Learning The Value of Teamwork to Form Efficient Teams

Ryan Beal and Sarvapali D. Ramchurn
Univeristy of Southampton

Abstract

In this paper we describe a novel approach to player valuation and team formation based on players influence within a team. Specifically, we propose a model of teamwork that considers outcomes of passages of plays (interactions of passes between players). Based on our model, we devise a number of metrics to capture the contribution of players and pairs of players. This is then used to learn the value of teamwork from historical team performance data. We apply our model to predict team performance and validate our approach using real-world team performance data provided by StatsBomb. Our model is shown to better predict the real-world performance of teams by up to 46% compared to models that ignore the strength of interactions between players.

1 Introduction

Due to the number of different positions and roles within a football team, it can be a challenge to value how each player contributes to the team performance. This normally leads to player ratings and tactical decisions being made due to subjective opinions through human decision making. There are techniques from Artificial Intelligence (AI) which can help us to value players and form teams which could be used for manager suggestions or team contribution values for the media and team.

The Team Formation (TF) problem that is presented by football is similar to the concept that underpins many multi-agent systems where heterogeneous agents with individual properties (e.g., roles, capabilities, costs) come together to undertake tasks. TF involves the evaluation of different sets of agents in order to determine how well they will, individually or collectively, perform their tasks. By so doing, it is then possible to pick sets of agents that form the most effective teams. This means that we can use similar techniques from this domain to the football TF problem. A similar real-world example where these techniques are also used are shown in teams of emergency responders which are formed based on individual agent's abilities to navigate a difficult environment or address threats (Chalkiadakis and Boutilier 2012). Similarly, in ride-sharing settings, groups of riders can be efficiently formed to minimise travel time and costs (Bistaffa et al. 2017b).[1] In this paper we apply these types of techniques to the problem presented by football where we can value players and sets of players contribution to a team and then form a optimal team based on the values that we calculate.

Against this background, we propose a novel approach to valuing players contribution in a team and forming teams using patterns that appear in a network of interactions between players. We are then able to validate and evaluate our approach by applying our models and algorithms to a real-world team formation problem presented by football. We show that our teamwork-focused model outperforms other player-focused approaches at predicting the teams that would be chosen by human-expert managers across games across 2 seasons in the English Premier League (EPL) and from the 2018 FIFA World Cup. We also show that our model is better at predicting the performance of teams from real world data. Thus, this paper advances the state of the art for research in football in the following ways:

1. We propose a novel approach to team formation based on the value of players team-work. Specifically, we propose a model that considers the outcomes of passages of play.
2. Based on our model, we propose a number of metrics to capture the contributions of individual players and pair of players to the team and the game outcomes.
3. We show how the value of teamwork can be learnt from data and then applied to the prediction of team performance.

When taken together, our results establish benchmarks for team formation algorithms in football and gives a new method to value players within the team.

The rest of this paper is organised as follows. In Section 2 we review the literature, while Section 3 defines the basic definitions and how we apply these to the problem. Section 4 discusses how we model the passage of plays and Section 5 provides the detail of methods that we use to value the players in the team. Section 6 shows how we form optimal teams and then we perform a number of Experiments on our

[1]Note that TF is different to coalition formation in terms of its focus on inter-agent interactions and non-selfish agents.

model in Section 7 and the findings are discussed in Section 8. Finally, Section 9 concludes.

2 Related Work

Boon and Sierksma (2003) discuss the design of optimal teams and calculates the value-added from new team members. Following on from this, (Vilar et al. 2013) look to understand how players' and teams' strategies result in successful and unsuccessful relationships between teammates and opponents in the area of play. There have also been applications to form optimal teams for fantasy football using an MIP and performance predictions in (Matthews, Ramchurn, and Chalkiadakis 2012). Our models differ from the previous work as we model the team as a weighted-directed network of agents and value players based on their influence on the team and the teamwork between players. We then form teams using a novel algorithm with MIP techniques.

