

An investigation of the intellectual structure of opinion mining research

Yongjun Zhu, Meen Chul Kim, and Chaomei Chen

Introduction. *Opinion mining has been receiving increasing attention from a broad range of scientific communities since early 2000s. The present study aims to systematically investigate the intellectual structure of opinion mining research.*

Method. *Using topic search, citation expansion, and patent search, we collected 5,596 bibliographic records of opinion mining research. Then, intellectual landscapes, emerging trends, and recent developments were identified. We also captured domain-level citation trends, subject category assignment, keyword co-occurrence, document co-citation network, and landmark articles.*

Analysis. *Our study was guided by scientometric approaches implemented in CiteSpace, a visual analytic system based on networks of co-cited documents. We also employed a dual-map overlay technique to investigate epistemological characteristics of the domain.*

Results. *We found that the investigation of algorithmic and linguistic aspects of opinion mining has been of the community's greatest interest to understand, quantify, and apply the sentiment orientation of texts. Recent thematic trends reveal that practical applications of opinion mining such as the prediction of market value and investigation of social aspects of product feedback have received increasing attention from the community.*

Conclusion. *Opinion mining is fast-growing and still developing, exploring the refinements of related techniques and applications in a variety of domains. We plan to apply the proposed analytics to more diverse domains and comprehensive publication materials to gain more generalized understanding of the true structure of a science.*

Introduction

Opinion mining refers to the task of finding opinions of people about specific entities (Feldman, 2013). It employs computational linguistics and natural language processing to identify and extract subjective information from source texts. Sentiment analysis, which parallels opinion mining, aims to quantify and determine the polarity of an individual with regards to judgement, affective state, or emotion (Pang and Lee, 2008). Sentiment analysis is used to automatically analyse evaluative texts and materials. Opinion mining and sentiment analysis are often seen as interchangeable concepts and express a mutual meaning (Medhat, Hassan, and Korashy, 2014). Pang and Lee (2004) described the year of 2001 as the beginning of widespread awareness of opinion mining and sentiment analysis. Most of the early studies placed emphasis on proposing computational techniques to measure the sentiment orientation of product reviews in a binary way (Thet, Na, and Khoo, 2010). Recent studies have been working on in-depth examination of source texts with a variety of techniques and units of analysis such as document, sentence, feature, or aspect. Within this context, there has been an exponentially increasing number of articles written on this topic (See Figure 2). Hundreds of startup companies are also developing sentiment analysis solutions and software packages (Feldman, 2013).

As the volume of the literature has exponentially grown, a wide range of communities has also paid large attention to opinion mining. Thus, a systematic analysis of the intellectual structure of the field is needed to understand its main applications and current challenges. The systematic domain analysis will allow the

followings. First, it helps the opinion mining community to be more self-explanatory as it has a detailed bibliometric profile. Secondly, researchers in the field can benefit from the systematic domain analysis by better positioning their research, identifying research trends, and expanding research territories. Finally, it guides researchers who are interested in the field to learn about emerging trends and current challenges. To our knowledge, however, there has been little concerted effort which aims to systematically understand the intellectual landscapes in opinion mining. The motivation of the present study lies in our intention to identify the intellectual structure of opinion mining and sentiment analysis research in a quantitative and systematic manner. We explored thematic patterns, landmark articles, emerging trends, and new developments of the field. In particular, our study was guided by a computational approach implemented in CiteSpace, a visual analytic system for illustrating emerging trends and critical changes in scientific literature (Chen, 2006; Chen *et al.*, 2012). We also employed a dual-map overlay technique (Chen and Leydesdorff, 2014) to expand the present study to the investigation of citation patterns at a disciplinary level .

In what follows, three research questions frame our investigation:

- *Research question #1:* what are epistemological characteristics of opinion mining research?
- *Research question #2:* which thematic patterns of research do occur in the domain?
- *Research question #3:* what are emerging trends and recent developments in the field?

The rest of the study is organized as follows. First, we survey preceding work. Then, the scientometric methodology of the study is introduced. We describe intellectual structure of the field and key findings. This is followed by the conclusion and proposed future study.

Related work

To date, there has been little research that has systematically investigated comprehensive intellectual landscapes and emerging research trends in opinion mining. Therefore, we examined review articles on opinion mining since review papers aim to survey and synthesize key findings and research trends in a specific domain of interest. Pang and Lee (2008) presented a detailed review on broad aspects of opinion mining and sentiment analysis. This survey featured applications of opinion mining and sentiment analysis such as classification and extraction of textual units and summarization of sentiment information from a set of single and multi-documents. Based on the authors' expertise on the field, a large number of articles were surveyed. A list of classifiers and their performance in grouping opinions were discussed by Govindarajan and Romina (2013). They summarized the performance of naive Bayes, support vector machine, and genetic algorithm on a variety of data sources along with the selection of different features. Evolution of the field and its future research trends can be found at a work done by Cambria and colleagues (2013). They reviewed 1) opinion mining techniques ranging from heuristics to discourse structure, 2) different levels of analysis ranging from document-level to aspect-level, and 3) four approaches regarding the identification of implicit information associated with word tokens such as keyword spotting, lexical affinity, statistical methods, and concept-based approaches. The authors considered the multi-modal sentiment analysis as one of the promising approaches, laying stress on the importance of investigating a variety of sources such as textual, acoustic, and video features. Feldman (2013) discussed opinion mining techniques to solve five specific problems: 1) document-level sentiment analysis, 2) sentence-level sentiment analysis, 3) aspect-level sentiment analysis, 4) comparative sentiment analysis, and 5) sentiment lexicon

acquisition. He introduced major applications of sentiment analysis such as customer review mining towards products and services, preference mining on candidates running for political positions, and market value prediction. Medhat and colleagues (2014) surveyed 54 research papers on algorithms and applications of sentiment analysis. They categorized these papers into six groups: 1) sentiment analysis, 2) sentiment classification, 3) feature selection, 4) emotion detection, 5) building resource, and 6) transfer learning. This paper showed that naive Bayes and support vector machine are the most frequently used techniques for grouping opinions. In addition, WordNet was said to be the most commonly used lexical source to solve sentiment analysis problems.

The review articles discussed above examined a variety of aspects of opinion mining research. However, there has been little concerted effort to investigate emerging trends and recent developments in opinion mining in such a comprehensive, systematic, and computationally driven manner that we use in the present study. In addition, one commonality of these papers is that the surveyed articles were selected based on prior domain knowledge or without specific selection criteria. While aggregated and/or brief discussions on individual papers were provided, the implication and importance of individual papers to the field were missing. This may have helped readers approach individual papers more systematically and selectively for a deeper understanding. The absence of clear thematic categories is also one possible limitation of these reviews. Therefore, these surveys lack objectivity and clarity in describing the intellectual structure and emerging trends in the field. Since opinion mining is a developing, fast-growing domain, a systematic and comprehensive investigation of intellectual landscapes and recent developments is essential. To bridge these gaps, the present study aimed to explore the intellectual structure of the field through a bibliometric analysis of extensive scientific literature, taking a variety of units of analysis into account. We also took a close look at landmark articles of the field that are chosen by multi-faceted criteria.

Methods

In this section, we introduce the data collection method and analytical approaches of the present study. The research procedure is shown in Figure 1.

