

Turning Words Into Consumer Preferences: How Sentiment Analysis Is Framed in Research and the News Media

Social Media + Society
July-September 2018: 1–12
© The Author(s) 2018
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/2056305118797724
journals.sagepub.com/home/sms


Cornelius Puschmann¹ and Alison Powell²

Abstract

Sentiment analysis is an increasingly popular instrument for the analysis of social media discourse. Sentiment scores seemingly represent an objective means of assessing the mood of social media users, consumers, and the public at large. Similar to other computational tools, sentiment analysis promises to reduce complexity and mitigate information overload, and to inform the decisions of marketers, pollsters, and scholars with reliable data. This article argues that the assumptions encoded into sentiment analysis as a method are accompanied by a number of constraints, both regarding its technical limitations (in terms of what sentiment analysis can and cannot accomplish) and conceptually (in terms of what the notion of sentiment implicitly represents), constraints which are often de-emphasized in public discourse. After providing an overview of its history and development in computer science as well as psychology and the social sciences, we turn to the role of sentiment as a currency in the attention economy. We then present a brief study of common framing of sentiment analysis in the news media, highlighting the expectations that exist regarding its analytical capabilities. We close by discussing the kind of conceptual work that takes place around computational methods such as sentiment analysis in specific cultural environments, highlighting their influence on the public imaginary.

Keywords

sentiment analysis, computational methods, behavioral research, framing

Introduction

Sentiment analysis, also commonly referred to as opinion mining, is increasingly used in a wide range of research areas, including the social sciences and the media and communication field (Ceron, Curini, Iacus, & Porro, 2014; Driscoll, 2015; Murthy, 2015; Papacharissi & de Fatima Oliveira, 2012; Schwartz & Ungar, 2015; von Nordheim, Boczek, Koppers, & Erdmann, 2018; Wojcieszak & Azrout, 2016). With origins in computational linguistics and computer science, its proliferation in market research, public relations, and political forecasting, and increasingly also as a method in the social sciences, parallels the globally growing relevance of social media, from a minority pursuit to a ubiquitous activity. Its purpose is to represent emotional or affective tendencies, often within user-generated content (UGC), though frequently no clear distinction is made between *affect expressed in a text* and *the emotional state of a text's author*. The growing prominence of sentiment analysis can be seen as a response to both datafication and “post-factual” politics, emphasizing the role of emotion in areas of inquiry previously understood chiefly in rational and analytical

terms. However, there is a risk that the expansion of the method away from its original contexts has produced misunderstandings and misinterpretations about how the method works, which are detrimental both to its application and to our broader social understanding of social media. This article argues that the expansion of sentiment analysis from the advertising industry across the media and communications field has ushered in a new interest in the quantification of feeling, as well as a new orientation toward political turbulence and affect that also coalesces into methodological debates within the social sciences (Halford & Savage, 2017; Margetts et al., 2016; Papacharissi, 2016).

We begin by pointing to three related problems. First, sentiment analysis as currently practiced has become

¹Hans Bredow Institute for Media Research, Germany

²London School of Economics and Political Science, UK

Corresponding Author:

Cornelius Puschmann, Hans Bredow Institute for Media Research, Rothenbaumchaussee 36, 20148 Hamburg, Germany.
Email: c.puschmann@hans-bredow-institut.de



disconnected from the methodological assumptions and material features of its original design in computational linguistics and the original applications to online rating systems, becoming embedded in a range of algorithmic decision-making systems. Second, this disconnection has produced an application without a concept—measurement of something called “sentiment” frequently fails to establish what sentiment might actually mean (beginning with the distinction made above, between the expressed polarity of a text, and the presumed emotional state of a writer). Third, the way in which different domains (marketing, opinion research, social science) describe, discuss, and critique sentiment analysis reveals something about the particular pitfalls of the concept and allows observations about the broader cultural acceptance of and epistemic faith in computational tools such as sentiment analysis across different national and cultural environments. As a result of these problems, the public perception of sentiment analysis is misaligned with its function and may create epistemological expectations that the method cannot fulfill due to its technical properties and narrow (and well-defined) original application to product reviews. Reinterpreting sentiment as currency in the attention economy allows us to assess the power of sentiment analysis as an economic instrument.

We proceed as follows. After providing an overview of the history and development of sentiment analysis, with a focus on computer science/computational linguistics and psychology/social science, we describe the kind of conceptual work that takes place around sentiment analysis in specific environments, highlighting its influence in the public imaginary. We then present a study of sentiment analysis’ representation in the media in order to investigate whether public perceptions match with its capabilities. We close with critical observations on the potentials and dangers of computational methods for social inquiry.

Sentiment Analysis in Computational Linguistics and Computer Science

In everyday language, the term *sentiment* has at least two distinct meanings. It is commonly used to describe (a) a feeling, or something of emotional significance, and (b) a particular (usually subjective) point of view. Sentiment analysis, also referred to as opinion mining, describes a collection of approaches that address the problem of measuring opinion, sentiment, and subjectivity in texts (for overviews, see Liu, 2010; Pang & Lee, 2008; Taboada, Brooke, Tofiloski, Voll, & Stede, 2011). While the study of subjectivity, emotion, and varying viewpoints is clearly within the purview of several different fields, including literature, history, rhetoric, political science, and the arts, sentiment analysis is generally considered a subfield within computational linguistics and occasionally computer science. Furthermore, sentiment analysis is considered a branch within computational linguistics, whereas it is generally treated as a method in social science.

