

Complex networks untangle competitive advantage in Australian football

Calum Braham^{1, a)} and Michael Small^{1, 2, b)}

¹⁾*School of Mathematics and Statistics, The University of Western Australia - Crawley, WA, 6009, Australia*

²⁾*Mineral Resources, CSIRO - Kensington, WA, 6152, Australia.*

We construct player-based complex network models of Australian football teams for the 2014 Australian Football League season; modelling the passes between players as weighted, directed edges. We show that analysis of these measures can give an insight into the underlying structure and strategy of Australian football teams, quantitatively distinguishing different playing styles. The relationships observed between network properties and match outcomes suggest that successful teams exhibit well-connected passing networks in which the passes are distributed between all 22 players as evenly as possible. **Linear regression models of team scores and match margins show significant improvements in R^2 and BIC when network measures are included in addition to conventional measures, demonstrating network analysis contain useful, extra information. Several measures, particularly the mean betweenness centrality, were shown to be useful in predicting the outcomes of future matches, suggesting it measures some aspect of the intrinsic strength of teams.** In addition, several local centrality measures were shown to be useful in analysing individual players' differing contributions to the team's structure.

Australian football was developed in the late 1850's as a way for cricketers to keep fit in the winter. It is a fast-paced, free-flowing game that is currently one of the most popular sports in Australia. Its highest level of competition, the Australian Football League (AFL), is the most attended and viewed sports league in the country. An Australian football team can be thought of as a complex system, with the players interacting through passes. **The lack of rigid positioning or restrictions on player movement make it a good candidate for modelling passing structures with player-based complex networks. In this work, we develop a set of context-appropriate network properties for analysing such networks and demonstrate their relationship to team and player performance. At present, common statistical analysis of**

AFL matches is based around the number of times particular events occur. For example, the number of times a particular player passed the ball or the number of turnovers committed by each team. Network analysis has the potential to provide measures that give a more quantitative picture of the intrinsic dynamics and strategies that lead to these events occurring. By demonstrating a relationship between particular network properties and the outcomes of matches we can suggest network characteristics that are common to successful teams. Analysing a player's node within their team's network we create quantitative measures of a player's role within their team's passing structure; capturing interactions with other players rather than simply counting the number of effective actions performed. Although the results themselves are naturally specific to the 2014 AFL season, the relationships observed can be seen as a proof of concept that such network analysis methods could be used to provide

^{a)}Electronic mail: calum.braham@research.uwa.edu.au

^{b)}Electronic mail: michael.small@uwa.edu.au

useful information about Australian football, or indeed extended to other team settings.

I. INTRODUCTION

In a game of Australian football two teams of 18 players (with another four on the interchange bench) compete to score as many points as possible. Games are divided into four 20-minute quarters, with the clock stopping every time a goal is scored or the ball goes out of play. The playing fields vary in size but are approximately oval shaped with long axis 165 metres and short axis 135 metres. At each end of the field there are four equally spaced posts, with two larger central posts called ‘goal posts’ and the two outer posts called ‘behind posts’. If the ball is kicked between the two goal posts by the attacking team, a *goal*, worth 6 points, is awarded. If the ball passes through the goal posts in any other way, hits a goalpost or passes between a goal post and a behind post, a *behind*, worth 1 point, is awarded. The ball is approximately ellipsoidal in shape and can be passed between teammates by kicking or handballing (holding the ball with one hand and punching it with the other). Players may also run with the ball down the field. The opposition can gain possession of the ball in a number of ways, including intercepting a pass or tackling the player in possession of the ball. If a kicked ball is caught on the full this is known as a *mark*. The player who caught the ball receives a free kick from their current position and cannot be tackled.¹

Several authors have used complex networks previously as a tool to analyse sports. In 2004 Onody and de Castro² created a network modelling the interactions between Brazilian soccer players and their clubs over their careers. Tripathy, Bagchi and Jain³ and Mukherjee⁴ both model the interactions between cricket players, linking them as nodes on a network if they have scored runs in a partnership together. Radicchi⁵ used network methods to rank professional tennis players while Xu, Hu and Liu⁶ analysed the

interactions of Go players in Asia over several thousand tournaments. Considering sport as a form of improvised dramatic entertainment, Kiley et al.⁷ used network methods to classify the ‘game stories’ of Australian football, showing that AFL matches exhibit a wider variety of narrative motifs than random walks.

Networks similar to those presented here, representing players as nodes and the passes as edges, have been constructed several times previously and used to analyse basketball and association football⁸⁻¹² (henceforth referred to as soccer). Cotta et al.⁸ and Peña and Touchette¹¹, created networks modelling soccer teams during the latter stages of the 2010 FIFA World Cup. Cotta et al. analysed how network metrics such as the number of consecutive passes, passes per minute, clustering coefficient and degree centrality varied across the course of the match for the Spanish team while Peña and Touchette include the measures of betweenness and eigenvector centrality and compare network properties between teams.

Duch, Waitzman and Amaral⁹ analysed the passing networks of soccer teams in the 2008 European Cup as a means to demonstrate a method of quantifying individual contributions to a team activity. Specifically, they used the betweenness centrality of players with regard to their opponents goal as a measure of the performance of individual players. Fewell et al.¹⁰ created passing networks to analyse basketball matches, considering the 2010 NBA play-offs. They compared global network metrics between the teams, finding that clustering and network entropy were most strongly correlated with the team’s performance, and considered the local network metrics, such as degree centrality, with respect to playing position.

This paper is the first to use network models to analyse the passing structures in Australian football. Here we create similar types of network but consider a wider range of analysis measures, particularly focusing on developing those that make use of the weight and direction information of the edges. More significantly, in the former studies the authors had smaller data sets at their disposal; generally only a few matches

for each team. This meant that analysis was based partially on a heuristic comparison between the results of the network analysis and the observed characteristics of those matches. As our data set covers all 207 matches from the 2014 AFL season we are able to statistically justify the relevance of the network measures to Australian football, particularly through their relationships to match outcomes. In this way we are making a trade off by losing the ability to examine specific matches in detail in the hope of performing a more generally applicable and quantitative analysis.

