g. lower attacking threat than average. Interpolated heatmaps showing more likely transition probabilities given spatial context, with heatmap colours matching team colour.

As shown in these examples, the ball carrier is generally well marked, and most attacking options are being blocked by defenders. This suggests that it is difficult for the attacking team to progress the ball into a more dangerous position without losing the ball, and the defence is well set up to limit the threat of moving the ball. Therefore, we see low TAx values.

7.4 Highest TAx Examples

Examples are now shown for situations where TAx is much higher, suggesting that the attacking threat is higher than usual due to the location of attackers and defenders. These are shown in Figure 11.

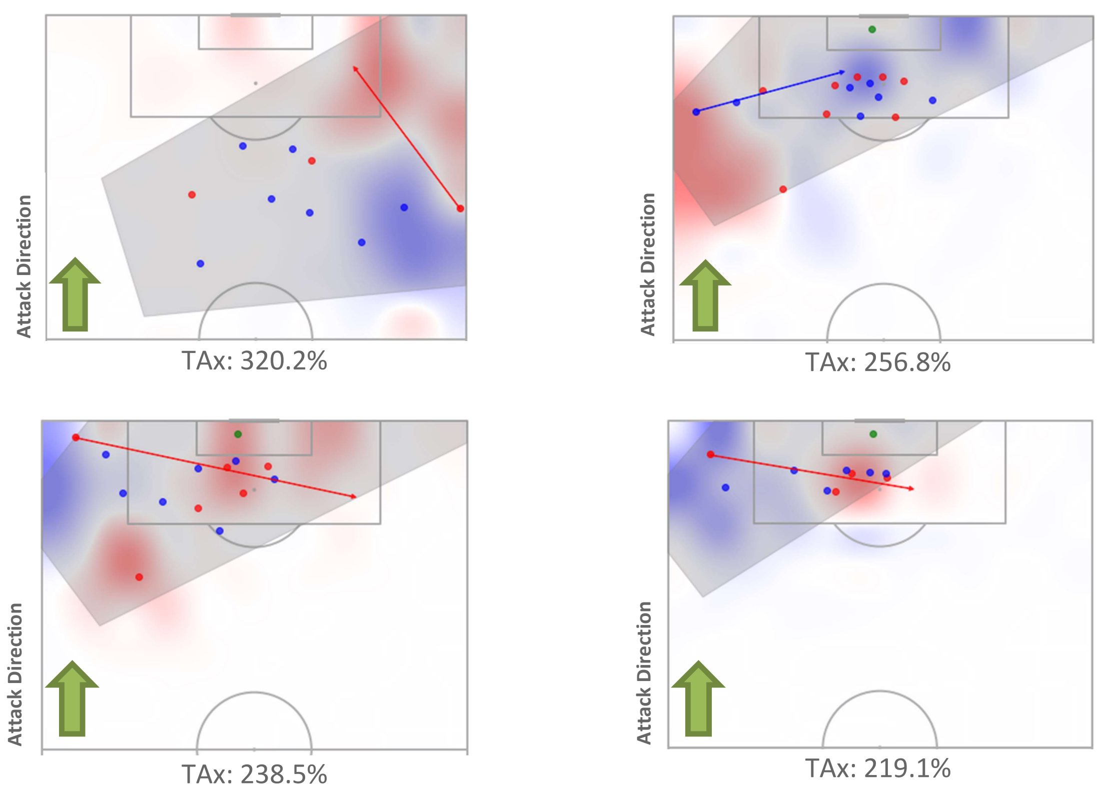

Figure 11 – Examples of highest TAx values e.g. higher attacking threat than average. Interpolated heatmaps showing more likely transition probabilities given spatial context, with heatmap colours matching team colour.

These heatmaps introduce an interesting aspect of the model. In situations where the ball carrier is far wide and it is clear that the ball is about to be crossed, the model has a high TAx value. This is likely because the model expects that the ball will reach a dangerous area in the next action, and if won by the attacking team, pose a high goal threat, particularly in comparison to the average wide-play.