There are examples in multi-agent systems literature which form teams using analysis of agents within a network. Gatson and DesJardins (2005), propose a number of strategies for agent-organised networks and evaluate their effectiveness for increasing organisational performance. The same authors also present an agent-based computational model of team formation, and analyse the theoretical performance of team formation in two simple classes of networks - ring and star topologies (Gaston and DesJardins 2008). Recently, (Bistaffa et al. 2017b) proposed a cooperative game theoretic approach to deal with the problem of social ridesharing. They first formed a social network representation of a set of commuters, then proposed a model to form the coalition and arrange one-time rides at short notice. The authors model their problem as a Graph-Constrained Coalition Formation (GCCF) (Bistaffa et al. 2017a). Their model is based on two principles, first they solve the optimisation problem for making coalitions while minimising the cost of the overall system. The set of feasible coalitions in their model is restricted by a graph (i.e., the social network). Secondly, they address the payment allocation aspect of ridesharing.

To our knowledge, none of the discussed approaches have looked at directed interactions between team members (such as passes) and how chains of interactions lead to different team outcome events. More importantly, these approaches have only been validated on synthetic data. Instead, our work is validated on granular data about team performance in real-world games involving teams of humans shown by football.

3 Basic Model Definitions

In football, a manager/coach selects a team of 11 players from a squad of 25 in the EPL or 23 in the World Cup. The objective is to select a team with the highest chance of winning a game. Against this background, we define the squad of players as our set of agents $\mathcal{A}$, the interactions $\mathcal{I}$ are the passes between the players in earlier games, and the graph G represents the network of passes between all the players in the squad. The passage of play $\mathcal{P}$ is made up from a number of passes which represent interactions between players. In football, a passage of play is ended by some event (e.g., tackle, shot, goal, miss and ball out of play). We characterise events into 4 possible outcomes, $\mathcal{E} = \{e_1, e_2, e_3, e_4\}$, where e_1 is a Goal, e_2 a shot on-target, e_3 a shot off-target and e_4 is a loss of possession. We are then able to learn the weights α_i for each outcome. In this case we assume $\alpha_1 > \alpha_2 > \alpha_3 > \alpha_4$. Using the model discussed in the Section 5 we calculate the value of each player $v(a_i)$ and form an optimal team based on the values considering the specific positional constraints of a football team. An example of a passage of play is shown in Figure 1 where the red arrows represent the passes between players and the blue arrow represents the outcome of the walk which in this case was a shot on target.

There are some positional constraints that are specific to football, making it more complex than a standard team formation problem. Each team in a game of football must have 1 goalkeeper and 10 outfield players which are formed from defenders, midfielders and strikers. In most positional formations in football there are between 3-5 defenders, 3-5 midfielders and 1-3 strikers. An example formation is 4-4-2 which is 4 defenders, 4 midfielders and 2 strikers.

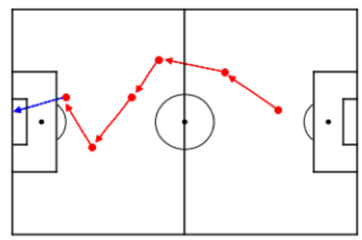

Figure 1: Example of a Walk in Football

4 Modelling Passage of Plays

Players will interact sequentially with each other (i.e., player a_x interacts with a_y who in turn interacts with a_z). To this end, we define a passage of play as a sequence of passes between players in the data. Another example of a similar sequence of interactions in a real-world application would be the movement of a data-packet through a mesh network where the packet moves from router to router until it reaches the destination.

A passage of play leads to an event of a specific type. For example, a player scoring a goal at the end of a sequence of passes or a data packet being used to complete a file download. There may be many different event types. Formally, the set of possible events $\mathcal{E}$ is defined as $\mathcal{E} = \{e_1, e_2, \ldots, e_k\}$ where e_κ is the event and k is the number of possible events from the walk.

