Design-weighted Regression Adjusted Plus-Minus (DWRAPM) for Evaluation of Player Impact

Michael Schuckers[a] (schuckers@stlawu.edu), Brian Macdonald[b], Seongwon Im[a], Jesse McNulty[c]
[a]St. Lawrence University, [b]University of Miami (FL), [c]Atlanta Blaze

February 8, 2019

Sports, particularly fluid team sports, present challenges for evaluating the impact of individual players. Regression Adjusted Plus-Minus (RAPM) is one widely accepted way to assess these impacts. RAPM tracks which players are involved, broadly defined, in an observation and then averages the effect of individual players conditional on the other players involved in that same observation. Below we will call these observations 'events'. Our definition of an event will be sport specific and can be quite variable because of the nature of individual sports. For example, [Macdonald, 2011] uses ice hockey shifts as events, while [Schuckers and Curro, 2013] uses ice hockey events from the National Hockey Leagues play by play files, [Kasan, 2008].

RAPM was first publicly introduced by [Rosenbaum, 2004] for basketball though he suggests that Jeff Sagarin and Wayne Winston had already developed a similar system. Rosenbaum used a formulation where a single parameter is used to assess the impact of a player. Results of this methodology are given by [Ilardi, 2007] for the 2006-07 National Basketball Association season. Separate offensive and defensive player impacts are modelled by [Ilardi and Barzilai, 2007]. Sill introduced regularization via out-of-sample prediction in [Sill, 2010]. This work was enlarged by Engelman who adds events in a team context for which a single individual is given credit in [Engelman, 2017] and [Engelman, 2015].

In this article, we outline the basic attributes of the RAPM approach and provide an extension of this methodology to allow for better assessment of individual player impact on an event.

$$x_{ij} = \begin{cases} -1 & \text{if player } j \text{ is on the field for the Away team for event } i, \\ 0 & \text{if player } j \text{ is not on the field for event } i, \text{and} \\ 1 & \text{if player } j \text{ is on the field for the Home team for event } i. \end{cases} \tag{1}$$

Shortly thereafter [Ilardi and Barzilai, 2007] modify the basic RAPM framework to distinguish between offensive impact and defensive impact. Effectively this is to double the number of columns of the design matrix, $\mathbf{X}$, so that each player is represented by a column for offense and defense. In this formulation, the x_{ij}'s only take the values 1 or 0 dependent on whether or not player$i \times$situation j where situation is whether or not the player is on offense or defense. Using regularization methods for the fitting of these player estimates as well as cross-validation for assessment of their predictive capabilities was presented by [Sill, 2010]. A further innovation comes from [Engelman, 2017] , and is explicated in [Engelman, 2015], whereby the design matrix is further expanded to address whether or not a player is directly involved in a specific type event, for example taking a shot in basketball.

The above developments of the RAPM were motivated by data from basketball. The RAPM has been applied to data from other sports. In particular, it has been applied to hockey by [Macdonald, 2011] and [Schuckers and Curro, 2013]. Some earlier work in soccer particularly, [scaryice, 2009] and [Hamilton, 2010b, Hamilton, 2010a, Hamilton, 2014] suggested that the RAPM was not a great fit for soccer in large part due to the infrequent substitutions of players on the field of play. Recently, [Francesca Matano, 2018] has used prior information from video game ratings to improve the performance of RAPM for soccer.