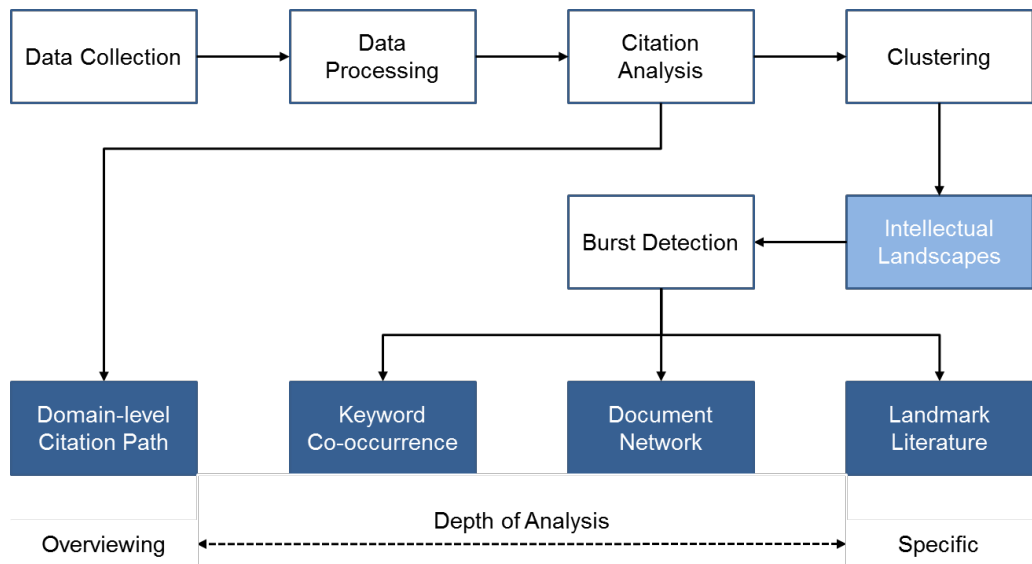


Figure 1: Research procedure

Data collection

The goal of the present study was to explore the intellectual structure of opinion

mining research. We aimed to identify thematic patterns, research networks, emerging trends and new developments, together with landmark articles in the domain. Toward that end, we retrieved bibliographic records from the Web of Science, using a topic search. There were two problems that we encountered with this data collection approach per se. First, it is acknowledged that the topic search does not include relevant records if querying terms do not appear in the targeted fields such as titles, abstracts, and keywords. In addition, despite the fact that conference proceedings are a common document type in computer and information-related sciences, the Web of Science under-represents publications from conferences (Kim and Chen, 2015). To remedy these issues, we employed a two-step complementary approach in data collection. First, we assumed that if an article cites at least one of the records retrieved by the topic search, then the article is topically relevant to opinion mining (Chen, Hu, Liu, and Tseng, 2012). As Garfield (1979) argues, citation indexing is an alternative strategy to capture a much broader context of a study. Therefore, we collected additional records, using a citation expansion implemented on the Web of Science. The dataset obtained by the topic search was regarded as the core dataset and the expanded one represents a broader context of the core. Using the citation expansion supports our goal to explore much broader landmark references influencing emerging trends and recent developments in opinion mining. Second, we triangulated our data collection by conducting patent and citation searches on Derwent Innovations Index. This database indexes technological inventions in chemical, electrical, electronic, and mechanical engineering.

The detailed procedure of our data collection is as follows. First, a data collection method proposed by preceding scientometric research (Chen, Dubin, and Kim, 2014a, b; Kim and Chen, 2015; Kim, Zhu, and Chen, 2016) was used. For this basic search, four query phrases identified as representative for the domain were used: “*opinion mining*” OR “*sentiment anal**” OR “*subjectivity anal**” OR “*review mining*”. We employed the wildcard * to capture relevant variations of a word, including no character, such as sentiment analysing and sentiment analytics. The use of double quotation marks helped queries considered to be clauses. Then, a record was regarded as relevant if any of the terms is found in the title, abstract or keyword fields of the record. The queries resulted in 2,010 records from 2002 through 2015 as of August 31 2016 when considering articles, proceeding papers, and review articles as representative record types. Then, the core records were cited by 3,418 articles, proceedings, and review articles on the Web of Science. Finally, the query set returned 168 patent records between 2005 and 2015 from Derwent Innovations Index. We integrated all of these data sets and used the merged one for the present scientometric investigation. Table 1 describes the brief statistics of the dataset. Figure 2 depicts the number of records over time. As depicted in the figure, opinion mining has received exponentially increasing attention from the community. Table 2 describes the query terms used to search records and the number of corresponding records to each term. It shows “*sentiment anal**” contributes the literature most.

Dataset	Duration	Results	Articles	Proceedings	Reviews	Authors	References	Keywords
Core	2002-2015	2,010	643	1,322	45	6,470	52,948	12,815
Expanded	2005-2016	3,418	2,170	1,222	126	11,354	150,503	31,393
Patent	2005-2015	168	-	-	-	577	936	1,047

Table 1: Data statistics

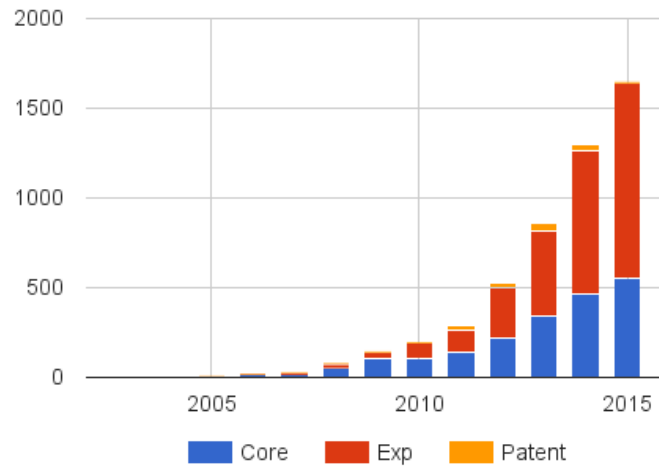


Figure 2: Records distribution over time

Query term	Core dataset	Expanded dataset	Patent dataset
sentiment anal*	1,615	2,972	149
opinion mining	767	1,591	17
subjectivity anal*	34	213	1
review mining	30	76	1

Table 2: Query terms and retrieved records (in descending order of the number of core records)

Investigation of the intellectual landscapes, emerging trends, and new developments in opinion mining

Scientometrics is the quantitative study of science. We employed scientometric approaches to analyse bibliographic records for the present domain analysis. There are several advantages in our study in comparison with a conventional domain analysis as follows (Chen, Dubin, and Kim, 2014b). Firstly, a much broader and more diverse range of relevant topics can be explored due to the use of citation expansion and patent records. Secondly, this kind of domain investigation can be conducted as frequently as needed, although an inquirer does not have prior expertise in a targeted domain. Finally, a scientometric analysis provides an additional point of reference.

CiteSpace, developed by Chen (2006), is a scientometric toolbox to generate and analyse networks of co-cited references using bibliographic records. The input is a collection of academic literature relevant to a specific topic. Given bibliographic records, this tool computationally detects and depicts thematic patterns and emerging trends in science. On top of this, CiteSpace provides a visual representation called a dual-map overlay which renders a domain-level view of the growth of the literature (Chen and Leydesdorff, 2014).

In a scientometric study, the intellectual landscapes of a scientific domain can be represented by a variety of networked entities such as cited references and keywords (Chen, Dubin, and Kim, 2014b). Specifically, we focused on document citation networks and networks of co-occurring keywords to explore emerging trends and recent developments in opinion mining research. The key features of the present investigation include: 1) domain-level citation paths, 2) frequently assigned subject categories, 3) keyword co-occurrence, 4) networks of highly cited references, and 5) influential articles selected by a variety of metrics.