II. DATA SET AND NETWORK CONSTRUCTION

The data set used was obtained from Champion Data, the official data provider to the AFL. It lists each successive event in the movement of the ball for each match in the 2014 AFL season. In essence, the data describes which players had possession of the ball in succession, where they were on the field and how the ball got to its next location. It is useful to partition the data into *possession chains*, which are unbroken strings of passes between members of the same team. A possession chain ends when the opposition takes possession of the ball or if play is stopped and the ball goes to the umpire, even if the original team regains possession immediately thereafter.

The basic type of network created will be referred to as a *passing network*. In such a network the nodes represent the players on one team and the passes between them are shown as the edges. The edges are directed, as a pass is specifically from one player to another, and weighted by the number of passes between the two players. Primarily, the analysis presented here was performed on networks created from the passes made by one team during one match, but it is naturally possible to consider networks from other subsets of the data, for example, all the passes made before a goal was scored.

Networks were also created including an extra node representing a shot at goal being taken. The weight of the edge from a player’s node to

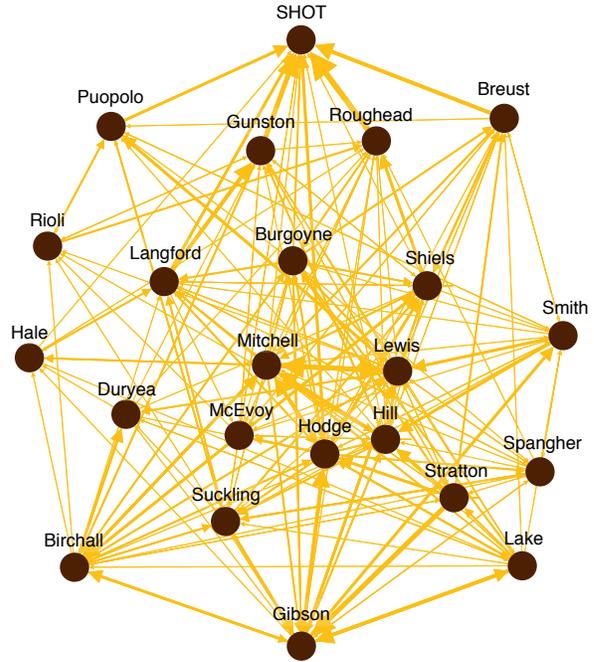


FIG. 1: Shot network for Hawthorn during the 2014 AFL grand final.

the ‘shot’ node represents the number of shots they took. These networks are labelled *shot network* and are intended to represent all intentional methods by which players release the ball. Note that the networks do not include information about how the team gains or loses possession of the ball, only complete passes between members of the one team and shots at goal. Figure 1 shows the shot network created from Hawthorn’s passes during the grand final.

Australian football has added layers of complexity which make it more difficult to model abstractly as a network than basketball or soccer. This is in part because there are more players on the field (18 players on each team with 4 on the interchange bench) and multiple ways the ball can be passed between teammates (kicked, handballed or tapped). There are also regularly several players in close proximity to the ball, with neither team properly gaining possession for any appreciable length



FIG. 2: Congested play in an AFL Match.

of time. This is known as *congested* play. Although the competition to gain clean possession is important to the development of the match, these situations are not captured well in our data set as it is difficult to record the rapid changes of possession. In general, game events when no player takes possession of the ball are not recorded in the data, for example, if multiple players compete to mark the ball unsuccessfully. However, we believe the data set still to be appropriate, as we are concerned with the intentional, strategic passes made by a team in open play and not the more chaotic, congested situations when constructing passing networks. In doing so, we are not intending to analyse the physical methods and techniques of Australian Football, rather how teams choose to distribute the ball through passes so as to best move it towards the goals. Nevertheless, this provides another motivation to compare network results over the full season, as an attempt to prevent the over-fitting of our results to data that imperfectly represents reality.

Considering these limitations, we have chosen to keep the network models simpler, not including specifics such as the method, length or location of pass made. As the significant variability in the match situation of each pass is not captured well in the data, any extra structure would necessarily be somewhat arbitrary. We have also chosen not to model both teams in a match on the same network as turnovers between teams often occur in congested situations.

In other ways, however, Australian football is more suited to network analysis. It is a fluid, free-flowing game without a rigid positional structure and with simple passing and positional rules. Players may move anywhere on the field and when in possession are subject to the same rules at any location. There are no defined roles such as a goalkeeper in soccer. Over the past decades it has developed into a game that values running endurance where a team's positional structure moves with the ball rather than remaining spread. A player's position relative to the other players and the ball is often of more importance than their field position. As such, we feel it is appropriate to use a network model which is centred around the players and treats them all equivalently.

III. NETWORK ANALYSIS METHODS

The network properties that were measured for analysis are described in Table I. Also listed is the type of network used for calculation (with or without the 'shot' node) and a description of how each measure is interpreted in the context of Australian football. Although most of the measures are calculated locally for a single player's node, we are often interested in the global characteristics of the team's network to analyse their overall passing structure and strategy. The final column lists the global properties that were collated for each local measure. Where the measures differ from conventional definitions they are discussed below.

The degree and strength measures are defined as is conventional for weighted, directed networks. Inward and outward directed paths are included equally when calculating the closeness centrality.

The betweenness was calculated using shot networks, such that it also considers shortest paths between other players and a shot being taken. Note that $1/(N-2)^2$ is used for normalisation rather than the conventional $2/((N-1)(N-2))$. The factor of 2 is cancelled as the networks are directed and, as the 'Shot' node has no outward edges, the number of ordered

TABLE I: Network analysis methods

Measure	Network type	Australian football interpretation	Global measure
Out-degree (In-degree)	Shot	Number of different players a particular player passed to (received passes from)	Mean out-degree (Mean in-degree)
Out-strength (In-strength)	Shot	Number of passes and shots made (received) by a player	Mean out-strength (Mean in-strength)
Closeness centrality	Passing	How well-connected and central a player is within their team’s passing structure	Mean closeness
Betweenness centrality	Shot	The extent to which the flow of passes through a team’s network depends on that particular player	Mean betweenness, Number of players with non-zero betweenness
Clustering coefficient	Passing	The extent to which the players a particular player passes to also pass with each other	Global clustering coefficient
Eigenvector centrality	Passing	The asymptotic probability that a particular player has the ball	Standard deviation of eigenvector centralities
Entropy	Shot	The unpredictability in who a particular player passes to	Team entropy (weighted mean of individual entropies by eigenvector centrality)

pairs of other nodes that could be connected by a path becomes $(N - 2)^2$. As the football networks are generally well-connected there is the potential for many of the shortest paths to be on the direct edge between two nodes, not contributing to betweenness values. As such it is possible that many nodes could have zero or low betweenness. To determine if this gave useful structural information the number of players with non-zero betweenness centralities was measured.