Sentiment analysis can be considered a subfield of information extraction, the research area within information and computer science that aims to condense, summarize, and draw inferences from collections of textual documents (Cowie & Lehnert, 1996; Sarawagi, 2008). Sentiment analysis in a computational sense rose to prominence within computational linguistics only in the early 2000s (Pang & Lee, 2008). While research on subjectivity within the broader paradigm of artificial intelligence reaches back to the 1970s and 1980s (Carbonell, 1981), it both approached the object of study in a different fashion from contemporary sentiment analysis and played a role of limited importance. To understand why this was the case, it is important to note the focus, both within computational linguistics and computer science, on rationality and objectivity, reflected by Liu’s (2010) statement in a handbook article that “textual information in the world can be broadly categorized into two main types: facts and opinions” (p. 627). For the first several decades from its inception in the 1950s, computational linguistics sought to solve a set of problems related to human language understanding, chief among them machine translation. The interpretation of human language was initially regarded as a problem of logic: a system equipped to parse the syntax of a natural language sentence and map the semantics of a natural language proposition would be able to correctly interpret sentence meaning and in turn produce meaningful sentences. In practice, a lack of both data and computing power severely limited early attempts to automate translation. When significant advances were made, this was through an approach that departed considerably from deductive logic and instead focused on exploiting statistical patterns within natural language data, that is, machine learning.

While machine learning provided the foundation to modern sentiment analysis, two other requirements had to be fulfilled. First, sufficient volumes of natural language data had to be accessible, a requirement that was met by the rise of the commercial Internet. Second, the textual data in question had to express opinions, emotions, and sentiment, rather than objective facts. For much of the 20th century, the canonical genres of textual expression studied revolved around a factual-objective style, from newspapers and scientific publishing to popular culture and fiction. Certain genres, such as fiction, expressed emotion, but not necessarily that of the author, and frequently artistic considerations influenced “natural” self-expression. Subjective writing, for example, in personal diaries, while widely practiced, often remained unpublished and accordingly unavailable to researchers. Pang and Lee (2008) acknowledge the dual academic and commercial appeal of sentiment analysis when noting “the fascinating intellectual challenges and commercial and intelligence applications that the area offers” (p. 5).

As a consequence of the popularization of social media, millions of people began to express their views online, not just about politics or current affairs but also about products and services. To add to the commercial importance of this

form of self-expression, the business models of fledgling Internet companies were (and still are) advertisement-driven, making it particularly important to provide a picture of consumers to advertisers that promised to be more granular than that offered by traditional opinion polls and focus groups.

The problem of synthesizing emotions and opinions quickly became a prominent research interest in computational linguistics. A common approach is to begin with a set of definitions of objects, features, opinion passages, and opinion holders (Liu, 2010, pp. 630-632) that together map the structure of, for example, a comment on a product. Once such a mapping has been developed and tested, the next step is to establish a scale through which the degree of sentiment expression can be measured for the entire document (document-level sentiment classification), though applying sentiment analysis on the level of sentences or phrases is also common (Feldman, 2013; Wilson, Wiebe, & Hoffmann, 2005). The aim in terms of the desired outcome is to develop a scale that measures sentiment within a particular posting (usually a product review on an e-commerce site such as Amazon.com). The simplest scale for this is a binary (positive-negative) or modal (positive-negative-neutral) categorization, or alternatively a score of between 1 (very positive) and -1 (very negative). While representing a separate problem in many respects, and being signaled differently, agreement and disagreement (e.g., with a particular political proposition) is often operationalized along the same lines (agree/disagree/no opinion) and treated an extension of the same principal problem.

A range of tactics, from precompiled sentiment lexicons to machine learning, are employed to determine an unambiguous polarity score for each document. The simplest approach is to employ a lexicon, that is, a precompiled word list of terms that indicate positive or negative expressions of sentiment. To score a real-life text, the occurrence of a term labeled as positive in the dictionary would increase the tally for classifying the posting as positive, while the reverse would hold true for terms labeled as negative. In some cases, a simple majority decides the final labeling of the post, while in others, more elaborate strategies—for example, through recognizing negation or the resolution of complex predicate structures—would be employed (see Liu, 2010, for an overview). The approaches for constructing the sentiment also differ, from hands-on approaches (human experts, Mechanical Turk) to more elaborate ones (e.g., examining units beyond individual words, such as part of speech, n-grams, clauses, or sentences). A problem with manually compiled dictionaries that is widely recognized within computational linguistics, but not always taken into account within the social sciences, is the highly context-dependent nature of human language. Issues also arise when sentiment dictionaries developed for one genre (e.g., product reviews) are suddenly applied to another (e.g., politics discourse) and context-dependent word meanings no longer fit with the original context (see González-Bailón & Paltoglou, 2015, for a methodological discussion of dictionary approaches

and their shortcomings). Put another way, this is a problem of perception linked to the tools used for generating particular visions of the social world. As art historian Jonathan Crary writes, referring to the new ways of seeing generated by photographic capture in the 19th century, “perception transformed alongside new technological forms of spectacle, display, projection” (p. 2). Similarly, the tools of sentiment analysis generate frameworks of perception for social media that change the way political discourse is described, even when the tools are not necessarily designed with this application in mind, and the theoretical justification appears far sounder for their original purpose than for the new one.

A solution widely employed to overcome precision issues is to inductively develop sentiment dictionaries that are tailor-made to particular genres, rather than universally applied. Since users in a sense annotate their review when awarding a five-star rating to an “outstanding restaurant,” it is fairly simple to infer that a review without any rating that also uses the adjective “outstanding” is positive, or that a review that awards a low rating but mentions many positive terms may be mislabeled by the user by mistake (by contrast, an “outstanding payment” should be awarded a different score). The importance of detecting such cases is obvious from the vantage point of commercial application, both encompassing the ability to improve review quality and enabling entirely new services, such as social media monitoring.