Terminology

In order to clearly communicate the technical approaches and findings of the present study to the audience, we introduce structural measures and text mining techniques used in this paper as follows:

g-index: The *g*-index, suggested by Egghe (2006), is an author-level metric to measure the scientific productivity of an individual. It considers the unique largest number of the top *g* highly cited articles received together more than or equal to *g* square citations. Examining the entire entities in our data, *i.e.* keywords and cited references, may be computationally challenging. It may not intuitively communicate the true structure of the domain to the audience as well. In addition, we argue that employing the *g*-index in selecting core entities is sounder than using the *h*-index since the *g* number of cores is always at least as big as the *h* cores. Thus, we used *g*-index to select a significant fraction of frequently occurring keywords and cited articles within a 1-year slice of time .

Pathfinder networks: Bibliographic networks can be highly dense with many links between entities. Network pruning or link reduction, the process in order to systematically remove excessive links, can address this issue. Based on the proximity of entity pairs, pathfinder networks capture the shortest paths, so links are eliminated when they are not on shortest paths (Schvaneveldt, Durso, and Dearholt, 1989). In this study, we employed pathfinder networks to eliminate redundant links between entities of analysis.

Betweenness centrality: Betweenness centrality is an indicator of a node considering the number of shortest paths from all vertices to all others that pass through the node (Brandes, 2001). A node with a high value of betweenness centrality has a large influence on the transfer of information through the network (Chen, 2011). If a node provides the only link between two large but previously unconnected groups of nodes, it would have a very high degree of betweenness centrality. In our study, we regarded this topological property as a significant sign of a bibliographic entity's influence.

Burstiness: Kleinberg (2002) proposed an algorithm called burst detection which captures the burstiness of events with certain features rising sharply in frequency. Based on this concept, an entity can be regarded as having bursting activities if it shows an intensive frequency of appearance during a specific duration of time. It overcomes the limitation of just considering the cumulative number of metrics such as citations as a measure of an entity's impact.

Sigma: Sigma is a measure identifying scientific publications with topological burstiness. It is defined as (betweenness centrality + 1) to the power of burstiness such that the temporal brokerage mechanism plays more prominent role than the rate of raw citations (Chen, *et al.*, 2009). We regarded this metric as another important sign of a bibliographic entity's structural burst.

Automatic cluster labelling: In order to automatically label clusters of cited references, we extracted candidate terms from titles and abstracts of citing articles. In CiteSpace, these terms are selected by three different algorithms: 1) latent semantic indexing (LSI) (Deerwester, *et al.*, 1990), 2) log-likelihood ratio (LLR) (Dunning, 1993), and 3) mutual information (MI). Labels extracted by LSI tend to capture implicit semantic relationships across records, whereas those chosen by LLR and MI tend to reflect a unique aspect of a cluster (Chen, Ibekwe-SanJuan, and Hou, 2010).

Results

Research trends at a disciplinary level

A dual-map overlay is an analytical approach that presents the domain-level concentration of citations through their reference paths (Chen and Leydesdorff, 2014). A base map depicts the interconnections of over 10,000 published journals and these journals are grouped into some regions that represent publications and citation activities at a domain level (Chen, Dubin, and Kim, 2014a). The dual-map overlay of opinion mining research is displayed in Figure 3. In the visual representation, the left clusters, *i.e.* research fronts, represent where the retrieved records publish, while the right clusters indicate where they cite. Regions are labelled by common terms found in the underlying journals. Citation trajectories are distinguished by citing regions' colours. The thickness of these trajectories is proportional to the z-score-scaled frequency of citations. Based on this map, we can identify patterns of how published articles in opinion mining refer to other intellectual bases (cited references). As rendered in the figure, there are four main citation paths in our dataset. Table 3 summarizes these paths with citing and cited region names.

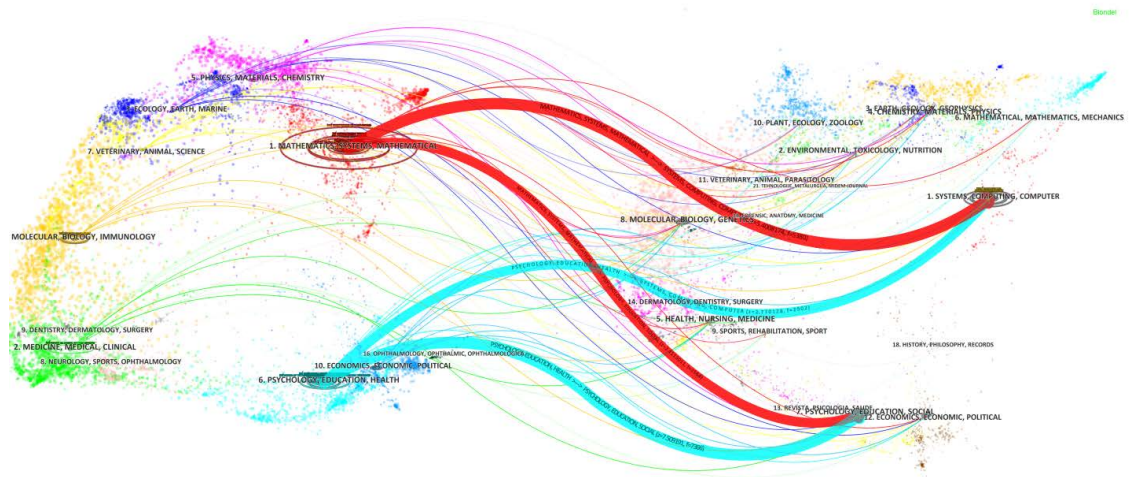


Figure 3: Domain-level citation patterns in opinion mining research

Citing region	Cited region	Z-score
psychology, education, health	psychology, education, social	7.729
mathematics, systems, mathematical	systems, computing, computer	5.803
psychology, education, health	systems, computing, computer	2.727
mathematics, systems, mathematical	psychology, education, social	2.468

Table 3: Citation trends at a domain level

The relationships are sorted by the z-scores in descending order where the values are rounded to the nearest thousandth. Each row is identified by the same colour of the corresponding path displayed in Figure 3. As described in Table 3, the domains the most frequently covering the records are: 1) *6. psychology, education, health*, and 2) *1. mathematics, systems, mathematical*. Then, this literature is mostly influenced by *7. psychology, education, social* and *1. systems, computing, computer*. On top of these domains, *5. physics, materials, chemistry* (rendered in purple lines), *2. molecular, biology, immunology* (depicted in yellow lines), and *4. medicine, medical, clinical* (illustrated in green lines) contribute to the domain-level citation trends in opinion mining. It indicates a multidisciplinary aspect of opinion mining since publications from multiple

domains contribute to the citation landscape of the domain. It also shows that social sciences study opinion mining and sentiment analysis within similar disciplinary contexts while the mathematical and algorithmic concepts of the domain are mostly influenced by the literature that investigates systems and computer (See the first and second rows of Table 3). Based on these observations, we argue that opinion mining research is partially multidisciplinary and partially monodisciplinary. We also identify an interdisciplinary characteristic as described at the third and fourth rows of Table 3. It is shown that psychological, educational, and biomedical investigations in opinion mining are based on systems perspectives of computing and vice versa.