In their 2007 paper, Saramäki et al.¹³ present a summary of several methods that generalise the clustering coefficient for weighted networks. **The generalisation we selected is similar to the method Peña and Touchette used in their analysis of the passing networks of soccer teams¹¹. It accounts for edge weights by considering the sum of the geometric means of the edge weights for all triangles a node participates in rather than simply the number of triangles. We generalise this for directed networks by taking the mean of the two edge weights between each pair of nodes. The global clustering coefficient is then defined, similarly to the unweighted case, as the sum of geometric means for all triangles**

in the network, divided by the total number of connected triples.

The eigenvector centrality is calculated by converting the adjacency matrix of the network to a transition matrix and calculating the steady state distribution (or equivalently the principal eigenvector). Note that if a player made no passes during a game their node will have an out-strength of 0 and act as a sink node. To avoid the long term probability being 1 at this node and 0 elsewhere we assume the player in question passed to all others with equal probability.

The transition matrices were also used to calculate Shannon’s entropy¹⁴ for each player, a method used by Fewell et al.¹⁰ to analyse basketball networks. In order to measure the unpredictability of a team’s overall ball movement, we calculate the *team entropy* as a weighted mean of the individual entropies over all players in the team. The weights are given by the eigenvector centralities in the same network with the ‘Shot’ node removed.

Several non-network measures are used for comparison to the network properties described above. Except for the mean chain length and

TABLE II: Non-network measures

Measure	Description
Disposal	A pass made or shot taken by a player
Disposal efficiency	Proportion of disposals made by a player that were effective (received by a teammate or scored a goal)
Shot efficiency	Proportion of shots taken that resulted in goals
Contested possession	When a player wins possession of the ball in a contest with an opposition player
Clearance	Clearing the ball out of congestion after a stoppage in play
Clanger	An unforced error made, e.g. kicking the ball to an opponent
Mean chain length	Mean number of passes in the team's possession chains
Passing rate	Number of passes per minute by a team when in possession of the ball
Passing rate (no shots)	Passing rate with shots at goal excluded as players are allowed extra time when taking a set shot at goal

passing rate, these are all conventional AFL statistics. They are recorded as the number of times each event occurred for the team in question. The measures not previously defined are described in table II. The conventional measures are recorded as the number of times each event occurred for a particular team.

IV. RESULTS & DISCUSSION

The variety of types of passing network that can be created and the number of properties that can be measured necessarily results in a large number of possible permutations for analysis. We divide this section into four avenues of analysis as a proof of concept that these network measures can be used to give useful information about Australian football.

A. Can we quantitatively characterise the styles and strategies of Australian football teams?

We first consider whether it is justifiable that the differences in network measures are quantifying different passing styles and structures. **Considering how the properties of individual networks compare to theoretical bounds may give insight into the network structure, but we have few reference points from other networks in similar contexts to determine what this would mean practically. As such, comparing properties between the networks of different teams or between one team's networks from different matches was considered to give more applicable results than analysing the networks in isolation.**

Table III shows each team's average value for several global measures across the 2014 season. Disposal efficiency, shot efficiency and passing rate are also included as non-network measures that may quantify differences in passing strategies. The teams are presented in the same order as the league table finished, that is, with the better performing teams at the top. By highlighting the upper quartile of each measure in red and the lower quartile in blue we notice that there is a trend for the stronger teams to have higher values for most network measures. This supports the assertion that these measures are distinguishing between the teams and measuring some tangible difference in passing structure. If the measures were not adequately relevant to Australian football we would not expect to see any trend.

The mean out-degree, mean closeness and global clustering all measure how well-connected a team's network is, but in different ways. In teams with a high mean out-degree players pass to more of their teammates during a particular match on average, potentially implying that the team's offensive strategy is more varied. **As this is measured on shot networks it is also influenced by the average number of players who take shots at goal. However, the mean in-degree is unaffected by shots and showed a similar ordering, suggesting this is not the primary effect.** The mean closeness,

TABLE III: Selection of network measures (and some other relevant characteristics) for each team. The teams are ordered as they were in the league table at the season’s end. The top quartile of scores are shown in red and the bottom quartile in blue.

Team	Mean out-degree	Mean betweenness	Team entropy	Mean closeness	Global clustering ($\times 10^{-2}$)	Disposal efficiency (%)	Shot efficiency (%)	Passing rate (no shots)
Sydney	8.69	11.0	3.03	0.880	9.17	71.6	52.3	13.8
Hawthorn	9.34	6.62	3.13	0.932	9.39	74.3	59.7	14.4
Geelong	8.57	9.94	3.02	0.882	8.88	73.1	51.8	13.7
Fremantle	8.63	6.15	3.02	0.853	8.66	72.2	54.7	13.5
Port Adelaide	8.68	6.76	3.04	0.881	8.27	71.2	51.9	13.5
North Melbourne	8.77	5.87	3.04	0.892	8.43	74.0	51.7	13.8
Essendon	9.13	7.69	3.12	0.926	9.55	74.2	51.5	13.9
Richmond	8.64	5.30	3.05	0.876	8.67	73.1	53.7	13.9
West Coast	8.27	6.87	2.95	0.856	8.12	71.7	55.6	13.8
Adelaide	8.63	5.60	3.01	0.863	9.12	72.0	53.0	13.5
Collingwood	8.22	6.92	2.98	0.853	8.80	70.6	50.6	13.6
Gold Coast	7.91	6.91	2.93	0.806	8.22	70.1	53.6	12.6
Carlton	8.09	7.68	2.97	0.826	8.44	71.7	54.5	12.4
Western Bulldogs	8.57	4.72	3.06	0.860	8.99	72.2	51.8	13.9
Brisbane	8.05	3.88	2.98	0.861	7.99	71.6	53.0	12.9
GWS	8.36	4.85	2.90	0.861	7.99	72.3	51.9	13.3
Melbourne	8.49	4.82	3.03	0.853	9.10	73.8	49.9	12.6
St Kilda	8.20	4.04	3.00	0.846	8.72	71.8	51.8	13.1

however, is more strongly influenced by the total number of passes between players in a network and whether all players are connected strongly in the network. The global clustering will be high if there are sets of players who pass regularly between each other and minimal variation in edge weights across the network.

