

# Microblogs Data Management: A Survey

Amr Magdy<sup>1</sup> · Laila Abdelhafeez<sup>\*1</sup> · Yunfan Kang<sup>\*1</sup> · Eric Ong<sup>1</sup> · Mohamed F. Mokbel<sup>2</sup>

the date of receipt and acceptance should be inserted later

**Abstract** Microblogs data is the micro-length user-generated data that is posted on the web, e.g., tweets, online reviews, comments on news and social media. It has gained considerable attention in recent years due to its widespread popularity, rich content, and value in several societal applications. Nowadays, microblogs applications span a wide spectrum of interests including targeted advertising, market reports, news delivery, political campaigns, rescue services, and public health. Consequently, major research efforts have been spent to manage, analyze, and visualize microblogs to support different applications. This paper gives a comprehensive review of major research and system work in microblogs data management. The paper reviews core components that enable large-scale querying and indexing for microblogs data. A dedicated part gives particular focus for discussing system-level issues and on-going effort on supporting microblogs through the rising wave of big data systems. In addition, we review the major research topics that exploit these core data management components to provide innovative and effective analysis and visualization for microblogs, such as event detection, recommendations, automatic geotagging, and user queries. Throughout the different parts, we highlight the challenges, innovations, and future opportunities in microblogs data research.

## 1 Introduction

Microblogs data, the micro-length user-generated data that is posted on the web, such as tweets, online reviews, news comments, social media comments, and user check-ins, has become very popular in recent years. As micro-length data, it is easy and quick for users to generate plenty of them every day. In fact, every day, over one billion users post more than four billions microblogs [106, 332] on Facebook and Twitter. Such tremendous amounts of user-generated data have rich content, e.g., news, updates on on-going events, reviews, location information, language information, user information, discussions in politics, products, and many others. This richness has motivated researchers and developers worldwide to take advantage of microblogs to support a wide variety of practical applications [226, 249], including public health [270, 327], disaster response [69, 70, 105, 153, 154, 160, 172], public safety [324], education [355], real-time news delivery [49], geo-targeted advertising [256], and several disciplines of academic research such as social science [331], information modeling [156], human dynamics [307], engagement in education [329], political sciences [330], behavioral sciences [306], and even medical-related research [328]. The distinguished nature of microblogs data that combines large data sizes, high velocity, and short noisy text, has introduced new challenges, which motivated researchers to develop numerous novel techniques to support microblogs data management, analysis, and visualization at scale.

This paper provides a comprehensive review for existing major techniques and systems for microblogs data management since the inception of Twitter in 2006. The literature on microblogs is rich and includes several major research communities, e.g., data management, natural language processing, and information retrieval. However, this survey paper is addressed to the data management community that pro-

\* Authors have equal contributions and are ordered alphabetically.

<sup>1</sup>Department of Computer Science and Engineering, University of California, Riverside  
Riverside, CA, USA

E-mail: amr@cs.ucr.edu, labde005@ucr.edu, ykang040@ucr.edu, eong001@ucr.edu

<sup>2</sup>Department of Computer Science and Engineering, University of Minnesota, Twin Cities

Minneapolis, MN, USA

E-mail: mokbel@umn.edu

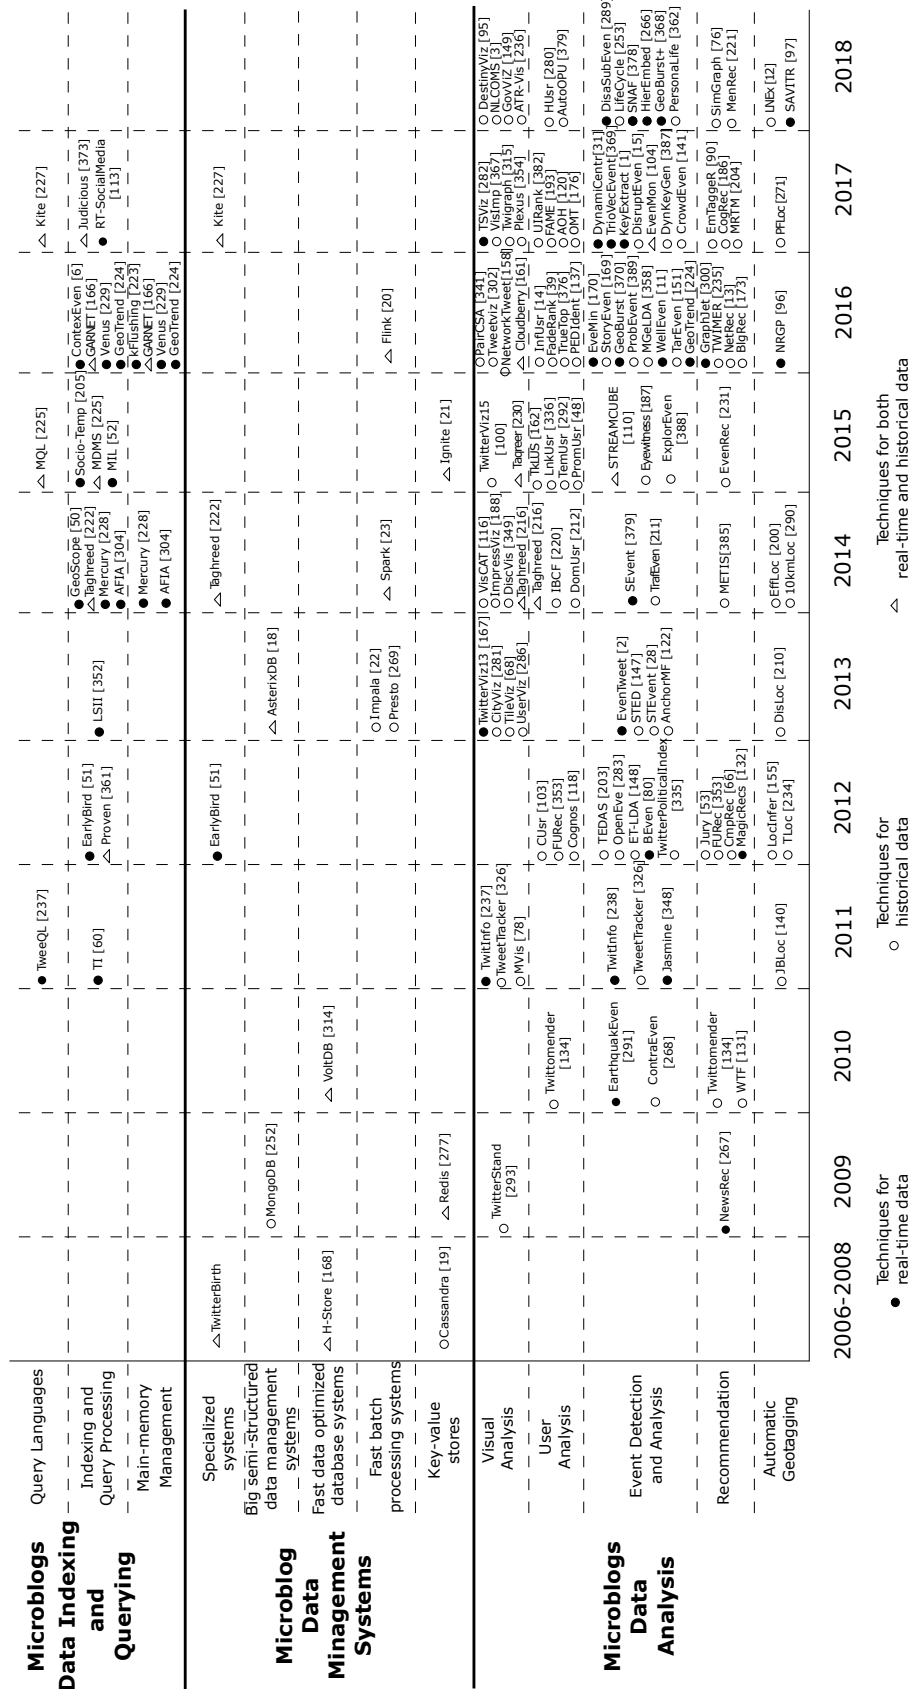


Fig. 1 Microblogs literature timeline

vides scalable infrastructures for indexing and querying microblogs and incorporate them in data management systems to enable managing this data at scale. The paper includes three main parts. The first part reviews core indexing and query processing components of microblogs data management, including their query languages and associated main-memory management techniques. The second part focuses on major genres of data management systems that are either designed for microblogs data or equipped with infrastructures to manage fast and large data, which are distinguishing characteristics for microblogs. The third part highlights major research topics that exploit data management infrastructures to build applications and analysis modules on top of microblogs, such as visual analysis, user analysis queries, and event detection. This part does not include other major research directions, e.g., natural language processing and information retrieval, as they are orthogonal to the data management research and out of the scope of this paper. In fact, dedicated survey papers review parts of their literature [82, 119].

Figure 1 depicts a summary of different parts and the research topics that will be covered in this survey paper in a timeline format. The horizontal axis in Figure 1 represents the year of publication or system release for each technique/system, while the vertical axis represents the research topic. The techniques are then classified into three categories: (1) techniques that deal with real-time data, i.e., very recent data, depicted by a filled black circle, (2) techniques that deal with historical data, depicted by a blank circle, and (3) techniques that deal with both real-time and historical data, depicted by a blank triangle. As the vertical axis of Figure 1 depicts, the paper is organized around three main parts: *indexing and querying*, *systems*, and *data analysis*, each part is outlined below:

**(1) Data Indexing and Querying:** this part covers existing work for indexing and querying microblogs data that is depicted in the first to third rows of Figure 1 and includes the following three topics:

- **Query languages:** this work provides generic query languages that support SQL-like queries on top of microblogs. This facilitates basic operators and advanced functions to express a variety of queries on microblogs.
  - **Indexing and query processing:** this work includes various indexing and their associated query processing techniques that have been proposed to index incoming microblogs either in main-memory [50, 51, 222, 228, 304, 352, 361] or in disk [60, 222]. This includes keyword search based on temporal ranking [51, 60], single-attribute search based on generic ranking functions [352], spatial-aware search that exploits location information in microblogs [228], personalized social-aware search that exploits the social graph and produces user-specific search results [205], and aggregate queries [50, 224, 304] that find trending keywords and correlated location-topic pairs instead of individual microblog items.
  - **Main-memory management:** this work includes techniques that optimize for main-memory consumption and utilization. Most microblogs indexing techniques depend on main-memory to manage microblogs in real time. Thus, some techniques are equipped for main-memory management such that memory resources are efficiently utilized, either for aggregate queries [224] or basic search queries that retrieve individual data items [223, 228].
- (2) Data Management Systems:** this part highlights the current state and the challenges of managing microblogs data through major types of big data systems [17, 20, 23, 25, 51, 222, 245, 314], depicted in the fourth to eighth rows of Figure 1. In specific, we give a briefing on system challenges and motivational case studies to provide system-level data management for microblogs. Then, we highlight the data management features that are related to managing microblogs in the following system genres:
- **Specialized systems:** such as Twitter EARLYBIRD [51, 245], TAGHREED [222], and KITE [227] that are designed considering the distinguishing characteristics of microblogs data and queries.
  - **Big semi-structured data management systems:** such as ASTERIXDB [17] that is a generic big data management system to support various data sources. Recently, ASTERIXDB has extended its components to support fast data [123], e.g., microblogs, natively in the system. We review the fast data support in ASTERIXDB, which shows the current challenges of persisting fast data.
  - **Fast data optimized database systems:** such as VOLTDB [314] that is mainly optimized for database transactions on fast data, e.g., microblogs. We review the challenges of supporting transactional applications on fast data and solutions at the system level.
  - **Fast batch processing systems:** such as Apache SPARK [23] and Apache FLINK [20] that are optimized to process high-throughput applications on fast data via batch processing models. We discuss viable use cases as well as challenges and limitations of such systems to support efficient management for different microblogs applications.
  - **Key-value stores:** such as Apache CASSANDRA [19] and REDIS [277] that store big datasets in key-value pairs. We discuss the adequacy of such systems to support certain microblogs applications as well as their limitations to support other applications.

- **Hybrid system architectures:** such as gluing stream processing engine, e.g., Apache STORM [25], with a persistent data store, e.g., MONGODB [252]. We discuss the challenges to manage real-time data in such setting showing the need to consider data velocity inherently in different system components.

