Beating the bookie: A look at statistical models for prediction of football matches

Helge LANGSETH

Department of Computer and Information Science,
Norwegian University of Science and Technology, Trondheim, Norway

Abstract. In this paper we look at statistical models for predicting the outcome of football matches in a league. That is, our aim is to find a statistical model which, based on the game-history so far in a season, can predict the outcome of next round's matches. Many such models exist, but we are not aware of a thorough comparison of the models' merits as *betting models*. In this paper we look at some classical models, extract their key ingredients, and use those as a basis to propose a new model. The different models are compared by simulating bets being made on matches in the English Premier League.

Keywords. Association football, statistical models, predictions, betting

1. Introduction

Association football (also known simply as "football", and sometimes even "soccer") is one of the most popular spectator sports in the world. Everyday, millions of people watch and discuss football, and one favorite pastime of the football fan is guessing on outcomes of future games. The motivation can be just for fun, or for making money in a betting market.

We can also see an increasing number of statistical models being made to perform game predictions. In this paper we will first review some of these models briefly. Secondly, since the models have been evaluated on separate datasets (and often with respect to different target objectives), we will compare the models on a dataset containing games played in Barclay's Premier League during the 2011-2012 and 2012-2013 seasons. The models are primarily intended as a means towards beating the bookmaker and make a profit, and we will therefore also discuss techniques to decide what amount of money to ideally put on each available bet. Finally, we will argue that football is a numbers game, and when traditional prediction models like [6,2,10] only use information about the final outcome of the game, they loose out on important information that can improve the predictive ability. As a first step towards more descriptive statistical models, we will propose a simple model that we discuss in some detail.

2. Predicting football matches

2.1. Classical models for match betting

One of the first thoroughly analyzed models for prediction of results in football matches in a league was made by Maher in 1982 [6]. He modeled the number of goals team i scores against team j when playing at home, X_{ij}, as a Poisson distributed variable. Similarly, the number of goals *conceded* by team i in the same game, denoted Y_{ij}, was also modeled as a Poisson variable, but with a different parameter. Crucially, he also assume X_{ij} and Y_{ij} to be independent.

Maher continued by assuming that each team i has an *attack strength* denoted α_i and a *defense* strength, β_i. A high attack-strength indicates that a team scores many goals; values are normalized so that the average value over the league equals one. On the other hand, a low value for the defense-strength indicates a tendency to concede only a few goals. Finally, there is a *home field advantage* denoted k, which is assumed to be the same for all teams. The full model now states that the outcome of the game between Team i and Team j follows the probability distribution

$$P(X_{ij} = x, Y_{ij} = y|\boldsymbol{\alpha}, \boldsymbol{\beta}, k) = \text{Poisson}(x|k \cdot \alpha_i\beta_j) \cdot \text{Poisson}(y|\alpha_j\beta_i).$$

Here we use $\text{Poisson}(x|\lambda)$ to denote the Poisson probability distribution function with parameter λ evaluated at x. Maher found maximum likelihood estimators for the parameters in his model.

Two fundamental objections have been raised against Maher's model. Firstly, if X follows the Poisson distribution, then the the two first central moments are equal, i.e., $\mathbb{E}[X] = \text{Var}(X)$. Karlis and Ntzoufras [4] found that for real game data, the variance is larger than the expectation, thereby violating the Poisson distribution assumption. We can compensate for this fact by casting Maher's model into a Bayesian setting. Let Λ be a random variable from a suitable distribution (e.g., the Gamma-distribution), and let $X|\{\Lambda = \lambda\} \sim \text{Poisson}(\lambda)$. Then, by the law of total variance, $\text{Var}(X) = \mathbb{E}[\Lambda] + \text{Var}(\Lambda) > \mathbb{E}(\Lambda) = \mathbb{E}(X)$, which therefore is in-line with Karlis and Ntzoufras observations.[1] Secondly, even if the independence assumption is crucial in the above mentioned model, it appears to go against the layman's understanding of football. A team that is vastly superior to its opponent will show that by both scoring many goals and at the same time concede very few. Still, thorough examination of results from many of the major football leagues in Europe indicate that the assumption is not strongly violated in real match-data [4].