Each of the possible events e_κ may have a different impact on the environment and therefore affect the overall performance of the team. Thus, for each $e \in \mathcal{E}$, the value function $\alpha : \mathcal{E} \rightarrow [0, 1]$ determines an associated value. For example, in a game of football if the event e_κ is a "goal" event, this will have a bigger impact on the overall outcome of the game and team performance in comparison to if e_κ is a "loss of possession" event .

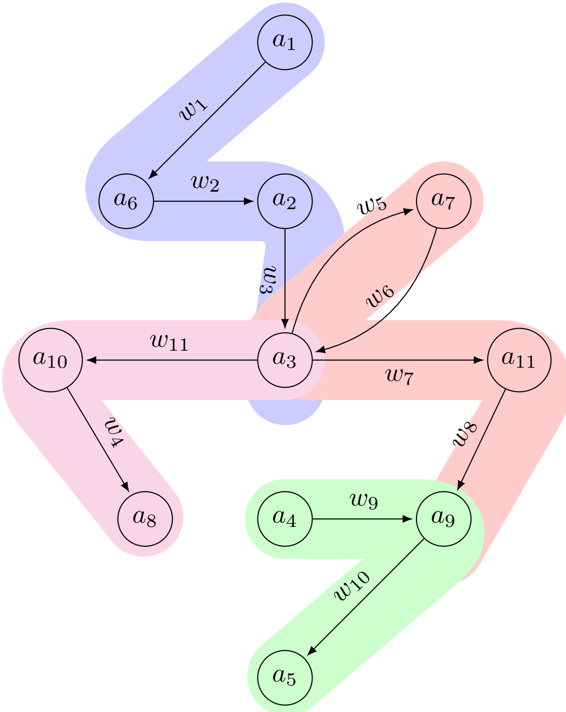

Figure 2: An example of 4 passage of plays through a sample graph of 11 players for an event e_k. The directed edge between two vertex represents the interaction between them and each highlighted colour represents a passage of play.

Note that each passage of play originates from one player and involves chains of directed interactions between pairs of players, resulting in an event. Hence, we next propose methods to extract the contribution of each player as well as sets of players to these individual events.

5 Valuing Players in the Team

Passage of plays $\mathcal{P}$ and associated events $\mathcal{E}$ can be used to infer the value of players and sub-sets of players within the team. We propose two metrics to value the contribution of individual players and sets of players as follows:

- Centrality: $v_{cent} : \mathcal{A} \to \mathcal{Z}$ refers to the sum of the weight of edges incident (incoming and outgoing) to a_i. This measures the influence of player a_i in the network. Thsi is used to value individual players.
- Frequency: $v_{freq} : 2^{\mathcal{A}} \to \mathcal{Z}$ refers to the number of times an player or a subset of players appears in all passage of plays. This represents the influence of an agent in the team. This is used to value sub-sets of players within the team.

It is important to note that these metrics attempt to summarise team performance in different ways, each with a different degree of information loss. Using centrality results in the most loss of information as it ignores who the interactions are made with. Using walk frequency considers all pairwise interactions that lead to specific events, and as we show later, is more representative of teamwork and can be used to predict the performance of teams more effectively.

Now, for each event, we will have different values for each metric for each player or sets of players (as for frequency). However, each event has a different impact on team performance (e.g., goals lead to a win, loss of possession likely to lead to a loss), and to determine the contribution of an player or subset of players to team performance, we need to learn the impact each metric has on the team's performance (discussed in Section 5.3). We assume that each event is independent and therefore use a weighted sum of the values for each of the possible events. This is shown in Equation 1.

$$v_m(a_i) = \sum_{k=1}^{K} \alpha_k v_m(a_i|e_k) \tag{1}$$

where, $v_m(a_i)$ is the value of a_i using the metric m, K is the number of possible events, α_k is the weight of the event e_k (which is learned from the data) and $v_m(a_i|e_k)$ is the value of a_i given the event e_k. We next expand on the above metrics in the following sub-sections.