7.5 Defensive Optimiser

We introduce here a proof-of-concept to identify the optimal defensive setup to minimise xT for a given situation. For now, we demonstrate examples of this approach, and we look to refine this algorithm and perform further defensive evaluations using this algorithm in future work. We approach this task by using a basin-hopping global optimiser with a sequential least squares programming algorithm where the spatial xT is looking to be minimised using the set of defender locations as a decision variable where each defender could be located within 5 metres of their actual location and the bounds of the pitch. This aims to find reachable optimal defensive positioning given where they were actually positioned during the event. A couple of examples can be seen below in Figure 12. Similarly to the TAx figures, these heatmaps visualise the increase and decrease in transition probability for T^* compared to the original transition model T to give an idea of most likely transitions in comparison to the average transitions from that current state.

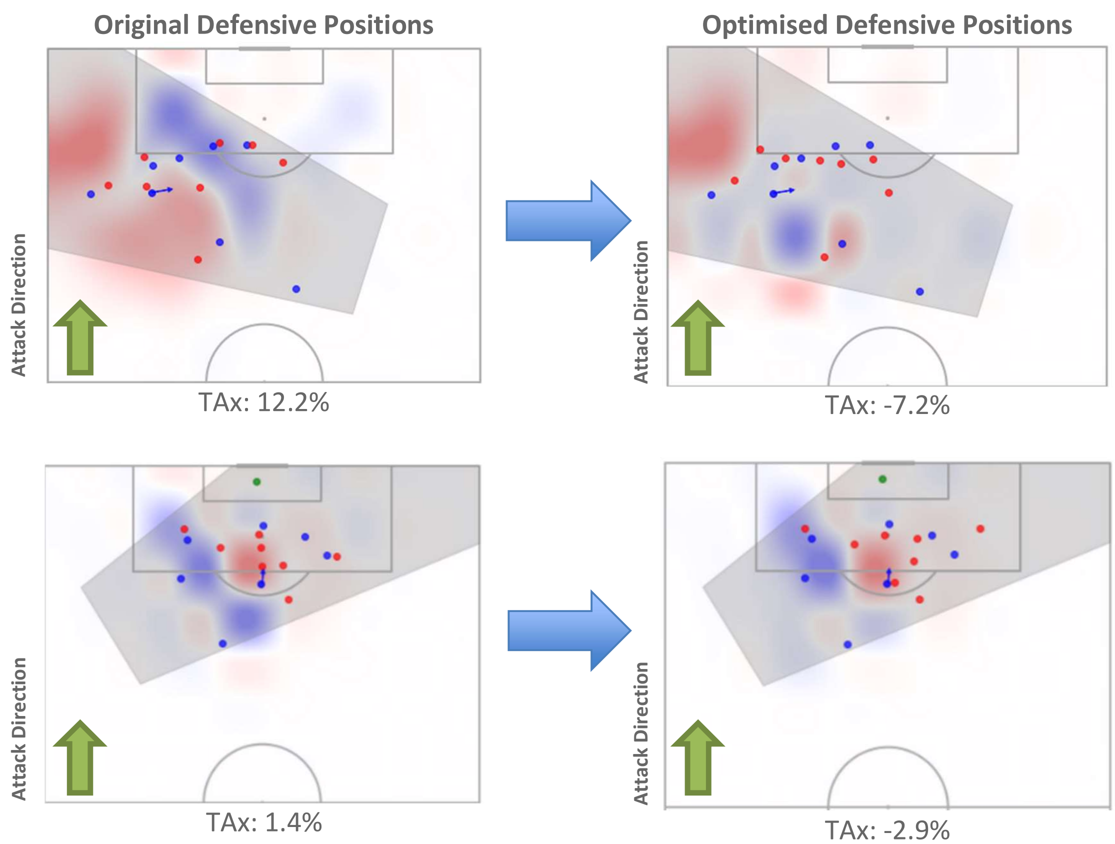

Figure 12 – Examples of the defensive optimiser, showing the actual player locations and an example with optimal defensive positioning computed from the optimiser.

As shown, this can be used to find defensive setups which reduce the attacking threat of situations given attacker locations. This could be used by teams and analysts to learn improved defensive structure and identify areas of the pitch that can be improved most. In particular, this could aid in post-match analysis to identify when the defence could have improved and how. Whilst this optimiser is a good initial step, improvements are needed to help find global optima, as due to long running times and discrete pitch zones, it is difficult to reach global optima and instead finds local optima due to the irregular optimization landscape, a product in part of the relatively coarse pitch grid. This leads to results that sometimes can be hard to interpret. However, a fully functional optimiser is a key area of future work which could be used to evaluate defences in rigorous detail.