In this paper, we propose to use a generalization of RAPM that considers values for x_{ij} other than one when a player is involved in event i. We continue to use $x_{ij} = 0$ when player j is not involved in event i. However, we propose the following values for x_{ij}

$$x_{ij} = \begin{cases} d_{ij} & \text{if player } j \text{ is on the field for the Away team for event } i \text{ with an impact weight of } d_{ij} \text{ on event } i, \\ 0, & \text{if player } j \text{ is not on the field for event } i, \text{and} \\ -d_{ij} & \text{if player } j \text{ is on the field for the Home team for event } i \text{ with an impact weight of } d_{ij} \text{ on event } i, \end{cases} \tag{2}$$

We refer to this formulation as design-weighted regression adjusted plus-minus (S. For the above equation we are allowing d_{ij} to vary by both player and event as opposed to being 1 for every player involved in each event. This might be utilized in soccer (association football) or hockey where the impact can be thought of as a function of the distance between player j and event i. In American football, d_{ij} might have different values among linemen on running plays, depending on whether or not the lineman was on the same side of the field of the running play. d_{ij} might be influenced by the position of the player. For example in ice hockey, the values for d_{ij} might be different when defencemen in the defensive zone as opposed to when they are in the offensive zone and *vice versa* for forwards. In this paper we illustrate how we might choose d_{ij} and discuss cross-validation for selection of these methods. We also recognize that changing our definitions of the x_{ij}'s is potentially a slippery slope and can lead to overfitting. Thus consideration of multiple methods for the evaluation of d_{ij} is critical as is cross-validation and out of sample prediction for assessing these approaches.

The way that we evaluate player distance from an event can be quite flexible and here the idea of distance is flexible. We do note that smaller values for d_{ij} denote larger values for the coefficients for the players assuming the the response is the same. To illustrate this, we consider least squares where the column for one player has been alter by multiplying by δ. Denote $\mathbf{I}_k(\delta)$ as an identity matrix with the k^{th} entry along the diagonal replaced by δ. Then

$$\begin{aligned}
\hat{\beta}_\delta &= \left([\mathbf{X}\mathbf{I}_k(\delta)]^T[\mathbf{X}\mathbf{I}_k(\delta)]\right)^{-1}[\mathbf{X}\mathbf{I}_k(\delta)]^T\mathbf{Y} &(3)\\
&= \left(\mathbf{I}_k(\delta)\mathbf{X}^T\mathbf{X}\mathbf{I}_k(\delta)\right)^{-1}\mathbf{I}_k(\delta)\mathbf{X}^T\mathbf{Y} &(4)\\
&= \mathbf{I}_k(\frac{1}{\delta})\left(\mathbf{X}^T\mathbf{X}\right)^{-1}\mathbf{I}_k(\frac{1}{\delta})\mathbf{I}_k(\delta)\mathbf{X}^T\mathbf{Y} &(5)\\
&= \mathbf{I}_k(\frac{1}{\delta})\left(\mathbf{X}^T\mathbf{X}\right)^{-1}\mathbf{X}^T\mathbf{Y} &(6)\\
&= \mathbf{I}_k(\frac{1}{\delta})\hat{\beta} &(7)\\
&= \left[\beta_1,\ldots,\frac{1}{\delta}\beta_k,\ldots,\hat{\beta}_p\right]^T &(8)
\end{aligned}$$

where $\hat{\beta} = \left(X^TX\right)^{-1}X^TY$. Consequently, by increasing δ we are lowering the value of the corresponding estimate of the impact of a player, $\hat{\beta}_k$. If we want to give a player more credit for a given play *ceteris paribus* then we need decrease their d_{ij} for a given event.

The notion that we have for how a player impacts a play has been specified by their 'distance' but we are thinking of this as a flexible concept that does not necessarily imply Euclidean distance. While field or court distance might be reasonable, it might also be possible to consider distance to be something like the average value of the real estate between a player and the event along the lines of [Cervone et al., 2016]. Figure **??** has some possible structure for ways that distance could be measured. We envision that these 'distances' will have a floor and a ceiling as denoted by d_{low} and d_{high} in that figure with the value of one included between these two values.