A subject category of an article can also be regarded as evidence showing an upper-level thematic concentration of the article. In order to specify the findings above, we examined the Web of Science subject categories assigned to the records. Table 4 describes 20 subject categories most frequently given to the dataset. The table shows each category's year of first occurrence, cumulative frequency, and betweenness centrality in descending order of the assignment frequency (third column). As described in the table, computer science, engineering, linguistics, and social sciences such as library and information science, business, economics, management, and psychology are among the leading subject categories in opinion mining. Among them, computer science, engineering, and their related fields have been assigned to the records most frequently. In addition, the category, interdisciplinary applications of computer science, shows to have the highest betweenness centrality of 0.21. It indicates that interdisciplinary applications of computer science have had the largest influence on the emergence, development, and diffusion of the ideas in opinion mining. In turn, psychology plays an important bridging role between domains that participate in opinion mining research (betweenness centrality: 0.20). It is also obvious that the publications from electrical engineering, artificial intelligence, linguistics, and library and information science have transferred scientific findings to opinion mining. Recently, opinion mining has received significant attention from operations research and management. It is also interesting to see that business and economics have published literature from the very early years of the domain. This may indicate that opinion mining started with an aim of understanding and scaling customers' subjective opinions toward products. Based on these observations, we assume that the researchers in the domain began by investigating algorithmic, linguistic, and psychological aspects of opinion mining in order to understand, quantify, and measure the sentimental orientation of texts. Successive literature may have explored the improvement and practical applications of the techniques to multiple domains. In the next subsections, we try to understand these findings deeper at the keyword and reference levels.

Category	First occurrence	Frequency	Centrality
cs	2003	3,453	0.03
cs, artificial intelligence	2003	1,769	0.10
cs, information systems	2005	1,671	0.03
cs, theory & methods	2007	1,303	0.03
engineering	2007	1,029	0.03
engineering, electrical	2007	888	0.14
cs, interdisciplinary apps.	2004	423	0.21
info. sci. & lib. sci.	2005	348	0.07
business & economics	2002	332	0.02
ops. res. & mgmt. sci.	2008	327	0.03
cs, software engineering	2006	326	0.01
linguistics	2004	251	0.08
language & linguistics	2004	227	0.02

business	2006	188	0.06
telecommunications	2008	182	0.06
cs, hardware & architecture	2007	182	0.01
management	2006	177	0.04
robotics	2008	145	0.00
cs, cybernetics	2006	108	0.04
psychology	2006	106	0.20

Table 4: Top 20 frequently assigned subject categories in the dataset

Keywords as indicators of emerging trends and new developments

The investigation of keywords can add richer interpretations to understanding the concentration of research themes since keywords represent underlying concepts of an article. Table 5 describes 20 keywords the most frequently given by authors and indexers to the records. It shows each keyword's year of first occurrence, cumulative frequency, and betweenness centrality. Based on its frequency of appearance and centrality, it is evident that *sentiment analysis* is the keyword considered as the most representative of the literature from the early years of the domain. This keyword is followed by *opinion mining* which is the application of sentiment analysis to a variety of textual materials such as *web* and *review*. Following *sentiment analysis* (betweenness centrality: 0.20), *emotion* is also located on the shortest path connecting pairs of other concepts in opinion mining research (betweenness centrality: 0.19). It indicates that the identification and extraction of subjective information in source texts have been regarded as the most important concepts in opinion mining and had the largest influence on the growth of the domain. Technique-wise, *text mining*, *machine learning*, and *natural language processing* have been frequently employed to quantify sentiment analysis. Classification is among the most frequently investigated techniques, also being a goal of understanding the sentiment orientation of text. As reflected in the recent keywords such as *social media*, *twitter*, *social network*, and *word of mouth*, the literature now focuses on the practical applications of opinion mining to marketing and social networking platforms.

Keyword	First occurrence	Frequency	Centrality
sentiment analysis	2006	1,458	0.20
opinion mining	2006	721	0.09
social media	2011	378	0.07
twitter	2011	340	0.06
classification	2009	303	0.07
text mining	2007	246	0.05
model	2010	242	0.03
social network	2009	207	0.03
machine learning	2006	194	0.07
emotion	2005	184	0.19
word of mouth	2011	182	0.07
web	2010	179	0.02
text	2009	178	0.11
information	2010	174	0.02
system	2009	165	0.03
network	2010	159	0.03
natural language processing	2009	154	0.04
review	2011	153	0.01
sentiment classification	2006	145	0.09
internet	2007	132	0.04

Table 5: Top 20 frequently occurring keywords in the dataset

Figure 4 (below) displays the co-occurrence network of these keywords. For the clarity of the figure, we only labelled keywords listed in Table 5. The size of a node is proportional to its frequency and each node is coloured by community membership (Blondel, Guillaume, Lambiotte, and Lefebvre, 2008). Finally, the more frequently keywords co-occur, the more closely they are located. As depicted in the figure, the formulation of the network is consistent with the findings above.

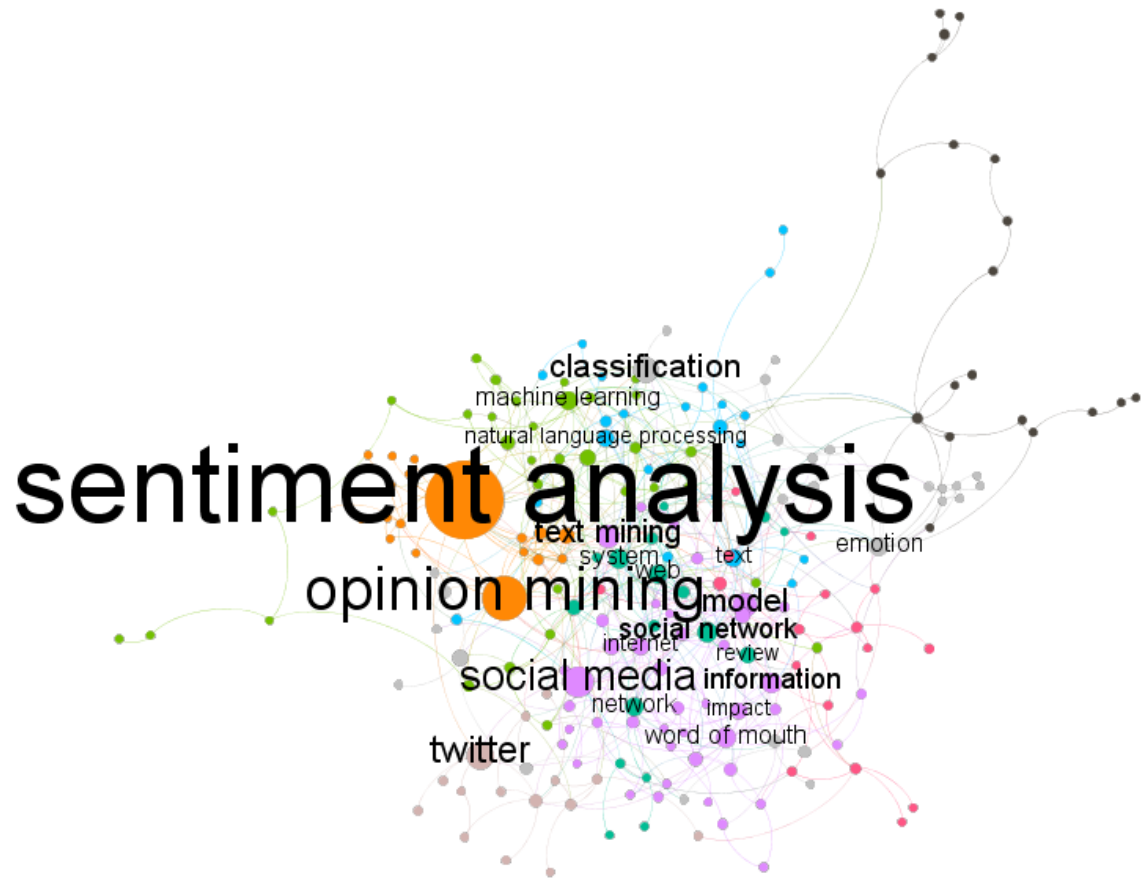


Figure 4: Keyword co-occurrence network

Table 6 (below) describes 20 keywords receiving the most intensive attention during a specific span of time. We sorted the keywords by the beginning years of the bursts so that we could identify temporal patterns. *Opinion mining* and *sentiment analysis* are among the most intensively bursting keywords across the entire time duration. *Sentiment classification* has been one of the most important goals from the earlier years. *Opinion extraction* is the task of automatically extracting structured opinions from unstructured data such as text. For *sentiment classification* and *opinion extraction*, *machine learning* and *text mining* have been employed from the community. In the previous section of examining trends in subject category assignment, we assumed that literature in an earlier stage may have explored linguistic aspects sentiment orientation. The burstiness of the keyword *language* complements this assumption. A variety of techniques to extract sentiment and to apply quantified subjectivity have shown bursting trends with *language* (See *information retrieval* and *information extraction*). Then, the automatic collection of a *corpus* from consumer contributed content such as *blog* has had large attention and *natural language processing* has been intensively employed to this end. Recently, there has been growing interest in applying *data mining* and *sentic computing* to a variety of textual genres such as *web* and *microblogging*. Thus, investigating the burstiness of keywords adds richer

interpretations to the understanding of the emerging trends in opinion mining than just considering the cumulative number of keyword occurrence.