A higher mean betweenness value means that more of the shortest paths between nodes on a network pass through other nodes. This means that a team’s passing network is structured such that certain players are important to connecting other players and maintaining the flow of the ball through the network. Of particular note is the wide range of mean betweenness centralities across the teams, suggesting that this strongly distinguishes between different passing strategies and structures.

The team entropy gives a measure of the unpredictability of a team’s offensive structure. A team that used the same passing chains to move the ball down the field, passing regularly between the same pairs of players, would have a lower team entropy than a team using a wider

variety of passing chains. Entropy will show some relationship to the mean out-degree, as the more players a particular player passes to, the more unpredictable their ball movement. Interestingly, the entropy results for the Australian football teams are quite similar to those calculated for basketball teams by Fewell et al¹⁰, which are in the range 2.93-3.23 Shannons. This may be explained by the fact that although basketball teams have fewer players, they are on a smaller court and so there is generally the possibility for any player to pass to any other player. In Australian football, despite the larger network and thus larger maximum possible entropy (4.52 for our shot networks compared to 3.32 for basketball), the field is bigger and consequently the number of disposal options available to a player may be similar to basketball.

The passing rate represents the speed with which the teams move the ball when in possession. This could be an indication of a tactical choice to ‘play on’ and move the ball more quickly or simply a result of how quickly the next passing option becomes available. **Two**

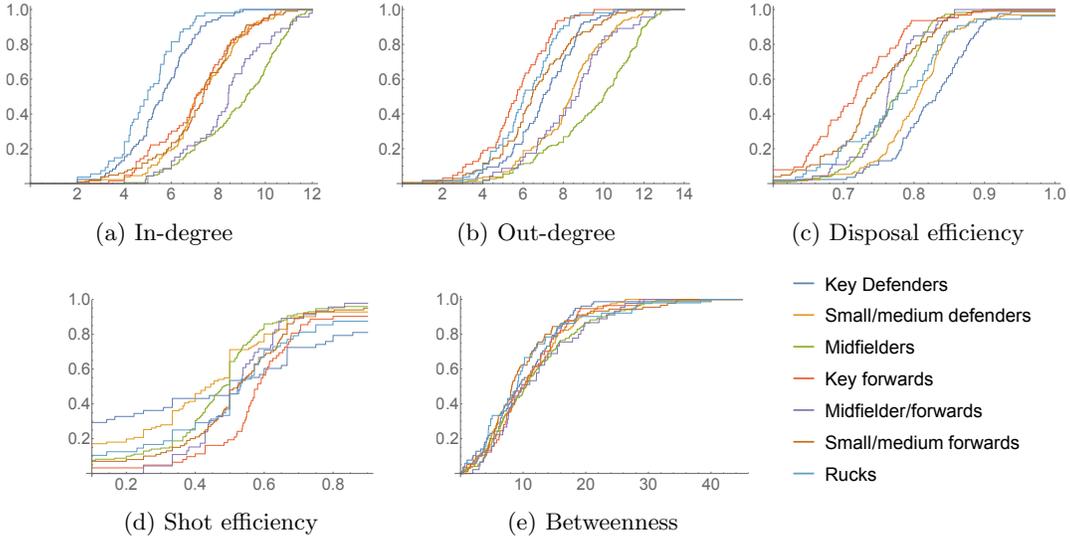


FIG. 3: Empirical cumulative distributions of the mean in-degree, mean out-degree, disposal efficiency, shot efficiency and mean betweenness for players of each nominal playing position.

teams known for their quick movement of the ball, Hawthorn and the Western Bulldogs, have high passing rates, suggesting the measured rates fit with the general perception of the tactics used by the teams. The disposal and shot efficiencies are related both to the skill of the players and the passing strategy used. For instance, a team playing directly and taking more risks would likely have lower disposal efficiency.

Figure 3 shows empirical cumulative distributions of a selection of local centrality measures for players of different nominal playing positions. The position of each player was taken from the AFL player ratings website¹⁵. We note that in modern AFL the role of each position is not rigidly defined and some players may be used in different areas throughout a match. For readers unfamiliar with the AFL, the ‘key’ forwards and defenders are generally tall and often are the focus of a team’s offensive or defensive structure. The ‘rucks’ have the designated role of competing to win possession of the ball after a stoppage in play.

It can be seen that these measures distinguish players of different positions, highlighting the

different roles played in the team’s structure. The in-degree, out-degree, disposal efficiency and shot efficiency all show an ordering that fits with the expected roles of each position. For example, midfielders have higher degree measures as they follow the passage of play around the ground and are more central to the structure, while key defenders and rucks have lower in-degree as they have other roles defined that do not involve receiving a wide variety of passes. Other network measures related to the number of passes received or given showed similar ordering to in-degree or out-degree. Comparing the measures, forwards have relatively higher in-degree and shot efficiency while defenders have relatively higher out-degree and disposal efficiency. This can be explained as defenders give more passes to a wider variety of other players in their role setting up the offensive play from the back of the field, while forwards receive the ball from more players but generally have more specific targets (either taking a shot or passing to another forward in a better position). Defenders may have higher disposal efficiency as they make more passes under less pressure, but

once the ball is in the forward line the opposition will be actively trying to defend each pass. It is interesting to note that the distributions of mean betweenness centralities do not distinguish between positions as significantly as the other centrality measures. This suggests that players in each position are reasonably equally utilised in connecting the team’s network and transporting the ball from the defensive line towards the goals, despite the midfielders having significantly higher scores in the simpler centrality measures.

B. Do any of the network measures show a relationship to match outcomes?

We now consider more directly whether there are characteristics of passing networks that exhibit relationships with match results, with the hope of uncovering the common network characteristics of successful teams. Table IV shows the proportion of winning teams that had higher values of each network measure. For each measure, the null hypothesis was taken that there is no relationship between the team that won each match and the team with the higher measure, i.e. the probability of the winning team having a higher network measure is 0.5. **The null distribution is thus a sample of 206 Bernoulli trials (there were 207 matches in the season with one draw ignored) with 0.5 probability of success.** Matches where both teams had equal values of a measure were also ignored, slightly reducing the number of data points for some measures. The bold line distinguishes measures at each tail of the distribution.

