



Delft University of Technology

The trade-off between model flexibility and accuracy of the Expected Threat model in football

van Arem, K.W.; Söhl, Jakob; Bruinsma, Mirjam; Jongbloed, Geurt

Publication date
2025

Document Version
Final published version

Published in
MathSport International 2025

Citation (APA)

van Arem, K. W., Söhl, J., Bruinsma, M., & Jongbloed, G. (2025). The trade-off between model flexibility and accuracy of the Expected Threat model in football. In D. Goossens (Ed.), *MathSport International 2025: 11th International Conference on Mathematics in Sport* (pp. 150-155).

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

The trade-off between model flexibility and accuracy of the Expected Threat model in football

Koen W. van Arem^{1*}, Jakob Söhl¹, Mirjam Bruinsma² and Geurt Jongbloed¹

¹ Delft University of Technology, The Netherlands

² AFC Ajax, The Netherlands

* k.w.vanarem@tudelft.nl

Abstract

With an average football (soccer) match recording over 3,000 on-ball events, effective use of this event data is essential for practitioners at football clubs to obtain meaningful insights. Models can extract more information from this data, and explainable methods can make them more accessible to practitioners. The Expected Threat model has been praised for its explainability and offers an accessible option. However, selecting the grid size is a challenging key design choice that has to be made when applying the Expected Threat model. Using a finer grid leads to a more flexible model that can better distinguish between different situations, but the accuracy of the estimates deteriorates with a more flexible model. Consequently, practitioners face challenges in balancing the trade-off between model flexibility and model accuracy. In this study, the Expected Threat model is analyzed from a theoretical perspective and simulations are performed based on the Markov chain of the model to examine its behavior in practice. Our theoretical results establish an upper bound on the error of the Expected Threat model for different flexibilities. Based on the simulations, a more accurate characterization of the model's error is provided, improving over the theoretical bound. Finally, these insights are converted into a practical rule of thumb to help practitioners choose the right balance between the model flexibility and the desired accuracy of the Expected Threat model.

1 Introduction

An average football (soccer) match contains around 3,000 on-ball events [1]. This data can provide new meaningful insights for football clubs. It can, for instance, be used to study players for scouting or opponent analysis. This can help practitioners make better informed decisions. Mathematical models can be used to extract more complex information from these on-ball events [2], which can give football clubs an advantage over the competition.

However, advanced mathematical models introduced by researchers might not be adaptable or practical enough for coaches and teams [3]. It is important to have models that can be explained in such a way that they can be understood by practitioners and utilized in practice [3]. To achieve this, models to distill information from the in-game data should also be assessed on how explainable and interpretable they are [4]. This means that explainable models are of added value because they offer an accessible option to retrieve complex information from in-game data.

The Expected Threat [5, 6] model quantifies the quality of an action, which gives detailed information about offensive player quality, and it is praised for its interpretability [7]. Nonetheless, practitioners still face one key design choice when applying the Expected Threat model. The Expected Threat model estimates the probability of scoring, and to do this, it divides the pitch into different in-game states. The number of in-game states, M , describes the model flexibility and can be chosen. A model with more game states can distinguish between more situations [8]. On the other hand, more states decrease the accuracy of the Expected Threat model because more probabilities are estimated with the same amount of data. Increasing the amount of data is often infeasible in practice, because of the additional costs. This trade-off between model flexibility and accuracy hinders the application of the otherwise accessible Expected Threat model in practice. The aim of this study is to provide a rule of thumb for practitioners to manage the trade-off between model flexibility and accuracy of the Expected Threat model.

2 Expected Threat

The Expected Threat model [5, 6] considers football as a Markov chain and is based on the idea that good actions increase the probability of scoring a goal within the possession chain. The pitch is divided into M squares, and the state of the game is defined as the square where the ball-possessing player is. For each state s , the model then calculates the probability of scoring, denoted by $xT(s)$. The quality of an action is defined as the difference before and after the action: $\Delta xT(s_{\text{before}}, s_{\text{after}}) = xT(s_{\text{after}}) - xT(s_{\text{before}})$.