5.1 Network Centrality

Here we value an player a_i based on their centrality in the network. This value is equal to the sum of the weights of the edges incident to node a_i (both incoming and outgoing edges). For example, in the graph shown in Figure 2, $v_{cent}(a_9|e_k) = w_8 + w_9 + w_{10}$. Equation 2 shows the value calculation using the centrality metric for any player a_i for the given event e_κ:

$$v_{cent}(a_i|e_\kappa) = \sum_{a_j \in \text{Adj }(a_i)} w(a_i, a_j) + w(a_j, a_i) \tag{2}$$

5.2 Frequency

The frequency of an ordered set of players $A' = [a_x, a_y, ..., a_z]$ based on the passage of play frequency is the number of times the player(s) A' appears in all the passage of player. For example, in Figure 2, suppose the passage of play $[a_4, a_9, a_5]$ occurs three times, the passage of play $[a_7, a_3, a_7, a_3, a_{11}, a_9]$ occurs four times, $[a_3, a_{10}, a_8]$ occurs twice and $[a_1, a_6, a_2, a_3]$ occurs once. In this case the value of $A' = [a_9]$ will be $v_{freq}(A'|e_k) = 3 + 4 = 7$. The same reasoning can be applied to subsets of players.

We can compute such a metric for all sub-sets of players in the passage of play. Given a passage of play $\mathcal{P}$ of length l, the number of sub-sets of players of length j constructed from the play $\mathcal{P}$ is calculated by picking the consecutive $j+1$ vertices in the play $\mathcal{P}$. The total number of sub-sets of length l is $\sum_{j=1}^{l-1}(l-j+1)$ and extracting such sub-sets from each play is relatively straightforward. We focus on sub-sets involving pairs of players as we combine such pairs in a combinatorial optimisation algorithm to consider chains of passes. We next describe how we will learn the weights of events α_k to compute Equation 1.

5.3 Learning Event Weights From Data

As shown in Equation 1 we must learn weights for the event outcomes. To learn the set of weights $\mathcal{W}$, which correspond to the impact of the possible walk events $\mathcal{E}$, we use a Logistic Regression algorithm (Hosmer Jr, Lemeshow, and Sturdivant 2013). This allows us to extract the coefficient weights of each of the input features i.e., the weight α_κ (which corresponds to an event e_i) which is used to calculate the final value $v_m(a_i)$ for each agent or sub-team of agents.

Hence, for an outcome y (e.g., a team wins a match), the probability that an player a_i contributes to this outcome is dependent on the individual events (e_κ) to which an agent contributes, as captured by the metrics computed in the previous section. This can be summarised as per the linear combination in Equation 3.

$$p(y|a_i) = \alpha_0 + \alpha_1 v_m(a_i|e_1) + ... + \alpha_k v_m(a_i|e_k) \quad (3)$$

The result of running the logistic regression algorithm is the set of weights $\alpha_\kappa \forall e_\kappa \in \mathcal{E}$. Given this, we can now compute efficient teams according to the learnt measures.

6 Forming Teams

We use two methods to form efficient teams using values that are calculated in the previous section. Firstly, we form teams based on the values of singleton agents. Secondly, we form teams based on the value of pairs that the agents appear in, so that teams are formed between agents who communicate and work well together.

6.1 Individual Player Approach

To form the efficient team based on individual player values, we use the values $v(a_i)$ (depending on the value metrics m) for each player a_i. We use these values alongside constraints over players' positions to form the optimal team. The approach we use to solve this is an edited version or the MIP approach shown in Equation 4. Where we maximise $\Sigma_{n=1}^{N}(V(p_n) \cdot z_n)$ and do not consider the pair decision variable x_i. The other constraints remain the same. This gives a combinatorial optimisation problem (knapsack packing problem) that can be solved using standard mixed-integer programming (MIP) techniques. Similar methods are also used in (Pochet and Wolsey 2006; Fitzpatrick and Askin 2005; Matthews, Ramchurn, and Chalkiadakis 2012).