These results indicate that several of the measures are related to match outcomes. Based on conventional AFL analysis it would be expected that teams who had more disposals, higher disposal efficiency or shot efficiency would win more matches. However, the highlighted network measures give new information about the structure of the passing networks of successful teams. The measures in green are significant at the right tail of the distribution (suggesting teams with higher values win more matches)

TABLE IV: Proportion of winning teams with greater value of each network measure. Matches where measures were equal are excluded, $n = 206$ otherwise.

Measure	Winning teams (%)	p-value
Mean out-strength	74.4 ($n = 203$)	$< 10^{-4***}$
Disposals	72.5 ($n = 204$)	$< 10^{-4***}$
Mean out-degree	71.8 ($n = 202$)	$< 10^{-4***}$
Mean in-degree	71.7 ($n = 205$)	$< 10^{-4***}$
Disposal efficiency	69.3 ($n = 205$)	$< 10^{-4***}$
Shot efficiency	68.3 ($n = 205$)	$< 10^{-4***}$
Mean in-strength	67.8 ($n = 202$)	$< 10^{-4***}$
Team entropy	64.6	$< 10^{-4***}$
Mean closeness	64.1	$< 10^{-4***}$
Players with non-zero betweenness	60.8 ($n = 176$)	0.0026**
Mean chain length	60.2	0.0021**
Global clustering	56.3	0.041*
Mean betweenness	54.9	0.093
Passing rate (no shots)	46.6	0.15
Eigenvector std. dev.	40.3	0.0032**
Passing rate	33.5	$< 10^{-4***}$

while the eigenvector standard deviation in blue is significant at the left tail. These measures all relate to how well connected a team’s passing network is and the evenness of the passing distribution. The direction of the relationships suggests that successful teams have better connected networks and a more even distribution of passes between the players. For example, winning teams were more likely to have a lower eigenvector standard deviation (implying a more even distribution of probabilities that any player has the ball) and higher team entropy (implying players passed to a wide variety of teammates at similar rates).

Table V shows two different stepwise linear regression models used to analyse whether the network measures add extra information to conventional AFL statistics. The first model uses the values of each measure for all teams as predictor variables for a team’s score. All 414 data points from the 207 matches in the season are used. The second model uses the difference in measures between the two teams in a match (home team - away) as predictors for the dif-

ference in score between the two teams. Only intrinsic measures are used, i.e. those not directly related to the score in a match. For example, the number of shots taken is not considered. The Bayesian information criterion (BIC) has been used to compare the models. In both cases a stepwise regression process is used, first determining the best model using conventional AFL measures (with respect to BIC), before including the network measures. The increase in R^2 and decrease in BIC with network measures are estimated. BIC is used over AIC as it penalises model complexity more heavily thus providing a better test of whether adding network measures adds extra information to the models. In all cases the measures are first standardised to allow easier interpretation of effect sizes.

In both frameworks the network measures of mean betweenness, mean out-degree and team entropy improve the models, increasing the correlation coefficient and decreasing BIC. They are all utilised to a statistically significant degree in the combined models. The BIC decrease in both cases is greater than 9.2, implying that the improvements over the models using only conventional AFL measures are significant at the 1% level. This suggests that these network measures contain information relevant to match outcomes that is not present in the conventional AFL measures and offers some evidence for their use in quantitatively measuring team performance.

The effect directions in the models are all as would be expected. The only exception is the team entropy, which showed a positive relationship in table IV and a negative relationship here. We suggest this occurs as team entropy is related to mean out-degree and thus it is possible the mean out-degree is masking its influence. Indeed, if the mean out-degree is removed from the models the team entropy shows a positive effect. Also note that the mean in-degree was not included in the models, although it significantly improved the fits. The only difference between mean in-degree and out-degree is the edges connecting to the ‘Shot’ node, so including them together allows their linear combination to become a proxy measure for the number

of shots taken, which we wish to avoid.

C. Are team network characteristics consistent enough to assist in predicting the outcome of future matches?

Given that some network characteristics appear to be related to the outcomes of matches, it is natural to consider whether they represent some measure of team strength that is consistent enough between matches to be used for predicting future match results. To analyse this we perform a similar test to Table IV but instead of comparing the network measures of each team in the current match, we compare their average value of each measure across all previous matches. In this way, we are predicting each match using all information that would be available on each team’s network measures prior to the match. We label this a team’s *previous match mean* of a measure. The team with the highest previous match mean was predicted to win for the measures where more than 50% of teams won with higher values in the current match (see Table IV). The opposite was predicted for the other measures. Under the null hypothesis each match again represents a Bernoulli trial with 0.5 probability of success. However, the first round cannot be predicted using this method so the number of matches used is reduced to 197. The results of this test are shown in Table VI. Predicting using the bookmakers’ favourite was included for comparison and the net profit or loss that would have been made betting \$1 on each match is calculated for each measure.

Notably, several of the network measures did show some predictive capacity using this method. Predicting using some of the measures would have made a net profit, even given the bookmakers margins. The most successful network measure is the mean betweenness, predicting 134 winners correctly out of 197. Despite predicting less matches correctly than choosing the favourite, this was enough to make a profit as some underdog teams are correctly predicted, giving higher returns. To test the robustness

TABLE V: Linear regression models of team score against their intrinsic measures and match score margins against measure margins. Stepwise regressions are performed, first creating the best model possible using conventional AFL statistics, then showing the improvement when network measures are added. Effect sizes and p-values for the estimated models are shown.