In this paper we present three illustrations of this new approach. The first illustration is an analysis of the impact of offensive lineman on rushing yardage in the National Football League. Using NFL play by play data, we can assess the yardage gained but also the location relative to the offensive line of the run, ie. was the play run 'up the middle' or 'off right end'. We define our design matrix entries as

$$x_{ij} = \begin{cases} d_{ij} = f(\delta_{ij}) & \text{if offensive lineman } i \text{ is on the field for event } j, and \\ 0 & \text{if offensive lineman } i \text{ is not on the field for event } j, \end{cases} \tag{9}$$

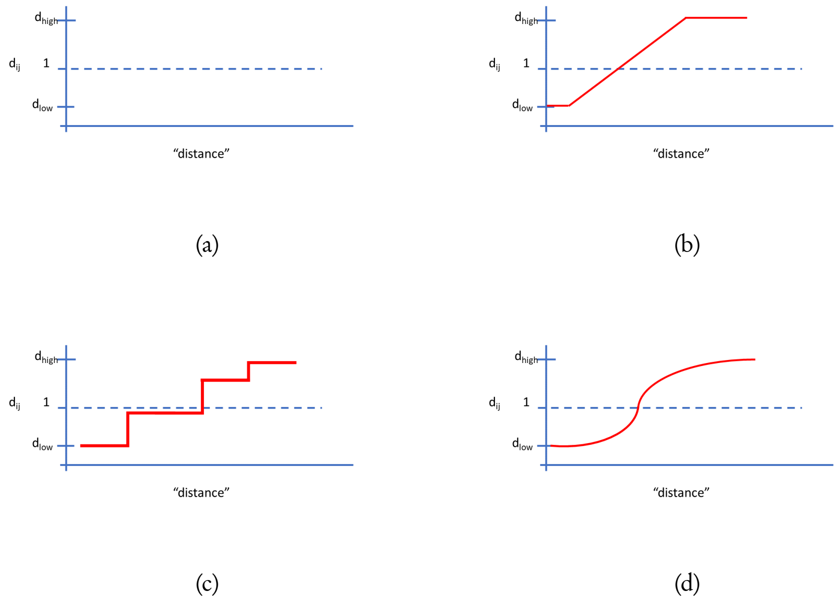

Figure 1: Four example structures for calculation of d_{ij}

where we are interested in only 'rushing' events, we define δ_{ij} as the number of gaps that offensive lineman i is from where the event was recorded and $f(\delta_{ij})$ is some measure of distance between the lineman and the gap of the play. More generally, we can think of δ_{ij} as some measure of 'distance'.

Here we focus on the elements of assigning d_{ij} for offensive lineman but thought would also have to be given to values of d_{ij} for quarterbacks, running backs and wide receivers. Our illustration below focuses on a formation with 2 tight ends one to each side of the formation with those players aligned next to the offensive tackle on their side of the formation. This formation can be found in Figure **??**.

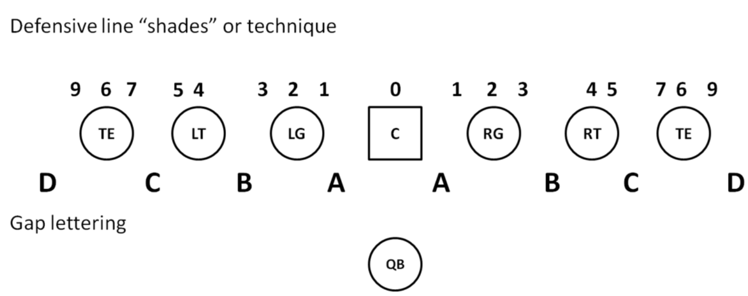

Figure 2: Football formation with two tight ends (TEs) [1]

For a run play from this formation we will use a design matrix entries d_{ij} dependent on location of the run. We will denote the value for d_{ij} for each player in **??**. The weight assigned to the left and right tight ends, respectively, will be denoted as d_{LTE} and d_{RTE}. The rest of the design matrix entries for the offensive line from left to right will be specified as $d_{LT}, d_{LG}, d_C, d_{RG}, d_{RT}$. For this illustration we will drop the i from notation for now.