The selection of linguistic features also requires choices regarding the cues used for classification. For example, term frequency has been found to be a poorer predictor than term presence (Pang & Lee, 2008). The usage of highly emotionally charged terms is more significant than their exact frequency, as is the usage of terms more suggestive of objective statements. Word class, multi-word phrases, syntax, and negation have all been used as features, as have been the use of exclamation marks, all caps, and character repetition (Brody & Diakopoulos, 2011). Misclassification can occur for a number of reasons, one being domain specificity. As Pang and Lee (2008) point out, “simply applying the classifier learned on data from one domain barely outperforms the baseline for another domain” (p. 25)—the baseline being random guessing. Appropriate usage of sentiment analysis therefore generally presumes detailed knowledge of the domain of application.

In addition to the described approaches to the measurement of sentiment, the detection of fraudulent reviews and evaluations—so-called opinion spam—has also become a growing area of research in computational linguistics and computer science in recent years (Jindal & Liu, 2008; Liu, 2010; Ott, Choi, Cardie, & Hancock, 2011). Liu (2010) characterizes opinion spam as

fake or bogus opinions that try to deliberately mislead readers or automated systems by giving undeserving positive opinions to some target objects in order to promote the objects and/or by giving malicious negative opinions to some other objects in order to damage their reputations. (p. 629)

Opinions are taken to be spam if they are deceptive, are related to a brand only rather than a product, or are entirely unrelated to the product under review. Some commonalities to email spam exist: a high degree of similarity in the wording of multiple reviews can identify opinion spam, alongside indicators taken from the user profile. Relying on Linguistic Inquiry and Word Count (LIWC; see the next section), Ott et al. (2011) find commonalities between opinion spam and imaginative writing based on part-of-speech distributional similarities, returning in a sense to the original association made by Liu (2010) and others regarding sentiment and subjectivity.

Sentiment Analysis in Psychology and Social Science

In parallel to the development in computer science and computational linguistics, relying on language as an indicator of psychological well-being and a conduit for emotions more generally also has a long tradition in clinical psychology. In 1969, Gottschalk and Gleser introduced an influential method for applying content analysis to affective language, commonly referred to as the Gottschalk–Gleser approach (Gottschalk, Winget, & Gleser, 1979). The method derived its measures from the grammatical clause and the agent and recipient of the meaning of the clause's verb, with scores obtained by calculating a content-derived score per 100 words. The scales identified by Gottschalk and Gleser included the Anxiety, Hostility, Social Alienation, Cognitive Impairment, Depression, and Hope scales (Gottschalk, 1997). Importantly, the data used in this approach were collected by elicitation (asking the subjects questions, often superficially unrelated to their well-being). Inspired by this scientifically objective approach to measuring subjective states, Tausczik and Pennebaker (2009) emphatically argue that “[l]anguage is the most common and reliable way for people to translate their internal thoughts and emotions into a form that others can understand” (p. 25) and claim that “for the first time, researchers can link daily word use to a broad array of real-world behaviors” (p. 24). Noting that “[t]he roots of modern text analysis go back to the earliest days of psychology” (p. 25), the authors establish a parallel history of the study of emotion in text that differs significantly from the one told by Pang and Lee about the development of computational linguistics, their narrative focusing not on the polarity of a text but on the emotions of its author.

LIWC (Pennebaker, Booth, & Francis, 2007), a software developed by James Pennebaker and colleagues at the University of Texas at Austin and translated into a number of languages, represents the pinnacle of this development within psychology. Pennebaker's research was initially centered on psychological well-being and the emotional response to traumatic experiences. Studying individuals who had undergone such events, he discovered that the well-being of subjects improved when they wrote down accounts of their

experiences. This was generally interpreted to be the result of cognitively working through the event, thereby making the experience coherent through narrative. Pennebaker became interested in the content of the patients' trauma diaries and in the correlation between their writing and their self-reported and observed psychological well-being. This interest developed into LIWC, a program that assigns a particular word a range of emotive dimensions and scores users according to their use. LIWC encodes a number of assumptions about language, some of which can be challenged on the grounds of their conflation of words with discrete units of meaning. For example, some of the words in LIWC's category index are polysemous, taking on very different distinct meanings depending on the text surrounding them, resulting in potential false positives. Further issues arise when examining the categories and their psychological validity. It is one thing to assume that neuroticism is a human psychological trait, but another to argue that use of the first-person pronoun signals it, because this assumption is difficult to test empirically. Genre differences, for example, between a diary, a newspaper article, a student textbook, and a party program, may shape the distribution of the features that LIWC uses as the basis of its scoring, along with regional and social differences (Danescu-Niculescu-Mizil, Gamon, & Dumais, 2011). This furthermore extends to non-Western languages whose differences in morphology, syntax, writing system, and other areas may lead to errors. LIWC presupposes English or a language quite similar to it, and a genre of text in which the author freely expresses their feelings rather than acting under particular genre constraints or communicating strategically. LIWC is hugely popular in psychology and beyond, with countless social media studies applying it (De Choudhury et al., 2013; Gilbert & Karahalios, 2009; Kramer, Guillory, & Hancock, 2014; Pfeil, Arjan, & Zaphiris, 2009). This is in spite of considerable reservations regarding its validity that have been articulated by scholars who have sought to improve the approach (Danescu-Niculescu-Mizil, Gamon, & Dumais, 2011; González-Bailón & Paltoglou, 2015; Schwartz et al., 2013; Su et al., 2017).