**(3) Data Analysis:** this part covers the *major* types of analysis on microblogs data that are depicted in the ninth to thirteenth rows of Figure 1. The selected types of analysis are the ones that exploit the data management infrastructure to pose queries of massive number of microblogs and popular in the research community. This does not include either ad-hoc non-research applications, such as web applications that exploit microblogs data, or orthogonal research directions, such as linguistic analysis or information retrieval, which are intractable and dedicated surveys review only portions of them [82, 119]. This part includes the following five types of analysis:

- **Visual analysis:** this work covers existing microblogs data visualization techniques that make use of the underlying scalable queries to enable visual analysis for excessive number of microblog records. This work use both aggregate queries, for aggregation-based visualization [95, 116, 282, 315], non-aggregate queries for sampling-based visualization [222, 293], or a combination of both [161, 236, 326].
- **User analysis:** this work is mainly interested in querying user information for different purposes, such as identifying top influential users in certain regions or topics [162, 222, 336] or discovering users with similar interests [33, 134]. Such users, or group of users, can be used in several scenarios, including posting ads and enhancing their social graph.
- **Event detection and analysis:** this work exploits the fact that microblogs users post many updates on ongoing events. Such updates are queried, grouped, and analyzed to discover events in real time [2, 291] or analyze long term events [238, 326], e.g., revolutions.
- **Recommendation:** this work exploits microblogs user-generated content as means for catching user preferences to support diverse recommendation tasks, such as recommending content to follow [13], real-time news to read [267], authority users to follow [53], products [385], or users who share similar interests [134].
- **Automatic geotagging:** this work tries to attach geo-locations to microblogs data that are not geotagged based on analyzing their different attributes. This is mainly motivated by the small percentage of geotagged microblogs, e.g., less than 4% of tweets, that is faced by the need of many location-aware applications on top of microblogs, e.g., [2, 228, 256, 293].

Other sporadic analysis tasks are addressed on microblogs data in both research community, e.g., news extraction [267, 293], and topic extraction [142, 275], and industrial community, e.g., geo-targeted advertising [256] and generic social media analysis [323, 383]. However, we outline the major analysis that exploit the data management infrastructure and include a wide variety of research techniques, which is of interest for the data management research community.

The rest of this paper details each of the three parts highlighting existing challenges, innovations, and future opportunities in microblogs data management research. Section 2 gives details of the data indexing and querying part. Section 3 gives details of the data management systems part. Section 4 gives details of the data analysis part. Finally, Section 5 concludes the paper and discusses different open problems in microblogs research.

## 2 Microblogs Data Indexing and Querying

This section gives a comprehensive review for data management techniques that support large-scale indexing and querying for microblogs data. We first introduce microblogs query languages, in Section 2.1, that enable high-level declarative interfaces for querying microblogs. Then, Section 2.2 reviews the core indexing and query processing techniques. Finally, Section 2.3 outlines main-memory management techniques that are used in association with in-memory index structures.

### 2.1 Query Languages

There are few attempts in the literature to standardize query languages tailored for the needs of microblogs, and inspired by SQL query language: TWEEQL [237] and MQL [225, 227], each outlined below.

TWEEQL [237] is a wrapper over Twitter APIs<sup>1</sup> so the user can post SQL-like queries on top of Twitter data and the underlying query processing is performed through accessing her Twitter developer account. TWEEQL supports *Select-Project-Join-Aggregate* queries, recognizing aggregation as a major part of querying microblogs in several applications, e.g., trend discovery. In addition, TWEEQL allows two additional constructs. First, built-in filters for the three major microblog attributes: keywords, spatial, and temporal attributes. Second, user-defined functions that allows higher-level analysis of tweets, such as automatic geotagging and sentiment analysis.

Unlike TWEEQL, MQL [225, 227], stands for *Microblogs Query Language*, is proposed as an inherent part

<sup>1</sup> <https://developer.twitter.com/en/docs/api-reference-index.html>

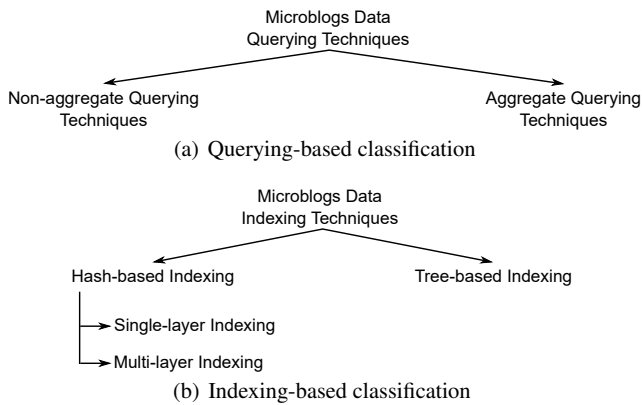


Fig. 2 An overview of microblogs data management literature.

of data management systems that support microblogs. MQL allows *Select-Project-Join-Count* queries, focusing on *count* as the only useful aggregate measure on microblogs. The major distinction of MQL is promoting *top-k* and *temporal* aspects as mandatory in all queries, arguing that there is no practical microblog query that can avoid these two aspects. Even if the user does not explicitly provide a top-k ranking function and temporal horizon for the query, MQL beefs up the query with default values. In addition, MQL supports filtering data based on arbitrary attributes, including spatial boundaries and keywords, and continuous queries similar to traditional data streams [58, 127, 250, 343, 357, 359, 392].

## 2.2 Indexing and Query Processing

This section reviews indexing and query processing techniques that are proposed to support large-scale querying on microblogs. Figure 2 depicts a high-level overview of this literature classified based on both query type (Figure 2(a)) and index type (Figure 2(b)). Based on query type, existing techniques are classified into *non-aggregate querying techniques* (detailed in Section 2.2.1) and *aggregate querying techniques* (detailed in Section 2.2.2). Based on index type, microblogs are indexed using either *tree-based indexing* or *hash-based indexing* that could employ a single or multiple layers of hash-based indexes. Table 1 provides more details summarizing these techniques in terms of the query attribute(s), index structure, index cell content order, and top-k ranking function. As the table shows, all existing queries on microblog include both temporal and top-k aspects regardless their other details. This is attributed to the nature of microblogs as they come in large numbers around the clock. This large number mandates retrieving the most useful k microblogs based on a certain top-k ranking function, otherwise, many useless data will be reported. In addition, being a kind of streaming data, the data is real-time by nature and many users and applications are interested in recent

microblogs. This inspired almost all the existing techniques to embed the time aspect by default in the query signature, unless it is disabled by the user. In fact, without using the time aspect, a query might retrieve data from several years ago, which leads to a significant querying overhead. So, by disabling this default option, users become aware of the implications on querying performance if they consider data of long temporal periods.

A generic query signature that represents all queries in Table 1 is: “*Find top-k microblogs/keywords ranked based on a ranking function F.*” In non-aggregate queries that retrieve individual microblogs, the ranking function  $F$  can be *temporal* [6, 51, 60], *spatio-temporal* [228, 229], *significance-temporal* [352], or *socio-temporal* [110] as shown in Table 1. In aggregate queries [50, 166, 224, 304], the temporal aspect is used as a filter for queried data and the ranking functions mostly depend on keyword counts and their derived measure, e.g., trendline slope, except GEO-SCOPE that employs a correlation measure.

Almost all indexing techniques of microblogs are optimized for high digestion rates in a main-memory index for real-time data indexing, and secondary storage indexing is assumed to have older data to query historical microblogs. The only exception is TI [60] that primarily uses a disk-based index. In addition, the query processing techniques are optimized for top-k and temporal queries. The rest of this section briefly outlines each technique that is shown in Table 1, for both non-aggregate and aggregate querying.

### 2.2.1 Non-aggregate Indexing and Querying

This section reviews non-aggregate querying techniques that “*Find top-k microblogs ranked based on a ranking function F.*”, and retrieve individual microblog records.

TI [60] employs a disk-based inverted index structure where microblogs are sorted based on their timestamp. The main idea in TI is to defer indexing unimportant microblogs to reduce the number of microblogs that are indexed immediately and cope up with the large number of incoming data records. So, it keeps in memory a set of *recent and popular queries and topics*. Then, it categorizes each incoming microblog and decides whether it should be indexed immediately or deferred. The categorization considers the microblog recency, the user’s page rank, popularity of the topic, and the textual relevance. The unindexed microblogs are written into log file and an offline batch indexing is performed periodically to reduce real-time indexing latency. This is the first work to consider temporality in optimizing for microblogs, but it uses disk-based index solution which cannot scale for high microblogs arrival rate. The following techniques use in-memory structures that can digest fast data rates as well as providing low query latency. Eventhough, TI achieves higher indexing throughput compared to tradi-

**Table 1** Summary of indexing and top-k querying of microblogs.

	Query Attribute(s)	Index Structure	Cell Content Order	Top-k Ranking Function
TI [60]	Keyword, Temporal	Inverted index	Temporal	Temporal recency
EARLYBIRD [51]	Keyword, Temporal	Inverted index	Temporal	Temporal recency
CONTEXEVEN [6]	Keyword, Temporal	Inverted index	Temporal	Popularity and temporal recency
MIL [52]	Keyword, Temporal	Multi-layer inverted index	Temporal	Keyword similarity
MERCURY [228], VENUS [229]	Spatial, Temporal	Partial quad-tree	Temporal	Spatial proximity and temporal recency
LSII [352]	Arbitrary, Temporal	Log-structured inverted index	Temporal	Significance, keyword similarity, and temporal recency
SOCIO-TEMP [205]	Social, Keyword, Temporal	3D inverted index	Temporal	Social relevance, keyword similarity, and temporal recency
PROVEN [361]	Keyword, Temporal	Inverted index	Temporal	Provenance similarity
RT-SOCIALMEDIA [113]	Keyword, Spatial, Temporal, Social	Inverted index	Temporal	BM25, Spatial proximity, Temporal recency, Social distance
JUDICIOUS [373]	Keyword	Inverted index	TF-IDF	Keyword similarity

**Top-k Non-aggregate Queries**

	Query Attribute(s)	Index Structure	Top-k Ranking Function
AFIA [304]	Keyword, Spatial, Temporal	Spatio-temporal grid index	Keyword count
GEO TREND [224]	Keyword, Spatial, Temporal	Partial quad-tree	Keyword count and trendline slope
GARNET [166]	Arbitrary, Keyword, Temporal	Multi-dimensional grid index	Keyword count and trendline slope
GEO SCOPE [50]	Keyword, Spatial, Temporal	Hash index	Correlation between locations and topics

**Top-k Aggregate Queries**

tional techniques. Indexing time ranges from 0.1 seconds to 1 seconds based on index parameter settings, whereas for the traditional index the indexing time is a constant 1.6 seconds. Query processing time ranges from 30 milliseconds to 90 milliseconds as the number of involved microblogs increases with growing answer size value. Query accuracy also increases with minimum of 90% for all settings.

EARLYBIRD [51] — *the core retrieval engine that powers Twitter’s real-time search service* — is a distributed system where each node manages multiple inverted index segments to index keywords in real time. Incoming data first goes to a partitioner that divides tweets over nodes. In each node, ingested tweets first fill up the active segment before proceeding to the next one. Therefore, there is at most one index segment actively being modified, whereas the remaining segments are read-only. Each index segment is a traditional inverted index, however, postings for each term are maintained in reverse chronological order. It is worth mentioning that EARLYBIRD reduces the concurrency management overhead by adopting a single-writer multiple-readers model to eliminate any contention and race conditions. When a query comes, a blender receives it and determines which nodes should be accessed. Then, the query is posted to these nodes, the partial answers are retrieved and compiled by the blender to return the final answer. The

experimental evaluation shows that EARLYBIRD achieves 7,000 tweet/second indexing rate at latency of 180 milliseconds.

CONTEXEVEN [6] also supports keyword search queries on real-time microblogs in favor of finding real-time content of the top-k relevant events. It defines the event context with a set of keywords and organizes incoming data in an inverted index based on these keywords. Each index entry maintains a list of event ids that correspond to a certain keyword ordered by a hybrid score that combines popularity and time recency, while an event is represented with a temporal tree that shows the chronological order of data within the same event [7]. To cope up with high velocity data, each index entry divides its posting list into buckets of exponentially growing sizes to reduce the insertion overhead in real time. In addition, CONTEXEVEN adapts a lazy update strategy for the index that defers updating the event id order until it is moved to the  $(2 \times k)^{th}$  position, sacrificing a slight query accuracy with real-time efficiency. The query processor then iterates over all index entries that correspond to the query keywords and aggregate the event final ranking score from all entries to return the final top-k events.

MIL [52] is another event-based real-time search system that employs multi-layer inverted index that organizes event data based on keywords. The index has  $m$  layers, each

layer maintains a separate inverted index. The index key at the  $i^{th}$  layer is a set of  $i$  keywords that co-occur in certain events, while the posting list stores a list of event ids that correspond to these keywords. So, layer 1 key has a single keyword while layer 2 key has a pair of co-occurring keywords and so on. A new microblog is inserted into all layers that correspond to different combinations of its keywords. Incoming queries also access all index layers to perform a nearest neighbor search based on cosine similarity. Experiments show that MIL outperforms variants of its competitor IL in search time, pruning power, and index update time. MIL search time is below 2 milliseconds with different data sizes and query length, where pruning power is constant and it is almost 1. Index update time is less than 0.1 milliseconds for up to 10 millions records.