Consider a play from this formation where the run goes through the center- guard gap on the left side of the formation. Figure ?? has a red arrow where that run would take place. The values that we might utilize for this

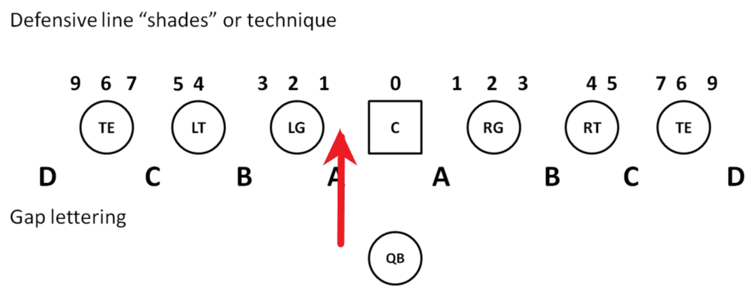

Figure 3: Run between the left guard (LG) and the center (C)

play are the following: $d_{LG} = 0.8$, $d_C = 0.8$, $d_{RG} = 1$, $d_{LT} = 1$, $d_{LTE} = 1.1$, $d_{RT} = 1.1$, $d_{RTE} = 1.15$. These values are chosen so that the blockers who are closest to where the run is designated have the smallest d_{ij}'s and, thus, the largest impact.

We consider another running play. This time with the run designated between the right guard (RG) and the right tackle (RT). Thus those players and the others nearby are given the smallest values for that event in the design matrix. Those values might be something like: $d_{RG} = 0.8$, $d_{RT} = 0.8$, $d_{RTE} = 1$, $d_C = 1$, $d_{LG} = 1.1$, $d_{LT} = 1.15$, and $d_{LTE} = 1.2$

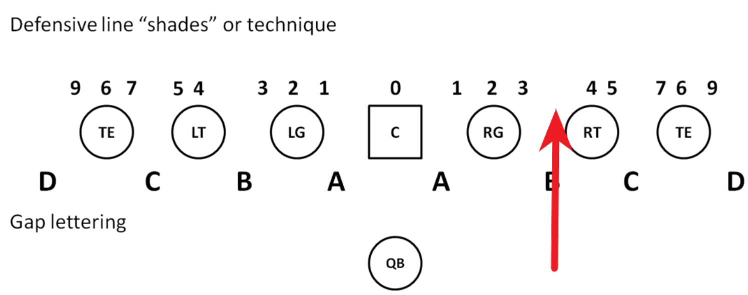

Figure 4: Run between the right guard (RG) and the right tackle (RT)

Our second illustration extends the work of [Schuckers and Curro, 2013] to evaluate forwards and defencemen in hockey using the following framework where we given additional impact to forwards in the offensive zone and to defencemen in the defensive zone.

$$
x_{ij} = \begin{cases} d_{ij} = a & \text{if player } j \text{ is on the home team, on the ice for event } i \text{ and } S_i = T_j, \\ d_{ij} = b & \text{if player } j \text{ is on the home team, on the ice for event } i \text{ and } S_i \neq T_j, \\ 0 & \text{if player } j \text{ is not on the ice for event } i, \\ d_{ij} = -b & \text{if if player } j \text{ is on the away team, on the ice for event } i \text{ and } S_i \neq T_j, \\ d_{ij} = -a & \text{if player } j \text{ is on the away team, on the ice for event } i \text{ and } S_i = T_j \end{cases} \quad (10)
$$

where $S_i \in O, D$ is the location of event i either offensive (O) or defensive(D) zone and $T_j \in O, D$ is the players position either center or wing (O) or defencemen (D).

The selection of the values of d_{ij} (or equivalently, $f(\delta_{ij})$) in each of our illustrations are chosen based upon the context of the sport but also based upon out-of-sample validation of these values using regularization as in [Macdonald et al., 2012] or [Sill, 2010].