Measure	Team's score against measures		Score margin against measure margins (Home - Away)	
	Conventional	Combined	Conventional	Combined
Conventional AFL measures				
Disposals	7.6 ($< 10^{-4***}$)	-0.96 (0.70)	-	-
Disposal efficiency	12.4 ($< 10^{-4***}$)	9.6 ($< 10^{-4***}$)	-	-
Effective disposals	-	-	29.2 ($< 10^{-4***}$)	18.1 (0.00039***)
Handballs	-7.7 ($< 10^{-4***}$)	-6.8 ($< 10^{-4***}$)	-16.9 ($< 10^{-4***}$)	-14.7 ($< 10^{-4***}$)
Contested possessions	10.3 ($< 10^{-4***}$)	9.3 ($< 10^{-4***}$)	17.2 ($< 10^{-4***}$)	15.2 ($< 10^{-4***}$)
Shot efficiency	9.5 ($< 10^{-4***}$)	10.2 ($< 10^{-4***}$)	10.8 ($< 10^{-4***}$)	11.9 ($< 10^{-4***}$)
Clangers	-4.0 ($< 10^{-4***}$)	-3.6 (0.00015***)	-10.5 ($< 10^{-4***}$)	-11.5 ($< 10^{-4***}$)
Clearances	2.8 (0.010**)	4.1 (0.00018***)	-	-
Free kicks received	-	-	-8.0 (0.00015***)	-9.4 ($< 10^{-4***}$)
Network measures				
Mean betweenness	-	3.7 ($< 10^{-4***}$)	-	8.3 ($< 10^{-4***}$)
Mean out-degree	-	12.3 ($< 10^{-4***}$)	-	16.1 (0.0011**)
Team entropy	-	-3.0 (0.020*)	-	-7.6 (0.0018**)
R²	0.588	0.622	0.732	0.777
R² increase		0.034		0.045
BIC	3597.9	3580.4	1932.9	1910.9
BIC decrease		17.5		22.0

of this conclusion two other prediction methods were considered: using only the measures from each team's previous match and using the mean of their measures from their past four matches. The results for mean betweenness are shown at the bottom of Table VI showing that it predicts a statistically significant number of matches correctly and makes a net profit in both cases.

Table VII shows stepwise linear regression models in a similar vein to Table V. In this instance the differences between the two team's previous match means for each measure are used as predictor variables in a model of the score margins. Similarly, the best model with respect to BIC using conventional AFL measures is created. Network measures are then included in a combined model and the improvement in R^2

and BIC estimated. The first test allows only intrinsic measures to be included, as with the previous analysis. As a second test, any measures are allowed (including, for example, the mean number of shots taken or points scored) as no measures from previous matches would be expected to be directly related to the score in the current match. A model using only network measures is included for comparison.

Including the network measures improves the models in both cases, increasing R^2 and decreasing BIC. When only intrinsic measures are considered, R^2 and BIC improve significantly with the addition of network measures. Furthermore, the model with only network measures performs similarly to the combined model, suggesting that the previous match means of the intrinsic network measures are more closely

TABLE VI: Winners predicted using the means of network measures from all previous matches in the season. Measures where a profit is made are shown in bold. Prediction using the mean betweenness over one or four previous matches is shown below.

	Winners predicted ($n = 197$)	p-value	Net profit
Favourite	143	$< 10^{-4***}$	-7.65
Mean betweenness	134	$< 10^{-4***}$	11.07
Possessions	131	$< 10^{-4***}$	5.67
Mean out-strength	129	$< 10^{-4***}$	3.56
Mean in-strength	127	$< 10^{-4***}$	4.21
Players with non-zero betweenness	127	$< 10^{-4***}$	-4.79
Mean closeness	126	$< 10^{-4***}$	4.09
Mean out-degree	124	0.00017***	-0.61
Mean in-degree	122	0.00050***	-3.34
Passing rate (no shots)	119	0.0021**	8.58
Team entropy	119	0.0021**	-8.36
Mean chain length	104	0.24	-8.24
Passing rate	102	0.33	12.17
Disposal efficiency	101	0.39	-23.13
Shot efficiency	96	0.67	-21.40
Global clustering	96	0.67	-14.05
Eigenvector std. dev.	86	0.97	-2.46
Mean betweenness (1 match)	130	$< 10^{-4***}$	26.46
Mean betweenness (4 matches)	113 ($n = 170$)	$< 10^{-4***}$	14.36

related to the margin of the current match than that of the intrinsic conventional measures. Thus, the network measures can be more easily interpreted to represent consistent characteristics of the team's playing style.

When any measures are permitted the best predictor of the current match was the mean number of shots taken in previous matches. However, a model including the eigenvector centrality standard deviation and number of players with non-zero betweenness centrality still improved R^2 and BIC over the best model using conventional measures. A model replacing the number of players with non-zero betweenness with the mean betweenness performed similarly (with a BIC of 2029.1). The decrease in BIC is not as large as when only intrinsic measures are included, but in conjunction with the other

formulation still suggests the network methods are adding useful information.

Note that, for this analysis, measures that no longer showed significant effects were removed from the combined models such that the comparisons were between the best conventional measure models and the best combined models. The fact that this was necessary implies that the network measures and conventional measures share some information. However, as the combined models have at least as many predictors as the conventional measure models and R^2 shows improvement in both cases, we can still infer that the network measures contain extra information not present in any of the conventional measures.

We are not suggesting, however, that these models optimise the prediction of AFL matches. Indeed, the best models in all formulations tested were simply to choose the bookmakers' favourite. Instead, we suggest that some network measures are successful in predicting future matches as their relationship with match outcomes is based on intrinsic differences between team's passing strategies rather than the specific events that occur in each match. Interestingly, significantly more matches were predicted correctly by summing the mean betweenness over all previous matches than were won by the team with the highest mean betweenness during each actual match. It also showed larger effects in the linear model using previous match means than the model using current match values. This suggests that the mean betweenness distinguishes teams well when averaged over previous matches, but shows too much variation between individual matches to show a direct relationship on this basis. It may be sensitive to events in each match, but on average be measuring underlying differences between team's network structures or passing strategy with a predictive capacity strong enough to consider it related to a team's strength.

TABLE VII: Linear regression models of match score margins against margins of previous match means for each measure. Stepwise regressions are performed, first creating the best model possible using conventional AFL statistics, then showing the improvement when network measures are added. Effect sizes and p-values for the estimated models are shown.