6.2 Team-Centred Approach

Using the values of the player pairs we form teams using the MIP formula presented in Equation 4. When forming teams we ensure that all the pairs of players are part of the same squad and can be selected together. We also consider the positions of the players so that we pick a team in a reasonable positional formation. This is represented by position range constraints.

$$\begin{aligned}
\underset{x}{\text{maximise}} \quad & \Sigma_{i=1}^{I}(V_1(p_i) \cdot x_i + \beta V_2(p_i) \cdot x_i) \\
\text{subject to} \quad & \Sigma_{n=1}^{N}(z_n) = 11 \\
& z_a = x_i, z_b = x_i \\
& \Sigma_{i=1}^{N}(gk_n \cdot z_n) = 1 \\
& 3 \leq \Sigma_{n=1}^{N}(def_n \cdot z_n) \leq 5 \\
& 3 \leq \Sigma_{n=1}^{N}(mid_n \cdot z_n) \leq 5 \\
& 1 \leq \Sigma_{n=1}^{N}(str_n \cdot z_n) \leq 3
\end{aligned} \quad (4)$$

where a binary decision variable x_i represents the selected pairs and z_n represents whether a player is picked or not. There is then a number of binary sets for each position (gk, def, mid and str) containing if a player plays in the corresponding position or not and we aim to maximise the pair values V_1 and V_2 in the selected team; β is set to 0.05.[2]

In more detail, we aim to form a team of N agents using two binary decision variables: x_i which represents if a pair of agents is selected and z_j represents if an agent is picked or not. The agent decision variables z_a and z_b represent the two agents in pair i and must equal the decision variable x_i. We aim to maximise the sum of $V_1(p_i)$ (the value for p_i using the agent pair values we have calculated) and $V_2(p_i)$ which represents the *interactional alignment* (the value of pair p_i which is calculated by Equation 5). This value is weighted by β which can be learned form the data.

$$V_2(p_i) = \sum_{K}^{k=1}(V_1(p_k) \cdot x_k)_{\{p_i \cap p_k\}} \quad (5)$$

Equation 5 represents the sum of all pair values where there is an overlap with pair p_i. By maximising this, it allows us to increase the strong links between selected pairs while decreasing the weak links.

7 Evaluation

To evaluate our model for team formation we use a dataset collected from the past two seasons in the English Premier League (2017-19) as well as comparing results to the 2018 FIFA World Cup.[3] The dataset breaks down each game in the tournaments into an event-by-event analysis where each event gives a number of different metrics including: event type (e.g., pass, shot, tackle etc.), the pitch coordinates of the event and the event outcome. This type of dataset is industry-leading in football and used by top professional teams. Thus, we believe that this is a good, real-world, dataset with the richness and challenge appropriate to rigorously assess the value of our model.

7.1 Experiment 1: Player Value Outputs

In this experiments we show the top 10 players/pair of players which have been calculated using our new metrics. Figures 3 and 3 shows the top individual players from the past 2 EPL seasons. Following on from this Figures 5 and 6 show the top pairs of players.

These results show that using our valuations the top players and pairs are dominated by Manchester City players (who won the league in both of the leagues we tested in). This is likely due to the number of players involved in the build up play which is fundamental to their style of play. This may suggest that these metrics are best to compare the players in the same team and play in the same style so could be used by the teams to help form teams and identify weaker areas that could be improved in the transfer window.

[2]This is calculated using a brute force method.

[3]All data provided by StatsBomb - www.statsbomb.com.

