Decision Support Systems 88 (2016) 76–84

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

Predicting wins and spread in the Premier League using a sentiment analysis of twitter

Robert P. Schumaker [a], A. Tomasz Jarmoszko [b], Chester S. Labedz Jr. [c]

[a] Computer Science Dept., University of Texas at Tyler, Tyler, TX 75799, USA
[b] Management Information Systems Dept., Central Connecticut State University, New Britain, CT 06050, USA
[c] Management & Organization Dept., Central Connecticut State University, New Britain, CT 06050, USA

ARTICLE INFO

Article history:
Received 11 January 2016
Received in revised form 30 May 2016
Accepted 31 May 2016
Available online 8 June 2016

Keywords:
Business intelligence
Topic:
Sentiment analysis
Twitter
Sports analytics
English Premier League
Crowdsourcing

ABSTRACT

Can the sentiment contained in tweets serve as a meaningful proxy to predict match outcomes and if so, can the magnitude of outcomes be predicted based on a degree of sentiment?

To answer these questions we constructed the CentralSport system to gather tweets related to the twenty clubs of the English Premier League and analyze their sentiment content, not only to predict match outcomes, but also to use as a wagering decision system. From our analysis, tweet sentiment outperformed wagering on odds-favorites, with higher payout returns (best $2704.63 versus odds-only $1887.88) but lower accuracy, a trade-off from non-favorite wagering. This result may suggest a performance degradation that arises from conservatism in the odds-setting process, especially when three match results are possible outcomes. We found that leveraging a positive tweet sentiment surge over club average could net a payout of $3011.20. Lastly, we found that as the magnitude of positive sentiment between two clubs increased, so too did the point spread; 0.42 goal difference for clubs with a slight positive edge versus 0.90 goal difference for an overwhelming difference in positive sentiment. In both these cases, the cultural expectancy of positive tweet dominance within the twitter-base may be realistic. These outcomes may suggest that professional odds-making excessively predicts non-positive match outcomes and tighter goal spreads. These results demonstrate the power of hidden information contained within tweet sentiment and has predictive implications on the design of automated wagering systems.

1. Introduction

Predicting the outcomes of sporting events has a long and rich tradition. Since ancient times people have designed methods to divine natural and physical events. Today the urge to successfully predict still grips gamblers and academics alike. Prediction is no longer an art and probability is now considered a complex science. The most difficult aspect of prediction rests with identifying the relevant parameters and separating them from the noise of the event. Critical parameters are sometimes difficult to identify or measure, are constantly changing, or are not yet fully explored. The inability to correctly identify the most relevant parameters can sometimes lead to crippled systems relying on unimportant data or, worse, may create forecasts not based on sound science (e.g., basing predictions on the color of a uniform).

One way to simplify this problem of choosing and weighting parameters is to implement crowdsourcing as a forecasting tool. In James Surowiecki's seminal book, *The Wisdom of Crowds* [1], he made claim that large groups of individuals are better at making forecasts in conditions of uncertainty than are domain experts. This stems from collective intelligences, on the whole, being better able to properly sift through and analyze data than an individual. About the same time, another milestone book, *Moneyball* [2], popularized the use of statistics and sabermetric techniques (a quasi-scientific methodology of identifying relevant sports metrics, applying and refining them) in sports. Academic focus was not too far behind as the field of sports analytics gained popularity [3].

Twitter has been a boon to academic research with rich crowd-based datasets that can be easily collected and analyzed. Its data have been used to make predictions on phenomena as diverse as crime [4], the stock market [5], political elections [6], public opinion polls [7], public health [8] and movie sales [9]. The lure of twitter for academic research is two-fold. It provides a rich topical memory in the form of author-annotated hashtags, and provides a record of trends in public perception. Coupled together, academics can mine the twitterverse (i.e., universe of twitter data) and identify valuable insights.

Our research aims to demonstrate a crowdsourced system that can extract sentiment information from twitter to make match and point spread predictions in the English Premier League. Further, we analyze specific sentiment components such as tone and polarity, use them to calculate the degree changes in club-level sentiment and predict the magnitude of match goal differentials.

E-mail addresses: rob.schumaker@gmail.com (R.P. Schumaker), jarmoszko@ccsu.edu (A.T. Jarmoszko), clabedz@gmail.com (C.S. Labedz).

http://dx.doi.org/10.1016/j.dss.2016.05.010
0167-9236/

The rest of this paper is framed as follows. Section 2 provides an overview of literature concerning crowdsourcing, sentiment analysis and relevant studies. Section 3 presents our research questions. Section 4 introduces the CentralSport system and explains its various components. Section 5 sets up the Experimental design. Section 6 details the Experimental results and discussion. Finally Section 7 presents study conclusions and suggests further extensions of this stream of research.