Another interesting trajectory of computational methods for the analysis of text in the social sciences is that of software developed to augment and in some cases even supplant traditional content analysis (for an overview, see Grimmer & Stewart, 2013). The description provided above may lead to the conclusion that sentiment analysis came to prominence exclusively within computational linguistics and psychology, but a rich parallel history within the social sciences, particularly within political science exists (see Ceron, 2015; Ceron & Negri, 2016; Young & Soroka, 2012, for contemporary examples). Early systems for computer-aided text analysis (CATA) and automated content analysis (ACA) reach back to the 1950s, particularly the General Inquirer (Stone, Dunphy, & Smith, 1966) and DICTION (Hart, 2000). As Stone's (1997) overview attests, the links between linguistics and social science were stronger in the pioneering age of

computational text analysis, when fewer out of the box solutions existed and interdisciplinary collaboration was the norm (see also Carley, 1990). By contrast to early software development for academic research, the establishing of sentiment as a commercially relevant measurement in the Web 2.0 era was strongly driven by the interests of the early Internet industry and its business model and led to a renewed interest in the field. Proposals to code texts manually according to their expressed support for or opposition to a particular theme are noted by Roberts (1997, p. 277), who points to the work of Ithiel de Sola Pool in the 1950s and Ole Holsti in the 1960s. Yet, sentiment generally took the backseat in text-based research in political science, sociology, and media and communication, mostly because the genres under study were public discourse (newswriting, political speeches and documents) where semantics usually had precedence over expressed emotions. Furthermore, as Roberts (1997) outlines, diverging philosophies of content analysis made it difficult to reconcile strong claims about manifest versus inferential approaches to the subject. This changed only when public discourse in genres of spontaneous, informal, and widely accessible channels such as social media platforms became widespread, redefining what types of phenomena could be studied with content analysis in the social sciences.

Not in all cases are applications of sentiment analysis based on a long-standing research trajectory, as is the case in the political science tradition of studying institutional textual data such as parliamentary transcripts. Examples such as Driscoll (2015), Lindgren (2012), and Wojcieszak and Azrout (2016), which rely on *SentiStrength* (Thelwall, 2017), an out-of-the-box software tool for measuring sentiment, illustrate the trajectory of the method into less specialized research communities.

To summarize, the academic history of sentiment analysis (and of similar approaches by other names) can be characterized as one of continued innovation, both in computer science and computational linguistics, and in psychology and the social sciences. Researchers have sought to “extract” emotions, subjectivity, and sentiment from textual data in the same fashion that semantic information is derived from natural language. Their objectives have ranged from correctly assigning the sentiment of product reviews in computer science and computational linguistics to inferring the psychological health of subjects from their writing in psychology. While applications in political science and sociology have not been per se concerned with sentiment, assigning sentiment has at times also been a goal of traditional content analysis.

Sentiment as Currency in the Attention Economy

With these complex histories in mind, we turn to the public perception of sentiment analysis, particular in relation to its commercial applications. Sentiment analysis’ unique

contribution is that it makes affect quantifiable, thus paving the way for treating it as a form of capital in the attention economy (Marwick, 2015; Tufekci, 2013; Turow, 2012). Only by first establishing sentiment as a concept and by providing instruments for its precise measurement does a sentiment economy become viable, in which sentiment is an essential tool for marketers (Rambocas & Gama, 2013). The ability to measure how people feel toward a product or service and claim that this accurately reflects their views is paramount to the relevance of sentiment analysis in marketing, as is its ability to address the challenge of information overload posed to companies by UGC.

The dominant framing of sentiment analysis in marketing illustrates why this is an important prerequisite to the technique’s success: *there is an imperative need to standardize the measurement of human emotion in social media in order to efficiently monetize it*. Users’ preferential signaling is a vital component of the “Like economy” that underpins the business models of all major social media companies (Gerlitz & Helmond, 2013). This view is reflected in a number of recent critical scholarly accounts of sentiment analysis and similar computational methods in media and communication research. Kennedy (2012) argues that sentiment analysis is increasingly a key site of cultural production. She suggests that this is a result of the rising trust in peer recommendations and the expansion of marketing messages, as well as the pervasive quality of social media messages, arguing that “data gathered through sentiment analysis are believed to provide detailed information about something to which direct access did not previously exist: public opinion and feeling” (p. 435). Similarly, Hearn (2010) argues that measurement of feeling is one means of effectively enrolling in online economies of attention and affect: “online reputation measurement and management systems are new sites of cultural production; they are places where the expression of feeling is ostensibly constituted as ‘reputation’ and then mined for value.” She describes the rise of “feeling-intermediaries”—social listeners and other information measurers—and tracks the expansion of this group of intermediaries. Hearn (2010) then identifies the expansion of a political economics of attention and affect, where the very stuff of experience is transformed into value. Katherine Hayles examines these differences in terms of the different ways that they construct narrative. She identifies machine-enabled “hyper reading” that combines multiple forms and strands of text as if they were one, as a form of reading attuned to “an information intensive environment” (Hayles, 2012, p. 12). This kind of reading makes a far different kind of meaning than narrative construction might. As Hayles (2012) writes, “narratives gesture toward the inexplicable, the unspeakable, the ineffable, whereas databases rely on enumeration, requiring explicit articulation of attributes and data values” (p. 179). For our discussion of feeling and sentiment, the shift in this narrative construction also accompanies a fundamental shift in how ideas about worth, goodness, or value are constructed. From being part of other

kinds of social experiences, which in a narrative form might be indeterminate or difficult to resolve, under the alternative narrative of computational sentiment analysis, these values can be enumerated on a scale in relation to a specific dictionary definition that is largely oblivious to context. This is not so much a bug as it is a feature: the success of sentiment analysis as a metric depends on its claim to objectify what is otherwise hidden and subjective.

Broadening the view on this potential of analytics, Andrejevic, Hearn, and Kennedy (2015) see them as instrumental for the development of new strategies for forecasting, targeting, and decision making. They also can usher in opaque forms of discrimination based on “incomprehensibly large, and continually growing, networks of interconnections” (p. 379). Therefore, sentiment analysis as a technique adds to processes through which categorization and quantification emerge as key features of “logistical media” (Peters, 2015) and “social media logic” (Van Dijck & Poell, 2013). Andrejevic et al. (2015) note that in general this form of media arranges information “not to understand or interpret their intended or received content but to arrange and sort people and their interactions according to the priorities set by [digital capitalism’s] business model” (p. 381). They identify an analytic turn in media and cultural studies that decenters the human and meaning-making processes within communication, and introduces various “flat ontologies” for analysis (new materialism, object-oriented ontology, and others) that primarily engage with the circulation of affects and effects, or the data-processing capacities of media. This unfolds as a parallel process of “meaning-resistant” theorization to the expansion of computational possibilities for mapping the social world. Operationally, it also muddles the distinction between “structured” and “unstructured” data, speaking to the constructive power of calculative methods. As Amoore and Piotukh (2015, p. 347–348) point out,

because text analytics and sentiment analysis conduct their reading by a process of reduction to bases and stems, their work exposes something of the fiction of a clear distinction between structured and unstructured data. Through processes of parsing and stemming, everything can be recognized and read as though it were structured text.

