OUT OF GAS: QUANTIFYING FATIGUE IN MLB RELIEVERS

Burris, Kyle; Coleman, Jake

Duke University, Durham, NC, USA

† E-mail: *kyle.burris@duke.edu*

The study of reliever usage is becoming increasingly important, now that managers are pulling their starting pitchers earlier than ever before. To address this need, we develop a framework for measuring the effect of fatigue for MLB relief pitchers by estimating how their effectiveness responds to workload. We predict the probability of a swinging strike using random forests to obtain a metric for the quality of a given pitch. Using this metric as a dependent variable, we fit a hierarchical Bayesian model with an autoregressive-like structure that models game-day pitch quality as a function of prior workload.

We find a strong signal that fastball quality is negatively affected by workload and that some components of curveball quality may be diminished as well. Moreover, we estimate that pitches thrown more than three days in the past have no effect on pitch quality and that a pitch thrown two days ago has between 25-45% of the effect as a pitch thrown yesterday. Finally, we find no substantial evidence that fatigue effects differ significantly between relievers.

NO "I" IN TEAM: A NETWORK ANALYSIS OF DIVISION I MEN'S BASKETBALL OFFENSES

Dickman, Grace

Macalester College, Saint Paul, MN, USA

† E-mail: *gdickman@macalester.edu*

Sports contain interactions between players that are measurable, such as passing the ball through the offense or screening for teammates. These interactions result in outcomes which dictate the quality of the relationship based on productivity. The networks can be used as ways to look at who is the heart of the team – the player that makes everyone else go. This is recognized by coaches' analysis of film when scouting, but is often neglected by statistics. We used four men's Division I basketball teams and their game film from the 2015-2016 season to create networks that characterize their offense. We use weights based on significance of statistical categories to create an algorithm for weighted, directed edges between players. The resulting networks reveal patterns in teams' offensive schemes and highlight star players

as centers of the networks. There are distinct communities for each team based on who plays with whom as well as who interacts on the floor the most. Weighted degree measures directly correlate to offensive production, whereas other centrality measures tell a story not seen in the box score. PageRank is the most telling centrality measure, revealing who the most important player on each team really is. The networks and the analysis confirm coaches' intuitions of how the offense flows and who should be defended. They could be used as reliable methods for scouting and preparing for opponents. Similarly, they could be used for player acquisition and termination by showing who truly fits well into the offense.

FAIR STANDINGS IN SOCCER AND OTHER ROUND-ROBIN LEAGUES

Firth, David

University of Warwick and the Alan Turing Institute, Coventry, UK

† E-mail: *d.firth@warwick.ac.uk*

A specialized extension of the Zermelo-Bradley-Terry model is developed for balanced round-robin leagues with simple win-draw-loss points systems such as the 3-1-0 system used in association football. The model allows schedule-strength differences to be eliminated coherently, to produce fully "retrodictive" match-by-match league standings that are more informative than the usual ranking based on accumulated points. Results from several seasons of major European soccer leagues are used to assess the model, and to calibrate it in aspects such as the frequency of draws and the "home advantage" effect. We consider also how to present the resulting fair-standings tables as transparently as possible.

A POINT-BASED BAYESIAN HIERARCHICAL MODEL TO PREDICT THE OUTCOME OF TENNIS MATCHES

Ingram, Martin

Silverpond, Melbourne, AUS

† E-mail: *martin.ingram@gmail.com*

A well-established assumption in tennis is that points in a match are independent and identically distributed (iid). With this assumption, it is enough to specify the serve probabilities for both players to derive a wide variety of event distributions, such as the expected winner, number of sets, and number of games. However, models using this assumption, which I will refer to as "point-based", have typically performed worse than other models in the literature at predicting the match winner.

I present a Bayesian hierarchical model for predicting the outcome of tennis matches. The model predicts the probability of winning a point on serve given surface, tournament and match date. Each player is given a serve and return skill which is assumed to follow a Gaussian random walk over time. In addition, each player's skill is assumed to vary by surface, and variable conditions at different tournaments are modeled as tournament-specific intercepts.

The serve probabilities estimated by the model are then used as inputs to the iid model to predict match outcomes. When evaluated on over 2,000 ATP matches in 2014, the model outperforms other models using the point-based approach, predicting serve probabilities with lower root mean squared error (0.077 vs. 0.082) and match outcomes with greater accuracy (68% vs. 66%) and lower log loss (0.60 vs. 0.64). The results are competitive with approaches modeling the match outcome directly, demonstrating the forecasting potential of the point-based modeling approach.