Spatio-temporal ranking functions can be depicted in MERCURY [228] and its successor VENUS [229]. MERCURY employs a partial quad-tree structure, where each cell contains a list of microblogs that have arrived within the cell boundary in the last  $T$  time units, ordered chronologically. As traditional data insertions, expiration, and index structuring are very inefficient for real-time data, MERCURY employs bulk data insertion, speculative index cell splitting, piggybacked deletion, and lazy cell merging to significantly reduce the overall indexing overhead and scale for real-time microblogs. The bulk insertion buffers incoming data and insert them every  $t$  seconds, where  $t$  is 1-2 seconds, to navigate different index levels once for several thousands microblogs. In addition, deletion and index structuring operations are piggybacked on the insertion navigation. For the index structuring, cell splitting is performed if and only if the cell exceeds maximum capacity and the microblogs in that cell will span at least two quadrants of the quad-tree node. Cell merging is deferred until at least three out of the four quadrant siblings are empty to reduce redundant splits and merges in real time. The query processing in MERCURY have two phases, namely the initialization phase and the pruning phase. In the first phase, cells lying within the query range are ordered based on a spatio-temporal proximity score, and microblogs are retrieved from these cells based on their score. The pruning phase, tightens the original search boundaries where microblogs outside the new boundaries are early pruned. This significantly reduces the total number of processed microblogs to get the final answer. Experimental results shows that VENUS supports high arrival rates up to 64,000 microblogs/second and average query latency of 4 milliseconds.

LSII [352] supports top-k queries based on combining three ranking scores for a microblog: its significance, its keyword similarity with the query, and its temporal freshness. A microblog is more significant if it is posted, for example, by an authority user or has high popularity with large number of forwards and replies. High keyword similarity in-

dicates a high relevance to the query and freshness measures the temporal recency of the microblog. LSII consists of a sequence of  $m$  inverted indexes where each index  $I_i$  is double in size its predecessor index  $I_{i-1}$ . The first index  $I_0$  is a read-write structure to which new data is appended, and the microblog list of each keyword is ordered chronologically. The indexes from  $I_1$  to  $I_{m-1}$  are read-only indexes and each keyword has three microblog lists sorted with the three ranking scores. The small size and simple organization of  $I_0$  enables high digestion rates of real-time data, while the three sorted lists of subsequent indexes enables efficient query processing. When index  $I_{i-1}$  size reaches a certain threshold, a merge operation with index  $I_i$  is triggered and index  $I_{i-1}$  is flushed. To process an incoming query, LSII first scans  $I_0$  to get the initial set of top-k microblogs, then it proceeds in scanning other indexes. If the upper-bound of index  $I_i$  is no more than the scores of the top-k candidates, index traversal is stopped and proceeds to the next one. Since each index  $I_i$  is less recent than index  $I_{i-1}$ , the search is more likely to get pruned at earlier indexes since they have higher fresh scores. Extensions to LSII includes personalized search, when a user is only interested in microblogs from specific users. Performance of LSII is compared to append only approach and Triple-Posting-List approach. The query processing time for LSII is less for both number of microblogs the query asks for and for the number of queries in the mixed stream of queries and updates. The query time is between 1 and 10 seconds for varying number of microblogs asked by a query and increases from this range linearly with increasing the number of queries. The total processing time is almost 10 seconds and does not vary with changing the weights of the ranking function.

Another type of ranking is considering the social relevance as well as the textual relevance along with microblog freshness. A 3-D inverted index structure is proposed in [205] where each index cell is a three-dimensional data cube, a dimension for term frequency, a dimension for social relevance, and a dimension for time freshness. Each dimension is partitioned into intervals, the social graph is partitioned with  $k$ -way partitioning using minimum cut utility. The time and textual dimensions are sorted at indexing time whereas social dimension is sorted at query time. Data is first partitioned by time, then cubes in each time interval are indexed with a B+ tree to avoid maintaining many empty cubes. New data records are added to the last time interval. When the size of data in latest time interval exceeds a threshold, it is concluded and a new time interval is introduced. For query processing, cubes are first sorted by their estimated total score. Then, the query processor iterates over neighbouring cubes and gets actual scores for microblogs. When the existing top-k records are more relevant than the next unseen cube, the query processing terminates and prunes all remaining data cubes to ensure efficient query

latency. The 3-D index outperforms both time pruning and frequency pruning, the two state-of-the-art techniques, with an average of 4-8x speedups for most of the parameter settings.

PROVEN [361] optimizes keyword search on microblogs for a unique similarity measure that depends on data provenance, measured through microblog content such as hashtags, URLs, and keywords. Incoming microblogs are grouped into bundles based on their provenance similarities and ordered based on their temporal evolution. An inverted index organizes bundles that are continuously updated with incoming microblogs. The inverted index has provenance elements, such as hashtags, URLs, and keywords as index keys and bundles as values. Through this index, incoming queries retrieve whole connected bundles of microblogs, which improves the search result relevance.

RT-SOCIALMEDIA [113] proposes a generic index structure for generic query function that can be extended to support temporal, spatial and/or social aspects. It proposes using the inverted index structure with a space-partitioning strategy in which the documents ids are partitioned into intervals, and each interval partitions documents based on keywords into blocks. To facilitate the top-k retrieval, meta-data is stored within each block. The metadata includes an interval id, a maximum score, and a bit map signature to determine which documents are present in this block. The maximum score is an upper bound for all documents in the block, so if the current top  $k^{th}$  score exceeds this bound, the block is safely pruned. The signature field also provides a tighter bound to fasten the search process as absent documents are not included in the upper bound value. Documents are sorted in the inverted index by the document id, so newer documents are appended to the end of the list to naturally support the temporal aspect. To support the spatial aspect, the index is extended with a uniform grid where in each cell we store the interval ids present in this cell, which helps to demote absent documents that do not appear the query cell. To support social aspect, the index metadata is extended with a friendship bitmap, which helps to determine quickly if a user is a friend of another user. Experiments show that RT-SOCIALMEDIA reports better query latency compared to competitors in keyword search and spatial keyword search, and better query latency in most cases as compared to LSII [352] in temporal keyword search.

JUDICIOUS [373] is the only microblog querying technique that does not consider temporality in their indexing. It offers a compact inverted index structure that treats rare terms, that are not frequently present in the data, differently from common terms, that are frequently present. For rare terms, a traditional inverted index is used. For common terms, a compact inverted index is proposed that uses block partitioning schemes, where microblogs are hashed into intervals, each interval is stored in a block with maximum

score as meta-data for the block to facilitate early pruning in query processing. Thus, whole blocks are pruned if their maximum scores are not within the current query upper bound. Incoming queries has two types, singular queries that ask for one type of terms, either rare terms or common terms, and mixed queries that ask for both rare and common terms. Singular queries are answered from their corresponding index. In mixed queries, the rare item lists are retrieved first and used as fancy lists that tighten the query upper bound score and speed up pruning the search space. Experiments have shown that JUDICIOUS achieves 2-3 times query speedup over the state-of-the-art approaches with much smaller index size. For the same dataset, JUDICIOUS maintains an index of 35GB, whereas competitors BM-OPT and BMW-LB-PB maintain indexes with sizes 49GB and 50GB, respectively. Average response time on TREC queries ranges from 9 to 130 milliseconds with increasing the number of keywords in JUDICIOUS, whereas it ranges from 25 to 290 milliseconds in the other two techniques. With increasing answer size, JUDICIOUS average response time ranges from 21 to 30 milliseconds, whereas the other two techniques ranges from 70 to 110 milliseconds.

### 2.2.2 Aggregate Indexing and Querying

This section reviews aggregate querying techniques that “*Find top-k keywords ranked based on a ranking function F.*” These techniques retrieve keywords, rather than individual microblog records, ranked based on aggregate information, e.g., frequency or frequency growth over time.

AFIA [304] retrieves top-k frequent keywords that lie within any arbitrary spatial range and temporal interval. To support this at scale, AFIA maintains in main-memory a set of spatial grid indexes at different spatial and temporal granularities. Each grid cell keeps track of a summary of top-k keywords that lie within its spatial and temporal ranges, using a modified version of the SpaceSaving algorithm [243]. At query time, the query range is mapped to the corresponding grid cells, summaries from all cells are merged together to get the top-k keywords for the query spatio-temporal range. Despite using the SpaceSaving algorithm the consumes small memory footprint, AFIA is consuming significant memory resources when supporting fine spatial and temporal granularities, as shown in [166, 224], due to maintaining a huge number of summaries without supporting deletions or data expiration.

Unlike AFIA, GEOTREND [224] limits its search scope to recent microblogs and retrieves top-k *trending* keywords that lie within any arbitrary spatial range within the last T time units. GEOTREND accommodates various trending measures including trendline slope, which gauges the keyword frequency growth over time, and keyword frequency. To support this efficiently, GEOTREND maintains a partial



quad-tree structure where each cell contains aggregate information about keywords that arrive within its spatial boundaries. A list of top-k keywords is materialized in each cell at indexing time. At query time, GEOTREND first gets local top-k trending keywords within cells that intersect with the query boundaries. Then, to get the global top-k trending keywords, the global trending value of each keyword is aggregated from local values, using Fagin’s algorithm [107], and final top-k keywords are returned. Experimental evaluation shows that GEOTREND supports arrival rates up to 50K microblogs/second, average query latency of 3 milliseconds, and 90% query accuracy under limited memory resources.

GARNET [166] generalizes trend discovery to any arbitrary user-defined context instead of being limited to the spatial space. In specific, GARNET finds top-k trending keywords within: (a) a  $d$ -dimensional context that is defined on arbitrary  $d$  microblog attributes, and (b) an arbitrary time interval. For example, it could find trending keywords that are posted by teenagers in Spanish during July 2018. In this example, the context is two-dimensional and defined over age and language attributes. Each of the contextual attribute is divided into a set of discrete values or disjoint intervals, e.g., age attribute can be divided to *child*, *teenager*, and *elder*, while the language attribute can be categorized into *English*, *Spanish*, *French*, and *Others*. Then, a  $d$ -dimensional grid index is employed to map incoming data to the corresponding context grid cells. An in-memory grid index is maintained for recent data and in-disk grid index is maintained for historical data. Each in-memory grid cell maintains a list of top-k trending keywords over the last  $T$  time units, while each in-disk grid cell maintains a temporal tree that maintains top-k trending keywords for multiple temporal granularities over extended periods. At query time, top-k keywords are aggregated from corresponding grid cells and a final top-k list is compiled in a similar way to [224]. Experimental evaluation has conducted to show index scalability and query performance with different numbers of grid cells. The comparison with AFIA [304] has shown the superiority of GARNET. GARNET in-memory insertion time is below 400 milliseconds for up to 24,000 microblog/second rate, and reaches up to 1 second for higher rates. For varying grid cells, query latency ranges from 0.1 to 1 millisecond for both frequent and trending queries. The naive scanning alternative is not a competitor and increases query latency up to 1 second.

Unlike all other techniques, GEOSCOPE [50] measures localized trending topics based on correlation between topics and a predefined set of locations, e.g., list of cities. The main idea of GEOSCOPE to discover localized trending topics rather than topics that are popular all over the space. For example, a presidential election campaign is trending in many cities all over the country while a city council election campaign is trending only within a

specific city. To this end, GEOSCOPE limits the number of monitored locations to the  $\theta$ -frequent locations, keeps track of topics that is only  $\phi$ -frequent at least in one location, and then only tracks  $\psi$ -frequent locations of this topic. GEOSCOPE has two main data structures: *Location-StreamSummary-Table* and *Topic-StreamSummary-Table*. *Location-StreamSummary-Table* maintains top frequent topics for each location while *Topic-StreamSummary-Table* maintains top frequent locations for each topic. At query time, these aggregate information are processed to retrieve topics that correlated only to the query location, distinguishing them from topics that are popular in all locations. Experiments show that GEOSCOPE consumes almost constant amount of memory and reports constant amount of time with increasing window size. Also, it reports perfect recall and near-perfect precision.