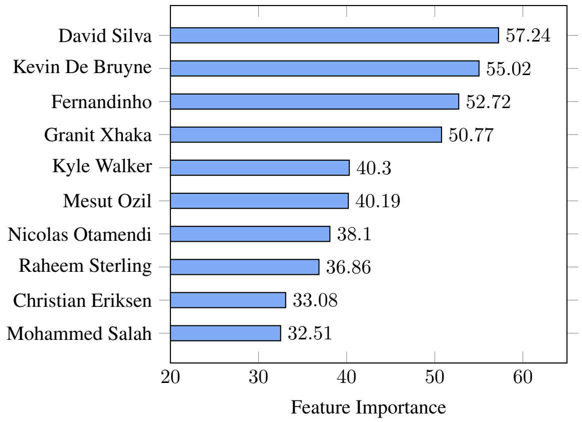

Figure 3: 2017/18 EPL Individual Players.

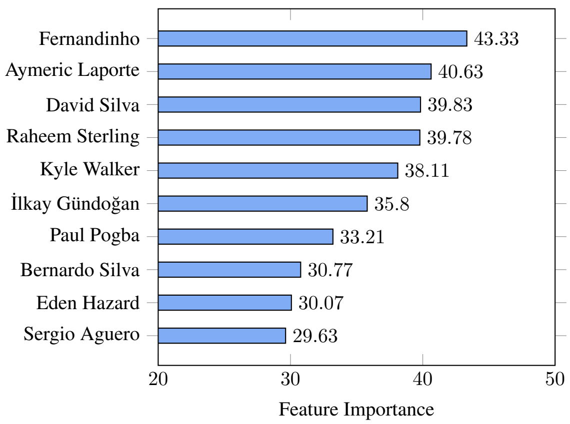

Figure 4: 2018/19 EPL Individual Players.

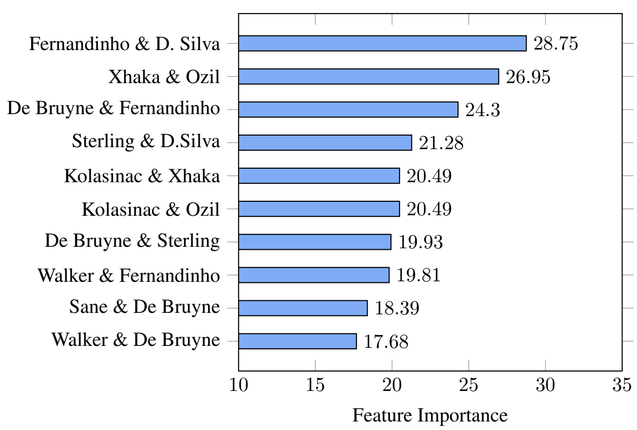

Figure 5: 2017/18 EPL Pairs of Players.

7.2 Experiment 2: Performance Comparison to Human Formed Teams

We evaluate our model using all games from the 2017/18 and 2018/19 EPL Season as well as the 2018 FIFA World Cup. We compare both the individual player value approach and

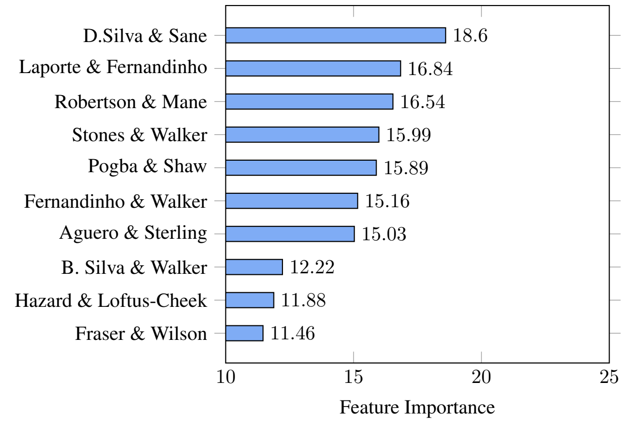

Figure 6: 2018/19 EPL Pairs of Players.

the pairs approach with the teams selected by the human-expert manager (focusing on both the starting 11 player). The results are presented in Figure 7 (where error bars represent a 95% confidence interval).