Keyword	Burst	Begin	End	2002-2015
opinion mining	52.860	2006	2012	
sentiment analysis	8.549	2006	2011	
sentiment classification	7.197	2006	2011	
machine learning	4.859	2006	2009	
text mining	8.486	2007	2011	
opinion extraction	3.621	2007	2011	
language	9.148	2008	2012	
information retrieval	4.798	2008	2011	
information extraction	3.619	2008	2010	
corpus	2.729	2009	2012	
blog	5.115	2009	2013	
natural language processing	3.434	2009	2010	
opinion analysis	3.684	2009	2011	
data mining	3.700	2010	2012	
document	3.796	2010	2013	
sentiment analysis	7.981	2010	2012	
web	7.511	2010	2012	
text analysis	3.220	2011	2013	
microblogging	4.547	2011	2012	
sentiment computing	3.922	2011	2012	

Table 6: Top 20 bursting keywords in the dataset (bursts rounded to the nearest thousandth)

Document co-citation network

Figure 5 (below) visualises the document co-citation network in the dataset. The network consists of g highly cited references within a given slice of time. Each node represents a cited article and the size of a node is proportional to its cited frequency. References with citation bursts are rendered with rings in red whereas nodes with purple rings have high betweenness centrality values. Landmark articles cited more than or equal to 133 times are labelled in black. Nodes are grouped in same lines by a clustering technique called smart local moving (Waltman and van Eck, 2013). Clusters are numbered in such a way that higher rankings are given to the clusters containing more references. The colour legend at the top indicates that links and citations in cooler colours happen closely to the year of 2002 whereas hotter ones occur in close years to 2015. Referring to the legend, we can keep track of the thematic trends and temporal developments in opinion mining research.

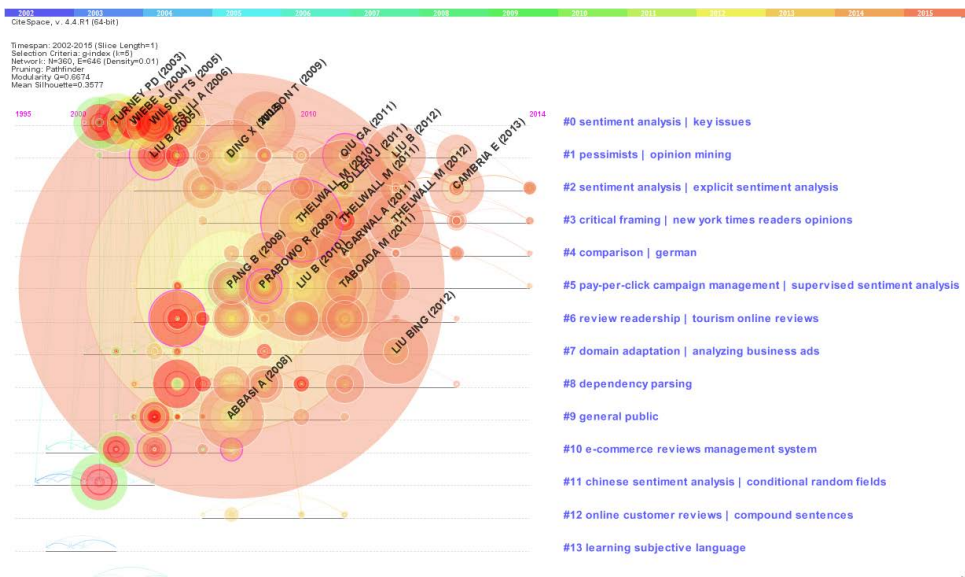


Figure 5: Timeline visualization of the network (labelled by LSI)

Considering both the size and recency of member nodes, we consider clusters #0 through #4 as representing emerging research themes in the domain. Table 7 summarizes these clusters in terms of number of member articles, three types of labels extracted from titles and abstracts, and average published year of cites. Among these clusters, cluster #0 is the largest and oldest one in terms of the cluster size and the mean year of cited references. The references in this cluster influenced successive research applying sentiment analysis to the investigations of social media and user requirements. It is interesting that even though social media is one of the recently appearing keywords (See Table 5), relatively old papers influenced research under this theme. Cluster #1 is the second largest and second oldest one. Literature citing references in this group focuses on mining opinion words and analysing customer preferences. Cluster #2, one of the most recent ones, consists of references which influenced successive research on developing a fuzzy system for the extraction of affective information. Cluster #3 includes references affecting later literature on understanding newspaper readers' opinions and sharing behaviour on social media. Cluster #4 studies the application of sentiment analysis to comparing specific products. It coheres with the findings from examining subject categories as reflected in *management* and *business* (See Table 4) and keyword assignment (See *word of mouth* in Table 5).

Cluster	Size	Cluster labelling technique			Mean year
		LSI	Log-likelihood ratio	Mutual information	
0	44	sentiment analysis key issues	social media	user requirement	2004
1	37	Pessimists opinion mining	opinion word	customer preferences analysis	2008
2	35	sentiment analysis explicit sentiment analysis	affective information	fuzzy system	2011
3	35	critical framing new york times readers opinions	social media	sharing behavior	2011
4	26	comparison german	electronic cigarette	comprehensive empirical comparison	2010

Table 7: Emerging clusters in the dataset

Landmark literature

Emerging trends and recent developments in a domain can be investigated in detail by surveying influential manuscripts of the discipline. For measuring the importance of an article, we considered a variety of indicators such as cumulative citation counts, intensity of citations during a specific span of time, topological importance, and burstiness of structural importance. Tables 8, 9, 10, and 11 show landmark articles identified by these criteria. Landmark articles are extracted from cited references. In each table, manuscripts are coloured as follows: 1) green: articles that appear three times across all the tables, 2) yellow: articles that appear in two tables, and 3) white: articles that only appear in one table. In reviewing them, we chose only the articles that appear at least twice across the tables. In addition, some articles are not considered for this discussion if they are review papers, books, and book chapters. Topically irrelevant articles are also omitted. We argue that they do not influence thematic patterns and emerging trends in a domain. Eight papers selected in this way are summarized in Table 12. In the following paragraphs, we discuss each of eight papers in chronological order.

