

Visual Analysis of Explicit Opinion and News Bias in German Soccer Articles

D. Oelke, B. Geißelmann and D. A. Keim

University of Konstanz, Germany

Abstract

Most state-of-the-art opinion and sentiment analysis techniques were developed for customer feedback data or reviews. Applying them to another domain is often not possible because the algorithms are based on the assumption that the opinions are expressed explicitly in the text. However, news articles, for instance, convey an opinion in a more subtle manner. In this work we analyze German soccer articles with respect to the sentiment that is expressed in them. Besides adapting conventional sentiment analysis algorithms to the specific domain, we also investigate what can be measured with these techniques and what should be measured on news articles according to communication scientists. We suggest to bridge the existing gap with visual analytics methods and demonstrate the usability of the techniques on a concrete application example.

Categories and Subject Descriptors (according to ACM CCS): Document and Text Processing [I.7.5]: Document Capture—Document Analysis;

1. Challenges in News Opinion Analysis

Most sentiment and opinion analysis algorithms assume that subjective statements are explicitly expressed. However, in some domains such as news articles, opinions are conveyed in a more subtle manner. This is a challenge for automatic algorithms because detecting such subjective statements requires much background knowledge and a thorough understanding of the subtleties of natural language. To clarify the challenges, we juxtapose the view of communication scientists on how opinion is expressed in news articles and previous work in computer science in this section.

1.1. What should be measured?

While editorials or letters to the editor contain statements that are explicitly marked as subjective, this is rarely the case for other categories of news. Instead, they often convey opinions in a more subtle manner. Below we list some examples for rhetorical techniques that can be used to insert subjectivity into an otherwise factual report:

- **Framing:** By carefully selecting, highlighting, stressing, or excluding certain information, some aspects of reality are put on focus and the news is viewed from a certain

perspective. As a consequence this suggests certain interpretations or assessments. [Sch07]

- **Word choice:** Word Choice can make a big difference in terms of how an issue is perceived by the reader. Does the text talk about “freedom fighters” or “terrorists”? Was a country entered by “coalition forces” or “invasion forces”? [NB12]. And what associations are made? Are waterboarding practices referred to as “torture” or is this association avoided by the article? [DPRF10]
- **Upvaluation or degradation:** If an article reports on the actions of a certain person or cites someone’s statements, this person can be upvalued and made more trustworthy by simple linguistic means such as adding a title (Professor of ...) or calling someone an expert. Similarly, a person can be degraded, e.g., by mentioning the employer (which may suggest that the person is biased towards a certain position in an argument) or by indicating that former actions of this person were unsuccessful (which will likely degrade the persons trustworthiness even if the actions are not related to the currently discussed issue). [Frü07]
- **Stating opinions as facts:** If an issue is controversial, an opinion can be propagated by reporting on a disputed fact as if this would be generally accepted (e.g., “Nuclear power plants provide a safe and clean energy source.”).

1.2. What is already measured?

Much research has been conducted on the development of sentiment analysis techniques for customer feedback data (cf. [LZ12, PL08]). But comparatively few approaches deal with the analysis of newspaper articles. Examples include Balahur et al. [BSG*09] that aim at mining explicit opinions from quotations found in newspaper articles. Similarly, Godbole et al. [GSS07] use sentiment detection algorithms to classify an entity as subjective / objective and to determine its polarity. Park et al. [PLS11] raise the issue that a classification into positive / negative might not be appropriate for news articles. Instead, they propose to classify articles according to conflicting views on the issue, being thus close to the idea of news frames that was discussed in the previous section. The semi-automatic approach presented in [Mii97] is based on the idea that news articles carefully select what they report on to stress certain aspects of a controversial issue. They determine term clusters that are manually inspected and assigned to one of the conflicting parties and can then be used to classify the articles.

2. Analysis of Explicit Opinions in Soccer News

We chose soccer articles as a domain to apply the sentiment algorithms to because sports news are known to be more emotional and subjective than other section of the newspaper. Most newspapers are (frankly) biased towards one of the teams - usually the one that is geographically closer. Therefore, we can expect that soccer news do not only contain hidden but also explicit sentiments that allow us to investigate how well a common sentiment algorithm that was developed for customer feedback data works on such a domain.