When a player has ball possession, there are two ways to score a goal: either directly score a goal or move the ball to another state with a dribble or pass and score from there. To score a goal, the player has to decide to shoot, and the player has to score the shot. If the current game state is denoted as s , the probability of this happening is $P(\text{shot}|s) \cdot P(\text{goal}|\text{shot}, s)$. Because $P(\text{goal}|\text{shot}, s)$ is the quantity described by Expected Goal (xG) models, this can also be denoted as $P(\text{shot}|s) \cdot xG(s)$. If the player decides not to shoot, a goal can be scored by moving the ball to each other game state s' and scoring from there. The probability of scoring via the game state s' can be written as $T_{s \rightarrow s'} \cdot xT(s')$, where $T_{s \rightarrow s'}$ is the probability of transitioning from s to s' . This means that the probability of scoring from game state s is

$$xT(s) = P(\text{shot}|s) \cdot xG(s) + \sum_{s'} T_{s \rightarrow s'} \cdot xT(s'). \quad (1)$$

In practice, the values of $P(\text{shot}|s)$, $xG(s)$, and $T_{s \rightarrow s'}$ are estimated by counting the occurrences of these events in the data set. When these are estimated, the only unknowns in (1) are the values $xT(s)$. Because (1) holds for each state s , it gives a system of equations, which is generally solved using an iterative algorithm. In this way, the probability of scoring from state s is estimated.

Due to randomness in the training data, errors are made in the estimation of $P(\text{shot}|s)$, $xG(s)$, and $T_{s \rightarrow s'}$. These estimation errors cause errors in the estimated xT -values. For practitioners, it is important to have a bound on these errors. The model error in this study is defined as the maximal difference between the true and the estimated xT -values, denoted by $\|xT - \widehat{xT}\|_{\infty}$. The distribution of this error depends on the number of training points N and the number of game states M and can be used to describe the trade-off between model flexibility, described by M , and model accuracy.

Using the properties of the Expected Threat model, it is possible to derive probabilistic bounds on the error of the model. For this purpose, the transition matrix is the matrix with the transition probabilities $T_{s \rightarrow s'}$ as entries. A summary of our theoretical results is described in the following theorem.

Theorem 1. Let $g \in \mathbb{R}^M$ be defined by $g_s = P(\text{shot}|s) \cdot xG(s)$ and let T be the transition matrix. Assume that $\|T\|_\infty < 1$ and that for the estimated transition matrix $\|\hat{T}\|_\infty < 1$. Moreover, assume that the estimates of the quantities in the Expected Threat model are obtained by taking averages of N independent Bernoulli random variables. Then the following bounds hold with probability at least $1 - \alpha$:

$$\|\widehat{xT} - xT\|_\infty \leq \frac{1}{1 - \|T\|_\infty} \left(M \sqrt{\frac{\log(2M^2/\alpha)}{2N}} + \sqrt{\frac{\log(2M/\alpha)}{2N}} \right) \leq \frac{2}{1 - \|T\|_\infty} \left(M \sqrt{\frac{\log(2M^2/\alpha)}{2N}} \right). \quad (2)$$

More specifically, the term $M \sqrt{\frac{\log(2M^2/\alpha)}{2N}}$ corresponds to the error in estimating T and $\sqrt{\frac{\log(2M/\alpha)}{2N}}$ to the error in estimating g .

This theorem shows that the error in estimating the xT -values is of the order $O(M \sqrt{\log(M)}/\sqrt{N})$. However, the results also suggest that, in practice, implemented models might fall under a finite sample regime where the error in g is still larger, and where a faster decay of the error can be observed.

3 Methods

To find the error distribution of the Expected Threat model in practice, simulations were performed based on the Markov chain underlying the model. The data used for this simulation study is obtained from the openly available Statsbomb data set [9]. All available games from the Premier League, Ligue 1, Serie A, La Liga, and the Bundesliga were used. The events that did not describe passes, dribbles, errors, clearances, or shots were filtered out. This resulted in a data set of approximately 4,000,000 events, which is equivalent to around 6.5 full seasons of one league.

In this research, the maximal model error was studied for Expected Threat models with different discretization grids, and thus for different flexibilities M , which are described in Table 1. For each grid size M , one Expected Threat model was calculated with the Statsbomb data. These models were assumed to be the ground truth for elite male leagues within the scope of this study. Because the data set is relatively large, this is a reasonable assumption.

Each of these ground truth models describes a Markov chain, which can be used to resample a new data set. For each ground truth model, new data sets of different sizes N were resampled, which were then used to train resampled Expected Threat models. The model error was then obtained by calculating the maximal absolute difference between the ground truth xT -values and the xT -values based on the resampled data. This process was repeated 1,000 times for each combination of M and N as described in Table 1. This created 48,000 data points describing the model error, the grid size M and the number of data points N .

The goal of this research is to obtain more insight into the trade-off between the model flexibility governed by M and the model accuracy, described by the model error. Because situations with $M \log(M)/\sqrt{N} \geq 15$ resulted in errors too large for practice, these were filtered out. This makes it possible to study the distribution of the errors in this setting, which is interesting for practitioners.