8. Conclusion

In this paper, we proposed a novel modification of the Expected Threat model which considers player positioning in football. Our model utilises StatsBomb 360 data to consider a contextual approach to move transition probabilities within the Expected Threat architecture, with the use of a convolutional neural network model. Using these transition probabilities, we compute a modified Expected Threat value and validate our model by comparing its predictive accuracy to original Expected Threat. We also propose a new “Threat Above Expected” (TAx) metric for computing the success of attacking and defensive player positioning and compare positional effectiveness of 10 teams in the 2021/22 English Premier League season. Finally, we show how our model can be applied to perform a deeper analysis of team structures and used as an analytical tool in opposition analysis, with the inclusion of a defensive optimiser to find optimal defensive positioning. In future work, we look to add a contextual shot model to our modified Expected Threat. Furthermore, we aim to improve and streamline the defensive optimiser to evaluate teams and optimal defending over a large sample of games.

References

[1] K. Singh, "Introducing Expected Threat(xT)," 2019. [Online]. Available: https://karun.in/blog/expected-threat.html. [Accessed 08 August 2022].

[2] T. Decroos, L. Bransen, V. H. Jan and J. Davis, "VAEP: an objective approach to valuing on-the-ball actions in soccer.," in *Proceedings of the twenty-ninth international joint conference on artificial intelligence, IJCAI-20*, 2020.

[3] J. Fernandez, L. Bornn and D. Cervone, "A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions.," *Machine Learning 110.6,* pp. 1389-1427, 2021.

[4] W. Spearman, "Beyond expected goals," in *Proceedings of the 12th MIT sloan sports analytics conference.*, 2018.

[5] W. Spearman, B. Austin, G. Dick, R. Hotovy and P. Pop, "Physics-based modeling of pass probabilities in soccer.," in *Proceeding of the 11th MIT Sloan Sports Analytics Conference*, 2017.

[6] B. Burriel and J. Buldu, "The quest for the right pass: Quantifying player's decision making," in *Statsbomb Conference*, 2021.

[7] S. Rudd, "A Framework for Tactical Analysis and Individual Offensive Production Assessment in Soccer Using Markov Chains," in *New England Symposium on Statistics in Sports*, 2011.

[8] T. Decroos and J. Davis, "Valuing on-the-ball actions in soccer: a critical comparison of XT and VAEP," 2020.

[9] M. Van Roy, P. Robberechts, W.-C. Yang, L. De Raedt and J. Davis, "Leaving goals on the pitch: Evaluating decision making in soccer.," in *2021 MIT Sloan Sports Analytics Conference*, 2021.

[10] M. Van Roy, P. Robberechts and J. Davis, "Optimally Disrupting Opponent Build-ups.," in *StatsBomb Conference*, 2021.

[11] M. Stockl, T. Seidl, D. Marley and P. Power, "Making offensive play predictable-using a graph convolutional network to understand defensive performance in soccer.," in *Proceedings of the 15th MIT Sloan Sports Analytics Conference*, 2021.

[12] D. Link, S. Lang and P. Seidenschwarz, "Real time quantification of dangerousity in football using spatiotemporal tracking data.," *PloS one,* 2016.

[13] J. Fernandez and L. Bornn, "Soccermap: A deep learning architecture for visually-interpretable analysis in soccer.," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2020.

[14] H. Sotudeh, "Potential Penetrative Pass," in *StatsBomb Conference*, 2021.

[15] D. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *International Conference for Learning Representations*, 2014.

[16] J. Fernandez and L. Bornn, "Wide Open Spaces: A statistical technique for measuring space creation in professional soccer," in *Sloan sports analytics conference. Vol. 2018*, 2018.

[17] J. Fernández, L. Bornn and D. Cervone, "Decomposing the immeasurable sport: A deep learning expected possession value framework for soccer.," in *13th MIT Sloan Sports Analytics Conference.*, 2019.