Measure	Intrinsic measures only			Any measures	
	Network	Conventional	Combined	Conventional	Combined
Conventional AFL measures					
Disposals	-	9.5 (0.0035**)	-	8.2 (0.0062**)	-
Clearances	-	-8.3 (0.013*)	-	-6.9 (0.021*)	-
Clangers	-	-8.9 (0.010*)	-9.3 (0.0049**)	-	-
Contested Possessions	-	18.3 ($< 10^{-4***}$)	9.2 (0.0071**)	-	-
Shots taken	-	-	-	20.3 ($< 10^{-4***}$)	16.5 ($< 10^{-4***}$)
Network measures					
Mean betweenness	14.2 ($< 10^{-4***}$)	-	13.2 ($< 10^{-4***}$)	-	-
Mean out-degree	27.2 ($< 10^{-4***}$)	-	13.8 (0.00034***)	-	-
Global clustering	-9.8 (0.0078**)	-	-10.9 (0.0024**)	-	-
Eigenvector std. dev.	-15.6 (0.0010**)	-	-7.5 (0.018*)	-	-8.7 (0.0021**)
Team entropy	-17.4 (0.019*)	-	-	-	-
Players with non-zero betweenness	-	-	-	-	7.9 (0.011*)
R²	0.273	0.193	0.294	0.261	0.289
R² increase			0.101		0.028
BIC	2043.2	2058.6	2042.7	2035.9	2028.2
BIC decrease			15.9		7.7

D. Can we use centrality measures to show a player's importance to their team's structure?

We begin analysing this question from the premise that demonstrating a positive relationship between individual players' centrality scores and match outcomes would justify the use of that centrality measure in representing some kind of player contribution to their team. As discussed in Section IV A, different centrality measures can quantify the variety of different roles players can have. As such, we argue that different centrality measures could measure different types of influence on a team's network structure. **This information could be useful in strategic planning, with the centrality measures showing which players are more crucial to a particular style or the structure of a team, and thus, as an opposition team, important to de-**

fend.

For each centrality measure we perform a similar analysis to Table IV, except using the mean of the largest n centrality scores for each team in each match rather than the global measures. For example, taking $n = 1$ we are comparing teams based on the player with the maximum centrality score on each side. When setting $n = 22$ teams are compared based on the mean centrality score of all 22 players. Again, the null hypothesis is taken that the difference in means of centrality scores is unrelated to which team wins the match, **implying each match is a Bernoulli trial with probability 0.5.** Figure 4 shows these results for each centrality measure, with the solid black lines showing the critical values for significance at the 1% level (≥ 120 or ≤ 85) for scale. For the properties that are calculated using integers and integer ratios (such as disposal or shot efficiency) there is an approx-

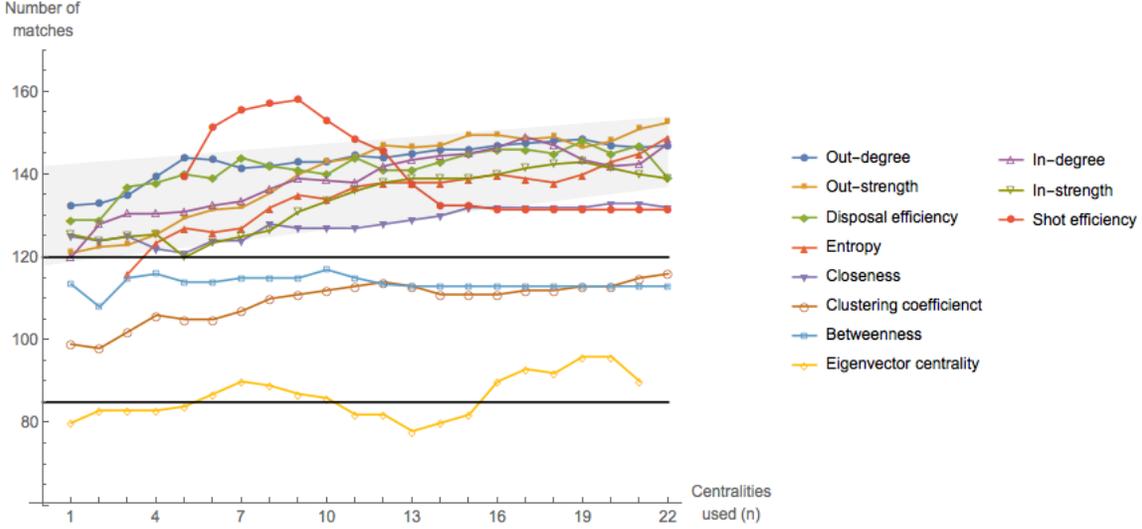


FIG. 4: Number of times (out of 206 matches) the winning team also had a greater mean value of each centrality measure. Results are computed using the largest n values of each measure, from 1 to 22. For example, the results for $n = 10$ represent the number winning teams whose largest 10 values of a centrality measure had a greater mean than their opposition's. The 1% bounds for statistical significance are shown in black.

cial chance that two opposing teams may have equal mean values, particularly for low n . In this case the number of matches with a positive result was calculated both by including and excluding teams with equal mean centrality measures and the average taken. If the difference between the two values was more than 30 we consider the uncertainty to be too significant and the result is not shown in Figure 4.

Firstly, several centrality measures produce significant results at the right tail for most values of n , suggesting teams with higher values of these measures won more matches. For low values of n the degree, strength, closeness and disposal efficiency measures show significant relationships (entropy and shot efficiency become significant when slightly more players are included). This suggests a relationship between the centrality scores of a small group of players (those with the highest centrality) and the match outcome. It then follows that such measures are capturing the contribution of individual players to their team's goal of winning the

match.

Most of these significant measures show an upward trend as the centralities of more players are included in the comparison. This lends weight to the argument that the performance of a team is dependent on all the players and not simply influenced by the performance of the top few, supporting the results from Section IV B that suggested spreading the passes as evenly as possible was advantageous. However, this is in contrast to the result obtained by Duch, Waitzman and Amaral⁹ for soccer, who made the argument that averaging the highest two 'match performances' was most closely related to the match result. Perhaps this is an indication of a difference in style between the sports - with less shots in soccer the outcome may depend more on individual brilliance to move past the defensive line. This is also potentially the result of the different methods used. Duch, Waitzman and Amaral based their definition of 'match performance' on the proportion of a players involvement in successful or unsuccessful passing

chains. As network methods are predominantly suited to showing the interactions it is possible that the methods used here, which maintain the full information of the network system, better show the contribution and interactions of all players.