Dixon and Coles [2] proposed two important extensions to Maher's independence-model. Firstly, they defined that

$$P(X_{ij} = x, Y_{ij} = y|\boldsymbol{\alpha}, \boldsymbol{\beta}, k) = \tau(x, y) \cdot \text{Poisson}(x|k \cdot \alpha_i\beta_j) \cdot \text{Poisson}(y|\alpha_j\beta_i),$$

where $\tau(x, y)$ is a function that makes low-scoring draws (0-0 and 1-1) slightly more probable, and the results 1-0 and 0-1 slightly less probable than in Maher's model. This

[1] Even though many football prediction models have been set in the Bayesian framework, we are unaware of over-dispersion being used to motivate the model. In the experiments reported in Section 4 we are using the Bayesian formulation (with inference using MCMC) instead of the likelihood-based formulation originally suggested.

was motivated by an investigation into the results from 6629 league and cup results in the period from 1992 to 1995. Secondly, it is argued that it seems unnatural to assume that a team's attacking and defensive strengths are constant throughout a season. Both an important player becoming unavailable for a part of the season either through an injury, a suspension, or changes made in the transfer window, as well as a fundamental changes in the playing style, e.g., due to the manager being sacked, can lead to fluctuations in a teams offensive and defensive abilities. Dixon and Coles [2] therefore proposed a time-dependent model, where team i's attacking ability at time t is modeled by a dedicated parameter α_i^t (and similar for the defensive strength). A smoothing function was used to incorporate a set of historical games when evaluating the game at time t.

Rue and Salvesen [10] extended this idea by defining a random-walk model, with $\alpha_i^t | \left\{ \alpha_i^{t'} \right\} \sim N(\alpha_i^{t'}, \frac{t-t'}{\tau}\sigma^2)$ and similar for the defensive strength. τ determines the "strength of memory", and is estimated from data.

2.2. *Looking behind the results*

If a team does not obtain results that are quite as good as had been expected, a mantra from managers and fans alike is that "*we need to look behind the results*". The idea is that often, a game "should" have been won, had it not been for a couple of unlucky incidences that totally changed the run of play.[2] The idea to look behind the results seems worthwhile when it comes to building statistical models. While the classical models we have covered so far were created in an age when information about a football game was limited to James Alexander Gordon reading the classified football results on BBC every Saturday evening, we are now in a situation where e.g., Opta[3] and Prozone[4] offer extremely detailed statistics from each player's kick on the ball in close to real-time.

A natural development is thus to extend the classical models for football betting by taking this vast amount of data into account. A first step along this path is described in the following, where the ideas of the model by Maher [6] first is extended by dynamic abilities (in the spirit of Rue and Salvesen [10]), then further enhanced by taking `FiredShots` and `ShotsOnTarget` into account.[5] The structure of the model is as follows: Firstly, when Team i plays Team j at home at time t, they will have a number of *chances* during a game. Their ability to create chances follows a Poisson distribution with parameter λ_i^H, $C_{i,j} \sim \text{Poisson}(\lambda_i^H)$. Now, each chance may lead to a shot being fired, unless the away team is able to break up the chance. Team j's defensive ability at time t, β_j^t is driven by a random walk over time, and is used to model $F_{i,j}$, the number of shots fired by Team i towards Team j's goal, as $F_{i,j} | \{\beta_j^t, C_{i,j}\} \sim \text{Binomial}(C_{i,j}, \text{probit}(\beta_j^t))$. Here we use $X \sim \text{Binomial}(N, p)$ to denote that the random variable X follows the Binomial distribution with N trials and success-probability p. Further, probit(ϕ) denotes the probability that a standard Gaussian variable takes a value of ϕ or larger. A similar model defines T_{ij}, the number of shots on target. For this quantity, it is the attacking strength of Team i at time t that comes into play, $T_{i,j} | \{\alpha_i^t, F_{i,j}\} \sim \text{Binomial}(F_{i,j}, \text{probit}(\alpha_i^t))$. Finally, X_{ij}, the number of goals that

[2] Often, this is said to counter the statement *"the league table doesn't lie"*, which is shouted from the rooftops by the fans who do not need to look *behind* the results to find something to cheer them up.

[3] `http://www.optasports.com/`

[4] `http://www.prozonesports.com/`

[5] Data has been downloaded from `http://football-data.co.uk/`.

