

# Sweet FA: Sentiment, Swearing and Soccer

Emma Byrne  
ebyrne.uk@googlemail.com

David Corney  
School of Computing & Digital Media  
Robert Gordon University, Aberdeen  
d.p.a.corney@rgu.ac.uk

## Abstract

The sentiments expressed by football fans in the stories that they tell are often intensified by the use of swear words. Football provides a useful test bed for sentiment analysis due to the symmetric nature of events in matches: what is good for one team is bad for the other. We can relate social media messages to the narrative that fans of a given team might be expected to construct. We use these features of football-related tweets to investigate some common assumptions about swearing as a sentiment marker on social networks. The results demonstrate that swearing and other sentiment markers depend heavily on context, and that understanding this context is essential if sentiment is to be detected faithfully. We also show that swearing is not always indicative of negative sentiment.

*Please note that strong language is used throughout this paper!*

## 1 Introduction

Football fans construct a shared narrative about their team's performance. Fanzines, terrace talk and pub conversations have always formed part of fan identity and communication. Online social networks have added another venue in which those shared identities can be constructed.

Football stories shared by fans are emotional, and narratives shared via Twitter are freighted with sentiment. The language used by fans is often scattered

---

*Copyright © by the paper's authors. Copying permitted only for private and academic purposes.*

In: S. Papadopoulos, P. Cesar, D. A. Shamma, A. Kelliher, R. Jain (eds.): Proceedings of the SoMuS ICMR 2014 Workshop, Glasgow, Scotland, 01-04-2014, published at <http://ceur-ws.org>

with obscenities which provide colour and also intensify expressions of sentiment. In this paper, we describe an initial study exploring the way that sentiment is intensified by swearing in Twitter messages from football fans during games. By limiting our analysis to messages containing swear words, we focus on the more intense expressions of sentiment.

Football provides a particularly interesting test bed for sentiment analysis due to the symmetric nature of events within matches: what is good for one team is (equally) bad for the other. For example, a goal has two opposing "valences" which makes it possible to search for and analyze approximately equal volumes of positive and negative sentiment per event. For simplicity, we restrict our analysis to FA English Premier League, which is the most-watched football league in the world.

## 2 Related work

Swearing is a feature of social networks and has been for a long time (by internet standards.) In his 2008 paper [9], Thelwall studied the occurrence of swearing in MySpace profiles. Swearing is not only widespread (67% of UK MySpace profiles of 16-19 year olds contained some form of swearing) but varied. Some of the swearing was self-directed ("tehe i am sorry.. i m such a sleep deprived twat alot of the time! lol") or clearly affectionate ("Chris you're slacking again !!! Get the fuck off myspace lol !! you good anyway?") It was frequently used in an approbatory fashion ("Thanks for the party last night it was fucking good and you are great hosts." "That 50's rock and roll weekender was fucking mint!") This paper demonstrates that, for young people at least, swearing is part of their performed identity online. Not only that, swearing is multipurpose, used to demonstrate amusement, affection and self-deprecation as well as negative sentiment.

Swearing is known to be a response to – and a mediator of – emotions [8], as are other forms of language that are *prima facie* abusive or insulting [1, 3, 7]. In

workplaces in the UK and New Zealand, for example, “jocular abuse” is part of team bonding. Swearing, piss-taking and other forms of abuse are not only tolerated, but form an essential part of the workplace group dynamic.

Even in a setting as replete with strong language as football, there is a strong imperative to keep the reality of the discourse of swearing off the airwaves. Consider this tweet for example:<sup>1</sup>

Ivanovic over kicked the ball and a Chelsea fan angrily swears “fucking cunt” and the commentator emotionally apologises Lmao #CFC #Setanta1.

However, as we demonstrate here, Twitter as a social network provides a much less filtered view of fans’ actual language. Swearing is used regularly to intensify sentiment and is, as such, a relevant marker in sentiment analysis.

Sentiment analysis is widely used by brand managers to monitor public perception of their products and services, including Amazon reviews and Ebay feedback. However, much of this work assumes that swearing typically represents negative sentiment. Hu and Liu have created lists of positive and negative words associated with opinion or sentiment [4]. In the current version of list list<sup>2</sup>, words such as “shit”, “fuck” and “damn” are included in the negative set and not in the positive set. While they argue that having such an opinion lexicon is not sufficient for sentiment analysis, it is nonetheless a useful tool [5].

A more sophisticated analysis of detecting sentiment from swear words has been presented by Maynard et al. [6]. They also use a gazetteer of opinion words, including swear words. They recognise that swear words may be used to intensify the expression of a sentiment, be it positive or negative. However, if a sentence contains swear words and no words recognised as implying a positive sentiment, then they assume the sentence is negative. While this may often be an effective approach, we would add that the wider context can often be used to interpret the sentiment even of isolated swear words as we discuss below (Section 4.8).

Twitter has been used to help automatically detect events during football matches [10]. That system collects tweets based on hashtags and then detects spikes in the volume of tweets collected which they associate with major game events. For each spike, they analyze the words of the tweets and classify the event using machine learning. They consider a fixed range of events (goals, own-goals, red cards, yellow cards and substitutions) and compare these classifications to the official match data to evaluate their system. They also clas-

sify individual tweeters as fans of one team or another by counting the number of mentions of each team over several matches, similar to our approach (Section 3.2).