The ability to standardize what is otherwise a much more idiosyncratic process represents the promise of sentiment analysis to both applied and academic social media research. Sentiment analysis, which began as a well-defined method within a particular application area, has significant consequences for the way that information is understood and acted upon—socially, economically, and politically, and its presentation in media reflects these consequences. As Karppi and Crawford (2016) note in their analysis of a 2013 Wall Street flash crash, “predictions based on social media are still a form of conjecture, often based on shaky assumptions regarding sentiment, meaning, and representativeness” (p. 77).

When such predictions are fully automated, the result—in this case a steep drop in the Dow Jones Industrial Average—can be grave.

Sentiment Analysis’ Portrayal in the Media

Having outlined the parallel academic histories of sentiment analysis and its relevance to marketing, we now take into consideration how sentiment analysis is framed in the media. Similar to other information technologies that have caught the public’s imagination in recent years, sentiment analysis has emerged from a niche subject at academic conferences and at marketing conventions to a topic of interest for mass media reporting, at least in conjunction with other technology trends, such as big data and artificial intelligence (Puschmann & Burgess, 2014). Our analysis is based on the assumption that the framing of sentiment analysis as a novel technology matters to the stakeholders engaged in its success as a metric, that is, social media companies, marketers, and consumer companies in particular but also society more broadly. We use this analysis to capture a broader discussion of the consequences of sentiment as currency in the attention economy and to examine the consequences of its move from a specific design and application into a (perceived) generalized paradigm for understanding the mediated social world. We examine how sentiment analysis is presented in the media—acknowledging that our own methods are also part of the same quantifying process that we describe above.

Data and Methods

In order to describe how sentiment analysis is depicted in the media, we collected a data set through Nexis Mass Media Germany/United Kingdom/United States consisting of 198 articles published between 2014 and 2017 in the German, British, and American press (Germany: 40, UK: 82, US: 76) that contained the term *sentiment analysis* or *Sentimentanalyse*. Our choice was underpinned by the assumption that sentiment analysis plays a role in the business discourse in these countries, which are economically similar in many respects, but differ in the relative speed of their social media uptake, with Germany lagging the United States and the United Kingdom. The selection includes national news sources, such as *USA Today*, *The New York Times*, *The Washington Post*, *The Guardian*, and *The Daily Telegraph*, as well as trade publications such as *Advertising Age*, *Bank Technology News*, and *Institutional Investor*.

After conducting an initial exploratory analysis, sentences mentioning sentiment analysis within the collection of articles were assigned one of eight thematic categories, resulting in a total of 1,062 codings. The categories were derived both inductively from surveying the material and

from comparable analysis of the framing of computational approaches in the media (Puschmann & Burgess, 2014; Van Dijck & Poell, 2013). Below, we focus on those three categories that reflect the issues previously raised in relation to the adoption of sentiment analysis in new societal areas, such as finance and health care, where its reliability is of particular importance, as well as the types of explanations provided for its capabilities. This allows us to provide an initial description of how the assumptions about and applications of sentiment analysis have begun to diverge from their original contexts. Table 1 summarizes the results and provides additional examples.

Explanations/Definitions

Explanations and definitions were used in articles mentioning sentiment analysis to clarify the concept to readers. Some of the descriptions are quite narrow and formal, for example, stating that in sentiment analysis words or phrases are assigned a polarity rating or a point score. Others are more functional, arguing that sentiment analysis enables new forms of statistical inquiry or makes it possible to process large volumes of data. The ability to process “not just gigabytes but petabytes” of data that “spans new channels,” and to do so “in real time,” is frequently characterized as a key aspect of sentiment analysis, although this hardly plays a role in methodological descriptions of the approach. Another set of descriptions emphasizes the role of sentiment analysis for marketing and branding, explicitly addressing entrepreneurs and relating the approach to “your products and brands.” In such accounts, “written words” are turned into “quantifiable consumer preferences.” These types of definitions are generally vague when it comes to how precisely the aim of generating useful knowledge for market research is realized.

A further group of definitions centers on individual psychological factors when describing sentiment analysis, claiming that it can “reliably capture mood” in a social media post, or “how whoever wrote it felt about the subject.” One article notes that “a chart in real time shows disgust, anger, fear, happiness, grief, surprise in percentages and curves,” pointing both to the dimension of real-time analysis and to data visualization as a new source of knowledge. Finally, some explanations emphasize the potential of sentiment analysis not on an individual but collective and societal level. Results are presumed to “depict online opinion,” “detect opinions and moods,” and “track how people feel about each candidate.” This is escalated further in depictions that (at least notionally) impart agency to the sentiment analysis software, for example, claiming that “computers around the world are studying comments and rating them as positive, negative or neutral” or that “a computer program, instead of asking people questions, surveys the Web on a large scale.” In some cases, the claim that sentiment analysis is concerned with the goal of natural language understanding more broadly

is made, although that area is clearly distinct in computational linguistics and concerns much harder problems than the relatively simple scoring mechanisms previously described. To summarize, these assessments vary greatly from basic descriptions of what sentiment analysis is on a technical level to macroscopic claims about their ability to capture social moods and even accurately reflect public opinion.