In this paper we have proposed an extension to the regression adjusted plus-minus model that is commonly used in some sports for the evaluation of players. This generalization allows for differential credit to be given to players on the court or field for each event. This credit is allocated via changes to the design matrix and so we call this methodology design weighted regression adjusted plus minus.

References

[Cervone et al., 2016] Cervone, D., Bornn, L., and Goldsberry, K. (2016). Nba court realty. http://www.sloansportsconference.com/wp-content/uploads/2016/02/1556-NBA-Court-Realty.pdf, note=Accessed February 8, 2019.

[Engelman, 2015] Engelman, J. (2015). Estimating a player's influence on his teammates' boxscore statistics using a modified rapm framework. https://www.youtube.com/watch?v=0uC0YZTADcE. Accessed April 18, 2018.

[Engelman, 2017] Engelman, J. (2017). Possession-based player performance analysis in basketball (adjusted +/- and related concepts). In Albert, J., Glickman, M. E., Swartz, T. B., and Koning, R. H., editors, *Handbook of Statistical Methods and Analyses in Sports*. Chapman and Hall/CRC.

[Francesca Matano, 2018] Francesca Matano, Lee F. Richardson, T. P. C. E. J. Q. (2018). Augmenting adjusted plus-minus in soccer with fifa ratings. https://arxiv.org/abs/1810.08032.

[Hamilton, 2010a] Hamilton, H. (2010a). Some starting documents on the plus/minus problem. http://www.soccermetrics.net/player-performance/some-starting-documents-on-the-plusminus-problem-2.

[Hamilton, 2010b] Hamilton, H. (2010b). Some work on plus/minus ratings in mls. http://www.soccermetrics.net/player-performance/some-work-on-plusminus-ratings-in-mls-2.

[Hamilton, 2014] Hamilton, H. (2014). Adjusted plus/minus in football – why it's hard, and why it's probably useless. http://www.soccermetrics.net/player-performance/adjusted-plus-minus-deep-analysis.

[Ilardi, 2007] Ilardi, S. (2007). Adjusted plus-minus: An idea whose time has come. http://www.82games.com/ilardi1.htm. Accessed April 18, 2018.

[Ilardi and Barzilai, 2007] Ilardi, S. and Barzilai, A. (2007). Adjusted plus-minus ratings: New and improved for 2007-2008. http://www.82games.com/ilardi2.htm. Accessed April 18, 2018.

[Kasan, 2008] Kasan, S. (2008). Off-ice officials are a fourth team at every game. http://www.nhl.com/ice/news.htm?id=388400. Accessed September 24,2013.

[Macdonald, 2011] Macdonald, B. (2011). A regression-based adjusted plus-minus statistic for nhl players. *Journal of Quantitative Analysis in Sports*, 7.

[Macdonald et al., 2012] Macdonald, B., Lennon, C., and Sturdivant, R. (2012). Evaluating nhl goalies, skaters, and teams using weighted shots. http://arxiv.org/abs/1205.1746. Accessed October 24, 2015.

[Rosenbaum, 2004] Rosenbaum, D. T. (2004). Measuring how nba players help their teams win. http://www.82games.com/comm30.htm. Accessed April 18, 2018.

[scaryice, 2009] scaryice (2009). Mls 2009 plus/minus data. http://usasoccer.blogspot.com/2009/12/mls-2009-plusminus-data.html.

[Schuckers and Curro, 2013] Schuckers, M. and Curro, J. (2013). Total Hockey Rating (THoR): A comprehensive statistical rating of National Hockey League forwards and defensemen based upon all on-ice events. Proceedings of the 2013 MIT Sloan Sports Analytics Conference http://www.statsportsconsulting.com/thor.

[Sill, 2010] Sill, J. (2010). Improved nba adjusted +/- using regularization and out-of- sample testing. 2010 MIT Sloan Sports Analytics Conference.