Team i ends up scoring, is assumed to be determined by a quantity γ_j^t, which is Team j's goal-keeper ability at time t: $X_{i,j}|\{\gamma_j^t, T_{i,j}\} \sim$ Binomial$(T_{i,j}, \text{probit}(\gamma_j^t))$. The process for modeling the number of goals scored by the away-team in the same match is an exact mirror-image of the process described above, except that the number of chances for the away team is determined by a parameter λ_i^A. Thus, the home-field advantage is modeled separately for each team, and for team i can be characterized by the fraction λ_i^H/λ_i^A. While the model is still extremely simple (the only addition to the parameter set of the previous models is the goalie's ability), we will see in Section 4 that we are making improvements over the previous models.

3. Money management

When it comes to football betting, the ultimate goal of the bettor is not "to be right", but rather to win money. Thus, he is looking for bets where the odds offered by a bookie is better than the calculated probability for a given outcome implies it should be. It is therefore rational to bet at an outcome with a low probability p, if the odds ω is "good enough". As an example, each of the described models allocated low probabilities (less than 10%) to Blackburn beating Manchester United at Old Trafford during the 2011-12 season. However, as the odds were as long as $\omega = 25$, a punt may still have been seen as having value. (As we know, Blackburn ended up winning that game. It was arguably the biggest upset of the season, and the loss eventually lead to Manchester United missing out on the title by goal difference.) The crucial quantity is, obviously, the expected gain per unit at stake, calculated as $p \cdot \omega - 1$.

During a round of games (typically ten games, played in the period from Saturday to Monday), there may be several bets that have a positive expected profit. The punter must then be able to balance how much money to put on each of these bets. Consider an example where two games are available:

Game A: Probability of an away win is $p = .30$, with odds $\omega = 3.6$.
Game B: Probability of a home win is $p = .90$, with odds $\omega = 1.2$.

The expected gain per unit stake is $p \cdot \omega - 1 = 0.08$ for both games, and they can therefore to some extent be seen as equivalent from the punter's perspective. One may still argue that the rational behavior is to put a larger percentage of the punter's bankroll C on Game B than on Game A, because the probability of winning is higher. On the other hand, more money must be wagered to potentially win, e.g., 100 credits, thus the downside for Game B is potentially larger. There are a number of different strategies to this *money management* problem, and we will briefly outline some of these solutions next. To fix ideas, we assume there is a set of possible bets $i = 1, \dots, n$ to wager, each having probability p_i and odds ω_i. Further, assume that the punter will allocate c_i on each bet, with $\sum_{i=1}^n c_i \leq C$. Let the profit from making bet j be denoted Δ_j. This is a random variable with $\mathbb{E}[\Delta_j] = c_j(p_j\omega_j - 1)$ and variance $\text{Var}[\Delta_j] = p_j(1-p_j)(\omega_j c_j)^2$. The task of the money management procedure is to find reasonable values for each c_i. Typically, one would want to find bets that give a high expected return and with low risk. We will therefore only consider bets with positive expected profit, $\mathbb{E}[\Delta_j] > 0$, and call these *feasible*. Below we assume all bets $i = 1, \dots, n$ have this property.

Fixed bet: Allocate the same amount to each feasible bet, $c_i \propto 1$.

Fixed return: Make sure the same *winnings* can be obtained from each bet. This will result in lower amounts staked on the high-gain/low-probability outcomes. $c_i \propto 1/\omega_i$.

Kelly ratio: Kelly [5] proposed to use a decision-theoretic approach to the money management problem. In this setup, the utility of having an amount C after a bet has been determined is defined to be $\ln C$, and the utility of going broke is therefore minus infinity. The expected utility of a bet c_i when the bankroll is C is thus $p_i \cdot \ln(C + \omega_i c_i) + (1 - p_i) \cdot \ln(C - c_i)$, which is maximized for $c_i \leftarrow C \cdot \frac{p_i \cdot \omega_i - 1}{\omega_i - 1}$. When this money management strategy is used in Section 4, it is modified so that the total bets during one round cannot exceed a predefined value C_0, which is chosen to be much smaller than the bankroll at the beginning of the simulation. This is to ensure that a system that looses heavily during the first few rounds is not excessively punished at a later stage.