AUTOMATIC EVENT DETECTION IN BASKETBALL GAMES

Keshri, Suraj; Oh, Min-hwan; Zhang, Sheng; Iyengar, Garud

Columbia University, New York, NY, USA

† E-mail: *skk2142@columbia.edu*

We propose a unsupervised machine learning framework for automatically tagging events in basketball games. Our framework uses the optical player tracking data made available for all NBA games since 2013. We take a two step approach: learn the time series of defensive assignment for each possession, which is then used as an input to detect events. We learn the defensive assignment using a novel player identity and court location dependent attraction based model which uses hidden Markov models (HMMs), Gaussian processes (GP), and "bond-breaking" based transition matrix. The GP uses a shared mean prior to pool information across both players and spatial locations on the court. The "bond energy" based transitions are used to learn a low-dimensional approximation for the full $(5^5) \times (5^5)$ dimensional transition matrix for the state space of all possible defensive assignments. The learned defensive assignments are used as an input to a set of HMMs that automatically detect events such as ball screens, drives and post-ups. We show that our models provide significant improvements over existing benchmarks both on defensive assignments and event detection.

A SHOT TAXONOMY IN THE ERA OF TRACKING DATA IN PROFESSIONAL TENNIS

Kovalchik, Stephanie

Tennis Australia, Melbourne, AUS

† E-mail: *s.a.kovalchik@gmail.com*

Shots are the vocabulary of tennis yet there has been limited quantitative study of the distinct types of shots at the professional level. In this study, we build a taxonomy of shots for elite tennis using tracking data from five years of men's and women's matches at the Australian Open. Our taxonomy is constructed using model-based multi-stage functional data clustering, an unsupervised machine learning approach. Among 284,035 men's and 194,091 women's shots, we found 15 and 13 distinct first serves to the Ad and Deuce court for male players; 13 and 16 for female players. The distinct number of groundstroke shot types were notably higher, men having 25 and 30 distinct forehand serve return and forehand rally shot types, while women had 31 and 29 types, respectively. Both male and female players showed less variety on the backhand, with men having 20 distinct backhand serve returns and 14 backhand rally shots; women having 20 and 15 types. This dictionary provides a framework for investigating the vocabulary of tennis, characterizing playing style, and advancing the modeling and analysis of elite player performance.

HOW OFTEN DOES THE BEST TEAM WIN? A UNIFIED APPROACH TO UNDERSTANDING RANDOMNESS IN NORTH AMERICAN SPORT

Lopez, Michael (1); Matthews, Gregory (2), Baumer, Benjamin (3)

(1) Skidmore College, Saratoga Springs, NY, USA; (2) Loyola University Chicago, Chicago, IL, USA; (3) Smith College, Northampton, MA, USA

† E-mail: *bbaumer@smith.edu*

Statistical applications in sports have long centered on how to best separate signal (e.g. team talent), from random noise. However, most of this work has concentrated on a single sport, and the development of meaningful cross-sport comparisons has been impeded by the difficulty of translating luck from one sport to another. In this manuscript, we develop Bayesian state-space models using betting market data that can be uniformly applied across sporting organizations to better understand the role of randomness in game outcomes. These models can be used to extract estimates of team strength, the between-season, within-season, and game-to-game variability of team strengths, as well each team's home advantage. Parameter estimates are validated by considering actual game outcomes. More generally, we use our framework to compare cumulative (across all weeks) and sequential (from all weeks prior) models. We implement our approach across a decade of play in each of the National Football League (NFL), National Hockey League (NHL), National Basketball Association (NBA), and Major League Baseball (MLB), finding that the NBA demonstrates both the

largest dispersion in talent and the largest home advantage, while the NHL and MLB stand out for their relative randomness in game outcomes. We conclude by proposing a new metric for judging competitiveness across sports leagues. While we focus on sports, we discuss a number of other situations in which our generalizable models might be usefully applied.

USING DATA ANALYSIS TO PREDICT ATTENDANCE FOR NHL REGULAR SEASON GAMES

Macdonald, Brian (1); Peterson, Michael (2); Cifu, James (1)

(1) Florida Panthers, Sunrise, FL, USA; (2) Tampa Bay Lightning, Tampa, FL, USA

† E-mail: *bmac@jhu.edu*

We will discuss a data-driven approach to predicting the popularity of an NHL game for the purposes of informing business decisions within an NHL organization. We first describe a model for forecasting the popularity of a game that uses only publicly available information that is known before tickets go on sale, when we have no sales data available for individual games. Once single game tickets go on sale and we do have internal sales data, we describe a second model which uses this new data to update and improve the predictions. We discuss how those prediction models not only use different data, but also have slightly different purposes in mind at different stages of the pricing and ticketing process. Finally, we give a brief overview of the other kinds of data analysis and visualization we use in business operations, and how we combine both team data and business data to inform our business decisions.