2.1. Measuring Explicit Opinions in Soccer News

The algorithm that we apply is a lexicon-based method for feature-based sentiment analysis. Given a list with positive and negative opinion signal words, the algorithm identifies all mentions of these terms in the document. The polarity is reversed if a negation signal word is found nearby the term. Next, the features that were commented on have to be identified. Finally, the algorithm determines which features are modified by which opinion signal words. Usually, such a mapping is made by syntactic analysis of the sentence or - as in our case - with a distance-based approach that assigns each feature the polarity of the closest opinion signal word in the sentence.

Adapting the algorithm to the new domain comes with several issues. In our case, features are the names of players, coaches, or teams. Fortunately, the set of possible features is restricted and a list of names can be fed into the system. Because the usage of anaphora is quite common in news articles, we use a simple heuristic to resolve anaphora that cannot be detailed here because of space limitations.

Furthermore, a domain-dependent sentiment lexicon is

needed. With the help of a bootstrapping algorithm and a general purpose sentiment lexicon (SentiWS [RQH10]), we compiled a dictionary with domain-specific opinion signal words. The input of the algorithm is a lemmatized and part-of-speech tagged corpus[†] of soccer news consisting of about 6000 articles. A short seed word list of 20 soccer-specific sentiment terms was used to identify patterns of part-of-speech tags that afterwards can be used to search for additional candidate terms. Candidate terms that were found in SentiWS were directly added to the lexicon. For the rest of the terms, the polarity of a candidate term was determined by looking up synonymous terms in a dictionary. If one of the synonyms was found in the SentiWS lexicon, the same polarity was assigned to the domain-specific term. Finally, the list was manually inspected to ensure a high quality.

Evaluation

To evaluate the algorithm, several articles that contained in total 132 features were manually annotated. This ground-truth was compared to the result of the automatic sentiment detection. Overall, we got an accuracy of 0.48 for the detection of explicit feature-based sentiment (classification as positive, neutral, or negative). Partly this could be further improved by putting additional manual effort into the compilation of the sentiment dictionary since about 35% of the errors were due to a missing sentiment term in the dictionary. In about 24% of the erroneous cases the polarity of the sentiment term in the dictionary was wrong. Failures in the feature-sentiment term mapping process caused 27% of the errors. A special challenge in such sports news seems to be that often several features are mentioned in one sentence and the opinion is expressed by setting them in relation to each other (e.g., in sentences saying that team A was superior to team B a simple distance-based mapping algorithm is often not enough). The remaining 13% were due to the lemmatization algorithm - a problem that most likely would be alleviated if English news articles were processed instead of German ones.

2.2. Visual comparison wrt. explicit opinions

We visualize the result of the sentiment detection algorithm in a tabular visualization to permit further analysis and interpretation of the results. Each row represents a team member, each column a newspaper that reported on the game. The coloring of the cell shows if the newspaper commented on a team member positively (blue) or negatively (red). In case of several opposing statements, the proportions of the two colors hint at the distribution of positive / negative comments. Because of the low accuracy of the sentiment detection, ratings must be considered as somewhat uncertain if only a single comment exists for a team member. Therefore, color sat-

[†] For lemmatization and part-of-speech tagging the TreeTagger of the University of Stuttgart was used (<http://www.ims.uni-stuttgart.de/projekte/complex/TreeTagger/>).

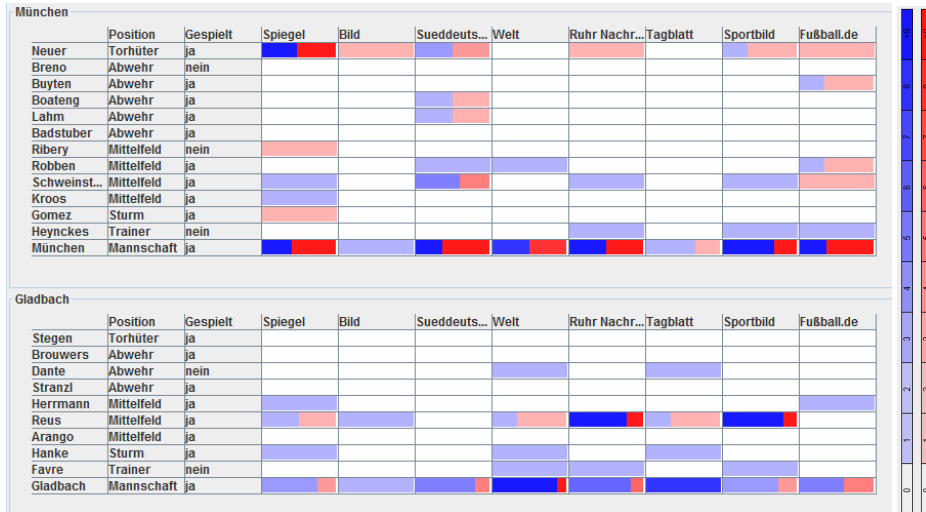


Figure 1: Visualization of the sentiment analysis results for a game between Bayern München and Mönchengladbach.

uration is mapped to the number of statements about a player that were found in an article, making cells whose value is based on more ratings visually more salient. Two columns next to the name of the team member provide further information about the person.