2. Literature review

Crowdsourcing is a tool through which the average of crowd forecasts is used to predict future events [1]. In sports, this forecasting behavior generally equates to wagering on favorites and has been found to be a fairly accurate and reliable indicator of expectations. In a study of UFC fights, crowds were better able to predict wins (85.7%) than were bookies (67.6%) [10]. In a study of the wagering on matches in the Bundesliga (Germany's premier football league), crowds were found to be more accurate in their forecasts than bookies [11]. In a study of the FIFA World Cup 2006 tournament, crowds were also better able to predict winners than were pre-tournament rankings or random chance [12]. All three studies suggest that crowds were able to collectively make more accurate forecasts by weighting the data, not scientifically, but naïvely. Their decision-making contrasts with the weighting schemes designed by experts, which are driven by experience and previously seen patterns of data and profit generation. While crowdsourcing has demonstrated itself as an effective prediction tool, critics observe that some bettors may simply select the crowd favorite rather than evaluate the data independently [13]. This reinforcing behavior could lead to over-valuing the crowd favorite and can have an impact on accuracy. However, empirical evidence has shown that this typically encompasses a minority of wagering activity [13].

2.1. Odds-makers and wagering

Before two clubs take to the soccer playing field, or pitch, odds-makers will set a betting line in an attempt to draw an equal currency amount of wagers on each club. By balancing the wagers, in effect the losing side of the wager pays the winning side minus the sportsbook's[1] commission. Should the line become unbalanced, the sportsbooks are responsible for the difference and this imbalance may cause them a monetary loss. If one club is heavily favored, the sportsbook will increase odds on the less favored club to give bettors an incentive to wager longshots and rebalance the line.

One type of popular wagering system is the Moneyline. In this system, clubs with negative values are favored and clubs with high positive values are longshots. Odds and payouts are based on a unit of £100. For example, Arsenal and Swansea may have a Moneyline of −220 and +550 respectively. For the bettor on Arsenal (the favorite) they would need to wager £220 to win £100. For the Swansea bettor, they would wager £100 in a bid to win £550. The odds-makers attempt to gauge betting interest on the match and adjust the Moneyline to balance the monetary amounts wagered on both clubs.

Once odds are initially set, odds will move in response to the currency amount of wagers to continually balance the odds-makers match balance sheet. Because there are a variety of odds-makers with which to place wagers, the amount of currency wagering between clubs may differ between sportsbooks. This will lead to differences in odds between books. Typically the sportsbook with the more favorable odds will attract more wagering and will thus force their odds to return to market equilibrium.

[1] We use the terms odds-makers and sportsbooks interchangeably.

2.2. Social media and prediction

There has been much academic interest in using social media to make predictions. These predictions have crossed a diverse number of domains because of social media's rich crowd-based datasets that can be easily collected and analyzed. These areas have included crime, movie sales, politics, the stock market and sports. In a study of social media prediction and crime, twitter content was topically clustered into distinct discussion areas, correlated with the geo-location of the tweet and fed into a crime prediction model to demonstrate better predictive performance in 19 of 25 crime types [4]. Even though the tweeters were not making predictions themselves of crimes, their topics of discussion were a decent predictor.

In a study that correlated social media attention to movie sales, twitter content was found to be a good predictor [9]. In particular, positive twitter content was associated with higher movie sales whereas negative content was associated with lower movie sales. The authors also noted that tweets expressing an intention to watch a particular movie had the strongest predictive effect. In this case, tweeters were expressing their intention to watch or not watch a particular movie. This differs from the crime prediction study where the topics of the tweets themselves were used for prediction.

In US politics, twitter tweet counts and sentiment have been used to predict voter outcomes. In a study of the German Federal elections, Tumasjan et al. used a simple and easy to implement method of counting tweets that mention a candidate or political party [14]. Their reasoning was that tweets mentioning a candidate or party indicated their voting intention. This method was fairly accurate when applied against German federal elections with an error rate of 1.65%. When more complex methods were investigated such as using a sentiment analyzer to further determine voter intention, the results were not as precise [15]. Although the results are dependent upon the methods of how sentiment was captured and analyzed. It was further noted that sentiment polarity methods, at the time, were not sophisticated enough to recognize political language nuances, had poor performance and produced unacceptable errors [16].

Another political study investigated using a moving average of candidate, or elected official, tweet sentiment as a replacement to traditional polling services [7]. This work noted that natural language processing techniques achieved an 80% correlation.

In a study of the sentiment of financial news articles and the stock market, Schumaker et al. used the article sentiment as a method for predicting the magnitude of stock price movements immediately following article release [17]. Their work found that articles with a negative sentiment were easiest to predict, netting a 3.04% trading return using a simple trading engine.

Fans post tweets in order to express their personal feelings, most fundamentally (as we collected them) about the strengths, weaknesses and prospects of the team they follow and its next opponent. Admittedly, the tone and polarity of fans' tweets likely do not affect match results, except in cases in which extraordinary fan base sentiment might exceptionally motivate or demotivate a team. Fans' tweets by themselves are not likely to influence betting lines offered to bettors, which of course potentially includes those fans. Nonetheless, the tones and polarities of opponent fans' tweets may modestly affect initial betting line odds or later adjustments thereto, as we note elsewhere.

As tweets suggest the expected outcome on the field (i.e., the full-time score line), the wisdom of crowds premise becomes more credible. Many thousands of fans, well versed through years as footballers themselves before advancing age and injuries transformed them into amateur pundits, bring considerable collective intelligence to sentiment crowdsourcing. This fan base is sufficiently diverse, decentralized through the reach of the Internet, able to be summarized and rabidly independent. In expressing their sentiments about real events on the turf pitch from odds-distorted results on the shadow field of wagering, fans' tweeted views contain useful raw information about future score lines.