Using the simulated data, the distribution of the maximal model error could be studied. To do this, the following lognormal model was assumed to describe the maximal model error:

$$\|\widehat{xT} - xT\|_\infty = C \frac{M^\alpha}{(\sqrt{N})^\beta} e^\varepsilon, \quad \text{where } \varepsilon \sim N(0, \sigma^2). \quad (3)$$

n_x	n_y	$M = n_x \cdot n_y$	N
16	12	192	100,000
32	24	768	130,000
40	30	1200	170,000
48	36	1728	240,000
56	42	2352	370,000
64	48	3072	630,000
			1,300,000
			4,000,000

Table 1: Grid configurations with corresponding M values and the sample sizes N of the simulations.

The model error, which is a maximal absolute difference, is known to be positive. Additionally, the theoretical results indicated that both the mean and spread of the error are small if M is small and N is large. This makes the lognormal model a reasonable assumption. Moreover, this model describes the powers of the variables M and N , which are unknown because the theoretical results only gave an upper bound.

If $c = \log(C)$, (3) is equivalent with

$$\log(\|\widehat{xT} - xT\|_\infty) = c + \alpha \log(M) - \beta \log(\sqrt{N}) + \varepsilon, \quad \text{where } \varepsilon \sim N(0, \sigma^2). \quad (4)$$

This formulation gives a linear model with normal residuals. Therefore, ordinary least squares (OLS) can be applied to estimate c, α, β , and σ^2 .

4 Results

Dep. Variable:	$ \widehat{xT} - xT _\infty$	R-squared:	0.835
Model:	OLS	Adj. R-squared:	0.835
No. Observations:	23000	Log-Likelihood:	-9297.8
Df Residuals:	22997	AIC:	1.860e+04
Df Model:	2	BIC:	1.863e+04

	coef	std err	t	P> t	[0.025	0.975]
c	-1.8758	0.026	-72.138	0.000	-1.927	-1.825
α	0.9898	0.003	330.569	0.000	0.984	0.996
β	1.0416	0.004	238.837	0.000	1.033	1.050
Variance residuals	0.1314					

Table 2: Summary of the OLS model fitted on the log maximal model error for the data points with $M \log(M) / \sqrt{N} < 15$.




Figure 1: A QQ plot of the residuals of the fitted OLS model.

The summary of the ordinary least squares applied to (4) is given in Table 2. It shows that the adjusted R^2 is 0.835, which indicates that the model is able to explain most variance of the errors with M and N . Additionally, Figure 1 shows the QQ plot of the residuals of the lognormal OLS model. It is visible that the residuals indeed seem to be well-described by the lognormal OLS model. Thus, it can be concluded that the model in (4) provides a good description of the distribution of the maximal error of the Expected Threat model. With the found values for α , β and σ^2 , the distribution of the model can be described by

$$\|\widehat{xT} - xT\|_{\infty} = e^{-1.8758} \frac{M^{0.9898}}{(\sqrt{N})^{1.0416}} e^{\varepsilon}, \quad \text{where } \varepsilon \sim N(0, 0.1314). \quad (5)$$

The values found for α and β are close to 1, although significantly different according to the confidence intervals in Table 2. This means that the maximal error of the model is of order $O(M^{0.9898}/(\sqrt{N})^{1.0416})$, which indicates that the error, in practice, is of a lower order than established by the theoretical results.

5 Rule of thumb

With the distribution of the model error, it is possible to give guidance to practitioners on how to use the Expected Threat model. If they have an existing model, the distribution of the error in (5) can be used to describe the distribution of the error of their model. For example, consider the Expected Threat model by Singh [6], which has a 16×12 grid and one season of the Premier League as training data. This corresponds to $M = 16 \cdot 12 = 192$ game states and roughly $N = 620,000$ training data points. In the experience of consulted experts, errors smaller than 0.03 are acceptable for scouting purposes. The distribution of the maximal error of this model is visualized in Figure 2. It indicates that there is a 62.09% chance of having an error that is lower than 0.03. This means that there is a reasonable chance of this model having an acceptable error.

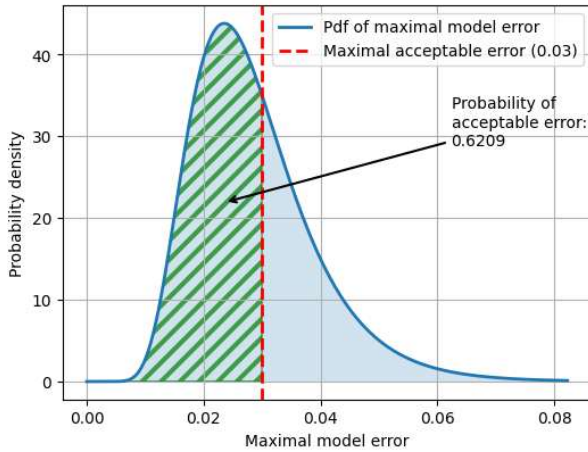


Figure 2: The distribution of the maximal model error of the Expected Threat model in [6], where $M = 192$ and $N = 620,000$.