Domains of Use

Sentiment analysis is used across a steadily growing range of societal domains, creating more opportunities for its methodological features to be interpreted (or misinterpreted) in new ways. In business, sentiment analysis is used in market research, public relations, customer service and support, and human resources in a range of industries, such as fashion, retailing, health, gastronomy, travel, finance, and media, as well as brands and companies more generically. The largest group of corporate customers unsurprisingly appears to be marketers, followed by financial firms. Non-technology companies include Schufa (German consumer credit rating agency), EDF (French energy provider), and Kia (Korean automotive company). A small number of articles mentions companies, including Deutsche Bank and Vodafone, that have chosen not to employ sentiment analysis, often on regulatory grounds. Actors listed as developers of sentiment analysis software fall into two broad categories: corporations, especially global Internet and IT companies, and research organizations, both public and private major corporate actors mentioned in conjunction with sentiment analysis include Google, Microsoft, IBM, SAP, Salesforce, and DataSift, and non-technology companies such as Thompson-Reuters (publishing) and BAE Systems (defense). Start-ups include Idibon, Qriously, Ocado, Emotient, and Recorded Future, with business models that are largely centered on market research. Non-profits such as Demos, a UK-based think tank, are also mentioned, along with academic institutes and a number of individual developers.

A distinction is often made between explicit feedback on a product or service (e.g., in customer reviews) and social media discourse that may contain relevant implicit clues about the interests and preferences of consumers, and sentiment analysis is expected to be able to assess market sentiment, that is, the ability to track a variety of data sources to predict changes in the stock market. Some commentators point to the relatedness of the two domains, for example, remarking that sentiment analysis can be used anywhere from “political events to product launches.” Another special area is health care more generally, where often the interest is less in a particular product or service, but in predicting mental health problems or in being able to detect the symptoms of particular illnesses based on social media postings. Less frequently cited areas include education, law enforcement, disaster relief, and the insurance industry.

Table 1. Types of Codes.

Type	Examples
Explanations/definitions	<p>"Words and phrases are assigned a polarity score of positive, neutral or negative" (US, 203)</p> <p>"Tries to mine opinion about particular products or brands" (US, 128)</p> <p>"Trawl social media for any mention of your product and analyse what people are saying about a brand to produce insights" (US, 107)</p> <p>"Sophisticated sentiment analysis would turn written words into quantifiable consumer preferences" (US, 160)</p> <p>"Reliably capture mood" (GER, 23)</p> <p>"Sentiment analysis tools break down text and read its tone and affect, determining what is being talked about and how whoever wrote it felt about the subject" (UK, 104)</p> <p>"Depicts online opinion" (GER, 27)</p> <p>"Detects opinions and moods" (GER, 47)</p> <p>"The index tracks how people feel about each candidate" (US, 220)</p> <p>"Computers around the world are studying comments and rating them as positive, negative or neutral" (US, 229)</p> <p>"A computer program, instead of asking people questions, surveys the Web on a large scale" (US, 226)</p> <p>"Programmed properly, news analytics algorithms can recognize implications and sentiment in context" (US, 119)</p> <p>"Commercial research" (US, 156)</p> <p>"Insurance processes" (US, 167)</p> <p>"Predict a user's behavior" (US, 196)</p> <p>"Look for threats" (US, 201)</p> <p>"Understand how the public feels about something and track how those opinions change over time" (US, 203)</p> <p>"Help organizations identify, integrate, and correlate all of their customer information" (US, 208)</p> <p>"Identify bullying-related tweets" (US, 223)</p> <p>"Mental health" (UK, 57)</p> <p>"Financial analytics models to help investors achieve higher returns" (UK, 59)</p> <p>"Industries like finance and medicine" (UK, 61)</p> <p>"Students in Yorkshire and Humber" (UK, 65)</p> <p>"Theranos's public standing" (UK, 74)</p> <p>"Market analysis" (UK, 104)</p> <p>"Trading" (UK, 113)</p> <p>"Pharmaceutical sector" (UK, 133)</p> <p>"The Singapore anti-smoking campaign, which used sentiment analysis to show how online conversation became more disposed to quitting" (UK, 141)</p> <p>"The relatively new science of sentiment analysis has been dogged by the difficulties of getting computers to understand the complexities of natural language" (UK, 95)</p> <p>"A computer's ability to make sense of intent is limited" (US 235)</p> <p>"Digital media might be the most measurable, but it is also the easiest to misinterpret" (UK, 141)</p> <p>"But, remember, if Twitter was a reliable guide to public opinion, Scotland would be independent" (UK, 86)</p> <p>"Such analysis underestimated, for example, the performance of France's hard-right National Front" (US, 168)</p> <p>"The U.K. is particularly fond of sarcasm, making it especially difficult to determine a poster's sentiment by hashtag or keyword" (US, 168)</p>
Domains of use	
Issues of reliability	

Issues of Reliability

Issues of reliability and the practical limitations of sentiment analysis are also mentioned. Interestingly enough, this aspect is often framed as one of making computers understand language, even though this is precisely not what sentiment analysis does, instead scoring words and leaving the remainder of interpretative work to the analyst. When journalists argue that “a computer’s ability to make sense of intent is limited” or that sentiment analysis is “dogged by the difficulties of getting computers to understand the complexities of natural language,” they conflate scoring and interpretation into the same process. Some commentators are skeptical, arguing that language is easy to misinterpret, not a reliable guide to public opinion and that social media users in particular are prone to sarcasm. In some cases, the amount of data is cited, usually to indicate either that having large quantities of data facilitates analysis (which is hardly true per se) or, that is, increases reliability (which is generally not true). Context is widely referred to as something that limits reliability, although a number of seemingly distinct concepts are all variously referred to as “context.” In some cases, the fact that sentiment analysis cannot determine why someone is unhappy is described as a limitation, though depending on the specific case, this problem is not easily solved by humans either. A common thread in these characterizations is the implicit assumption that sentiment analysis, rather than merely speeding up inference processes made by people and performing them more consistently, is able to make qualitatively *better* inferences.