### 2.3 Main-memory Management

All major indexing techniques of microblogs store data in main-memory to be able to support real-time indexing for fast data and provide low query response time. However, with the rapid increase in number of microblogs, it is infeasible to store all data in main-memory for extended periods. At certain point, the available memory becomes full and part of the memory content has to be moved to a secondary-storage structure to free up memory resources for incoming microblogs. To this end, different indexing techniques use, implicitly or explicitly, flushing policies that decide on which microblogs to flush from main-memory to secondary storage. Although the problem of selecting memory content to evict has been studied before for the buffer management in database systems [99], anti-caching in main-memory databases [87, 197, 375], and load shedding in data stream management systems [32, 114, 138], flushing in microblogs data management is different in terms of the optimization goals and the anticipated real-time overhead as detailed in [223]. This section reviews the major flushing policies that are proposed in the literature to manage main-memory for microblogs data management.

Many of the major microblogs indexing techniques implicitly depend on temporal-based flushing [51, 60, 110, 304, 352], where a chunk of the oldest data is flushed to disk to free up memory resources. The main intuition behind this simple policy is that: (a) recent microblogs are more important than old microblogs in several applications, and (b) incoming data, in these techniques, are indexed and ordered based on temporal recency, so flushing the oldest data will encounter very limited overhead in real time. This intuition is correct in a practical sense and gives the major advantage of the temporal flushing, which is its low overhead in real-time environments so its invocation does not limit the system

scalability. However, it encounters a major limitation that affects both main-memory utilization and query latency. It underutilizes memory resources and stores  $\sim 70\%$  of memory data that is never reported to any incoming query, as detailed in [223]. The main reason is that flushing decisions depend solely on data recency without accounting for what is actually needed for incoming queries. Subsequent techniques in the literature have addressed such limitation for different types of queries as outlined below. The main objective of all these techniques is better utilization for main-memory resources, as useless data are evicted and useful data accumulates in main-memory. This leads to increasing *memory hit ratio*, so more queries are answered from in-memory content without accessing disk content.

MERCURY [228], and its successor VENUS [229], provide flushing policies that decide on evicting non-aggregate data, i.e., individual microblogs. The flushing policy is optimized for top-k spatio-temporal queries that retrieve microblogs from a spatial boundary  $R$ , and temporal interval of the last  $T$  time units. By default, each index cell stores data from the last  $T$  time units. MERCURY flushing policies provide two tighter time bounds,  $T_c$  and  $T_{c,\beta}$ , both of them are no greater than  $T$ , where any data record outside  $T_c$  or  $T_{c,\beta}$  can be flushed to disk. The main observation behind finding such tighter bounds that highly populated areas, e.g., Downtown Chicago, has higher arrival rates than other areas. Then, top-k microblogs can be retrieved from a shorter time than that of areas of less arrival rates. Thus, values of  $T_c$  and  $T_{c,\beta}$  are derived based on the local arrival rate, ranking function, and query parameters.  $T_c$  ensures accurate query answers, which means any data record outside  $T_c$  is not reported to any incoming query. On the contrary,  $T_{c,\beta}$  employs a load shedding parameter  $\beta$ ,  $0 \leq \beta \leq 1$ , that allows to save up to  $100 \times \beta\%$  of the memory with probability  $\beta^3$  to miss a needed data record, trading off a slight decrease in query accuracy with a significant saving in memory resources.  $\beta$  in this case is an input parameter by the system administrator. Experimental results show that compared with the default case where data from the last  $T$  time units are stored, the policy consumes 65% less storage while achieving an accuracy of 98% to 99.5% when  $\beta = 0.3$ . At  $\beta = 0.7$ , 75% less memory are consumed and the accuracy is 97.5% to 99.3%. VENUS [229] extends this to provide an adaptive load shedding technique where the value of  $\beta$  is adaptively calculated and automatically adjusted with the distribution changes in incoming queries and data. This leads to different  $\beta$  value for each region, based on local data and query distributions, rather than a single global value for all regions. The strategy saves up to 80% of the storage while keeping an accuracy of more than 99% and is considered as significant enhancement over MERCURY.

KFLUSHING [223] is another flushing policy for non-aggregate data, . KFLUSHING accounts for a variety of top-k

queries for arbitrary attributes, ranking functions, and index structures. KFLUSHING performs flushing on three phases, a following phase is only invoked when the preceding phase cannot flush  $B\%$  of memory, where the default value of  $B$  is 10. The first phase keeps only  $k$  microblogs in each index cell and trims any records beyond  $k$ . The following phase removes the infrequent values of indexed attributes, e.g., keywords, with their associated microblogs in ascending order of their *latest arrival time*. If infrequent entries do not clear  $B\%$  of memory, the last phase removes data in least recently used order. The main idea in all three phases is evicting data on the level of index entry rather than the level of individual microblogs. This significantly reduces the real-time overhead and scale in highly dynamic data environments. Comparisons with the *first-in-first-out* and *least recently used* policies are made to demonstrate the superiority of KFLUSHING. The results show that KFLUSHING increases memory hit ratio by 26-330% when compared with the existing flushing schemes and saves up to 75% memory resources.

GEOTREND's flushing policy, TRENDMEM [224], depends on aggregate information to evict data from main-memory. GEOTREND queries find top-k trending keywords within an arbitrary spatial region and recent time, where different trend measures depend on keyword count. To effectively utilize memory resources, TRENDMEM evicts keywords that are consistently infrequent during all recent time periods, so they are unlikely to contribute to any top-k trending query answer. Targeting consistent infrequency ensures not to miss a rising keyword. Therefore, TRENDMEM periodically removes  $\epsilon$ -infrequent keywords every  $\frac{1}{\epsilon}$  insertions in each index cell, so dense spatial cells do not affect less populated cells. TRENDMEM achieves significant memory savings while maintaining highly accurate query answers.

GARNET [166] also provides a flushing policy that aims to use the minimal amount of memory rather than utilizing a fixed memory budget. The policy is tailored for its trending queries over arbitrary time periods. Each incoming microblog needs the past  $N+1$  index cells to calculate its trending measure. Thus, only these  $N+1$  cells are kept in memory and any older data is flushed to disk. If less than  $B\%$  of the memory is flushed, GARNET flushes from the least recently arrived keywords till it reaches  $B\%$ . Memory usages of TRENDMEM, GARNET, and AFIA are compared. By comparison, TRENDMEM consumes less than 10% of AFIA memory while GARNET consumes around 40% of AFIA memory. It is also shown that GARNET supports the highest arrival rate. The arrival rate supported by TRENDMEM is higher than AFIA and is also an order of magnitude higher than the current Twitter rate.

### 3 Microblogs Data Management Systems

In this section, we highlight the major data management systems that support either microblogs data in particular or similar characteristics so microblogs data can be one of their use cases. Due to the plethora of new systems that are emerging in the data management literature, our review gives representative examples for each major genre of systems. We identified the major genres based on the adequacy of systems features and components to handle microblogs data. In specific, microblogs combine both large volume and high velocity aspects, where major novel techniques on managing microblogs data give particular attention to its fast streaming nature. Managing fast data has been recently got attention in many data management systems, from both academia and industry, which makes some of microblogs queries manageable in different systems genres. This section reviews five genres of systems: specialized systems that are designed and developed for microblogs, semi-structured data management systems, fast-data-optimized database systems, fast batch processing systems, and key-value stores. In addition, we highlight hybrid architectures that combine two different types of systems to manage microblogs, showing the limitations of this approach.

Table 2 summarizes the microblog-related features for systems that are reviewed in this section. It summarizes their capabilities in terms of indexing, supported queries, and flushing policies, highlighting the minimum and ideal requirements for efficient management of microblogs data. The rest of this section outlines different genres of systems, discussing their challenges, solutions, and limitations.

**Specialized systems.** The literature has few systems that are specialized for microblogs data. A major example from industry is *Twitter* EARLYBIRD. As introduced in Section 2.2, EARLYBIRD system started as a distributed search system that powers real-time keyword search in Twitter [51]. However, Twitter added different functionalities [195, 208, 209, 239, 245, 246, 325] related to real-time data management, large-scale logging, and higher-level data analysis. We focus on one of such functionalities, which is real-time query suggestions, as it was a motivational use case to radically re-design the way Twitter is handling its real-time data and it shows the importance of radically re-thinking batch processing systems to support efficient queries on real-time data as detailed in [245]. When a user poses a keyword query, a query suggestion module finds potential related queries to suggest to the user. For example, a user who searches for *football* might receives suggestions such as *soccer*, *FIFA*, or *world cup*. Twitter was supporting query suggestions through a query analyzer that employs Hadoop MapReduce to analyze the query log of EARLYBIRD system and produce the suggestions. However, using Hadoop has led to significant overhead where an hourly data

is processed in fifteen minutes. This is much slower than the changes in Twitter queries distribution, which changes every few minutes [209, 239]. Thus, fifteen minutes latency to process one-hour data is way behind such fast changes and has led to producing inaccurate query suggestions. To overcome this, Twitter beefed up EARLYBIRD system with in-memory query analyzer modules that directly access user queries through EARLYBIRD blenders (see Section 2.2). Each in-memory query analyzer maintains statistics about incoming queries with a ranking module to filter top related query suggestions. Every five minutes, the suggestions are persisted to a distributed file system, that represents a data store from where query suggestions are retrieved to end users. Such addition to EARLYBIRD system was motivational for Twitter to add several latency-sensitive components to their internal systems and radically re-design solutions that depend on batch processing systems such as Hadoop.

Another two examples of specialized systems that come from academia are TAGHREED and KITE systems. TAGHREED [222], and its successor KITE [227], were early end-to-end holistic systems that focus on microblogs data management in academic systems groups. In particular, both system designs inherently consider microblogs characteristics of both *data* and *queries* in indexing, query processing, and main-memory management. For data, they support fast and large volume data requirements. To this end, they employ both in-memory and in-disk index structures as core components to store, index, and retrieve recent and historical data. Indexes at different storage tiers are optimized for different objectives. In-memory indexes are equipped with fast data ingestion through batching incoming data and segmenting the index into small segment sizes that is lightly updatable. In addition, in-memory indexes are equipped with flushing policies that are responsible for moving a portion of memory content to disk when the available main-memory budget is full. Flushing policies are optimized to sustain system real-time operations as well as careful selection of victim data to evict to utilize memory resources to store useful data that serve incoming queries. For microblogs queries, they promote temporal, spatial, textual, and top-k queries as first-class citizens through indexing and query processing. So, each of the two systems supports two families of index structures: a spatial index and a keyword index. Each index incorporates the temporal aspect in organizing its data, and in certain settings it incorporates the top-k ranking function. Moreover, index segmentation is based on the time dimension in both memory and disk indexes. Disk indexes are optimized for efficient queries over arbitrarily large temporal periods through a richer segmentation setting. Basically, the data are replicated over different temporal resolutions, e.g., day, week, and month, so that querying data over several months still access limited number of index segments and provide a relatively low query latency. Other than in-

**Table 2** Summary of systems features for supporting efficient management of microblogs data.

	<b>Indexed Attributes</b>	<b>Index/Data Storage</b>	<b>Supported Queries</b>	<b>Flushing Policies</b>
EARLYBIRD [51]	Keyword, Temporal	Main-memory	Top-k, Keyword, Temporal	Temporal flushing
ASTERIXDB [17]	Keyword, Spatial, Arbitrary	Disk	Keyword, Spatial, Aggregate, Range	N/A
VOLTDB [314]	Arbitrary	Main-memory	Aggregate, Range	N/A
TAGHREED [222]	Keyword, Spatial, Temporal	Main-memory and disk	Keyword, Spatial, Temporal	Temporal and top-k flushing
KITE [227]	Keyword, Spatial, Temporal, Arbitrary	Main-memory and disk	Top-k, Keyword, Spatial, Temporal, Aggregate	Temporal and top-k flushing
SPARK [23]	None	Main-memory and disk	Temporal, Aggregate	Least-recently-used flushing
FLINK [20]	None	Main-memory and disk	Temporal, Aggregate	Operator-dependant flushing
STORM [25]	None	Main-memory	Temporal, Aggregate	N/A
MONGODB [252]	Keyword, Spatial, Arbitrary	Disk	Temporal, Aggregate	N/A
Key-value stores [19, 277, 21]	Keyword, Arbitrary	Main-memory or disk	Aggregate	N/A

#### Existing systems features

Minimum	Keyword, Spatial, Temporal	Main-memory and disk	Top-k, Keyword, Spatial, Temporal, Aggregate	Temporal and top-k flushing
Ideal	Keyword, Spatial, Temporal, Arbitrary	Main-memory and disk	Top-k, Keyword, Spatial, Temporal, Aggregate, Social, Range, User, Personalized	Temporal, top-k, and customized flushing

#### Minimum requirements and ideal system features for microblogs data management

dexing and query processing, both TAGHREED and KITE give a particular attention to main-memory utilization as a core asset to manage hundreds of millions of microblogs. For this, they provide different optimization techniques in their flushing policies so that most useful data accumulates in main-memory and obsolete data is moved earlier to disk.