Similar methods have been used to classify events during American football (NFL) matches [11]. Tweets were collected based on team names and NFL terminology. Events were also detected by finding spikes in the volume of tweets and each event was assigned to one of a fixed number of classes, in this case using lexicographic analysis. Their system was very effective at detecting the most significant scoring events such as touchdowns, but was less effective at finding less significant events like interceptions and field goals.

## 3 Methods

### 3.1 Collecting Tweets

On a typical Saturday afternoon, 5 or 6 English Premier League matches kick off at 3pm, at the same time as many matches from lower divisions. For three consecutive weeks, we collected public tweets discussing the matches. For each match-day, we crawled Twitter using their standard Streaming API<sup>3</sup> and filtered using a total of 28 hashtags. These are the standard abbreviation hashtags of the 20 teams (e.g. #CFC, #MFC) along with certain widely-used hashtags that indicate team support (e.g. #KTBFH for Chelsea and #YNWA for Liverpool, both based on popular supporters’ chants). We collected tweets starting 30 minutes before kick-off and continued until 30 minutes after final whistles. On average, we collected 125,070 tweets per match-day. Our analysis focuses on period from 3pm to 5:00pm each Saturday. This includes 90 minutes of football, the half-time break (c. 15 minutes) and a few minutes of post-match response.

Throughout this analysis, we have made frequent use of the mainstream media accounts of matches, for example to verify when goals were scored, players were sent off and other noteworthy events occurred. In particular, we used the live-blogging “minute by minute” commentaries provided by the BBC for the three Saturday’s in question, namely 7/12/2013<sup>4</sup>, 14/12/2013<sup>5</sup> and 21/12/2013<sup>6</sup>. Knowing these match events and the team supported by each fan (see below), we can derive an “emotional ground truth” which we expect to then be reflected in the language used in fans’ tweets.

### 3.2 Linking Tweets to Teams

We identified the team that each Twitter user supports (if any). Initial manual inspection of a number of

<sup>1</sup>This, and all other tweets quoted and analysed in the paper are available on request from the authors.

<sup>2</sup><http://www.cs.uic.edu/~liub/FBS/sentiment-analysis.html#lexicon>, dated 12/3/2011

<sup>3</sup><https://dev.twitter.com/docs/api/1.1/post/statuses/filter>

<sup>4</sup><http://www.bbc.co.uk/sport/0/football/25264555>

<sup>5</sup><http://www.bbc.co.uk/sport/0/football/25365181>

<sup>6</sup><http://www.bbc.co.uk/sport/0/football/25463334>

tweets suggests that fans tend to use their team’s standard abbreviation hashtag greatly more often than any other teams’, irrespective of sentiment. This identification was then confirmed by inspection of the text of a sample of the tweets. We therefore define a fan’s degree of support for one team as how many more times that team’s abbreviation is mentioned by the user compared to their second-most mentioned team. For each user, we aggregated all their tweets and counted the total number of times they mention each team. Here, we include as “fans” any user with a degree of two or more and treat everyone else as neutral. Having assigned fans to teams, we can then associate specific tweets with specific teams even when no team is mentioned, if other tweets from the same person make their allegiance clear.

To evaluate this method, we randomly selected 100 tweeters that our algorithm had predicted to be fans of various teams. We then examined the tweets in our collection from each person to determine which team they expressed support for, if any. We manually labelled them as supporters of a specific team, neutral or unclear (e.g. due to non-English tweets). Of the 100 people thus analyzed, 93% were correctly assigned to teams by the algorithm; 7% appeared to be neutral commentators who showed no clear preference for any team and one was a spam account unrelated to football (but using a team hashtag to attract clicks). Our algorithm mis-assigned them to which specific team they happened to mention most often. In no case was a clear fan of one team assigned to any other, giving us a strong confidence in the rest of our analysis.

### 3.3 Filtering Tweets by Use of Swearing

The use of twitter by fans during matches results in an average of over 125,000 tweets per game containing the team’s hashtag. However, we are particularly interested in those tweets that contain not only indications of sentiment, but indications of a high-intensity sentiment. For this reason we filter the tweets that we collected using the stems of the two most common swearwords in the English language: “shit” and “fuck.”<sup>7</sup> Post filtering, our corpus consisted of an average of 6483 tweets per match-day, meaning that more than 1 in 20 tweets (5.36%) from football fans contain the words ‘shit’ or ‘fuck’ or their derivations.

We further filtered these tweets manually to remove messages that were not full-on “fannishness.” So, for example, tweets asking about television coverage or discussing betting on the match were not assessed for sentiment. We then manually assessed the remaining tweets, with both authors coding the sentiment(s) in

---

<sup>7</sup>Because we used these terms as stems, the filter also matched “shitting, batshit, shite, fucking, fucker, fucked...”

sets of tweets.