Issues of reliability are also often raised alongside with discussions of the data sources used in sentiment analyses. Data are widely described in broad terms, such as “online chatter,” “social media data,” or “big data,” or in terms of volume or type (“text,” “online comments”). However, the most common reference to a specific type of data source is to Twitter data or tweets, making up more than half of all references. While blog posts, social networks, and other online sources are also mentioned, this is mostly in combination with Twitter. Only a small number of examples are related to entirely different types of data, such as survey responses, novels, SMS, and email. Table 1 summarizes common claims about sentiment analysis.

Discussion

The results from our content analysis suggest the existence of an expanded set of assumptions about what sentiment analysis can do, where it should be used, and how reliable it is, ranging from the ability to “reliably capture mood” to “predicting a user’s behavior.” Although our analysis also contains references to the limitations of sentiment analysis, the references to mood tracking and behavior prediction suggest that sentiment analysis is attributed to qualities that were never part of its original design or application, with some

consequences for our understanding of the social world through the lens of this metric.

These assumptions contrast with how the method was conceived and is discussed in methodological literature in several ways. In computational linguistics, no claims are made regarding the validity of sentiment analysis as a method of empirical social research. In particular, representativeness is not important to model customer reactions within a particular e-commerce platform, and no causal inferences are made about the relation of a text’s sentiment and human behavior. Claims in psychology are often more ambitious though, positing that there exists a more immediate relationship between words and phrases and a subject’s emotions. However, the data used in hallmark studies in psychology have been elicited material whose production could be closely controlled by the researcher. Only in recent years has an expansion of LIWC to entirely other types of writing on social media platforms taken place. Similarly, computational approaches to sentiment generally assumed that a text had sentiment per the words used in it, but not necessarily that sentiment scores would reliably reveal the intentions of the writer or the effect on the reader. By focusing on very constrained types of texts, such as product reviews, the risk of misclassification could be kept reasonably low.

In computational linguistics and computer science, there are generally strong reservations against associating the sentiment of a text with the emotions of its author (Liu, 2010; Pang & Lee, 2008). Such claims are sometimes made by proponents of psychometric approaches such as Pennebaker and colleagues, who repeatedly argue that language use and behavior are linked (Pennebaker, Mehl, & Niederhoffer, 2003). Only through this linkage does language data become a legitimate data source for the social sciences. But because sentiment analysis is in many respects an exotic and arcane method in social science, a view that sometimes creeps into analysis is that the algorithm can interpret the same statement more accurately than a human could. Here, the direction of interpretational power becomes reversed: the algorithm does not replicate human behavior but becomes its motivation.

This shift becomes particularly apparent when examining sentiment analysis as a tool for market research. Our content analysis describes how public perceptions of sentiment analysis attribute qualities and capacities to sentiment analysis that diverge from how this method has been developed and refined in practice. Some perceptions—be they celebratory or critical—focus on the capacity of these methods to form predictions of mood from large repositories of data. This connects with a broader shift toward a perception of feeling as something to be computationally grasped.

Conclusion

This article has sought to describe the complex parallel histories of sentiment analysis in computational linguistics and