Citation count	Reference	Cluster #
888	Pang and Lee (2008)	5
257	Taboada, <i>et al.</i> (2011)	5
256	Liu (2010)	5
190	Bollen, Mao, and Zeng (2011)	2
174	Thelwall, <i>et al.</i> (2010)	3
166	Ding, Liu, and Yu (2008)	1
147	Abbasi, Chen, and Salem (2008)	9
144	Wilson, Wiebe, and Hoffmann (2009)	0
142	Liu (2012)	7
133	Cambria, Schuller, Xia, and Havasi (2013)	1

Table 8: Top 10 highly cited articles in the dataset

Burst	Reference	Cluster #
51.24	Turney (2002)	11
47.15	Pang, Lee, and Vaithyanathan (2002)	0
46.48	Turney and Littman (2003)	0
35.94	Wiebe, <i>et al.</i> (2004)	0
34.39	Pang and Lee (2004)	0
28.45	Dave, Lawrence, and Pennock (2003)	10
22.54	Kim and Hovy (2004)	0
21.49	Wilson (2005)	0
20.18	Liu, Hu, and Cheng (2005)	1
13.27	Pang and Lee (2005)	10

Table 9: Top 10 articles with the highest citation bursts

Centrality	Reference	Cluster #
0.14	Jindal and Liu (2008)	10
0.12	Pang and Lee (2005)	10
0.12	Chevalier and Mayzlin (2006)	6
0.12	Tang, Tan, and Cheng (2009)	5
0.11	Liu, Hu, and Cheng (2005)	1
0.10	Kanayama and Nasukawa (2006)	1
0.10	Thelwall, <i>et al.</i> (2010)	3
0.10	Qiu, Liu, Bu, and Chen (2011)	1

0.09	Wilson (2005)	0
0.09	Riloff, Wiebe, and Wilson (2003)	0

Table 10: Top 10 articles with the highest betweenness centrality values

Sigma	Reference	Cluster #
26.48	Pang, Lee, and Vaithyanathan (2002)	0
18.09	Turney (2002)	11
14.54	Turney and Littman (2003)	0
7.62	Liu, Hu, and Cheng (2005)	1
6.12	Wilson (2005)	0
4.41	Chevalier and Mayzlin (2006)	6
4.32	Pang and Lee (2005)	10
4.21	Kim and Hovy (2004)	0
3.90	Riloff, Wiebe, and Wilson (2003)	0
2.63	Yu and Hatzivassiloglou (2003)	0

Table 11: Top 10 articles with the highest sigma

Reference	Cluster #	Citation	Burst	Centrality	Sigma
Liu, Hu, and Cheng (2005)	1		X	X	X
Pang and Lee (2005)	10		X	X	X
Turney (2002)	11		X		X
Pang, Lee, and Vaithyanathan (2002)	0		X		X
Turney and Littman (2003)	0		X		X
Riloff, Wiebe, and Wilson (2003)	0			X	X
Kim and Hovy (2004)	0		X		X
Thelwall, <i>et al.</i> (2010)	3	X		X	

Table 12: Select landmark articles in the dataset

Turney (2002) proposed an unsupervised learning algorithm using mutual information to classify reviews into positive and negative groups. The method calculates the semantic orientation of a phrase, considering mutual information between the phrase and the sentiment groups, *excellent* and *poor*. The semantic orientation of a review is then classified compared to the average semantic orientation of included phrases. The proposed technique achieved an average accuracy of 74%, given a dataset consisting of documents on automobiles, banks, movies, and travels. Even though the method is simple and straightforward, the performance of the method needs to be further validated, due to the limited size of the dataset.

Pang, Lee, and Vaithyanathan (2002) tested the performance of three machine learning algorithms: Naïve Bayes, maximum entropy classification, and support vector machine on the sentiment classification of movie reviews. In testing these algorithms, they employed different feature sets such as unigrams, bigrams, and parts of speech. The support vector machine based on unigrams achieved the best performance, which reached 82.9%. The contribution of the work lies in its broad investigation of different features in classifying peoples' opinions. The authors laid stress on the importance of identifying plausible features that screen whether sentences are on-topic. The application of these approaches to other domains might be helpful for strengthening the generalizability of the study.

Turney and Littman (2003) introduced a method that infers the semantic orientation of words based on their statistical associations with other positive and negative words. Two co-occurrence based measures, pointwise mutual information and latent semantic analysis, were employed in the study. The evaluation showed an accuracy of 82% and it could be raised up to 95% when applied to mild words. The proposed approach can be applied to a variety of parts of speech such as adjectives, adverbs, nouns, and verbs. The limitation of the method is that it requires a large size of corpora to calculate significant statistical associations among words to achieve better performance.

Riloff, Wiebe, and Wilson (2003) proposed a system that classifies subjective and object sentences. The system uses extraction patterns to learn subjective nouns using bootstrapping algorithms. Then, a Naïve Bayes classifier was trained using the identified subjective nouns and other features such as discourse features and subjective clues. Results showed that the system achieved 77% recall and 81% precision. A contribution of the work is the use of bootstrapping methods for the identification of extraction patterns.

Kim and Hovy (2004) proposed a system that determines the sentiment of opinions by examining people who hold opinions toward a given a topic. The system has modules for determining sentiments of words as well as sentences that contain those words. To achieve this, they started with selecting less than 100 seed words (both positive and negative) manually and expanded the list by extracting additional words from WordNet, and finally obtained 20,000 words with the polarity information. They proposed three models with different ways of combining sentiments of individual words: a model that only considers polarities, a model that uses the harmonic mean of sentiment strengths, and a model that uses the geometric mean of sentiment strengths. Evaluation showed that the highest score the system achieved was 81% accuracy.

Liu, Hu, and Cheng (2005) proposed a system that analyses online customers' reviews toward products. The system supports a feature-by-feature comparison of customer opinions. A language pattern mining-based approach was proposed for the extraction of product features. This system provides a visual comparison of consumer opinions on competing products. A contribution of this work is the proposal of a supervised pattern discovery method that supports the automated identification of products features from Pros and Cons.

Pang and Lee (2005) proposed a meta-algorithm for rating-inference problem that take multi-point scale of rating into account. Instead of classifying a review as either positive or negative, they viewed the problem as a multi-category classification. They compared the performance of three approaches: multiclass SVM (OVA), regression, and metric labelling on the task. While there was no single approach that performed best in every cases, each approached showed its own strength in different tasks.

Thelwall and colleagues (2010) proposed a novel algorithm to extract sentiment strength from informal English text. They employed a machine learning approach to optimize sentiment term weightings. The authors extracted sentiment information from texts of nonstandard spelling. An evaluation was performed on MySpace comments and results showed that the proposed method achieved 60% and 72% accuracy in predicting positive and negative emotion respectively.

The literature discussed here has examined how to define and quantify the semantic orientation of texts, using lexicon- and/or machine learning-based approaches. Lexicon-based approaches examine words or phrases in a document and use metrics such as mutual information to judge their semantic orientation. Machine learning-based techniques employ widely accepted algorithms such as

Naive Bayes and support vector machines. These techniques train classifiers with labelled sentiment. On top of this, classification was among the most frequently employed techniques. In general, lexicon-based techniques show lower performance but higher applicability, *i.e.* can be applied to many domains, while machine learning-based techniques outperform, but are only applicable to limited domains. These approaches have been applied to a variety of units of analysis, ranging from individual words to complete reviews that are comprised of multiple sentences.

To sum up, most of the landmark papers focused on the computation and grouping of sentiment orientation. Even though often context-specific, they provide practical implications towards the understanding and applications of sentiment analysis. These findings are also consistent with the findings in previous sections.

Discussion

Research question #1: what are epistemological characteristics of opinion mining research?