Once we had collected our corpus of tweets from fans watching matches we could begin the process of relating them to a narrative. That narrative is multi-level and consists of discourse about events, games and the English Premier League competition as a whole. The competition and the game are easily defined: the English Premier League runs from August to May each year and each of the 20 teams plays the other twice, once at home and once away. The winner is decided on the number of points acquired from the matches and ties for position are decided using the number of goals a team has scored minus the number they have conceded.

Matches, too, are clearly bounded. We know which teams are involved, the ground at which they are playing, the start and end time of the match.

Events are more difficult to define. There are some canonical events that are noted as part of the statistical record: for example, goals, fouls, bookings, free kicks and penalties are all recorded with an associated time stamp. However, some events are not part of the official record, despite being matters of significance to the fans. Take the example, shown later, of the Liverpool captain Steven Gerrard suffering a recurrence of an old injury, which threatened to prevent him taking part in the next few matches. We can assign a timestamp to this event, but we need to mine the event commentary in detail to infer what has taken place.

Even more challenging to identify are the events that occupy a timeline rather than a timepoint. For the sake of clarity we will call states that persist “fluents” and reserve the term “events” for those actions that change the state of a match. Examples of fluents include “the run of play” - an informal definition of which team is currently dominating in terms of possession.

### 3.4 Complicating factors

A number of factors make analyzing a typical Saturday afternoon’s football especially challenging.

First, numerous games are played simultaneously. Here, we’ve been considering the five or six Premiership matches being played from 3pm, but at that same moment, up to 36 other matches may be being played in the three other professional football divisions in England. There are also matches played in the (separate) Scottish league, not to mention other matches in other countries.

Second, there is a great variety in the response to different classes of event. The emotional impact of being awarded free-kick expires more rapidly than the emotional impact of scoring a goal from that; and the emotional impact of the goal expires less rapidly if that

goal has changed who is winning. Furthermore, events within a single match typically overlap to some degree and have their own duration. They are not discrete, point-sources as may be assumed in theoretical analysis.

Thirdly, assigning fans to teams is made harder by the tendency of some fans to use the hashtags of their opposing team in order to get the attention of fans of that team, as opposed to being one of them. Insults and banter are only effective if their target is aware of them.

Finally, it requires some judgement to assess the degree of sincerity of messages, due to the use of sarcasm, irony etc. Both authors are native English speakers not unversed in such matters, but nonetheless, some messages may have been misinterpreted.

## 4 Case studies

The tweets that we examined for these case studies came from a number of matches in the English Premier League that took place in the first half of December. This window of time can roughly be considered as mid-season: fans have had around 15 league games to assess their team's performance this season, but with over 20 games still to play, there is still plenty to play for. This is the time of year that matches begin to be referred to as "real six-pointers"<sup>8</sup>.

The matches that we draw examples from are: Stoke City-Chelsea, Liverpool-West Ham (both 7/12/2013), Cardiff-West Bromwich Albion, Chelsea-Crystal Palace (14/12/2013) and Manchester United-West Ham United (21/12/2013). For context, the top four positions in the league on December 9<sup>th</sup> were occupied by Arsenal, Liverpool, Chelsea and Manchester City respectively. Manchester United were in (an unusually low) 9<sup>th</sup> place, just above Hull and Stoke. West Ham and Crystal Palace were both in the bottom four just below West Brom and Cardiff. It is also worth noting that different clubs have very different numbers of fans, with average home attendances varying between 20,000 and 75,000. This is likely to be reflected in the number of tweets in our collection related to each club.

From the analysis of only those tweets containing swearing we observed that sentiment is expressed in a complex and sometimes counter-intuitive manner, several aspects of which we now discuss.

---

<sup>8</sup>In the English football leagues, teams are awarded three points for a win, a point for a draw and no points if they lose. When teams are playing against opponents that are close to them in the league table it is important not only to win all three points, but also to deny the other team the opportunity to score any points. Hence the fact that these games are referred to – with more poetry than numeracy, perhaps – as "real six-pointers."

### 4.1 "We're Shit and We Know We Are"

In the English Premier League it is not uncommon for fans to be highly critical of their own team. This criticism is levelled at individual players, the team as a whole, the manager and – very occasionally – the fans themselves.

This is not an uncommon occurrence and – for some sets of fans in particular – accounts for the majority of swearing use in large segments of the match. For example, in the opening 42 minutes of the Liverpool vs West Ham United match (7<sup>th</sup> December), over half of the tweets we collected from Liverpool fans were gripes about their manager, Brendan Rodgers, and several team members. For example, the following tweets all arrived within a few minutes and all were sent by Liverpool fans:

15:17:57 @lfc Joe allen<sup>9</sup>, easily beaten un the middle of the park. Fuck Rodgers. #ynwa #lfc #LivWhu  
15:21:41 Mignolet<sup>10</sup> bails our shite defence out once again. #LFC  
15:25:31 This is like the hull game, creating fuck all atm and mignolet keeping us in the game #LFC  
15:26:17 Please learn to pass sterling<sup>11</sup>! Or fuck off #lfc  
15:27:17 This game is driving me mad. Just fucking score already damnit. #LFC  
15:30:07 Amazes me that people don't understand that we are shite, one player being world class doesn't make a quality team #LFC  
15:31:30 The fatc people on here are praising defenders and Allen means our attack is shite and nothing more. #LFC  
15:32:28 For a team that focuses so much on passing, passings been shite today #lfc  
15:33:28 Henderson<sup>12</sup>, sterling are shit #lfc  
15:34:18 Raheem Sterling knows fuck all. He should join Newcastle or Swansea. #lfc