However, the results for shot efficiency show anomalous behaviour, peaking when eight to ten players are included before decreasing as n is increased further. The relationships suggested between the shot efficiency and match results around this peak are some of the strongest across all measures considered. As around eight to ten players would be expected to take the majority of the shots this suggest that, as expected, the shot efficiencies of the forwards can significantly influence the outcomes of matches. With n increased further, more players who took minimal or no shots are included. We would expect their shot efficiencies to have little impact on the match outcome. As such, including this extraneous information may dilute the distinction between the teams, explaining the weaker relationship.

It has been demonstrated previously that diminishing returns in the shooting efficiency of basketball players leads to an analogy of Braess' paradox, where a team performs more efficiently without their 'star player'^{12,16}. Diminishing returns were not evident in either the disposal or shot efficiency of AFL players in our data set¹⁷ suggesting a similar analogy does not apply. As such, although there is evidence that spreading passes amongst as many players as possible is beneficial to the team, the 'star players' themselves are not performing worse by being overused.

V. CONCLUSIONS & APPLICATIONS

At a general level the results presented can be viewed as a proof of concept that network analysis provides additional insight into the structure and strategies of Australian football. Naturally, the specific results are only relevant to the 2014 AFL teams and season but the significant relationships shown provide evidence that

similar methods could give useful information when analysing more recent matches or applied in other team contexts.

It was seen that network measures could provide quantitative information about aspects of a team's style and structure that are not apparent in conventional Australian football analysis. Analysing the distribution of local centrality measures for each playing position showed how players contributed diversely to the team's structure. The similar distribution of betweenness centralities for players of each position suggested that, with regards to this metric, players of all positions were contributing relatively equally to connecting the team's passing network and transporting the ball toward the goals.

Some network measures, both global and local, showed statistically significant relationships to match outcomes. Analysis of the global measures suggested that teams who utilised all their players with a relatively even passing distribution were more successful. Furthermore, for each local centrality measure that gave significant results (other than the shot efficiency) a stronger relationship with the team's performance was observed when the centrality measures of the full team were included in the comparison. This suggested that the outcomes of matches were influenced by the full team and not primarily by a few 'star players'.

Team scores and match margins were modelled using stepwise linear regression, both using intrinsic measures from the current match and the means of measures from the each team's previous matches. Adding the network measures to models constructed using conventional AFL measures increased R^2 and decreased the BIC in all cases. The significance of the improvement observed gave evidence that the network measures contain useful information in addition to what is available from conventional AFL statistics.

Furthermore, the intrinsic network measures outperformed intrinsic conventional measures in predicting matches, suggesting that they are more closely related to consistent differences in style and strategy between the teams. In particular, betting on the team with the high-

est mean betweenness in previous weeks was shown to make a profit. This suggested that the mean betweenness differentiated teams consistently enough to be considered a representation of underlying differences in passing strategies and, furthermore, that it was related strongly enough to match outcomes to be associated with the team's strength.

We suggest that network measures could be used in the AFL as supplementary statistical measures. They provide information that is not easily quantifiable by conventional AFL statistics as they measure aspects of the intrinsic dynamics of game play and consider the interactions between players on a team. At a team level, for example, they could help elucidate situations in which a team was passing too predictably or underusing certain players. The local central measures could expose the different types of contributions that players make to their team's structure and strategy, which could be useful for player rating purposes. Analysing an opposition team's network characteristics could be useful in planning defensive strategy in order to best disrupt their passing structure.

¹For a more detailed explanation of Australian football the reader is referred to Harvey and East's discussion Available at www.footy.com.au/dags/FAQ1v1-5.html.

²R. N. Onody and P. A. de Castro, "Complex network study of Brazilian soccer players," *Phys Rev E Stat Nonlin Soft Matter Phys* **70** (2004).

³R. M. Tripathy, A. Bagchi, and M. Jain, "Complex network characteristics and team performance in the game of cricket," in *Big Data Analytics: Second International Conference, BDA 2013, Mysore, India, December 16-18, 2013, Proceedings*, edited by V. Bhatnagar and S. Srinivasa (2013) pp. 133–150.

⁴S. Mukherjee, "Complex network analysis in cricket: Community structure, player's role and performance index," *Advances in Complex Systems* **16** (2013).

⁵F. Radicchi, "Who is the best player ever? A complex network analysis of the history of professional tennis," *PLoS ONE* **6** (2011).

⁶X. Xu, J. Hu, and F. Liu, "Complex network study of asian Go players," *Chaos* **17** (2007).

⁷D. P. Kiley, A. J. Reagan, L. Mitchell, C. M. Danforth, and P. S. Dodds, "Game story space of professional sports: Australian rules football," *Phys. Rev. E* **93**, 052314 (2016).

⁸C. Cotta, A. M. Mora, J. J. Merelo, and C. Merelo-Molina, "A network analysis of the 2010 FIFA World

Cup champion team play," *Journal of Systems Science and Complexity* (2013).

⁹J. Duch, J. S. Waitzman, and L. A. Nunes Amaral, "Quantifying the performance of individual players in a team activity," *PLoS ONE* (2010).

¹⁰J. H. Fewell, D. Armbruster, J. Ingraham, A. Petersen, and J. S. Waters, "Basketball teams as strategic networks," *PLoS ONE* **7** (2012).

¹¹J. L. Peña and H. Touchette, "A network theory analysis of football strategies," in *Sports Physics: Proc. 2012 Euromech Physics of Sports Conference*, edited by C. Clanet (Ecole Polytechnique, Palaiseau, 2013) pp. 517–528.

¹²B. Skinner, "The price of anarchy in basketball," *Journal of Quantitative Analysis in Sports* **6** (2010).

¹³J. Saramäki, K. Mikko, J.-P. Onnel, K. Kaski, and J. Kertész, "Generalizations of the clustering coefficient to weighted complex networks," *Physical Review E* **75** (2007).

¹⁴C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal* **27**, 379–423,623–656 (1948).

¹⁵AFL player ratings can be found at www.afl.com.au/afl-player-ratings.

¹⁶D. Braess, "Über ein paradoxon aus der verkehrsplannung," *Unternehmensforschung* **12**, 258–268 (1968).

¹⁷C. Braham, *What can complex networks tell us about Australian Football?*, Undergraduate honors thesis, The University of Western Australia (2015).