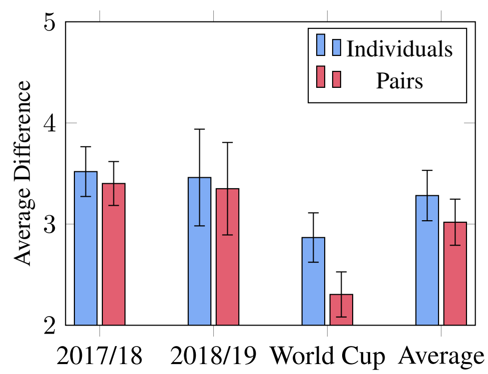

Figure 7: Average Difference Between Models and Human Manager (where lower is better).

This shows that the pair values optimisation method gives the closest teams to the human experts on average with a difference at the lowest of 2.3 per game for the starting team. This suggests that the human managers (either consciously or sub-consciously) consider the ability of players in the team to work together as the other methods only consider individual player values. At an average of 2.3 this could give managers suggestions of how changes could be made to the team that may give a better chance of winning the game. It is worth noting that the performance of the algorithms for the World Cup is better than over a whole EPL season. This may be due to the standard rotation which happens over a whole season, as well as there being less injuries leading to line-up changes in the World Cup as we do not incorporate injuries in our formation.

7.3 Experiment 3: Outcome Prediction Comparison

We see that in Figure 8 there is a positive correlation (p-value = 0.0011) between the teamwork values and the number of goals scored by the teams and we see similar results for the correlation of the team value for other metrics. Hence, to evaluate the strength of our methods, we use the valuations as a predictor of the actual real-world performance of the selected teams. We focus on match outcomes and other team performance metrics.

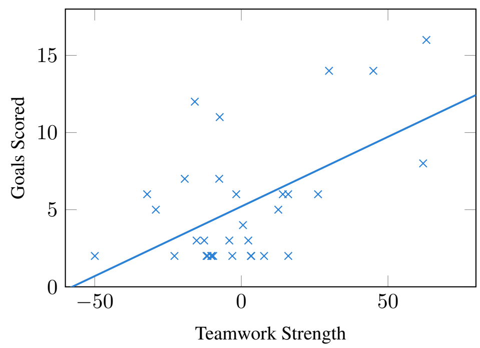

Figure 8: Correlation between team pair values and goals.

The results in Figure 9 and Table 1 suggest that the teamwork metric is a more accurate predictor of performance than individual player values, meaning that the teams with higher valued pairs are more likely to win the game and have better performance indicators. This is especially true when we predict the number of passes that a team will make in a game as this metric shows the strongest predictor when using the teamwork values and is a 46% increase on any other approach.

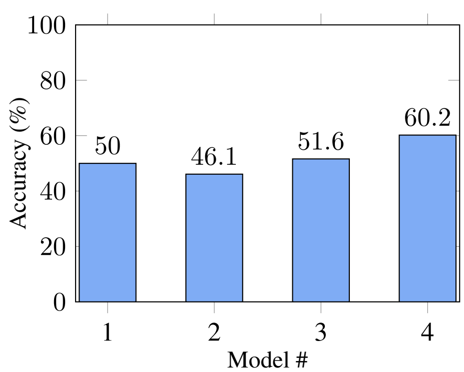

Figure 9: Accuracy of Valuation Methods For Outcome Predictions.

8 Discussion

We focus on forming teams based on individual player values and pair values. In further work this could be extended to look at larger sub-sets of player values and form teams based on those. However, we decided to focus on using the

Model #	Individuals	Pairs
Shots	4.33	**3.74**
Goals For	1.00	**0.87**
Goals Against	1.14	**1.06**
Passes	105.69	**57.07**

Table 1: Valuation Methods Root Mean Squared Errors for Performance Metrics (where lower is better).

pair values to show the teamwork so that we can easily identify the pairs of players that have a strong impact on team performance and the outcomes. This also allows us to calculate how the selected players will affect the rest of the team and therefore the overall teamwork of all players.