When criticism is sparked by a particular event fans may single out an individual for praise while criticising the team as a whole. For example, the tweet "Mignolet bails our shite defence out once again," is simultaneously a whinge about the team and praise for an individual. Inversely, when Chelsea keeper Petr Čech let in an equaliser in the Chelsea-Stoke game (7<sup>th</sup> December), he was singled out for criticism by the Chelsea fans, as this example illustrates:

15:43:53 What the fuck was you doing Cech. The team has played brilliant and then you go and do that. FFS!! Come on boys. Heads up.. #CFC

Such examples demonstrate that:

---

<sup>9</sup>Liverpool midfielder.

<sup>10</sup>Liverpool goalkeeper.

<sup>11</sup>Liverpool winger.

<sup>12</sup>Liverpool midfielder.

- Fans of a given team are likely to use swearing in disapprobation of their own team, players or manager.
- We also note that fans in these English Premier League games were much more likely to express a negative sentiment that was intensified by swearing about their own team than about an opposing team.
- Therefore negative sentiment intensified by swearing provide strong evidence of a tweeter’s affiliation but, perhaps counterintuitively to someone not familiar with “terrace culture”, they are most likely to be affiliated with the team that they are criticising.

#### 4.2 Off, Off, Off: on Bad Sportsmanship and Bad Players

During the Liverpool-West Ham game of 14<sup>th</sup> December, West Ham captain Kevin Nolan received a red card for deliberately stamping on Liverpool’s Jordan Henderson. This meant that he missed the rest of the match and – as it was his fifth red card of the season – he also received a three-match ban.

The tweets from both Liverpool and West Ham fans indicated strong disapprobation. Liverpool fans reacted predictably on Twitter:

16:39:29 Go on fuck off Nolan you prick #LFC  
 16:39:36 Kevin Nolan you fucking twat. #LFC  
 16:39:47 Fucking dirty bastard #LFC  
 16:40:02 NOLAN YOU FUCKING PRICK! #LFC  
 16:40:07 Fuck you nolan.#lfc  
 16:40:16 Nolan u dirty fuckin cunt! #LFC  
 16:40:42 Fuck off Nolan! Deserved red #LFC #PremierLeague

While we might usually expect a certain amount of dismay at the loss of the captain for three matches from a team’s fans, Kevin Nolan’s ban was greeted with an unexpected amount of positive sentiment by the West Ham fans.

16:39:18 FUCKING HAPPY DAYS KEVIN NOLAN IS GOING TO BE BANNED FOR 3 GAMES!!!! GET IN THERE HAHAHA #thereisagod #whufc  
 16:39:34 thank fuck Nolan banned for a while #coyi #WHUFC  
 16:40:08 Nolans last game for us? Let’s fucking hope so! #WHUFC  
 16:40:38 Yes come on now he is banned Wahoo!fuck off Nolan u prick #coyi #whufc #whu

As we will see when considering humour (Section 4.6), this may be more to do with Kevin Nolan’s performance as a player than his infraction of the rules.

From this we can conclude that:

- Fans rarely criticise players from the opposing teams for poor performances but they will criticise them for foul play
- Fans may criticise their own players for foul play and are apparently keen to do so if the player is performing poorly.

#### 4.3 Fuck’s Sake! v. Fuck, Yeah! Or Swearing Not Necessarily Considered Harmful

As noted earlier, there is a tendency to consider swearing in social media messages as a sign that the author of that message is expressing a negative sentiment [4, 6]. However, inspection of the content of these tweets demonstrates that this is sometimes wrong as these two tweets from Liverpool fans demonstrate:

16:04:57 I NEVER get to see #LFC and the #Oilers win on the same day but I’m feeling good about it today. Don’t let me down you fucks. #fucks  
 16:06:00 Fuckin get in reds, 2-0 #LFC

The tweet from the Oilers fan at 16:04 contains an affectionate mock threat, with the tongue-in-cheek nature of the tweet being reinforced by the use of a dafttag<sup>13</sup>. The more succinct “Fuckin get in reds, 2-0 #LFC” is a more straightforward expression of celebration.

However we can pair that second tweet from a Liverpool fan with a tweet from a West Ham United fan in response to the same event:

16:05:20 We are fucked... Hello championship!<sup>14</sup> Big Sam<sup>15</sup> out!! #whufc

Further analysis reveals the frequency of differing phrases using the word ‘fuck’ both positively and negatively. Table 1 summarises the total frequency with which each phrase is used across all matches for the three weeks shown. They include variant word forms (e.g. “fuck sake”, “fuck’s sake”, “fucks sake”)

From this we can infer that

- The same event, seen as positive by one tweeter and negative by another, can prompt tweets that contain swearing in both cases.
- While swearing may indicate that the sentiment of the tweeter is intense, it does not unambiguously demonstrate a positive or negative valence

<sup>13</sup>Here we define “dafttag” as a hashtag that is added, usually to the end of a tweet that communicates the tweeter’s sentiment in a manner that is self-parodying or an expression of the tweeter’s covert message content. These dafttags often indicate sarcasm, humour or other modes of communication and indicate to the reader that they should look beyond the tweet’s surface meaning.