Variance-adjusted: Rue and Salvesen [10] looked at the difference between the expected profit and the variance of that profit, and wanted to minimize this number. After betting c_i, the difference is $p_i \omega_i c_i - p_i(1 - p_i)(\omega_i c_i)^2$, which is minimized by choosing $c_i \leftarrow (2\omega_i(1 - p_i))^{-1}$.

Markowitz portfolio management: The variance-adjusted approach can be seen as a simplification of Markowitz portfolio management [7]. In this more general framework, one wants to find the allocation of bets c_i that in combination maximizes $\sum_{i=1}^{n} (\mathbb{E}[\Delta_i] - \nu \text{Var}[\Delta_i])$ under the constraint that the total bets during one round sums to a predefined value, C_0. Here, ν represents the level of risk acceptance by the punter. The dual representation of the optimization problem is to minimize $\sum_{i=1}^{n} \text{Var}[\Delta_i]$ under the constraints $\sum_{i=1}^{n} c_i = C_0$ and $\sum_{i=1}^{n} \mathbb{E}[\Delta_i] = \mu$, where now μ takes the role of the risk acceptance parameter. In the experiments reported in Section 4 we vary μ from $\mu_{\downarrow} = \left(\sum_{i=1}^{n} p_i \omega_i\right)/n - 1$ to $\mu^{\uparrow} = \max_i p_i \omega_i - 1$. The low value, $\mu = \mu_{\downarrow}$, models the risk-aversive approach, were we only require the expected return per unit stake of the combined bet to attain the average value of each bet. $\mu = \mu^{\uparrow}$ is the risk-seeking approach, which in practice will force all stakes to be placed on the single bet with the highest expected return.

4. Empirical results

In this section we compare three basic models for predicting football games, using data from the English Premier League 2011/2012 and 2012/2013. The prediction system is based on *JAGS: Just another Gibbs sampler* [9], which is a program for analyzing Bayesian hierarchical models using Markov chain Monte Carlo [3]. The probability distributions over the result of a game (home-win, draw or away-win) is found by sampling the number of goals scored by both the home-team and the away-team (using the models described above), then comparing the two to get a sample of the outcome of the match. This was repeated until we had generated 10.000 samples for each game.

The models are used to simulate betting over the last half of each season (i.e., starting from round 20), and the odds used are those given by William Hill, which is the biggest betting company in the UK. Each round (i.e., those during the same weekend or the same mid-week period) typically consists of ten games, and the betting system predicts the outcomes of each of them. Next, the money management system is used to choose how much to wager on each of the bets in that round, and collects the simulated winnings. Only data from the games that are prior to a round r is made available to the system when making predictions for round r. This repeats until the season is over, i.e., by letting r run from 20 until $r = 38$.

We compare three models: Firstly, we have implemented the Maher model [6], which is denoted STATIC below. Secondly, the Maher model is extended to incorporate the dynamic assessment of attack and defense [2,10]. Our implementation uses a Gaussian motion model [10], and is called DYNAMIC. Finally, the model described in Section 2.2 is denoted DATAINTENSIVE.

4.1. Results from the 2011/2012 Premier League

The results from the three models using different betting strategies are given in Table 1 and Table 2. The numbers give the bettor's gain over the full simulation as a percentage of the total amount betted. Positive numbers thus indicate that the system has made a profit over the season, and negative values means there has been a loss overall.

	Betting strategy			
Model	Fixed bet	Fixed return	Kelly	Variance Adjusted
STATIC	17.4%	17.4%	**23.2%**	**15.6%**
DYNAMIC	**22.7%**	14.3%	21.3%	12.0%
DATAINTENSIVE	20.3%	**24.2%**	23.0%	14.3%

Table 1. Results for the Premier League, 2011-2012 using standard betting strategies. The reported values give the total gains divided by the total bets made over the simulation.

	Markowitz strategy		
Model	Conservative	Intermediate	Aggressive
STATIC	**19.9%**	20.2%	-10.0%
DYNAMIC	17.9%	30.5%	**61.8%**
DATAINTENSIVE	17.4%	**37.5%**	10.5%

Table 2. Results for the Premier League 2011-2012 using Markowitz betting strategies. The *Conservative* strategy is to use $\mu = \mu_{\downarrow}$, the *Intermediate* strategy uses $\mu = (\mu_{\downarrow} + \mu^{\uparrow})/2$, and the *Aggressive* is defined by setting $\mu = \mu^{\uparrow}$.