PLAYER TRACKING IN AMERICAN FOOTBALL: SPATIO-TEMPORAL MODELING OF DEFENSIVE PLAYERS' INTENT

Pazdernik, Karl

North Carolina State University, Raleigh, NC, USA

† E-mail: *karl.pazdernik@gmail.com*

Gaining separation is paramount to the success of an individual playing offense in many team sports. Conversely, the ability to collectively and continuously limit separation for all offensive players is often what defines a great defense. To truly understand the complex spatio-temporal patterns of offensive separation and defensive coverage, identification of defensive intent is necessary. While such methodology has become a staple of the analytics

used in sports such as basketball, analogous methods for American football are limited. This presentation outlines methodology necessary to close this gap by estimating the probability that each defender is tracking each offensive player at predetermined intervals of time within a play using a hidden Markov model. Also, since trajectory differs between a defender in attack mode versus one in a surveillance motion, a secondary group of hidden states differentiating between these two behavioral patterns is included. The potential summary statistics that follow could revolutionize the sport, providing insight into aspects of player performance that have never been quantified, such as a player's instincts or their impact on the opposition's game plan. For illustration, both simulated data and data from NFL games obtained through the All-22 game film are used.

LINNET: PROBABILISTIC LINEUP EVALUATION THROUGH NETWORK EMBEDDING

Pelechrinis, Konstantinos

University of Pittsburgh, Pittsburgh, PA, USA

† E-mail: *kpele@pitt.edu*

Which of your team's possible lineups has the best chances against each of your opponents' possible lineups? In order to answer this question we develop LinNet. LinNet exploits the dynamics of a directed network that captures the performance of lineups at their matchups. The nodes of this network represent the different lineups, while an edge from node j to node i exists if lineup λ_i has outperformed lineup λ_j. We further annotate each edge with the corresponding performance margin (point margin per minute). We then utilize this structure to learn a set of latent features for each node (i.e., lineup) using the node2vec framework. Consequently, LinNet builds a model on this latent space for the probability of lineup λ_A beating lineup λ_B. We evaluate LinNet using NBA lineup data from the five seasons between 2007-08 and 2011-12. Our results indicate that our method has an out-of-sample accuracy of 69%. In comparison, utilizing the adjusted plus-minus of the players within a lineup for the same prediction problem provides an accuracy of 56%. More importantly, the probabilities are well-calibrated as shown by the probability validation curves. One of the benefits of LinNet – apart from its accuracy – is that it is generic and can be applied in different sports since the only input required is the lineups' matchup performances, i.e., not sport-specific features are needed.

OPPOSITION ANALYSIS: METHODS FOR FORECASTING STYLE OF PLAY IN SOCCER

Pleuler, Devin (1); Myers, Bret (1,2)

(1) Toronto Football Club, Toronto, ON, Canada; (2) Villanova University, Villanova, PA, U.S.

† E-mail: *dpleuler@torontofc.ca*

Despite the exponential growth of granular performance data being collected in professional athletics, soccer, like other dynamic team-invasion sports, has resisted efforts aimed at identifying and quantifying the underlying attributes that inform team and player success. While all soccer teams have the same literal goals in mind, their methods of reaching those goals are necessarily diverse.

The pass is the most important on-ball event in the game of soccer, as it represents the most economical method of moving the ball around the field and manipulating play. In a recently studied professional soccer league that included over 300,000 individual passes, each team averaged over 400 passes attempted per game.

By applying various data mining and machine learning techniques to on-field passing coordinate data provided by Opta Sports, we have identified passing patterns that persist on a team-by-team basis. These passes are found to be spatially distributed across the field in a structured and relatively symmetrical manner. However, in practice, this distribution can become quite asymmetrical and unique on a per-team basis. These team-level deviations from the standard league cartography correlate strongly on a game-to-game basis.

This methodology has helped identify factors that commonly vary between teams but remain consistent for a single team over the course of a season. In an applied setting, this information has proven to be remarkably useful for opposition analysis, helping to streamline various coaching processes.

STRATEGY IMPLICATIONS OF THE NFL'S NEW PAT RULE

Sackrowitz, Harold

Rutgers University, New Brunswick, NJ, USA

† E-mail: *sackrowi@stat.rutgers.edu*

The NFL has now completed two full seasons under the new PAT rule. The results of these last two seasons strongly suggest a phenomenon that could never have been firmly believed ever before. Namely, the probability of a successful 1-point conversion attempt is less than twice that of a successful 2-point conversion attempt. This dramatically impacts the results of any probabilistic approach to finding optimal PAT strategies. In the past, the 2-point attempt was essentially used to catch up when behind or protect the ability to reach overtime when ahead. Now the 2-point conversion attempt can be used as an offensive weapon to try to win in regulation time. This requires a different philosophy. We will discuss issues that need to be considered in this setting that were not of concern before the new rule.