<sup>14</sup>The division below the Premiership.

<sup>15</sup>Sam Allardyce, the West Ham manager.

Phrase	7/12	14/12	21/12	Totals
Fuck sake	80	43	54	Negative
Fuck this/fuck that	30	17	14	238
Fuck yeah	11	3	5	Positive
Fucking get in thank fuck	28 42	8 6	11 8	244

Table 1: Frequency and usual sentiment of various phrases collected during three Saturdays in Dec 2013.

to that sentiment, so we cannot universally classify swearing as positive or negative.

#### 4.4 And Another Thing... Multiple Sentiments in 140 Characters or Less

For a communication act that is limited to 140 characters, football fans’ tweets can display surprisingly rich and complex sentiments. For example, in the case of the Chelsea-Stoke City game (7<sup>th</sup> December), as Stoke equalized, one Chelsea fan managed to express disappointment, despair and hope in the same tweet:

15:45:00 Every fucking match something stupid gets #cfc in trouble. I’m hoping for quick response.

Likewise in the Liverpool–West Ham game, Liverpool fans were pleased with a West Ham own-goal that took Liverpool into the lead, but this didn’t prevent them from expressing their disappointment with their own team’s performance, particularly by the players Sterling and Allen:

15:42:39 Luck as fuck goal but I’ll take it after that Sterling miss. #LFC

15:44:18 Even though its an own goal we deserve the lead. Dominating the match but how shit is Joe Allen, seriously #LFC

Fans on Twitter also use simile, wordplay and allusion to comment on games. For example, this Chelsea fan is expressing displeasure at the scoreline against Crystal Palace, a team not generally thought of as strong opponents:

16:36:34 This is bollocks... we at home to Palace, not away to Barcelona, fucking painful just waiting for them to equalize. #CFC

This user alludes to the relative strengths of Chelsea, Palace and Barcelona<sup>16</sup>, the relative ease of playing at home versus away, the scoreline (a one goal difference between Chelsea and Palace) and the run of play, all within the space of a single tweet. The use of allusion – which relies on the reader’s understanding of

<sup>16</sup>Barcelona FC is widely regarded as playing some of the most beautiful football in the world, as well as being one of the most objectively successful teams. Crystal Palace, based in a London suburb, have been relegated from the English Premier League more often than any other team.

various elements of context – makes this a particularly information-dense message.

From these examples we can infer:

- Tweets, although brief, can contain multiple, sometimes oppositely-valenced sentiments.
- That individual tweets may therefore be too broad a unit of analysis if we wish to identify sentiment.

#### 4.5 Around the Grounds: The Story is not Restricted to the Match

While in the main, fans tweet about their own team and the match that they are currently engaged in, occasionally they will tweet about other concurrent games. This makes automatic event detection a challenging task. A Liverpool fan watching the game against West Ham may tweet about an event in the Chelsea-Stoke City match; for example, when Ousama Assaidi, on loan from Liverpool to Stoke, scored the final goal in a match that ended Stoke City 3 - 2 Chelsea, the Liverpool tweeters reacted exuberantly:

16:49:04 YES STOKE I FUCKING LOVE YOU! And an ex liverpool player does it ahahahahaahhaaha #LFC

16:49:33 Fucking Assiadi go on lad! Haha doing his club a favour #lfc

16:49:55 Assaidi you absolute beaut!! Fuck you Chelsea!!! #Assaidi #LFC

16:50:05 #Assaidi you fucking beauty!! #LFC #cheers #ChelseaKiller

16:50:13 ASSAIDI YOU FUCKING BEAUTY!!!!!!!!!!!! #LFC

At the beginning of the Chelsea-Crystal Palace game the following week (14<sup>th</sup> December), one Chelsea fan tweeted the following:

15:03:48 Get drunk as fuck & wake up to Arsenal losing 6-3 & a chance to cut the lead to 2 points. Fuck yes. #CFC

For context, Chelsea and Arsenal are both London teams with an historical rivalry. They also both occupied places in the top three of the table throughout December and were jockeying for the upper hand.

From these examples it is possible to infer that

- Fans tweet about games that their teams are not playing in.
- There appears to be a higher likelihood of this happening if there is either a positive link between the teams (e.g. a player is on loan to another team) or a negative link (e.g. a long-standing rivalry or a close position in the league table.)
- Thus the “story” from a particular tweeter’s point of view may not be restricted to a given game.

Rather it is anything that affects their team's position in the table, or that has some relationship to an affiliation or a rivalry with another club.

#### 4.6 Funny Old Game: Humour in Fan Tweets

The British are proud of their ready wit, even (perhaps especially) in times of emotional stress. It is unsurprising, then, that we find a lot of humour in the tweets sent by English Premier League football fans. Humour is interesting from a storytelling point of view as it often relies on ambiguity or wordplay for its effect; thus we need to parse humorous twitterances with the same care that we parse humorous utterances.