We choose to test our team formation methods by comparing the outputs to that of a human-expert team manager who selects the real-world team. Our results show that our model is able to form teams which are similar to the selections of human-experts and, that we are able to suggest a small number of changes that could improve the team. This comparison also suggests that human-experts consider the teamwork between players when selecting their teams (this may be sub-consciously).

Building on our models, in further work we would further evaluate the predictions of match-outcomes, based on our team-valuations of a starting 11 team, against other match-outcome prediction approaches such as (Dixon and Coles 1997; Constantinou, Fenton, and Neil 2012). We would also extend the models to address how the team formation could be improved by factoring in an opposition team (in games such as football this can have a significant difference to how a team is formed). Our results also suggest that this model could be applicable across a number of domains and, given a high-quality dataset, we could further validate the model performance to see if similar results are found (e.g., in emergency response or data transfers).

9 Conclusion

In this paper we have described a novel approach to team formation based on interactions between players. Our model of teamwork considers passage of play event outcomes. We defined and tested a number of metrics to value the contribution of players and sets of players and show how the value of teamwork can be learnt from data and then applied to predict the performance of teams. We tested and validated our models of valuing agents and forming teams by applying our models to problems posed by football and using StatsBomb data from the English Premier League and the 2018 FIFA World Cup. We showed that our model is able to produce similar team selections to an human-expert manager while also being able to suggest changes to the team. We also showed how our valuation methods are an effective predictor of the key team performance metrics in football.

References

Bistaffa, F.; Farinelli, A.; Cerquides, J.; Rodríguez-Aguilar, J.; and Ramchurn, S. D. 2017a. Algorithms for graph-constrained coalition formation in the real world. *ACM Transactions on Intelligent Systems and Technology (TIST)* 8(4):60.

Bistaffa, F.; Farinelli, A.; Chalkiadakis, G.; and Ramchurn, S. D. 2017b. A cooperative game-theoretic approach to the social ridesharing problem. *Artificial Intelligence* 246:86–117.

Boon, B. H., and Sierksma, G. 2003. Team formation: Matching quality supply and quality demand. *European Journal of Operational Research* 148(2):277–292.

Chalkiadakis, G., and Boutilier, C. 2012. Sequentially optimal repeated coalition formation under uncertainty. *Autonomous Agents and Multi-Agent Systems* 24(3):441–484.

Constantinou, A. C.; Fenton, N.; and Neil, M. 2012. pi-football: a bayesian network model for forecasting association football match outcomes. *Knowledge-Based Systems* 36:322–339.

Dixon, M. J., and Coles, S. G. 1997. Modelling association football scores and inefficiencies in the football betting market. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 46(2):265–280.

Fitzpatrick, E. L., and Askin, R. G. 2005. Forming effective worker teams with multi-functional skill requirements. *Computers & Industrial Engineering* 48(3):593–608.

Gaston, M. E., and DesJardins, M. 2005. Agent-organized networks for dynamic team formation. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, 230–237.

Gaston, M. E., and DesJardins, M. 2008. The effect of network structure on dynamic team formation in multi-agent systems. *Computational Intelligence* 24(2):122–157.

Hosmer Jr, D. W.; Lemeshow, S.; and Sturdivant, R. X. 2013. *Applied Logistic Regression*, volume 398.

Matthews, T.; Ramchurn, S. D.; and Chalkiadakis, G. 2012. Competing with humans at fantasy football: Team formation in large partially-observable domains. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*.

Pochet, Y., and Wolsey, L. A. 2006. *Production planning by mixed integer programming*.

Vilar, L.; Araújo, D.; Davids, K.; and Bar-Yam, Y. 2013. Science of winning soccer: Emergent pattern-forming dynamics in association football. *Journal of Systems Science and Complexity* 26(1):73–84.