Although TAGHREED [222] and KITE [227] share many characteristics in both objectives and system internals, TAGHREED is an earlier version of KITE that started to identify core components and requirements to support microblogs data and queries. Thus, TAGHREED focused in a single generic range query that allow to retrieve microblogs data within a spatio-temporal range and relevant to a set of keywords. Then, any further processing, e.g., top-k ranking, is performed on top of TAGHREED query processor. KITE generalized this to allow querying any arbitrary attribute, while still promoting temporal, spatial, and textual as the prime attributes. Also, KITE added support for more advanced queries in the system components, such as top-k queries and aggregate queries. Ideas in these systems are patented [251] and commercialized by a social media analysis startup company.

**Semi-structured data management systems.** A major example of such systems is *Apache* ASTERIXDB [17] that is a distributed big data management system that has been developed by academic research groups, and has been

recently incubated by Apache Foundation as a top-level Apache project [18]. ASTERIXDB is a general-purpose system that is designed to manage large volume, billion-scale, datasets that are limited to be managed efficiently in other systems. Recently, ASTERIXDB has introduced a core system component, called *data feeds*, to provide scalable ingestion and management for fast data [123], such as microblogs. A data feed digests and preprocesses raw data in main-memory. Then, data is forwarded to primary and secondary index structures. Each index is disk-based, however, it has in-memory components that aggregate data in main-memory before flushing them to disk-resident components. Data is accessible to the query processor when it is resident in the disk components. When data is congested, ASTERIXDB is equipped with different ingestion policies to select a portion of data to ingest promptly, while the rest of data is discarded or deferred. ASTERIXDB has achieved data digestion rates that are comparable to current Twitter peak rates with a cluster of five machines as experimented in [123]. Such performance is higher than what is reported by EARLYBIRD [51] in terms of data digestion per single machine. In terms of digestion latency (or searchability latency), i.e., average time between a microblog arrives to being indexed and available in search results, ASTERIXDB data feeds provide low latency that are appropriate for real-time applications with certain ingestion policies, and signif-

icantly high latency with other policies. So, it is crucial to configure the system carefully for the underlying application needs. As a general-purpose system that is not designed for microblogs data, ASTERIXDB provides common utilities that fit for general fast data use cases without focusing on particular microblogs characteristics, such as temporal and top-k query signatures.

**Fast-data-optimized database systems.** Although many of microblogs applications do not require transactional data management, database systems that are optimized for transactions on fast data are strong candidates to be used to handle some of microblogs queries, with optionally turning on or off the transactional features. This is due to their light weight management overhead with streaming data, while sustaining a high throughput of scalable queries. VOLTDB is an example for such systems. VOLTDB [314] is a distributed in-memory database management system (DBMS) that is designed and optimized to support high-throughput ACID database transactions on fast data. The system has started as an academic project, under the name of H-STORE [168], that is commercialized by VOLTDB [314, 338]. The main additions of VOLTDB to traditional disk-based database systems are driven by reducing the overhead of the database transaction manager. Particularly, VOLTDB identifies four major sources of overhead in transaction management: (1) multi-threading that is required to manage multiple transactions concurrently, (2) buffer manager that swaps in data pages from disk to a main-memory buffer and evicts pages to disk on full memory buffer, (3) locking that is used to manage data consistency in concurrency control, and (4) logging that is essential in recovery management of completed transactions and rolling back aborted transactions. So, the four main contributions of VOLTDB is to tackle such overhead sources to increase the throughput of transactions for fast data management. The multi-threading overhead is totally eliminated by assigning each transaction to a single dedicated CPU core. The buffer management overhead is totally eliminated by eliminating disk storage and storing all data in main-memory, so no buffer is managed in VOLTDB. The locking overhead is also eliminated through determining deterministic orders for executing transactions through introducing global and local serializer components. The global serializer is a component that is aware of different data replicas on different machines, while the local serializer has the transactions details on a single local machine. Both components exchange information so the global serializer is able to provide each local replica deterministic orders for transactions, which leads to eliminating the locking overhead. Finally, the logging overhead is significantly reduced through logging data images instead of logging single transaction commands. In particular, VOLTDB does not provide recovery management through

the traditional write-ahead logging that mandates to write each transaction step to the database log file. Instead, only transaction parameters are written to file proactively. Then, in lazy basis, a full image of current data is written to disk for recovery purposes. This significantly reduces disk access and increases the throughput to 16,000 transaction per core per second, with almost linear scalability when adding more cores. This light management overhead has significantly lifted up managing fast data. Thus, VOLTDB indexing and data management infrastructures are suitable to digest fast data efficiently and support important queries in real time, such as keyword queries. However, there are two major concerns for effectively supporting microblogs data end to end. First, VOLTDB and similar systems are not optimized for large volume datasets, as stated in their technical documentation, which will lead to limitations in handling historical microblogs, over several months, that are richly exploited in different use cases. Second, it has no support for prime attributes of microblogs, such as the spatial attribute, which makes it inadequate for several important queries even on fast microblogs data.

**Fast batch processing systems.** Recently, a new generation of distributed batch processing systems has been emerged, extending Hadoop-like systems with main-memory data management infrastructures for efficient processing of large and fast datasets. SPARK [23] and FLINK [20] are prime examples for these systems. Both systems primarily process data in main-memory with options to connect to popular file systems, such as HDFS, or store statuses in persistent data stores, such as RocksDB [284]. As in-memory systems that support fast data through streaming packages, e.g. SPARK STREAMING [24], some microblogs applications could fit as use cases for these systems. However, unlike all reviewed systems earlier, SPARK and FLINK do not inherently support data indexing. Instead, they provide an advanced generation of batch processing systems, similar in spirit to Hadoop, that provide efficient parallel scans over all data records using commodity hardware clusters. Batch processing has limitations in several applications that need inherent indexing for either large volume or high velocity data. Newer systems, e.g., Apache ASTERIXDB, have tackled these limitations and provide different types of indexing for large and fast data. Obviously, many of microblogs applications are among these applications that require data indexing of several types as detailed earlier. For that reason, any system that give particular attention to microblogs, e.g., EARLYBIRD, TAGHREED, or KITE, has provided different types of indexing for microblogs data. Furthermore, batch processing systems, such as SPARK and FLINK, do not consider query signatures that are popular in microblogs applications, e.g., top-k, spatial, and textual queries. This adds more overhead when powering large-scale microblogs applications on batch processing systems.

The pros and cons of SPARK and FLINK apply to other batch processing systems that share similar characteristics and architecture, e.g., Apache IMPALA [22] and PRESTO [269].

**Key-value stores.** A major genre of the emerging big data systems is key-value stores that work as massively distributed hashtables to store data in key-value pairs with various data models, e.g., Apache CASSANDRA [19], REDIS [277], and Apache IGNITE[21]. These systems are suitable for certain microblogs applications that require fast data ingestion with hash-based indexing, e.g., real-time keyword search. In fact, some of microblogs-oriented systems, e.g., EARLYBIRD [51] and KITE [227], are using the key-value store model to support in-memory keyword indexing. However, distributed key-value stores still lack other essential features that are needed in several microblogs applications, such as spatial indexing, temporal awareness, and top-k query processing. Such shortcomings limit them from being an end-to-end solution for managing microblogs, yet, they provide a solid foundation to build upon.

**Hybrid architectures.** An alternative way to handle fast and large data is gluing a streaming engine, such as Apache STORM [25], with a persistent data store, such as MONGODB [252]. In fact, MONGODB, a document-oriented database that provides several indexing and querying modules, has got a significant attention as a highly scalable database for persistent data, while Apache STORM has got similar attention for processing streaming data. However, each of them is designed and optimized for one aspect of big data, either large volume or high velocity, but not both. It has been experimented to glue these two systems in [123] to handle fast data that got persisted in large volumes. The comparison with Apache ASTERIXDB has shown up to two orders of magnitudes of higher digestion latency for the glued alternative, assuming that data is queried only when it is persisted to disk. Such significant overhead confirms the need of inherent support of fast data in the system components to provide scalable data indexing and querying. Similar conclusions are also drawn in other studies, e.g., [245], on the adequacy of adapting fast data management in systems that are optimized for large volumes. A major source of overhead is the incompatibility of system optimization goals, which leads to different decisions in different system components. For example, MONGODB is optimized for throughput, write concurrency, and durability, which leads to high wait time per single data write to disk and high ingestion latency in turn. Another source of overhead in such systems is the concurrency and transactions model that assume general-purpose applications with complex scenarios and requirements. This does not allow to use simple and scalable concurrency models, such as single-writer multiple-readers, that is adapted by several microblogs-oriented systems, e.g., [51, 228, 352].




Fig. 3 An overview of microblogs data visualization literature.

## 4 Microblogs Data Analysis

The reviewed data management techniques and systems on microblogs have enabled to power a variety of data analysis tasks at scale. This section highlights the major data analysis research for analysis tasks that exploit the scalable data management infrastructures on microblogs to provide high-level functionality. As microblogs data analysis is a broad literature and include several topics that are not related to the data management community, this section limits its scope only to the analysis tasks that lie in the intersection of two categories. First, they have novel research contributions, which excludes a plethora of development applications that analyze microblogs data without addressing novel problems. Second, they exploit the querying techniques that are developed by the data management community. This excludes major research directions that are orthogonal from the data management research, such as natural language processing and information retrieval. In fact, these research directions has a rich literature where dedicated survey papers review parts of it [82, 119]. The goal of this section is not discussing the details of various techniques. Instead, we present a high-level classification for techniques in the literature, and we summarize each topic through a generic framework that is induced from a variety of existing techniques when applicable. Then, we briefly highlight similarities or differences of each major technique in this topic compared with the induced framework. With such contributions, this section represents a road map for various microblogs data analysis that make use of the underlying data management infrastructures. We review major work in five main analysis tasks: *visual analysis* (Section 4.1), *user analysis* (Section 4.2), *event detection and analysis* (Section 4.3), *recommendations using microblogs* (Section 4.4), and *automatic geotagging* (Section 4.5). Finally, Section 4.6 briefly highlights other microblogs analysis tasks.

### 4.1 Visual Analysis

Visualizing microblogs data has gained a particular attention due to the importance of end users interactions with microblogs applications, e.g., political and disastrous event analysis, disease outbreaks detection, and user communities analysis. The challenges faced in visualizing microblogs data aligns with the general challenges in visualizing other types of big data [42, 59, 101, 102, 130, 178, 262]. So,

several pieces of the proposed research for big data visualization can be used for microblogs data as one type of big datasets. However, we review visualization work that targets a specific problem in microblogs datasets for different applications. In particular, microblogs have micro-length content, which makes them easy to be generated by users all the time, e.g., a user can easily generate a tweet in a few seconds or less. This leads to generating a large number of data records in relatively short times. Visualizing such large numbers is beyond the capacity of existing frontend technologies, such as mapping technologies, e.g., GoogleMaps. So, visualization techniques that focus on microblogs try to address this problem by either aggregation, sampling, or a combination of both. Figure 3 classifies the visualization literature into three categories of techniques: (1) aggregation-based techniques, (2) sampling-based techniques, or (3) hybrid techniques. The visualization modules in all these categories use underlying querying modules, both aggregate and non-aggregate queries, to retrieve the data to be visualized. Thus, they directly make use of the scalable data management infrastructures that are built for microblogs. The rest of this section outlines each category of techniques.

**Aggregation-based visualization.** Techniques in this category [3, 95, 116, 152, 158, 236, 282, 302, 315, 341, 349, 354, 367] reduce the amount of data to be visualized through visualizing aggregate summaries of microblogs at different levels of aggregation, e.g., different spatial levels or temporal levels, rather than visualizing individual microblogs. Such aggregation is application-dependent and is usually performed either based on major attributes, e.g., temporal aggregation [95, 152], spatial aggregation [116, 349], or keyword aggregation [95, 315], or based on derived attributes, e.g., sentiment [152, 282]. Thus, these techniques are lossless and present all available information in a summarized form without ignoring any portion of the data. Aggregation could be based on a single attribute (one-dimensional) or multiple attributes (multi-dimensional). Figure 4 shows an example of aggregation-based visualization based on a single attribute, the spatial attribute [258]. In Figure 4(a), spatial regions that have a large number of data points visualize a variable-size circle that shows the number of points in this region. On the contrary, regions that have sparse data, Arctic Ocean and Norwegian Sea in Figure 4(a), visualize the actual data points. On zooming on the map view, more detailed data is visualized up to the street level that shows detailed data points, as depicted in Figure 4(b) that shows street-level data in Riverside, California. Figure 5 shows an example of aggregation-based visualization based on two attributes, the spatial attribute and the language attribute [116]. In this case, number of microblogs is aggregated in each spatial region and the visualized circle categories data based on the language at-

tribute to show percentage of microblogs posted in English, Arabic, Indonesian, Persian, etc.