Examples include the following from the Liverpool–West Ham match where the first two goals came from West Ham own goals.

16:06:33 fuck sturridge<sup>17</sup>...this own goal is some player #LFC

And where West Ham captain Kevin Nolan was expected to turn in a disappointing performance, this tweet was very widely retweeted:

15:59:01 Kevin Nolan is very adaptable. He is equally shit in a number of positions #rubbishplayer #lfcshwhu #whufc

When West Ham pulled back a goal later in the match to bring the scoreline to 2-1, Liverpool fans offered the following:

16:35:58 #LFC in fuck this up "shock"

16:36:21 You know if Moses<sup>18</sup> was a horse he'd be a Pritt Stick by now #fuckingawful #LFC

Likewise in the Chelsea–Stoke City game, fans reported some on-terrace humour. In the last decade, Chelsea have won at least one major national or international competition in all but one season. Stoke City, however, have only managed to win the Autoglass Trophy – a competition for teams in the bottom two divisions of the English top-flight football league – since 1972<sup>19</sup>.

15:26:39 #scfc fans: Ur gonna win fuck all

#cfc fans: U've never won fuck all

#scfc fans: Autoglass trophy we've won it 2 times #clasicBanter

From these examples we can infer that:

- Humour is an important part of fans' storytelling both at games and online. Effective attempts at humour are picked up and passed around as either retweets or reports of stadium banter.

<sup>17</sup>Daniel Sturridge, striker for Liverpool and the English national team and seen in a generally positive light. Here, 'fuck' can be interpreted as 'forget' or 'never mind'.

<sup>18</sup>Victor Moses, Liverpool winger.

<sup>19</sup>This is not the kind of trophy that Premier League teams often boast of winning.

- Humour is not in and of itself a sign of positive-valenced sentiment. For example "LFC in fuck this up 'shock'" and "If Moses was a horse he'd be a Pritt Stick" are both jokes, but the originators of these jokes are not happy – as evidenced by the hashtag #fuckingawful.

- Jokes can rely on apparent or actual denigration of a player or team, usually one's own ("Fuck sturridge..this own goal is some player," "Autoglass trophy we've won it 2 times.") These are often examples of self-effacing humour from fans.

- Humorous tweets often carry a heavy freight of context and can be challenging to mine for sentiment without thorough understanding of that context.

#### 4.7 Sick as a Parrot, Fluffy as a Kitten: the use of Creative Language.

Football has a shared lexicon of formal and informal terminology<sup>20</sup> that has reached the level of cliché through widespread use on terraces and in newspaper coverage.

However, despite a fairly stable football cliché lexicon, football fans on social networks are inventive in their use of language. We see several examples of tmesis – the insertion of a word into the middle of another word or phrase.

16:48:46 Un-fucking-believable! #CFCLive

16:49:03 Assi-fucking-idi<sup>21</sup>!!!! what a strike son #scfc

We also see the use of relatively uncommon forms of swearing use, for example, the imperative form "motherfuck":

15:53:22 Mother FUCK this scoreline...

And the noun "fuckery":

16:49:47 This is some major fuckery... #cfc

As well as deliberate misspellings such as "bollox" – possibly used to turn a swear word into a "minced" variant<sup>22</sup>.

This Liverpool fan uses an unusual but evocative pair of similes to convey the change in emotional state that they have experienced through the recent portions of the game:

16:40:58 From being 2-1 and shitting our pants to being 4-1 and feeling as fluffy as kittens. #LFC

We also see this example of ironic litotes from West Ham fan who is commenting on another Twitter user's assessment of the performance of Manchester United and comparing it with their own team's performance:

<sup>20</sup>"A real six-pointer," "A game of two halves," "Sick as a parrot," etc.

<sup>21</sup>A misspelling of Oussama "Assa-fucking-idi", Stoke City winger on loan from Liverpool.

<sup>22</sup>A minced oath is a deliberate misspelling, mispronunciation or other mis-rendering of a word in order to render a euphemism.

15:47:51 "[USER\_REDACTED]: Haha man united lost again! They are so shite" wish #WHUFC was as shit as them

From these examples we must infer that:

- While tweets are short and ephemeral, they nevertheless contain rhetorical figures such as simile, litotes, irony, metaphor and allusion. They also contain wordplay in forms that include punning, euphemism and the use of uncommon parts of speech.
- As a result, we must be aware that tweets are not always simple utterances, and that they should be assessed with care.

#### 4.8 "I Fucking Hate Football Sometimes": the Importance of Assumed Context in Matchday Tweets

Football inspires a devoted following partly because it creates shared experiences on multiple timescales. For example, an event in a match ("Did you see Assaidi's goal?"), a match itself ("Can you believe we're 3-2 up against Chelsea?") and a competition as a whole ("I still think we'll be facing relegation come April") provide many opportunities for bonding over shared joys and sorrows.

However that intense sharing can baffle newcomers because so much common context is taken for granted. Both at the grounds themselves and on Twitter, fans assume that fellow fans will be aware of a team's recent performance history, a player's injury worries, or the likelihood of winning a competition. This may in part be deliberate obscurantism in order to highlight in-group vs. out-group differences. By demanding deep knowledge of a team's history before admitting someone to an inner social group of 'true fans', that group claims a stronger social identity.