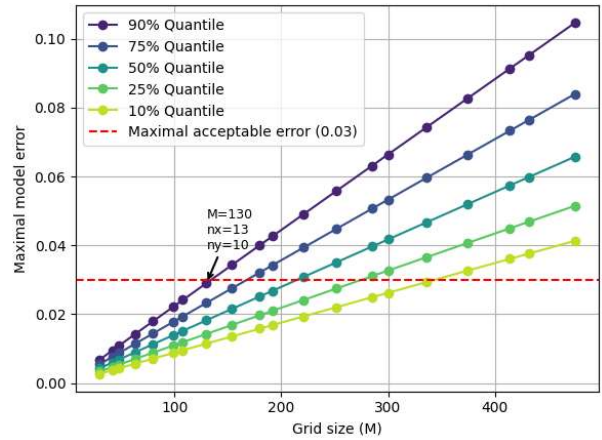


Figure 3: The quantiles of the maximal model error of an Expected Threat model with $N = 2,480,000$ training data points for different grid sizes M .

When training a new model, the number of game states M has to be chosen. It is desirable to have a low model error with a high probability. On the other hand, a more flexible model can better distinguish between in-game situations. To balance this trade-off, a reasonable idea would be to choose the most flexible model with an acceptable statistical error. Through consultation with experts, it was established that the Expected Threat model is sufficiently reliable for scouting purposes when the maximal model error is smaller than 0.03 with a 90% probability. This can be reformulated as the following rule of thumb: *select the most flexible model (highest M) such that the maximal model error is smaller than 0.03 with a 90% probability.*

To illustrate this rule of thumb, suppose a practitioner wants to train an Expected Threat model on a data set of 2,480,000 data points. This corresponds to data of 4 league seasons. Figure 3 shows different quantiles of the error for values of M for this number of data points N based on (5). The results show that the maximal M with a 90% quantile smaller than 0.03 is $M = 130$, which corresponds to a 13×10 grid. This means that the rule of thumb gives that a 13×10 grid yields the most flexible model with an acceptable model error. In this way, the rule of thumb provides guidance to practitioners on how to deal with the trade-off between model flexibility and accuracy, which makes the accessible Expected Threat model even more easily applicable for practitioners.

Acknowledgements

The authors would like to thank StatsBomb for making the StatsBomb Open Data publicly available.



References

- [1] Yorke, J. (2022) *StatsBomb 360: Exploring Line Breaking Passes*. Accessed: 14-5-2025 at <https://statsbomb.com/articles/soccer/statsbomb-360-exploring-line-breaking-passes/>.
- [2] Rein, R. and Memmert, D (2016) *Big data and tactical analysis in elite soccer: future challenges and opportunities for sports science*. SpringerPlus **5**, 1410.
- [3] Herold, M., Goes, F., Nopp, S., Bauer, P., Thompson, C., and Meyer, T. (2019) *Machine learning in men's professional football: Current applications and future directions for improving attacking play*. International Journal of Sports Science & Coaching, **14**(6), 798-817.
- [4] Davis, J., Bransen, L., Devos, L., Jaspers, A., Meert, W., Robberechts, P., Van Haaren, J. and Van Roy, M. (2024) *Methodology and evaluation in sports analytics: challenges, approaches, and lessons learned*. Machine Learning **113**, 6977–7010.
- [5] Rudd, S. (2011) *A Framework for Tactical Analysis and Individual Offensive Production Assessment in Soccer Using Markov Chains*. In Proc. New England Symposium on Statistics in Sports.
- [6] Sing, K. (2018) *Introducing Expected Threat xT* . Accessed: 5-11-2024 at <https://karun.in/blog/expected-threat.html>.
- [7] Van Roy, M., Robberechts, P., Decroos, T., and Davis, J. (2020) *Valuing On-the-Ball Actions in Soccer: A Critical Comparison of xT and VAEP*. In Proc. of the AAAI-20 Workshop on Artificial Intelligence in Team Sports.
- [8] Van Arem, K., Bruinsma, M. (2024) *Extended $xThreat$: an explainable quality assessment method for actions in football using game context*. In Proc. ISEA 2024 - The Engineering of Sport 15.
- [9] StatsBomb (2024) *Open Data*. Accessed: 15-10-2024 at <https://github.com/statsbomb/open-data>.