The investigation at the domain level reveals the following epistemological characteristics of opinion mining research. First, most of the studies came from two domain groups, *psychology, education, health* and *mathematics, systems, mathematical*, and each of these disciplinary groups cites research within similar fields, *psychology, education, social* and *systems, computing, computer*, respectively. The literature from other domains such as *physics, materials, chemistry, molecular, biology, immunology*, and *medicine, medical, clinical* also contribute to the domain-level citation trends. Second, we also identified the interdisciplinary nature of opinion mining: the two groups of domains frequently cross-reference to each other. Next, the advent and growth of opinion mining was mainly led by a few subject categories. Interdisciplinary applications of computer science and psychology were the driving forces that have had the largest influence on the emergence, development, and diffusion of the ideas in opinion mining. Moreover, opinion mining is relatively new and still developing. There are a few emerging research themes identified from keyword co-occurrence and document co-citation network. Finally, opinion mining has received large attention from multidisciplinary domains recently.

Research question #2: which thematic patterns of research do occur in the domain?

First, the subject category assignment shows that computer science, engineering, linguistics, and social sciences such as library and information science, business, economics, management, and psychology are among the leading categories in opinion mining research. In particular, computer science, psychology, electrical engineering, artificial intelligence, linguistics, and library and information have transferred important knowledge to the domain. Second, the exploration of keywords adds a richer interpretation: the identification and extraction of subjective information in source texts and the applications of sentiment analysis to a variety of textual materials have been regarded as one of the most important themes of the domain. The detection of emotion from source texts was the concept located on the shortest path connecting pairs of other concepts in opinion mining research. *Machine learning, text mining, data mining*, and *natural language processing* have been employed to this end. *Classification* is among the most frequently investigated techniques, also being a goal of understanding the sentiment orientation of text. The automatic collection of a *corpus* from a variety of textual genres such as *web, blog*, and *microblogging* has received large

attention for opinion mining purposes. Finally, the examination of the landmark articles reveals that investigating algorithmic and linguistic aspects of sentiment analysis has been of the community's greatest interest: they have put stress on understanding, quantifying, and applying the sentimental orientation of texts.

Research question #3: what are emerging trends and recent developments of the field?

Recently, opinion mining has received large attention from many multidisciplinary domains and recent literature has explored the practical applications of opinion mining to marketing and social networking platforms. Bursting keywords such as *word of mouth*, *social media*, *social network*, and *twitter* evidence this trend. The analysis of document co-citation network identifies that emerging clusters of research include 1) understanding consumer attitudes for effective online marketing and prediction of a product's market value, 2) developing a system for extracting affective information from text, 3) investigating newspaper readers' opinions and information sharing behaviour on social media. In recent years, landmark manuscripts also have explored the improvement and applications of opinion mining to a variety of domains such as informal English text.

Conclusion and future work

In the present study, we aimed to explore the intellectual structure of opinion mining research in a systematic and comprehensive way. Toward that end, we investigated domain-level citation trends, assignment of subject categories, keyword co-occurrence, document co-citation network, and landmark articles. Based on the findings from the study, we articulate that the field of opinion mining is still emerging. For example, a few domains have participated in publication while they mainly cite similar kinds of research. They less frequently cross-cite each other. Next, the advent, development, and diffusion of opinion mining have been largely led by a few domains such as computer science, engineering, linguistics, psychology, and library and information science. These domains have explored algorithmic and linguistic aspects of sentiment analysis to understand, quantify, and apply the sentiment orientation of texts. Recently, multidisciplinary domains have participated in studying opinion mining and this body of literature has explored the practical applications of opinion mining such as to marketing and analysing social networking platforms.

The approaches of the present study provide advantages in investigating intellectual structure of a science as follows. First, we systematically triangulated data collection. Conventional studies of domain analysis often cover only a fraction of published literature. The combination of topic search, citation indexing, and patent search in our method provides a systematic way to broaden the coverage of a knowledge domain, thus provides a much broader and more integrated context. Second, we investigated the domain from a multi-faceted point of view. Bursting keywords, document co-citation networks, emerging thematic patterns, and citation trends were identified in this study. In particular, we employed the concept of burstiness for measuring the importance of units of analysis. An entity's frequency of occurrence is a significant indicator reflecting its impact. However, we cannot identify its influence or density of that impact during a particular span of time if only considering the cumulative measure. Instead, we argue that using the concept of citation and keyword bursts is much more convenient in understanding the advent and decline of scientific trends. Finally, the analytical procedure and tool used in the present study enabled us to explore time-aware research trends in the domain. The scientometric approach and tool employed in this study can support comprehensive data collection and scalable

analytics. In addition, one can conduct domain analysis of his or her concern as frequently as needed without prior knowledge. Thus, the proposed approaches have advantages such as having a relatively higher reproducibility and lower cost for conducting studies at a larger scale, especially as the number of publications in sciences is fast-growing.

There are still several challenges in our study despite the fact that we have a variety of methodological advantages described above. First, the topic search along with citation indexing still may not be able to capture some relevant records. Generally, the vocabulary mismatch is said to present a challenge for keyword-based search ([Deerwester, Dumais, Furnas, Landauer, and Harshman, 1990](#)). Second, the Web of Science as our source of data may have underrepresented conference proceedings, which is also reported to be an issue for disciplines such as social sciences and arts and humanities ([Mongeon and Paul-Hus, 2016](#)). In addition, at the time of data collection, we could only collect records from the core collection of the Web of Science due to the institutional subscription. It might be possible that the number of records collected in each organization from the core collection of Web of Science varies due to different subscription policies. Some relevant records might have been omitted from our data sets accordingly. Additional sources such as Scopus are recommended for future refinements of this type of analysis. Unfortunately, at the time of the study, we did not have access to Scopus. Third, we selected g highly cited references for generating the intellectual landscapes. In spite of this metric's authority, we might be too strict in extracting an important proportion of records. It may be worth conducting a separate study of the theoretical implications of using a variety of conceivable selection criteria. We also plan to apply this scientometric approach to much more comprehensive records that cover a various type of publication materials.

About the authors

Yongjun Zhu is a doctoral candidate in Information Science at Drexel University, Philadelphia, PA. His research interests include information retrieval, network science, and text mining. He can be contacted at zhu@drexel.edu.

Meen Chul Kim is a doctoral candidate in Information Science at Drexel University, Philadelphia, PA. His research interests include learning analytics and data science. He can be contacted at meenchul.kim@drexel.edu.

Chaomei Chen is a Professor at the College of Computing and Informatics at Drexel University, Philadelphia, PA. His research interests include information visualization, visual analytics, knowledge domain visualization, network analysis and modelling, scientific discovery, science mapping, scientometrics, citation analysis, and human-computer interaction. He can be contacted at chaomei.chen@drexel.edu.

References

- Abbasi, A., Chen, H. & Salem, A. (2008). Sentiment analysis in multiple languages: feature selection for opinion classification in Web forums. *ACM Transactions on Information Systems*, 26(3), 12.
- Blondel, V.D., Guillaume, J.-L., Lambiotte, R. & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 10, P10008.
- Bollen, J., Mao, H. & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1), 1-8.
- Brandes, U. (2001). A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2), 163-177.
- Cambria, E., Schuller, B., Xia, Y. & Havasi C. (2013). New avenues in opinion mining and sentiment analysis. *IEEE Intelligent Systems*, 28(2), 15-21.