This can lead to tweets that are impossible to assess for sentiment unless the assessor is aware of this same context. For example, from the Liverpool vs West Ham United game, these tweets come from Liverpool fans:

16:12:27 ahh fuck Gerrard #lfc

16:37:51 luis suarez you lil shit #lfc

16:40:40 Luis. Fucking. Suarez. Again. #LFC

16:42:54 Suarez Is UnFuckingBelievable #LFC #YNWA

In the absence of context it may appear that the Liverpool fans are venting their displeasure at Steven Gerrard and Luis Suárez. However, examining the context shows us that at around 16:12, Steven Gerrard picked up a hamstring injury – a recurrent problem that has plagued his career – and that Suárez scored twice, at 16:37 and 16:40.

There are also many tweets where the sentiment is not hard to infer, but where the context has been stripped by the tweeter. For example, during the Chelsea vs Stoke City game, shortly after Stoke scored to take the game to 3-2 in their favour, Chelsea fans responded:

16:49:59 I fucking hate football sometimes #CFC

16:50:34 Fuck Fuck Fuck #CFC #STKvCHE

In the Crystal Palace game, it is possible to infer that this fan is responding to some form of commentary:

15:16:38 "Dominated possession"?? Fuck off. #CPFC

But it is impossible to know who made the original comment, or what they were commenting upon.

From these examples we can infer that:

- Fans assume that their audiences share their contextual information.
- It is not possible to provide a thorough analysis of sentiment on Twitter if we do not also mine for context.

#### 4.9 Speed is of the essence: Urgency and entropy

When a significant event happens, such as a goal being scored, fans tend to communicate their responses with great urgency. This is shown by an increase in the number of tweets sent in the following minutes and with a simultaneous shortening of those tweets. To investigate this, we examined the goals from three randomly-selected matches from our collection. We selected the tweets sent by both sets of fans for five minutes immediately before each goal, and for five minutes immediately after. Although there is considerable variation between the matches and the goals, the pattern is clear: in the relatively uneventful period before a goal, few tweets are sent (mean = 128.58 per minute in these cases) and they are relatively long (mean = 78.04 characters). In the minutes following a goal, fans from both teams respond rapidly with a surge in the number of tweets (mean 448.1 per minute, an increase of 248%) that are typically short (down to 62.8 characters, a drop of 19.5%). The details of these events are shown in Table 2.

Typical examples of longer tweets sent when no goal has been scored recently include (from West Ham-Manchester Utd, 21<sup>st</sup> December):

15:32:35 Stoke, Southampton, West Brom etc all go to Old Trafford and look decent. We go there and still look shit! #WHUFC #GoingDown

16:10:00 This is the Manchester United we have been watching all these years. Scaring shit out of their opponents when they attack.#Mufc.#GGMU

16:13:57 Fuck sakes, Welbeck injured! What sort of



training is Moyes putting these boys to, getting injuries like picking cherries? #MUFC

16:04:08 Why the fuck is Taylor not on the wing?! Fat Sam aint got a bloody clue. Useless fat fuck #COYI #whufc

Typical examples of shorter tweets sent just after goal include:

15:36:18 FUCK YEAH ADNAN!! #MUFC

16:40:12 That was fucking offside..#MUFC

and the terse:

16:31:06 3-0 fuckers #MUFC

From these results we can infer that:

- Fans respond rapidly to significant events by sending short focussed messages.
- Fans send many messages when a goal is scored.

## 5 Conclusions

Through a careful analysis of a collection of tweets about football that contain swearing, we have shown that bad language is not always negative; that the wider context is often crucial to interpret meaning; and that perhaps counter-intuitively, some of the strongest sentiments expressed are self-critical.

The examples in this paper demonstrate that swear words are used to intensify both positive and negative sentiments ('fucking beauty' v 'fucking painful'). However, even when the sentiment seems apparent, widespread irony makes it necessary to consider context before interpreting the valence of the sentiment. As noted earlier, the tweet "Luis. Fucking. Suarez. Again." is actually a positive sentiment by a Liverpool fan in response to Suárez scoring, but the automated sentiment analysis systems discussed earlier [5, 6] would see "fucking" as the only sentiment-carrying word of the message and so denote the whole message as negative.

This is one case where fans implicitly assume that their audiences share their context, meaning that apparently ambiguous expressions of sentiment will in fact be correctly interpreted by their intended audience. We suggest therefore, that it is often impossible to accurately analyze sentiment on Twitter if the context of utterances is not also analyzed.

The narratives that football fans tell are simpler to interpret than many of the stories played out on social media. For example, there is a very low incentive for a fan to give a false signal, and we have access to a ground truth – in the form of a match report – that indicates the likely valence of the sentiments expressed at different points in the narrative.

However, the narrative itself is open-ended – even in the close season fans talk about their hopes for next

year and their memories of previous triumphs and disasters. The events are also complex and overlapping – a goal within a match within a season for example.