The first observation is that each of the three betting models produce a profit, and the only combination that looses money is the simplest model (STATIC) combined with the Aggressive Markowitz strategy for money management. The best result is obtained by the DYNAMIC model combined with the Aggressive Markowitz strategy, with a gain of 61.8%. Looking at the details we find that most of the profit was won from two freak results (Liverpool loosing at home to both WBA and Wigan), and without these games a considerable loss would have had to be endured. Markowitz' intermediate strategy is more robust, and the DATAINTENSIVE model obtains a profit of 37.5% using this strategy. Again, Liverpool loosing to Wigan is the biggest contributor to the profit, but now it is accompanied by a total of 36 other winning bets (and 95 loosing bets). Figure 1 shows the robustness of the results on a round-by-round basis. We can see that the best round gave a profit of approximately 450%, and that all bets are lost during five rounds.

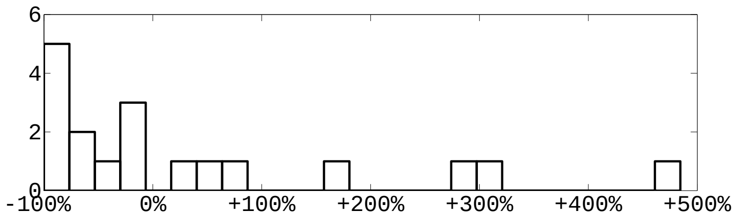

Figure 1. The histogram shows the spread of gain per round. A gain of -100% means all stakes that round are lost (this happens 5 times). The system lost money during 11 rounds of the 19 rounds, won money in 7 rounds, and preferred not to bet in one round.

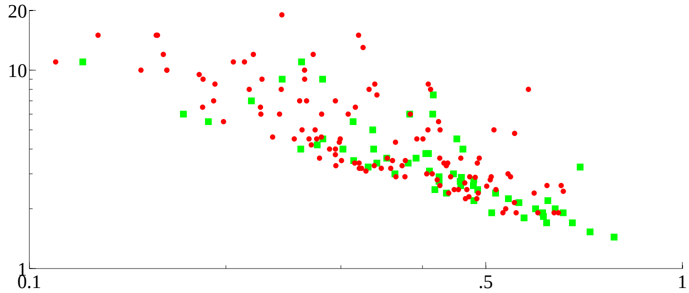

Figure 2. The odds of each game vs the calculated probability of winning. Only games that lead to betts being made are shown. Winning bets are shown as green squares, loosing bets are red circles.

Figure 2 shows how each game is evaluated by DATAINTENSIVE. The calculated probability of a particular outcome is given on the horizontal axis, and the corresponding odds offered by the book-maker is given on the vertical axis. Note the log-scale. Only the feasible bets are included in the figure, thus each datapoint is above the diagonal. Winning bets are given as green squares, loosing bets are shown as red circles. We can see that the system is able to pick out winning bets also amongst the low-probability results, but apparently is prone to over-estimating the probability of the events of low-to-medium probability.

Figure 3 shows how the money management affects the results. All bets being made by the DATAINTENSIVE model using the intermediate Markowitz strategy are shown in the plot; calculated probability of a particular outcome is given on the horizontal axis and the gained credits are given on the vertical axis. All bets above the x-axis thus correspond to winning bets, and those below the x-axis are loosing bets.

Finally, Figure 4 shows the accumulated winnings by the DATAINTENSIVE model using the intermediate Markowitz strategy as a function of the game-round. The system has a loosing behavior until Round 26, then accumulates a profit over the last 12 rounds of the season.

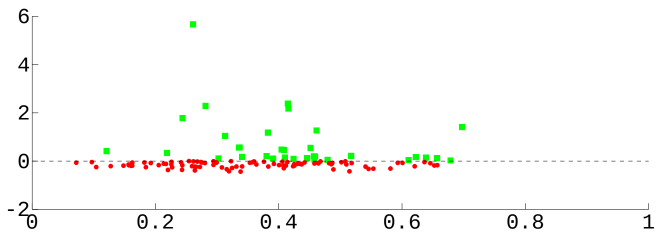

Figure 3. The gain of each bet vs the calculated probability of winning. Only games that lead to betts being made are shown. Winning bets are shown as green squares (positive values), loosing bets are red circles (negative gains).