Figure 1 shows the results for a game between the FC Bayern München and Borussia Mönchengladbach which Gladbach won 3:1. Unsurprisingly, Gladbach’s team got overall much better ratings. What is more interesting is who the newspapers blamed for the loss of Bayern that is generally known to be the better team and was expected to win. Apparently, this was Neuer, Bayern’s keeper, whose performance most newspapers discussed and rated negatively.

3. Analysis of News Bias in Soccer Articles

In terms of measuring inexplicit sentiment we concentrate on the aspect of framing. Our basic assumption is that the different perspectives that the newspapers take can be detected by comparing several newspaper articles to each other that report on the same topic. We put a focus on measuring news bias with respect to the persons quoted and teams that were reported on.

3.1. News bias analysis wrt. the teams reported on

Our dashboard (for space reasons not shown) provides the user with bar charts that show statistical data about the frequencies that the players are mentioned with. In general, a bias towards certain players can be observed. Among the most frequently mentioned players are usually the most famous members of the team but also the players that were involved in actions that determined the result of the game.

Besides this general trend the detailed visualization shows the bias between different newspapers.

3.2. News bias analysis wrt. interview partners

We are also interested in the question which newspapers resemble each other in terms of the interview partners that are quoted in direct speech in the articles. The underlying basic assumption is that all newspapers had the chance to talk to all team members. This may not always be true, but what we can expect is that newspapers choose their interview partners carefully to substantiate their views and beliefs.

To be able to analyze news bias with respect to interview partners, we first extracted all quotes together with the names of the quoted persons by means of regular expressions. We experimented with two different visualization techniques to make the data accessible: A thumbnail view which takes newspapers as entities (Figure 2) and a bubble view that puts a focus on the interviewed team members (Figure 3).

For building the thumbnail view, each newspaper is represented with a feature vector that contains one element per team member that was interviewed. An element of the vector is set to 1 if the respective player (or coach / manager) was quoted in the article. A euclidean distance function is used to calculate a distance matrix that is the input for an MDS algorithm [Alg]. In the 2D space, each newspaper is represented with a thumbnail that provides the user with further insight and reveals why certain newspapers were clustered together or set apart. The thumbnails could be considered a pixel-based, tabular representation of the feature vector with one column per team and one line per team member (see Figure 2). On demand, the user can access the extracted quotes of a certain article which permits a semantic analysis.




Figure 2: Thumbnail view for quotation analysis (FC Bayern München against FC Schalke 04).

The bubble view puts the team members in focus (see Figure 3) and consequently, feature vectors represent interview partners instead of newspapers which now constitute the elements of the vector. Again, the high-dimensional vectors are projected into 2D, but this time the names are visually grouped with the help of the Bubble Sets technique [CPC09]. A bubble encloses all team members that were quoted by a certain newspaper. To make the overlap between bubbles visible, semi-transparent colors are used. Unfortunately, the resulting blending paint can lead to artifacts that hamper the perception of the visual patterns. To ameliorate this problem, we implemented interaction facilities that allow the user to hide or display certain bubbles.

Figures 2 and 3 show the reports on a game between the FC Bayern München and FC Schalke 04. In the thumbnail view two clusters and an outlier (the *Sueddeutsche*) are visible. While the cluster at the right mainly consists of articles that interviewed players of Schalke, the cluster in the middle groups articles that put a focus on the team members of Bayern. The thumbnails provide insight with respect to the question why two newspapers were grouped together. However, it takes some effort to compare the thumbnails to find out which interview partners are shared. In the bubble view this is easier to see. Here it is readily apparent that the *Sueddeutsche* (our outlier) only shares one quote with another newspaper, the *Münchner Abendzeitung*. The two clusters split the display into an upper half and lower half. While the clusters were easier to see in the thumbnail view, the bubble view clearly shows that they are not entirely separated but that the coach of Bayern (*Heynckes*) and the player *Ribery* are quoted by newspapers from both clusters. Interestingly, a bias towards one team cannot only be constituted for local newspapers but also for the nationwide ones - an observation that might be worth further investigation.