We have seen that fans of English Premier League teams often swear about or at their own team, and relatively rarely about an opposing team or match officials. It may be that swearing is being used as a means of demonstrating affiliation to a particular group, or to demonstrate greater passion about their own team and an apparent indifference to all others. In either case, intense expressions of negative sentiment actually provide strong evidence of a tweeter's affiliation *towards* the target of their criticism. Note, however, that an event which is seen as positive by one tweeter and negative by another can prompt tweets that contain swearing from both. This means the storytellers can be regarded as "unreliable narrators" – likely to disparage what outsiders may see as neutral or positive in that team's performance.

Humour is an important part of fans' storytelling both at games and online. Effective attempts at humour are picked up and passed around as either retweets or reports of stadium banter. However, humour is not in and of itself a sign of positive-valenced sentiment. English fans demonstrate a striking ability to joke even (perhaps especially) when things are going against their wishes. Humorous tweets often carry a heavy freight of context and can be challenging to mine for sentiment without thorough understanding of that context.

Even very complex information, including mixed sentiments, can be expressed in fewer than 140 characters. In response to significant events, such as goals, fans tend to respond very rapidly with shorter, more focused messages than usual; but at other times, tweets can be packed densely with information and contain rhetorical figures, deliberate ambiguities and novel wordplay. It is important therefore to be aware that tweets are not always simple utterances despite their brevity, and that they should be analyzed and assessed with care.

In light of these observations, we would recommend that any automated sentiment analysis system should explicitly consider the nature of the evidence of sentiment, including the wider context. If the only evidence is from syntax or a lexicon, great care should be taken. The simple assignment of fans to teams used here is sufficient, but this process could be improved by considering tweets from each account over a longer period of time as fans tend not to change loyalties. Match analysis may be simpler if only one match is being played. This is more often true for internationals and cup final matches [2].

Date/time	Event	Prior count	Prior length	Post count	Post length
<b>Chelsea vs Stoke City, 7/12/13</b>					
15:10:00	Chelsea goal	898	82.346	3752	52.215
15:42:00	Stoke goal	392	81.719	1402	73.932
16:08:00	Stoke goal	646	80.517	2962	54.520
16:11:00	Chelsea goal	1155	66.762	3342	53.621
16:47:00	Stoke goal	655	79.988	2033	64.483
<b>Southampton v Newcastle, 14/12/2013</b>					
15:27:00	Newcastle goal	70	94.900	601	64.057
16:23:00	Southampton goal	214	61.243	458	67.590
<b>Manchester United v West Ham United, 21/12/2013</b>					
15:25:00	Manchester goal	400	83.933	2198	55.720
15:36:00	Manchester goal	709	79.609	3640	63.106
16:29:00	Manchester goal	546	70.495	3136	65.466
16:39:00	West Ham goal	1387	76.958	1122	76.647
<i>Mean</i>		642.9	78.04	2240.5	62.85

Table 2: Number and lengths of tweets collected for five minutes before and after goals in three matches

## References

- [1] Y. Baruch and S. Jenkins. Swearing at work and permissive leadership culture: When anti-social becomes social and incivility is acceptable. *Leadership & Organization Development Journal*, 28(6):492–507, Apr. 2007.
- [2] D. Corney, C. Martin, and A. Goker. Spot the ball: Detecting sports events on Twitter. In *European Conference on Information Retrieval ECIR2014*, pages 449–454, Amsterdam, Holland, 2014.
- [3] N. Daly, J. Holmes, J. Newton, and M. Stubbe. Expletives as solidarity signals in FTAs on the factory floor. *Journal of Pragmatics*, 36(5):945–964, May 2004.
- [4] M. Hu and B. Liu. Mining and summarizing customer reviews. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, Seattle, Washington, USA, Aug. 2004.
- [5] B. Liu. Sentiment analysis and subjectivity. In N. Indurkha and F. J. Damerau, editors, *Handbook of Natural Language Processing*. Chapman & Hall, 2nd edition, 2010.
- [6] D. Maynard, K. Bontcheva, and D. Rout. Challenges in developing opinion mining tools for social media. In *Proceedings of @NLP can u tag #usergeneratedcontent?! Workshop at LREC 2012*, Turkey, 2012.
- [7] B. A. Plester and J. Sayers. “Taking the piss”: Functions of banter in the IT industry. *Humor: International Journal of Humor Research*, 20(2), Jan. 2007.
- [8] R. Stephens and C. Allsop. Effect of manipulated state aggression on pain tolerance. *Psychological Reports*, 111(1):311–321, Aug. 2012. PMID: 23045874.
- [9] M. Thelwall. Fk yea i swear: Cursing and gender in MySpace. *Corpora*, 3(1):83–107, 2008.
- [10] G. van Oorschot, M. van Erp, and C. Dijkshoorn. Automatic extraction of soccer game events from Twitter. In *Proc. of the Workshop on Detection, Representation, and Exploitation of Events in the Semantic Web*, 2012.
- [11] S. Zhao, L. Zhong, J. Wickramasuriya, and V. Vasudevan. Human as real-time sensors of social and physical events: A case study of Twitter and sports games. *arXiv preprint arXiv:1106.4300*, 2011.