The literature currently has seventeen visualization modules that employ only data aggregation based on microblogs queries. We next briefly outline each of them, highlighting their aggregation attributes and visualization format. **VISCAT** [116] aggregates data based on categorical attributes, e.g., language, and spatial and temporal ranges. **DISCVIS** [349] aggregates tweets based on spatial region, language, and topics. **DESTINYVIZ** [95] aggregates tweets related to certain games based on time, sentiment, and keywords. **NLCOMS** [3] aggregates tweets based on user communities and visualize them in a graph form. **GOV-VIZ** [152] aggregates data based on time, country, topic, keywords, sentiment, and content objects, e.g., links, images, and videos. **VISIMP** [367] aggregates data based on communities and social interactions. **TWIGRAPH** [315] aggregates data based on keywords and visualizes it in a graph form. **PLEXUS** [354] aggregates data based on topics and emoji objects in the textual content. **TSVIZ** [282] aggregates data based on time, sentiment, and hashtags. **PAIRCSA** [341] aggregates data based on their location stamps or location mentions, to get relation between users locations and the locations they mention. **TWEETVIZ** [302] aggregates data based on sentiment for business intelligence. **NETWORKTWEET** [158] aggregates external passenger flow and unusual phenomena based on spatio-temporal attributes, and uses trending keywords from microblogs to understand users' behaviour. **TWITTERVIZ13** [167] aggregates data based on tweets' intensity (tweet/second) and tweets sentiment. **CITYVIZ** [281] aggregates data based on user behavior in cities to visualize periods intense/sparse of user activity. **TILEVIZ** [68] generates summary statistics of the data for each tile for exploring the raw data set. **TWITTERVIZ15** [100] provides two visualization views for Twitter data: (a) spatio-temporal analysis view, and (b) graph analysis view. The spatio-temporal view aggregates data based on spatial regions, sentiment, social bonds combined with spatio-temporal information, temporal evolution, and real-time statistics. The second view aggregates data based on social graph and real-time graph statistics. **IMPRESSVIZ** [188] aggregates textual and meta-data information to quantify user impression and visualize data in a six-dimensional impression space.

**Sampling-based visualization.** Techniques in this category [222, 293, 346] reduce the amount of visualized data through sampling. A sample of data is selected and visualized as a representative for the whole dataset, while the rest of data is not visualized. The sampling technique can be classified based on different dimensions. A sample could be a query-guided sample or an arbitrary sample. An example for a query-guided sample is OmniSci TweetMap<sup>2</sup>

<sup>2</sup> <https://www.omnisci.com/demos/tweetmap/>

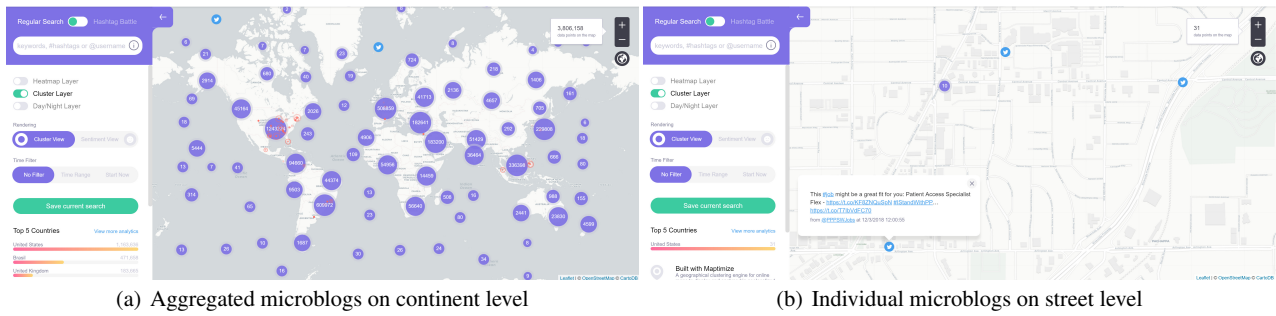


Fig. 4 Example of aggregation-based visualization based on the spatial dimension.

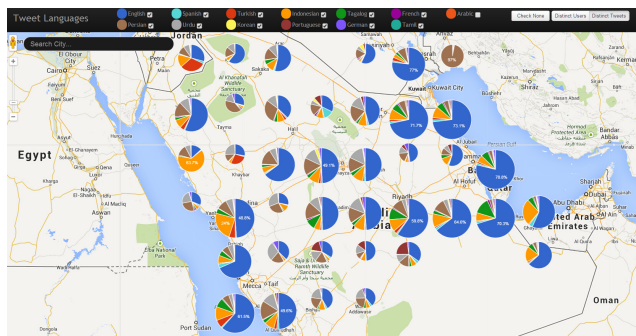


Fig. 5 Example of aggregation-based visualization based on both spatial and language dimensions.

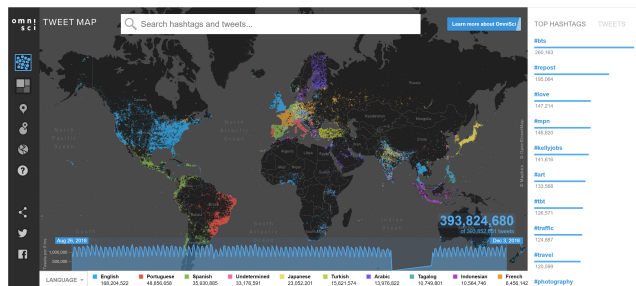


Fig. 6 Example of sampling-based visualization for tweets with different languages.

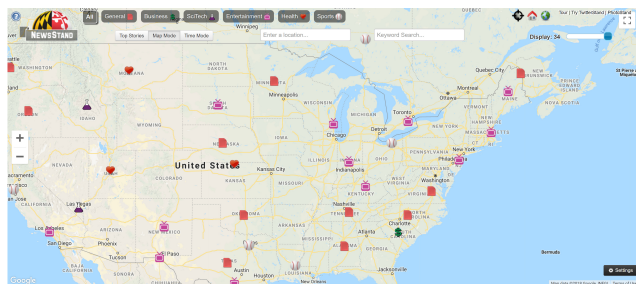


Fig. 7 Example of sampling-based visualization for news tweets.

(Figure 6) that samples tweets based their language as the query predicate filters data based on the language attribute. Another example is TwitterStand [293] (Figure 7) that samples tweets based on textual content that have news stories. For certain queries, the query predicate is generating a lot of

data that still cannot be visualized efficiently. In this case, applications, e.g., [222], selects an arbitrary data sample to reduce the data size. Another classification of the way of sampling is based on the amount of data in the sample. The sample is either fixed or interactive. For example, TwitterStand [293] takes a fixed sample of data that contains new stories. Any interaction for end users with the map view, in Figure 7, will not change the content of this sample. User interactions only change the subset of this sample that is shown on the map. On the contrary, an interactive sample changes the sample content based on user interactions. At the beginning, an initial sample of 100K, for example, is visualized from all languages including 30K English microblogs. When the user filters out data to show only English microblogs, the visualized English microblogs can be increased to 100K as it is solely visualized. Such interactive technique is exploiting the whole capacity of frontend technologies while increasing the overall amount of data visualized to users. Such technique is not heavily used and has several research challenges to support large-scale data.

Unlike aggregation-based techniques that are lossless, sampling-based techniques might be lossy or lossless depending on the application and the size of query result. If certain application queries are generating a reasonable sample size, then all data points are considered. Otherwise, such as in arbitrary sampling, a subset of data points are ignored and the sampling is lossy.

The literature currently has three visualization modules that employ only data sampling based on microblogs non-aggregate queries. We briefly outline each module highlighting the sampling attributes and stages. **CULTWEET** [346] samples data based on language, country, and topic. **TAGHREED** [222] performs two-step sampling. First, it samples data guided by query predicates based on spatial, temporal, and keywords. Then, if the sample size is still excessive, it performs an arbitrary sampling. **TWITTERSTAND** [293] samples data based on textual content and spatial extent.

**Hybrid visualization.** Some applications allow to use both aggregation and sampling to reduce the amount of data to be visualized [78, 161, 236, 238, 240, 326, 366]. For ex-




ample, event analysis applications [238, 326] sample microblogs based on their relevance to specific events. Then, event data need to be aggregated to summarize the event highlights to users, e.g., showing changes over time, space, users, or topics. Such applications usually do not encounter challenges in visualizing their data as the data size is reduced over two different phases, sampling and aggregation, which leads to significant reduction in their size and ease the visualization task. We highlight examples of such applications.

We highlight nine visualization modules that employ both data aggregation and sampling based on microblogs queries. We briefly outline each of them highlighting it different stages. **TWEETTRACKER** [326] samples tweets that are relevant to a set of tracked long-term events, then it aggregates them based on location, time, and keywords. **TWITINFO** [238] samples event-related data and aggregates them based on sentiment and spatial attributes. **ATRVIS** [236] samples tweets that are relevant to a set of input debates, then it aggregates and label tweets based on mentioned hashtags and the corresponding debate. **CLOUDBERRY** [161] samples data based on keywords and aggregates it based on space and time. **TWEETDESK** [240] provides a sample of top tweets of an event, along with a summary about the event. **CHINESENTIMENT** [366] visualizes sentiment distribution based on temporal, spatial and hot events. **EMOTIONWATCH** [174] visualizes sentiment summary of public reactions towards events. It allows visualization of intense emotional reactions (peaks), controversial reactions and emotional anomalies. **USERVIZ** [286] analyzes users' connections and the frequency of tweets sent by one or a group of users, classifies these tweets, generates a tag cloud, and visualizes the most popular users. **TAQREER** [230] samples microblogs based on user-defined categories, e.g., different car models, defined by a set of keywords, then data for each category is aggregated based on spatial and temporal ranges and visualized on map and aggregate views.

## 4.2 User Analysis

The importance of microblogs in different applications originates from its user-generated nature, where hundreds of millions of users worldwide are posting around the clock. Among the major analysis directions is analyzing the user behaviour related to different topics, locations, and communities based on their profiles and content of their microblogs. In fact, such kind of user analysis is highly overlapping in microblogs, i.e., the micro-length user-generated data, and social media in general that include both short and long posts and objects, e.g., images and videos. This section limits its scope to analyzing microblogs users where excessive numbers of data records are generated compared to regular social media data due to its micro-length.



**Fig. 8** An overview of microblogs user analysis literature.

Figure 8 classifies the literature of user analysis techniques on microblogs into techniques that either (1) find top-k users according to a certain ranking criteria, e.g., top-k influential users for a certain topic or top-k active users in a certain location, to provide useful answers for higher-level applications, or (2) classify users based on certain characteristics. Top-k users queries directly benefit from the data indexing and query processing techniques that are introduced in the data management literature to support different types of scalable top-k queries based on various ranking functions. In fact, usernames are used interchangeably with keywords as string keys for the index structures, which makes many of the proposed data management techniques applicable to user queries. Figure 9 depicts a high-level framework for user analysis in microblogs that is induced from the existing literature. The framework consists of three main stages. First, microblogs of each user are fed into a feature extraction module to profile the user behaviour through different pieces of information, such as keywords, followers/friends, timestmaps, and locations. The actual extracted information is different in different applications. Then, the extracted user information is forwarded to an indexing (or modelling) module that produces a relevant index/model for users based on their information. Finally, a query processor accesses the index/model to answer application-level queries. As the description shows, the last two stages of this user analysis framework significantly make use of the data management techniques, and hence new advancements in indexing and query processing techniques would positively affect the performance of user analysis tasks. Following the described framework, major techniques in the literature serves different applications with diverse purposes. We outline each category of techniques below.