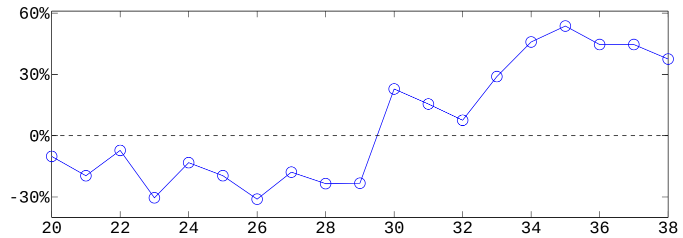

Figure 4. The accumulated gain per round of the Premier League 2011-2012 for the DATAINTENSIVE equipped with the Intermediate Markowitz betting strategy. The values on the y-axis give accumulated winnings divided by the accumulated bets being made. Values above zero indicates a winning strategy. At the end of the season, the system has a profit of 37.5%.

4.2. Results from the 2012/2013 Premier League

Results for the betting season 2012/2013 are given in Table 3.[6] These results are unfortunately far less impressive than those of the previous season, and only DATAINTENSIVE is able to generate a profit (and that is only in combination with the *Variance Adjusted* money management strategy). Note that William Hill had an average profit margin of 6.1% during this season, thus loosing less than 6.1% is "decent". Most of the results obtained by DATAINTENSIVE in Table 3 are at this level, albeit without making a profit.

	Betting strategy			
Model	Fixed bet	Fixed return	Kelly	Variance Adjusted
STATIC	-23.7%	-24.9%	-27.8%	-21.2%
DYNAMIC	-17.1%	-20.0%	-22.9 %	-15.9%
DATAINTENSIVE	**-6.3%**	**-0.7%**	**-3.4%**	**0.4%**

Table 3. Results for the Premier League, 2012-2013

A bookmaker's goal is to make sure that he will have a profit whatever the outcome of a game. This is obtained by ensuring that sufficient amounts of bets are put on all results, thereby "balancing the books". If one result is over-sold, the bookie will react by reducing the odds of that outcome and increasing the odds of the under-sold ones.

[6]Results using the Markowitz strategies, which are similar but even poorer, were left out due to lack of space.

Therefore, bookmakers sometimes value a game differently, and by shopping around for the best odds (instead of only using a single bookie) will improve the results.[7] Utilizing the odds differences, DATAINTENSIVE with the Variance Adjusted money management strategy obtains a profit of 9.1% (which should be compared to the average "profit margin" of the combined bookmaker at 0.3%).

5. Conclusions and further work

In this paper we have considered the use of different money management techniques to evaluate and compare some classical methods for predicting the results of football games. Additionally, we argued that one should try to make models that incorporate more of the available information describing the games to improve predictions. We have made a first step in this direction. The effect of using information beyond match results to characterize a team's performance is indicated in Figure 5, which shows how the two models DYNAMIC (solid line) and DATAINTENSIVE (dashed line) assess Aston Villa's attacking quality over the season 2011-2012.[8] The curves indicate that DYNAMIC is a "smoothed" version of the estimates by DATAINTENSIVE, and that the latter model is therefore better at quickly reacting to, e.g., a team's loss of form.

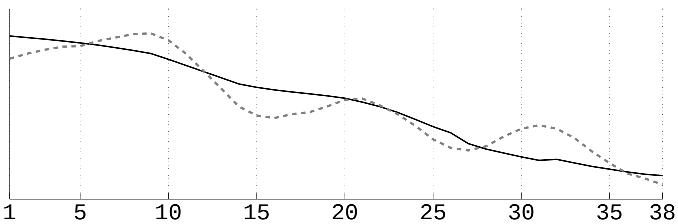

Figure 5. The attach strength for Aston Villa over the season 2011-2012. The results of DYNAMIC is shown with a solid black line, whereas the results of DATAINTENSIVE is shown with a dashed grey line.