- Chen, C. (2006). CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. *Journal of the American Society for Information Science and Technology*, 57(3), 359-377.
- Chen, C. (2011). *Turning points: The nature of creativity*. Beijing: Higher Education Press.
- Chen, C., Dubin, R. & Kim, M.C. (2014a). Emerging trends and new developments in regenerative medicine: A scientometric update (2000-2014). *Expert Opinion on Biological Therapy*, 14(9), 1295-1317.
- Chen, C., Dubin, R. & Kim, M.C. (2014b). Orphan drugs and rare diseases: A scientometric review (2000-2014). *Expert Opinion on Orphan Drugs*, 2(7), 709-724.
- Chen, C., Chen, Y., Horowitz, M., Hou, H., Liu, Z. & Pellegrino, D. (2009). Towards an explanatory and computational theory of scientific discovery. *Journal of Informetrics*, 3(3), 191-209.
- Chen, C., Hu, Z., Liu, S. & Tseng, H. (2012). Emerging trends in regenerative medicine: a scientometric analysis in CiteSpace. *Expert Opinions on Biological Therapy*, 12(5), 593-608.
- Chen, C., Ibekwe-SanJuan, F. & Hou, J. (2010). The structure and dynamics of co-citation clusters: a multiple-perspective co-citation analysis. *Journal of the American Society for Information Science and Technology*, 61(7), 1386-1409.
- Chen, C. & Leydesdorff, L. (2014). Patterns of connections and movements in dual-map overlays: A new method of publication portfolio analysis. *Journal of the American Society for Information Science and Technology*, 65(2), 334-351.
- Chevalier, J.A. & Mayzlin, D. (2006). The effect of word of mouth on sales: online book reviews. *Journal of marketing research*, 43(3), 345-354.
- Dave, K., Lawrence, S. & Pennock, D.M. (2003). Mining the peanut gallery: opinion extraction and semantic classification of product reviews. *Proceedings of the 12th International Conference on World Wide Web*, 519-528.
- Deerwester, S., Dumais, S.T., Furnas, G.W., Landauer, T.K. & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6), 391-407.
- Ding, X., Liu, B. & Yu, P. (2008). A holistic lexicon-based approach to opinion mining. *Proceedings of the 2008 International Conference on Web Search and Data Mining*, 231-240.
- Dunning, T. (1993). Accurate methods for the statistics of surprise and coincidence. *Computational Linguistics*, 19(1), 61-74.
- Egghe, L. (2006). Theory and practise of the g-index. *Scientometrics*, 69(1), 131-152.
- Feldman, R. (2013). Techniques and applications for sentiment analysis. *Communications of the ACM*, 56(4), 82-89.
- Garfield, E. (1979). *Citation indexing: its theory and applications in science, technology, and humanities*. New York, NY: Wiley.
- Govindarajan, M. & Romina, M. (2013). A Survey of Classification Methods and Applications for Sentiment Analysis. *The International Journal of Engineering and Science*, 2(12), 11-15.
- Jindal, N. & Liu, B. (2008). Opinion spam and analysis. *Proceedings of the 2008 International Conference on Web Search and Data Mining*, 219-230.
- Kanayama, H. & Nasukawa, T. (2006). Fully automatic lexicon expansion for domain-oriented sentiment analysis. *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, 355-363.
- Kim, M.C. & Chen, C. (2015). A scientometric review of emerging trends and new developments in recommendation systems. *Scientometrics*, 104(1), 239-263.
- Kim, M.C., Zhu, Y. & Chen, C. (2016). How are they different? A quantitative domain comparison of information visualization and data visualization

- (2000-2014). *Scientometrics*, 107(1), 123-165.
- Kim, S.M. & Hovy, E. (2004). Determining the sentiment of opinions. *Proceedings of the 20th International Conference on Computational Linguistics*, 1367.
- Kleinberg, J. (2002). Bursty and hierarchical structure in streams. *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 91-101.
- Liu, B. (2010). Sentiment Analysis and Subjectivity. *Handbook of Natural Language Processing*, 2, 627-666.
- Liu, B. (2012). *Sentiment Analysis and Opinion Mining: Synthesis Lectures on Human Language Technologies*. San Rafael, CA: Morgan & Claypool Publishers.
- Liu, B., Hu, M. & Cheng, J. (2005). Opinion observer: Analyzing and comparing opinions on the web. *Proceedings of the 14th International Conference on World Wide Web*, 342-351.
- Medhat, W., Hassan, A. & Korashy, H. (2014). Sentiment analysis algorithms and applications: a survey. *Ain Shams Engineering Journal*, 5(4), 1093-1113.
- Mongeon, P. & Paul-Hus, A. (2016). The journal coverage of web of science and Scopus: a comparative analysis. *Scientometrics*, 106(1), 213-228.
- Pang, B. & Lee, L. (2004). A sentimental education: sentiment analysis using subjectivity summarization based on minimum cuts. *Proceedings of the 42th Annual Meeting on Association for Computational Linguistics*, 271.
- Pang, B. & Lee, L. (2005). Seeing stars: exploiting class relationships for sentiment categorization with respect to rating scales. *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, 115-124.
- Pang, B. & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2), 1-135.
- Pang, B., Lee, L. & Vaithyanathan, S. (2002). Thumbs up? Sentiment classification using machine learning techniques. *Proceedings of the 7th ACL-02 Conference on Empirical Methods in Natural Language Processing*, 10, 79-86.
- Qiu, G., Liu, B., Bu, J. & Chen, C. (2011). Opinion word expansion and target extraction through double propagation. *Computational Linguistics*, 37(1), 9-27.
- Riloff, E., Wiebe, J. & Wilson, T. (2003). Learning subjective nouns using extraction pattern bootstrapping. *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003*, 4, 25-32.
- Schvaneveldt, R.W., Durso, F.T. & Dearholt, D.W. (1989). Network Structures in Proximity Data. In G. Bower (Ed.) *The psychology of learning and motivation: Advances in research and theory*, 24, 249-284. New York, NY: Academic Press.
- Taboada, M., Brooke, J., Tofiloski, M., Voll, K. & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2), 267-307.
- Tang, H., Tan, S. & Cheng, X. (2009). A survey on sentiment detection of reviews. *Expert Systems with Applications*, 36(7), 10760-10773.
- Thelwall, M., Buckley, K., Paltoglou, G., Cai, D. & Kappas, A. (2010). Sentiment strength detection in short informal text. *Journal of the Association for Information Science and Technology*, 61(12), 2544-2558.
- Thet, T.T., Na, J.-C. & Khoo, C.S.G. (2010). Aspect-based sentiment analysis of movie reviews on discussion boards. *Journal of Information Science*, 36(6), 823-848.
- Turney, P.D. (2002). Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, 417-424.
- Turney, P.D. & Littman, M. L. (2003). Measuring praise and criticism: inference of semantic orientation from association. *ACM Transactions on*

Information Systems, 21(4), 315-346.

- Waltman, L. & van Eck, N.J. (2013). A smart local moving algorithm for large-scale modularity-based community detection. *European Physical Journal B*, 86(11), 471.
- Wiebe, J., Wilson, T., Bruce, R., Bell, M. & Martin M. (2004). Learning subjective language. *Computational Linguistics*, 30(3), 277-308.
- Wilson, T., Wiebe, J. & Hoffmann, P. (2005). Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language*, 347-354.
- Wilson, T., Wiebe, J. & Hoffmann, P. (2009). Recognizing Contextual Polarity: an Exploration of Features for Phrase-Level Sentiment Analysis. *Computational Linguistics*, 35(3), 399-433.
- Yu, H. & Hatzivassiloglou, V. (2003). Towards answering opinion questions: separating facts from opinions and identifying the polarity of opinion sentences. *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, 129-136.

How to cite this paper

Zhu, Y., Kim, M.C. & Chen, C. (2017). An investigation of the intellectual structure of opinion mining research *Information Research*, 22(1), paper 739. Retrieved from <http://InformationR.net/ir/22-1/paper739.html> (Archived by WebCite® at <http://www.webcitation.org/6oGbzTjxJ>)

Find other papers on this subject

Check for citations, [using Google Scholar](#)

Facebook

Twitter

LinkedIn

Delicious

More

© the authors, 2017.

104 Last updated: 10 February, 2017

[Contents](#) | [Author index](#) | [Subject index](#) | [Search](#) | [Home](#)