**Top-k user queries.** Section 4.4 has reviewed several techniques that recommend top-k users as potential friends, which overlaps with top-k user queries. **DOMUSR** [212] finds most influential users based on nine features that are aggregated through different models to calculate a final influence score. The used models are both aggregation and SVM classification models. **TKLUS** [162] finds top-k local users who are most active for a certain topic in certain location. TKLUS uses textual, social, and spatial relevance in hybrid spatial-keyword index to organize and retrieve top-k users efficiently. **PROMUSR** [48] finds prominent users for certain event through a probabilistic model that analyzes their temporal and textual information. **LNKUSR** [336] identifies top candidate user entities with limited informa-

tion on microblogging platforms that can be linked to user entities on other platforms. It extends graph matching techniques with two heuristics to overcome the limited available information. **TWITMENDER** [134] finds top users with similar interest to the querying user to expand homogeneous communities of similar interests. It profiles user posts content and use collaborative filtering techniques to find similar users. **TEMUSR** [292] models users temporal behaviour for different short-term and long-term topics. **CUSR** [103] samples microblogs data records, rather than sampling  $k$  users, for efficient user community reconstruction based on strongly connected components. **IBCF** [220] uses dynamic user interactions in different topics to model the dynamics of relationship strength between users and topics over time. Then, the modeled relationships are used in matrix factorization recommendation model to improve social-based recommendation quality. **FUREC** [353] predicts the top- $k$  users who will retweet or mention a focal user in the future by formalizing the prediction problem as a link prediction problem over an evolving hybrid network. **INFUSR** [14] finds most influential users in a certain topic. A nodal feature called focus rate is introduced to measure how focal users are on specific topics. Then they incorporate nodal features into network features and use modified PageRank algorithm to analyze topical influence of users. **FADERANK** [39] evaluates the reputation of Twitter users. It summaries the past history in a bounded number of values and combines them with the raw reputation computed from the most recent behaviour to assign a final ranking score. **TRUETOP** [376] outputs top- $k$  influential non-sybil users among a given set of Twitter users. The system constructs an interaction graph and then performs iterative credit distribution using weighted eigenvector centrality as the metric to make the influential non-sybil users stand out. **UIRANK** [382] identifies influential users whose tweets can cause the readers to change emotion, opinion or behaviour. The algorithm is based on random walk and measures the user's tweet influence and information dissemination ability to evaluate the influence of the user. **FAME** [193] finds topical authorities on Twitter for a given topic. The algorithm adopts a focused crawling strategy to collect a high-quality graph and applies a query-dependent personalized PageRank to find important nodes that represent authorities. **COGNOS** [118] identifies expert users for a certain topic through mining the meta-data of Twitter user lists that are created by the crowd. Lexical techniques are used to infer user expertise, then experts in the same topic are ranked based on cover density ranking.

**User classification.** In addition to top- $k$  queries, user analysis is also performed to do user classification. **PE-DIDENT** [137] identifies pro-eating disorder (ED) Tumblr posts and Twitter users. They use the associative classification algorithm CMAR to generate classification rules and train a classifier to identify pro-ED posts and users.




Fig. 9 A framework for microblogs user analysis.

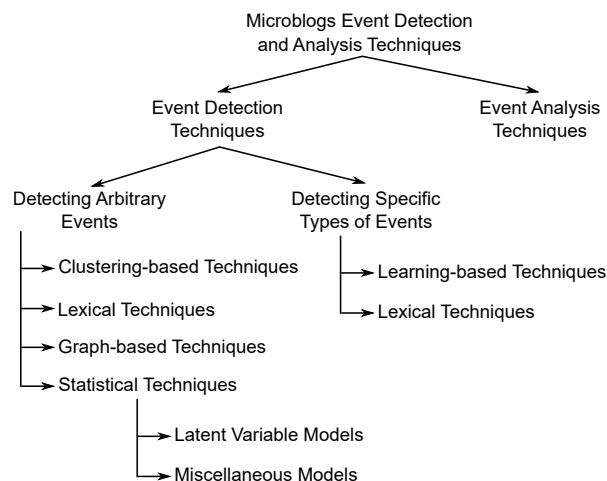


Fig. 10 An overview of microblogs event detection and analysis literature.

**AOH** [120] classifies users into automated agents and human users using a Random Forest classifier. **OMT** [176] identifies the orientation of a user by analyzing tweets which mention more than one orientation using a logistic regression model. **HUSR** [280] identifies hateful users from twitter. They first sample users using a diffusion process based on DeGroot's learning model. Then, a crowd-sourcing service was adopted to manually annotate the samples. **AUTOOPU** [379] detects the opioid users through a multi-kernel learning model based on meta-structures over heterogeneous information network.

### 4.3 Event Detection and Analysis

Event detection and analysis has gained tremendous attention with the rise of microblogging platforms [1, 2, 11, 15, 16, 31, 80, 93, 104, 110, 124, 148, 150, 151, 169, 170, 187, 203, 211, 217, 253, 266, 268, 283, 289, 291, 348, 358, 369, 370, 378, 384, 388, 389, 390]. The reason is the popularity of event-related updates that are posted by users through microblogs around the clock. This includes a wide variety of both short-term and long-term events, such as concerts, crimes, sports matches, accidents, natural disasters, social unrest, festivals, traffic jams, elections, and conflicts. Analyzing the event-related microblogs enabled several applications at different levels of importance, including crucial

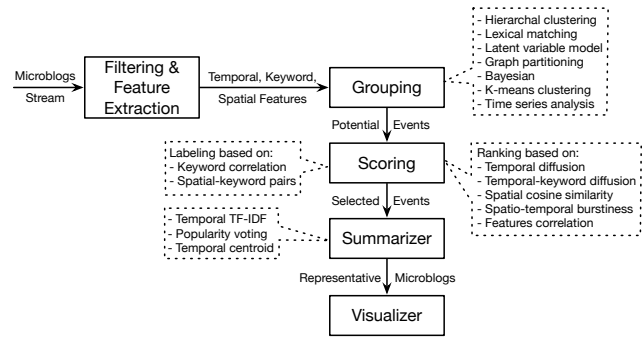
applications, leisure applications, and in-between applications. An example for crucial applications is rescue services and emergency response that have used microblogs to save hundreds of souls in different natural disasters since 2012 across the world [69, 70, 105, 153, 154, 160, 172]. An example of leisure applications is detecting surrounding entertainment events that are not collected in a single calendar, e.g., concerts, light shows, and special museum exhibitions in Los Angeles area. In between both types, other types of applications have become popular, such as news extraction based on events [49], event-driven advertising [256], public opinion analysis for political campaigns [334, 335], and analyzing protests and social unrest [27, 157, 333].

The advancements in microblogs data management enable significant performance enhancements in both tasks of events detection and analysis. As noted in the data management section, there are several state-of-the-art indexing and query processing techniques that are tailored for organizing and retrieving event data, such as CONTEXEVEN [6] and MIL [52] that are reviewed in Section 2.2.1. In a more general context, event detection makes use of indexing data based on temporal attributes that enables efficient retrieval of recent and temporally compact data, which is a major characteristics for grouping relevant data of a single event. In addition, indexing data based on spatial attributes gives an edge for discovering local events in geographic neighborhoods.


Figure 10 depicts an overview for the literature of event detection and analysis on microblogs. The rich literature is categorized into three main categories: (1) detecting arbitrary events, (2) detecting specific types of events, and (3) analyzing events. We summarize each category with a generic framework that is induced based on major work in the literature. Figure 11 shows three frameworks that correspond to the three categories. In the rest of this section, we review each category describing the different components of its framework and mapping existing literature to this framework highlighting similarities and deviations.

#### 4.3.1 Detecting Arbitrary Events


A major direction of event detection research focuses on detecting arbitrary events that have either no predefined or at most very high-level characteristics. For example, finding coherent discussions on Twitter [31, 80] without having a prior idea about what could be such discussions about. Another example is looking for local events in a certain city [2, 110] without determining any specific characteristics of such events. These events are arbitrary events as the user does not provide a prior detailed description for the event characteristics. Figure 11(a) depicts a framework that is followed by most arbitrary event detection techniques. The framework consists of five main stages: (a) *filtering*



(a) Detecting Arbitrary Events



(b) Detecting Specific Types of Events



(c) Event Analysis

**Fig. 11** Frameworks for microblogs event detection and analysis.

& feature extraction, (b) grouping, (c) scoring, (d) summarization, and (e) visualization. A microblog dataset, either streaming or stored dataset, is processed through the filtering and feature extraction stage to identify potentially relevant microblogs and extract their temporal [2, 15, 110, 187, 283, 348, 369, 384, 390], textual [15, 31, 93, 348, 369, 378, 390], spatial [2, 15, 110, 170, 187, 348, 369, 370, 378, 384, 390], and semantic (part-of-speech (POS) tags/named entities) [217, 283, 289, 358, 369, 388, 389] features. These four types of features are the main drivers for detecting new events. Then, microblogs are forwarded to a grouping stage that assembles microblogs with similar features into groups, each group represents an event candidate. The grouping stage uses different types of techniques, including clustering [2, 15, 80, 110, 169, 217, 348, 368, 369, 370], lexical matching [289, 378], graph partitioning [31, 93], and statistical techniques such as bayesian [369], latent variable models [283, 388, 389], and regression models [187], as depicted in Figure 10. The set of candidate events are then forwarded to a scoring module that gives a score (or a label) for each candidate to distinguish actual events from noisy groups. Scoring is performed in different ways, including labeling [170, 187, 368, 369, 384, 388] or ranking candidates based on diffusion [2, 80, 283], similarity [15, 31, 80, 169, 217, 283, 378, 390], correlation [93, 289, 348, 389],

**Table 3** Summary of clustering-based and lexical techniques that detect arbitrary events from microblogs data.

	Features	Grouping	Scoring	Summarization
EVENTTWEET [2]	Textual, Spatial, Temporal	Clustering: spatial cosine similarity	Ranking: Keyword burstiness, temporal diffusion	Number of keywords, spatio-temporal signature
STREAMCUBE [110]	Hashtag, Spatial, Temporal	Clustering: hierarchical spatio-temporal clustering	Ranking: Hashtag burstiness, local features	None
EVEMIN [170]	Textual, Spatial, Images	Clustering: word bursts using n-gram model, DCNN for images	Labeling: DCNN	None
REUTERSTRACER [217]	Named entities	Novel clustering technique	Ranking: newsworthiness probabilistic model	Cluster centroid
TRIOVECEVENT [369]	Semantic, Textual, Temporal, Spatial	Clustering: bayesian mixture model	Labeling: multimodal embeddings classifier	None
DISRUPTEVENT [15]	Textual, Spatial, Temporal	Online clustering based on cosine similarity	Replaced with naive bayes filtering	Top keywords ranked on novel temporal TF-IDF
MGE LDA [358]	Semantic, LDA micro topics	Clustering: k-means	None	Top frequent hashtags
STORYEVENT [169]	Textual, Temporal	Clustering: non-negative matrix factorization	Ranking: probabilistic model	Sequence of sub-events
JASMINE [348]	Textual, Spatial, Temporal	Lexical: based on location and time	Ranking: co-occurring words	None
DISASUBEVENT [289]	POS tags	Lexical: noun-verb pairs	Ranking: noun-verb co-occurrences	Integer linear programming

and/or burstiness of different combinations of temporal, keyword, and spatial features [2, 110, 370, 384]. For example, in scoring based on keyword correlation, if the group of microblogs has scattered keywords that are not related to each other based on statistical co-occurrences of words, then this group is discarded as a noisy group that does not reflect an event. On the contrary, if the set of keywords are cohesive and with high co-occurrence likelihood in real topics, it is assigned a high score as an actual event. Several scoring techniques also consider temporal and spatial similarities besides textual-based measures. Then, top scored candidates are selected as actual events while the rest of groups are considered noisy events. The events are then fed to an optional summarization module that identifies the most important microblog posts to represent a certain event using different signals of importance, such as popularity of the post, its temporal position, etc. Finally, the events are forwarded to a visualizer that displays representative microblogs along with their labels, content, locations, and temporal details to end users. The visualizer uses microblogs visual analysis techniques that are presented in Section 4.1, so the details of this are not duplicated in this section.

The described framework drives the major techniques in the literature. Tables 3 and 4 summarize the different stages of each technique. This literature can be categorized into four categories based on the grouping technique as the major stage that generates event candidates. Figure 10 depicts the four-category classification, namely, *clustering-based tech-*

*niques, lexical techniques, graph-based techniques, and statistical techniques.* Table 3 summarizes clustering-based and lexical techniques and Table 4 summarizes graph-based and statistical techniques. The rest of this section briefly outlines techniques of each category.