REPLAYING THE NBA: USING MARKOV DECISION PROCESSES TO TEST DECISION-MAKING FROM THE 2015-2016 REGULAR SEASON

Sandholtz, Nathan (1); Bornn, Luke (1,2)

(1) Simon Fraser University, Burnaby, BC, Canada; (2) Sacramento Kings, Sacramento, CA, USA

† E-mail: *nathan.sandholtz@gmail.com*

Last year, the Cleveland Cavaliers took 329 contested mid-range jump shots with over 10 seconds remaining on the shot clock. What could have happened if they had taken these shots 20% less frequently over the season? We attempt to answer these types of questions by modeling possessions from the 2015-2016 NBA regular season as Markov chains realized from team-specific Markov decision processes. To account for the dynamic nature of a basketball possession over the shot clock, we model the transition probabilities as a tensor exhibiting correlation in time. We assume the observed transition counts are multinomially distributed, governed by latent multivariate Gaussian distributions in order to explicitly impose a temporal covariance structure. We fit our model with STAN, using STATS SportVU optical tracking data. The draws from the transition probability tensor posterior distribution then serve as inputs in our regular season simulator.

We validate our simulation method by showing that we accurately recover the 2015-2016 transition counts for all intermediary and terminal states when simulating under the Cavs observed shot policy. To culminate, we simulate seasons under "alternate" shot policies proposed within the basketball analytics community and explore the net changes in production under these alternative policies.

UTILIZING ANALYTICS TO REDUCE NFL OFFENSIVE PLAY TYPE PREDICTIONS TO A MENTAL CHECKLIST

Yakubov, Ronen; Fernandes, Craig; Li, Richard; Prasad, Amrit; Chan, Timothy

University of Toronto, Toronto, Ontario, Canada

† E-mail: *ronen.yakubov@mail.utoronto.ca*

A great deal of past work and analyses have been conducted to try and predict with high accuracy whether an NFL offense will pass or rush the ball. Despite this past research, little has been done to effectively help defensive coordinators since all existing models require the use of a computer to perform the prediction. Since the NFL does not allow coaches to use any

external technological devices, we set out to find the best way to help defensive coordinators gauge their opponent's offense without the use of any such devices. We first used dozens of intricate attributes to train a neural network model on data from the 2012-2016 NFL regular seasons to achieve a prediction accuracy of 73.8%. We then analyzed passing predictability for each NFL team to inform coaches which teams deviate from expectation and which are more predictable. Finally, we created a more simplified model, using only a couple of feature variables that are easily observable by the defensive coach in real time, that predicts the opponent's offensive play with 72.3% accuracy. Doing so, we have created a model that does not require a computer, but can instead be executed through a mental checklist of what a defensive coach can visualize on the field. Thus allowing for predictive analytics to leave the domain of theoretical research, to a more practical application in the NFL.

NFLWAR: A REPRODUCIBLE METHOD FOR OFFENSIVE PLAYER EVALUATION IN FOOTBALL

Yurko, Ronald; Ventura, Samuel; Horowitz, Maksim

Department of Statistics, Carnegie Mellon University, Pittsburgh, PA, USA

† E-mail: *ryurko21@gmail.com*

The NFL lacks comprehensive statistics for evaluating player performance. One answer to this need was Total Quarterback Rating (Total QBR; Oliver, 2014). However, Total QBR is built on proprietary data, is not defined on a scale convertible to wins, and is only available for the QB position. We introduce our reproducible method for calculating Wins Above Replacement (WAR) for offensive skill positions in football, nflWAR, based on publicly available NFL play-by-play data from 2009-2016 accessible with nflscrapR. First, using our novel multinomial logistic regression expected points model, we estimate the "true" value for each play with expected points added (EPA; Burke et al., 2015). Then, similar to work measuring pitcher and catcher value in baseball (Judge et al., 2015), we extend a mixed model approach to isolate the EPA contribution made by individual offensive players and the opponent's defense as random effects while accounting for variables relating to the game situation as fixed effects. Next, we establish a robust way to define "replacement" level for each offensive skill given the historical play-by-play data. Finally, the expected points above replacement is converted to WAR (which we provide for all skill position players dating back to 2009) based on the observed relationship between points scored and wins. We emphasize that our reproducible nflWAR framework can be extended to estimate WAR for non-skill position players (e.g. linemen, linebackers, etc.) if provided with data specifying which players are on the field for both teams every play.