3.3. The SoccerNewsAnalyzer

Our SoccerNewsAnalyzer combines all visualizations into one tool. Interaction facilities permit further investigation of the data. For instance, for each player statistical data about the game can be viewed such as how many goals the player

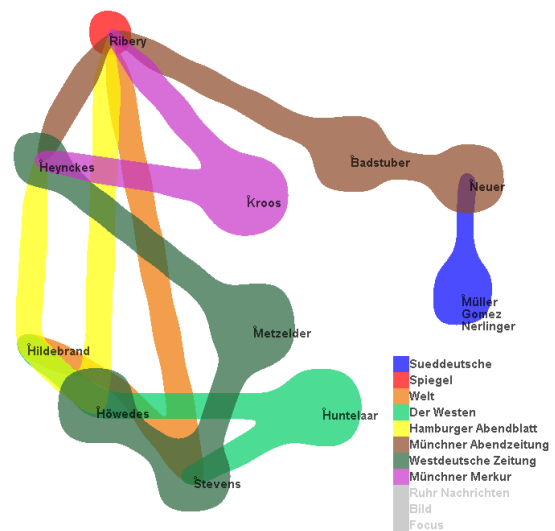


Figure 3: Bubbles view for quotation analysis.

hit, number of fouls, number of tackles with other players etc. This kind of information often helps to explain findings that were detected with one of the visualizations. Furthermore, maps of Germany with the main locations of the newspapers as well as the competing teams are displayed on demand which helps to investigate whether a newspaper's bias might be due to its geographical proximity to the team.

4. Conclusions

Automatic analysis of news in terms of the opinions expressed in them is challenging. In this paper we first outlined what should be measured according to communication scientists. Afterwards we presented a visual analytics tool that combines the results of an analysis of explicit sentiment with news bias analysis. In the future, the applicability of the approach to other news domains such as politics or economics needs to be tested. Furthermore, techniques that measure the other rhetorical techniques need to be developed.

References

- [Alg] Algorithmics Group. MDSJ: Java Library for Multidimensional Scaling (Version 0.2). Available at <http://www.inf.uni-konstanz.de/algo/software/mdsj/>. University of Konstanz, 2009. 3
- [BSG*09] BALAHUR A., STEINBERGER R., GOOT E. V. D., POULIQUEN B., KABADJOV M. A.: Opinion Mining on Newspaper Quotations. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 03* (Washington, DC, USA, 2009), WI-IAT '09, IEEE Computer Society, pp. 523–526. 2
- [CPC09] COLLINS C., PENN G., CARPENDALE S.: Bubble sets: revealing set relations with isocontours over existing visualizations. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1009–1016. 4
- [DPRF10] DESAI N., PINEDA A., RUNQUIST M., FUSUNYAN M.: *Torture at Times: Waterboarding in the Media*. Harvard student paper, Joan Shorenstein Center on the Press, Politics, and Public Policy, 2010. 1
- [Frü07] FRÜH W.: *Inhaltsanalyse - Theorie und Praxis*. UVK Verlagsgesellschaft mbH, Konstanz, 2007. 1
- [GSS07] GODBOLE N., SRINIVASIAH M., SKIENA S.: Large-Scale Sentiment Analysis for News and Blogs. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM)* (2007). 2
- [LZ12] LIU B., ZHANG L.: A Survey of Opinion Mining and Sentiment Analysis. In *Mining Text Data*, Aggarwal C. C., Zhai C., (Eds.). Springer US, 2012, pp. 415–463. 2
- [Mil97] MILLER M. M.: Frame Mapping and Analysis of News Coverage of Contentious Issues. *Social Science Computer Review* 15, 4 (1997), 367–378. 2
- [NB12] News Bias Explored - The art of reading the news, last accessed February 2012. <http://www.umich.edu/newsbias/>. 1
- [PL08] PANG B., LEE L.: Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval* 2, 1-2 (2008), 1–135. 2
- [PLS11] PARK S., LEE K., SONG J.: Contrasting opposing views of news articles on contentious issues. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1* (Stroudsburg, PA, USA, 2011), HLT '11, Association for Computational Linguistics, pp. 340–349. 2
- [RQH10] REMUS R., QUASTHOFF U., HEYER G.: SentiWS – a Publicly Available German-language Resource for Sentiment Analysis. In *Proceedings of the 7th International Language Resources and Evaluation (LREC'10)* (2010). 2
- [Sch07] SCHENK M.: *Medienwirkungsforschung*. Mohr Siebeck, 2007. 1