Although the picture is somewhat unclear, the DATAINTENSIVE model seems to find better bets than the simpler ones do (Table 1, Table 2, and Table 3). The results are even more debatable when the models are evaluated using standard statistical techniques, see Table 4. There is also a discrepancy between the results in Table 4, where DATAINTENSIVE is not clearly superior to the others for the the 2012-2013 season and Table 3 where it is. Similarly, Table 4, picks DATAINTENSIVE as the best statistical model for the 2011-2012 season, but this is not as clearly visible from the simulated bets reported in Table 1 and Table 2. Further investigation is therefore needed to better gauge the effect of using data-intensive models.

There are also several other directions for future research that we will consider. We plan to start harvesting more match-data to build even richer models. An excellent source for information is `whoscored.com`, but a careful analysis of which data dimensions

[7] Information about which bookmaker offers the best odds for each game outcome is available from a number of internet-sites. The reported results uses odds captured by `betbrain.com` on the evening prior to the start of a round of games.

[8] As the attack strengths have slightly different interpretations in the two models, we have transformed the results from DATAINTENSIVE linearly to ensure that both curves have same average and variation.

Model	2011/12 season 0/1-loss	Log-loss	Brier	2012/13 season 0/1-loss	Log-loss	Brier
STATIC	0.5105	-0.9815	0.5846	0.5211	-0.9883	0.5874
DYNAMIC	0.5053	-0.9804	0.5849	0.5263	**-0.9872**	**0.5862**
DATAINTENSIVE	**0.5368**	**-0.9685**	**0.5791**	**0.5316**	-1.0041	0.5983

Table 4. Model selection criteria for the Premier League 2011-2012 and 2012-2013. The numbers are 0/1-loss averaged over all games (i.e., accuracy), average log-probability of outcome, and the Brier score.

that are informative must be performed, see for instance [8]. A reasonable next step is to incorporate other less quantifiable effects. What is, e.g., the psychological effect of being in a relegation-battle towards the end-of-season? Our analysis will be in the spirit of [1], but completely data-driven. Similarly, one may consider to quantify the effect of an injury or suspension of a particular player. A more fundamental problem is to use historical data to understand and recognize mismatches that can be utilized tactically by one of the teams, and thereby improve the predictions further. We would also like to look into how the models can be bootstrapped to be useful also at the *beginning* of a new season. The relevance of the historical data is questionable at that time, e.g., due to ins and out during the transfer window, and other information must be used to enhance the predictions. Finally, season-bets is an interesting topic: While the result of a single game is somewhat unpredictable, the results of a full season is potentially easier to foresee. Placing a bet on, for instance, who will win the league a couple of months before it is decided may therefore be more beneficial than putting money a single game.

Acknowledgements

Thomas Rene Andresen, Damian Dubicki and Martin Belgau Ellefsrød contributed to this work through their master theses.

References

[1] Anthony C. Constantinou, Norman E. Fenton, and Martin Neil. pi-football: A Bayesian network model for forecasting association football match outcomes. *Knowledge-Based Systems*, 36:322 – 339, 2012.

[2] Mark J. Dixon and Stuart G. Coles. Modeling association football scores and inefficiencies in the football betting market. *Applied Statistics*, 46:265–280, 1997.

[3] Wally R. Gilks, Sylvia Richardson, and David Spiegelhalter. *Markov Chain Monte Carlo In Practice*. Chapman and Hall/CRC, 1995.

[4] Dimitris Karlis and Iovannis Ntzoufras. Statistical modelling for soccer games. In *Proceedings of the fourth Hellenic-European Research on Computer Maths and its Applications, HERCMA'98*, pages 541–548, 1998.

[5] John Larry Kelly. A new interpretation of information rate. *IRE Transactions on Information Theory*, 2(3):185–189, 1956.

[6] Mike J. Maher. Modelling association football scores. *Statistica Neerlandica*, 36(3):109–118, 1982.

[7] Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

[8] John Oberstone. Differentiating the top English premier league football clubs from the rest of the pack: Identifying the keys to success. *Journal of Qualitative Analysis in Sports*, 5(33):Article 10, 2009.

[9] Martyn Plummer. JAGS: A program for analysis of bayesian graphical models using gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, 2003.

[10] Håvard Rue and Øyvind Salvesen. Prediction and retrospective analysis of soccer matches in a league. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 49(3):399–418, 2000.