- disorganization*. Berkeley; Los Angeles: University of California Press.
- Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis, 21*, 267–297. doi:10.1093/pan/mps028
- Halford, S., & Savage, M. (2017). Speaking Sociologically with Big Data: Symphonic social Science and the future for big data research. *Sociology, 51*(6), 1132–1148. <https://doi.org/10.1177/0038038517698639>
- Hart, R. P. (2000). *Diction 5.0 User's Manual*. London: Scolaris Software, Sage Press.
- Hayles, N. K. (2012). *How we think: Digital media and contemporary technogenesis*. Chicago, IL: The University of Chicago Press.
- Hearn, A. (2010). Structuring feeling: Web 2.0, online ranking and rating, and the digital “reputation” economy. *Ephemera: Theory & Politics in Organization, 10*, 421–438.
- Jindal, N., & Liu, B. (2008). Opinion spam and analysis. In *Proceedings of the International Conference on Web Search and Web Data Mining (WSDM '08)* (pp. 219–229). New York, NY: ACM Press. doi:10.1145/1341531.1341560
- Karppi, T., & Crawford, K. (2016). Social media, financial algorithms and the hack crash. *Theory, Culture & Society, 33*, 73–92. doi:10.1177/0263276415583139
- Kennedy, H. (2012). Perspectives on sentiment analysis. *Journal of Broadcasting & Electronic Media, 56*, 435–450. doi:10.1080/08838151.2012.732141
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences of the United States of America, 111*, 8788–8790. doi:10.1073/pnas.1320040111
- Lindgren, S. (2012). “It took me about half an hour, but I did it!” Media circuits and affinity spaces around how-to videos on YouTube. *European Journal of Communication, 27*, 152–170. doi:10.1177/0267323112443461
- Liu, B. (2010). Sentiment analysis and subjectivity. In N. Indurkha & F. J. Damerau (Eds.), *Handbook of natural language processing* (2nd ed., pp. 627–666). New York, NY: Chapman & Hall.
- Margetts, H., John, P., Hale, S., & Yasseri, T. (2016). *Political turbulence: How social media shape collective action*. Princeton, NJ: Princeton University Press.
- Marwick, A. E. (2015). Instafame: Luxury selfies in the attention economy. *Public Culture, 27*, 137–160. doi:10.1215/08992363-2798379
- Murthy, D. (2015). Twitter and elections: Are tweets, predictive, reactive, or a form of buzz? *Information, Communication & Society, 18*, 816–831. doi:10.1080/1369118X.2015.1006659
- Ott, M., Choi, Y., Cardie, C., & Hancock, J. T. (2011). Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (HLT '11)* (pp. 309–319). New York, NY: ACM Press.
- Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval, 2*, 1–135. doi:10.1561/1500000011
- Papacharissi, Z. (2016). Affective publics and structures of storytelling: Sentiment, events and mediality. *Information, Communication & Society, 19*, 307–324. doi:10.1080/1369118X.2015.1109697
- Papacharissi, Z., & de Fatima Oliveira, M. (2012). Affective news and networked publics: The Rhythms of news storytelling on #Egypt. *Journal of Communication, 62*, 266–282. doi:10.1111/j.1460-2466.2012.01630.x
- Pennebaker, J. W., Booth, R. J., & Francis, M. E. (2007). *Linguistic Inquiry and Word Count (LIWC)*. Austin, TX. Retrieved from http://homepage.psy.utexas.edu/HomePage/Faculty/Pennebaker/Reprints/LIWC2007_OperatorManual.pdf
- Pennebaker, J. W., Mehl, M. R., & Niederhoffer, K. G. (2003). Psychological aspects of natural language. Use: Our words, our selves. *Annual Review of Psychology, 54*, 547–577. doi:10.1146/annurev.psych.54.101601.145041
- Peters, J. D. (2015). *The marvelous clouds: Toward a philosophy of elemental media*. Chicago, IL: University of Chicago Press.
- Puschmann, C., & Burgess, J. (2014). Metaphors of big data. *International Journal of Communication, 8*, 1690–1709. Retrieved from <http://ijoc.org/index.php/ijoc/article/view/2169>
- Pfeil, U., Arjan, R., & Zaphiris, P. (2009). Age differences in online social networking—A study of user profiles and the social capital divide among teenagers and older users in MySpace. *Computers in Human Behavior, 25*, 643–654. doi:10.1016/j.chb.2008.08.015
- Rambocas, M., & Gama, J. (2013, April). *Marketing research: The role of sentiment analysis* (FEP Working Paper no. 489). Porto, Portugal: University of Porto.
- Roberts, C. W. (1997). A theoretical map for selecting among text analysis methods. In C. W. Roberts (Ed.), *Text analysis for the social sciences: Methods for drawing statistical inferences from texts and transcripts* (pp. 275–283). Mahwah, NJ: Lawrence Erlbaum Associates.
- Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., & Ungar, L. H. (2013). Personality, gender, and age in the language of social media: The open-vocabulary approach. *PLoS ONE, 8*, e73791. doi:10.1371/journal.pone.0073791
- Sarawagi, S. (2008). Information Extraction. *Foundations and Trends in Databases, 1*(3), 261–377. <https://doi.org/10.1561/1900000003>
- Schwartz, H. A., & Ungar, L. H. (2015). Data-driven content analysis of social media: A systematic overview of automated methods. *The Annals of the American Academy of Political and Social Science, 659*, 78–94. doi:10.1177/0002716215569197
- Stone, P. J. (1997). Thematic text analysis: new agendas for analyzing text content. In C. W. Roberts (Ed.), *Text analysis for the social sciences: Methods for drawing statistical inferences from texts and transcripts* (pp. 35–54). Mahwah, NJ: Lawrence Erlbaum Associates.
- Stone, P. J., Dunphy, D. C., Smith, M. S., & Ogilvie, D. M. (1966). *General Inquirer: A computer approach to content analysis*. Cambridge, MA: MIT Press.
- Su, Y.-F., Cacciatore, M. A., Liang, X., Brossard, D., Scheufele, D. A., & Xenos, M. A. (2017). Analyzing public sentiments online: Combining human- and computer-based content analysis. *Information, Communication & Society, 20*, 406–427. doi:10.1080/1369118X.2016.1182197

- Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37, 267–307. doi:10.1162/COLI_a_00049
- Tausczik, Y. R., & Pennebaker, J. W. (2009). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, 29(1), 24–54. https://doi.org/10.1177/0261927X09351676
- Thelwall, M. (2017). The heart and soul of the web? Sentiment strength detection in the social web with SentiStrength. In J. A. Holyst (Ed.), *Cyberemotions, understanding complex systems* (pp. 119–134). Berlin, Germany: Springer. doi:10.1007/978-3-319-43639-5_7.
- Tufekci, Z. (2013). “Not this one”: Social movements, the attention economy, and microcelebrity networked activism. *American Behavioral Scientist*, 57, 848–870. doi:10.1177/0002764213479369
- Turow, J. (2012). *The daily you: How the new advertising industry is defining your identity and your worth*. New Haven, CT: Yale University Press.
- Van Dijck, J., & Poell, T. (2013). Understanding social media logic. *Media and Communication*, 1(1), 2–14. doi:10.17645/mac.v1i1.70
- von Nordheim, G., Boczek, K., Koppers, L., & Erdmann, E. (2018). Reuniting a divided public? Tracing the TTIP debate on Twitter and in traditional media. *International Journal of Communication*, 2, 548–569.
- Wilson, T., Wiebe, J., & Hoffmann, P. (2005). Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT)* (pp. 347–354). Morristown, NJ: The Association for Computational Linguistics. doi:10.3115/1220575.1220619
- Wojcieszak, M., & Azrout, R. (2016). I saw you in the news: Mediated and direct intergroup contact improve outgroup attitudes. *Journal of Communication*, 66, 1032–1060. doi:10.1111/jcom.12266
- Young, L., & Soroka, S. (2012). Affective news: The automated coding of sentiment in political texts. *Political Communication*, 29, 205–231. doi:10.1080/10584609.2012.671234

Author Biographies

Cornelius Puschmann is a senior researcher at the Hans Bredow Institute for Media Research in Hamburg. His research interests include online hate speech, the role of algorithms for the selection of media content, and methodological aspects of computational social science.

Alison Powell is an assistant professor in the Department of Media and Communications at the London School of Economics and Political Science. She researches how people’s values influence the way technology is built, and how technological systems in turn change the way we work and live together.