**Clustering-based techniques.** **EVENTTWEET** [2] proposes a framework to detect localized events in real time from a Twitter stream and track their evolution over time by adopting a continuous analysis of the most recent tweets within a time-based sliding window. Event candidate extraction is based on clustering keywords according to their cosine similarity of their spatial signature. Scoring the events is based on keyword burstiness and time diffusion from the cluster. Detected localized events are summarized by the number of related keywords and spatio-temporal characteristics. **STREAMCUBE** [110] system extracts microblogs hashtags along with spatio-temporal attributes. Then, hashtags are grouped through a single-pass hierarchical spatio-temporal clustering to detect event candidates, that are scored and ranked based on burstiness and local features. The system provides ways to explore events with different granularities in both time and location. **EVEMIN** [170] detects visual events based on photos and locations. Feature extraction calculates area weights and commonness score of words, grouping depends on word bursts using n-gram model and image clustering based on deep convolutional neural network (DCNN), and labeling uses another DCNN. **REUTERSTRACER** [217] extracts features based

**Table 4** Summary of graph-based and statistical techniques that detect arbitrary events from microblogs data.

	Features	Grouping	Scoring	Summarization
DYNAMICENTR [31]	Textual	Graph-based: dynamic semantic graph	Ranking: dynamic eigenvector centrality	Min weighted set cover
SNAF [378]	Textual, Spatial	Graph-based: spatial connected components (CC)	Ranking: number of nodes in CC	None
GEOBURST [370], GEOBURST+ [368]	Semantic, Textual, Spatial	Graph-based: geo-topics based on spatial and semantic features	Ranking: spatio-temporal bursts [370], Labeling: supervised model [368]	None
EVENDETECTTWITTER [93]	Textual, Temporal	Graph-based: multi-assignment graph partitioning	Ranking: cross-correlation similarity (short-time events), Riemannian distance (long-time events)	None
EXPLOREVEN [388]	Semantic, Temporal, Named entities, POS tags	Statistical: latent event and category model (LECM)	Labeling: matching entities with a semantic class	None
PROBEVENT [389]	Named entities, POS tags	Statistical: unsupervised latent variable model (LEEV)	Ranking: correlation of named entities, dates, locations, and words	None
OPENEVE [283]	Temporal, Named entities, POS tags	Statistical: latent variable model	Ranking: number of microblogs	Top phrases and entities based on a probabilistic model
EYEWITNESS [187]	Spatial, Temporal, Aggregate	Statistical: regression model	Labeling: anomaly threshold	Top microblogs based on text summarization
SPATIALEVENT [384]	Spatial, Temporal	Statistical: hidden Markov model (HMM)	Labeling: predefined taxonomy	None
SEVENT [390]	Textual, Spatial, Temporal	Statistical: location-time constrained model	Ranking: similarity join	None
BEVEN [80]	Hashtags, Temporal, User	Statistical: 3D probabilistic model	Ranking: probabilistic model	None

on named entities, while grouping uses a novel clustering algorithm that counts for microblogs features. **TRIOVECEVENT** [369] detects local events through extracting semantic textual, temporal, and spatial features that are used by a multimodal embedding learner to map correlated microblogs to the same latent space. Then, a novel bayesian mixture clustering model finds geo-topic candidate events. These candidate events are then passed by a classifier that relies on the multimodal embeddings to label whether an event is a local event. **DISRUPTEVENT** [15] framework has both classification and clustering. The classification phase is used for filtering event-related posts from noisy posts and based on a naive bayes model. Then, an online clustering is performed using temporal, spatial and textual set of features. After clustering, the framework offers event summarization using a novel temporal Term FrequencyInverse Document Frequency (TF-IDF) that generates a summary of top terms without the need of prior knowledge of the entire dataset. **MGELDA** [358] is a sub-event detection technique that extracts semantic features based on micro topics. The micro

topics are identified by a novel mutually generative Latent Dirichlet Allocation (LDA) model for microblogs hashtags. Then, k-means clustering is used to group related topics and discover events. **STORYEVENT** [169] also introduces a model that summarizes each event as a sequence of sub-events on a timeline based on non-negative matrix factorization (NMF) clustering.

**Lexical techniques.** **JASMINE** [348] extracts co-occurring words as well as geolocation and timestamp of microblogs. Then, microblogs that are generated within a short time and a small geographic area are grouped to form event candidates. Co-occurring words of each candidate are analyzed to distinguish noisy candidates from local events. **DISASUBEVENT** [289] extracts sub events from a bigger event, e.g. a disastrous event has a series of small-scale emergencies such as a bridge collapsing, airport getting shut, and medical aid reaching an area. Feature extraction is based on POS tagging, grouping of sub-events is based on noun-verb pairs, and ranking is based on the frequency of co-occurrence of their constituent nouns and verbs in the cor-

pus. For summarization, DISASUBEVENT uses an integer linear programming (ILP) technique that considers the maximum occurrence of nouns, verbs, and numerals.

**Graph-based techniques.** DYNAMICENTR [31] combines the first three stages of the framework depicted in Figure 11(a) through extracting emergent keywords from incoming data streams based on analyzing the dynamic semantic graphs, where nodes represent the keywords and the edges are the co-occurrence of the keywords. Then, events are summarized based on the minimum weighted set cover applied on the semantic graph of the dynamically highly ranked keywords. SNAF [378] detects local events based on spatial and textual features of microblogs. It first filters event-relevant microblogs based on lexical analysis and statistical user profiling. Then, relevant microblogs are geotagged based on a large gazetteer and distance-based data cleaning algorithms. The cleaned data is then grouped into spatial connected components that represent events. GEOBURST [370] uses spatial and keyword features to build a keyword co-occurrence graph that is used to infer semantic features through random walks. Then, geotopic clusters are formed as candidate events by combining both spatial and semantic features. A set of pivot microblogs are identified for each cluster, and then they are ranked based on spatio-temporal bursts and top-k are selected. GEOBURST+ [368] differs from GEOBURST by employing a new supervised framework for selecting the local events, instead of burst ranking. In addition, it performs keyword embedding to capture the subtle semantics of microblogs. EVENTDETECTTWITTER [93] framework identifies both short-term and long-term events. It first extracts temporal and textual features that include word frequency, conditional word frequency, inverse trend word frequency, fuzzy representation, and scale time modeling. The features are used to connect data in a graph model. Then a multi-assignment graph partitioning scheme is employed so that each microblog can belong to multiple events. The similarity measure differs based on event type, for short-term events a cross-correlation similarity measure is used whereas for long-term events Riemannian distance is used.

**Statistical techniques.** This category can be divided into two sub-categories. First, **latent variable models.** EXPLOREVENT [388] proposes a pipeline process of event filtering, extraction and categorization. The filtering is based on lexicon matching and binary classification to opt only event-relevant microblogs. Feature extraction then processes relevant microblogs for time expression resolution, named entity recognition, part-of-speech (POS) tagging and stemming, and the mapping of named entities to semantic concepts. The event candidate extraction and grouping phase is based on an unsupervised latent variable model, called Latent Event and Category Model (LECM). For labeling a cluster, the most prominent semantic class obtained

based on the event entities is employed as the event type. PROBEVENT [389] extracts features through POS tagging and named entity recognition, groups microblogs based on a novel unsupervised latent variable model, called LEEV model, which simultaneously extracts events and generates visualizations, and scores candidate events based on the correlation between named entities, dates, locations, and words. OPENEVE [283] extracts temporal, named entities, and POS tags, that are used to filter irrelevant microblogs through an event tagger module based on conditional random fields (CRF). The microblogs are then grouped based on latent variable model and ranked based on the association with event and time. Second, **miscellaneous models** that use different statistical methods, including regression, Markov models, graphical models, and temporal analysis. EYEWITNESS [187] extracts local events and summarizes them using time series analysis of geotagged tweet volumes from localized regions. The framework identifies features as count of data records based on spatial and temporal localities. Then, for a given region, a regression model is learned to predict volume of data versus data spikes as a function of time. Local events are identified when the actual volume exceeds the prediction by a significant amount. SPATIALEVENT [384] forecasts spatio-temporal events using an enhanced Hidden Markov Model (HMM) that characterizes the transitional process of event development by jointly considering the time-evolving context and space-time burstiness of Twitter streams. To handle the noisy nature of tweet content, words that are exclusive to a single event are identified by a language model that has been optimized by a dynamic programming algorithm to achieve an accurate sequence likelihood calculation. SEVENT [390] detects related events, e.g., a sinking boat and an on-going flood in same spatial region. It first extracts textual, spatial, and temporal features. Then, a novel graphical model-based framework, called Location-Time constrained Topic (LTT), is used to express each microblog as a probability distribution over a number of topics. To group related microblogs, a KL divergence-based measure is employed to gauge the similarity between two microblogs. Then, another longest common subsequence (LCS)-based measure is used for the link similarity between two sequences of user microblogs. Sequences are grouped based on spatial, temporal, and topical similarities. BEVEN [80] focuses on discovering breaking events and distinguishing real-life events from virtual events that happen only in the online community. Therefore, it categorizes microblogs based on three features extracted from the hashtags: (1) instability for temporal analysis, (2) meme possibility to distinguish social events from virtual topics or memes, and (3) authorship entropy for mining the most contributed authors. Based on these attributes, an unsupervised technique is used to categorize microblogs into advertisements, memes, breaking events, or miscellaneous.

The rich literature of event detection on microblogs not only contains holistic frameworks that start with raw data and output events to end users, but also specialized pieces of work that are not proposing holistic frameworks, however, it either focuses on one or more of the stages or studies a problem that is utility for event detection. We outline examples for such work in the rest of this section.

**HIEREMBED** [266] focuses on mining topics that are related to events in microblog streams. It presents an unsupervised multi-view hierarchical embedding (UMHE) framework that generates topics with a high accordance to the events from a microblog stream. The framework applies LDA to extract the feed-topic and topic-word distributions. Therefore, for each latent topic, there are two different view features namely the latent word distribution and the relevant feed collection. Then, it applies a novel multi-view bayesian rose tree (Mv-BRT) to refactor the latent topics into a hierarchy. A translation-based hierarchical embedding is formulated to encode the topics and relations in low dense vectors to better capture their semantic coherence. **ET-LDA** [148] proposes a joint model based on LDA to extract the topics covered in the event and tweets, and segmenting the event into topically coherent segments. **ANCHORMF** [124] solves the event context identification problem using a matrix factorization technique by leveraging a prevalent feature in social networks, namely the anchor information. A probabilistic model is built to consider users, events, and anchors as latent factors. An anchor selection algorithm is proposed to automatically identify informative anchors for the model. A Gibbs sampler and a maximum a posteriori (MAP) estimator are proposed to estimate the model parameters. **KEYEXTRACT** [1] focuses on extracting real-time local keywords through a time sliding window approach. For each keyword, a probability distribution over co-occurring places is estimated and used to eliminate spatial outliers. The spatial distribution is updated based on inserting new content and removing old content that is expired from the sliding window. **AUTOSUMMARIZE** [16] focuses on automatic summarization of Twitter posts using three methods, namely, temporal TF-IDF, retweet voting, and temporal centroid representation. The temporal TF-IDF is based on extracting the highest weighted terms as determined by the TF-IDF weights for two successive time frames. The voting method considers the highest number of retweets a post received in the time window. The temporal centroid method selects posts that correspond to each cluster centroid.

#### 4.3.2 Detecting Specific Types of Events

Another major direction of event detection research focuses on detecting specific types of events that have a set of distinguishing information to characterize the event, e.g., keywords. Examples of such events are crime events, earth-

quakes, or traffic jams. Crime events can be described by a set of keywords, while earthquakes are characterized by labeled training data, for example. In general, each event type is described by a set of event-related information. Figure 11(b) shows a framework that utilizes the event-related information along with incoming microblogs data to detect events of a specific type. The framework consists of three main stages: (a) *feature extraction*, (b) *event classification*, and (c) *visualization*. The incoming microblog data is processed to extract temporal [151], textual [151, 203, 211, 387], spatial [151], and sentiment features [362]. Then, the processed data is forwarded to a classification model that uses the event-related information to distinguish relevant data to the event type of interest from irrelevant data. The classification can be performed through two different types techniques, as depicted in Figure 10: (1) learning-based techniques [11, 141, 150, 151, 268, 291, 362], such as support vector machines (SVM) [141, 150, 291] and regression models [268], and (2) lexical techniques [203, 211, 387]. The type of classification is also coupled with the type of provided event-related information that might be keywords or labeled training data. The classified relevant microblogs are directly fed to a visualizer that displays events to end users. The visualizer still uses one of the visualization techniques that exploit aggregation, sampling, or both as presented in Section 4.1. Compared to arbitrary event detection (in Section 4.3.1), this framework replaces the clustering and scoring modules with a classification model that exploits the event-related information to directly group and filter relevant data and reduce noisy output.

This framework drives the major existing work on detecting different types of events. Table 5 summarizes the different stages of each technique. The literature includes two categories of techniques based on the classification stage, as depicted in Figure 10: *learning-based techniques* and *lexical techniques*. We briefly outline techniques of each category.

**Learning-based techniques.** This category includes both supervised and semi-supervised techniques. **EARTHQUAKEVEN** [291] detects earthquake events through Twitter. It uses SVM classifiers and labeled training earthquake data to classify earthquake-related tweets. **CONTRA EVEN** [268] detects controversial events through a regression classification model along with labeled training data on well-known controversial topics, such as Obama Nobel Peace Prize. **WELLEVEN** [11] extracts wellness events from tweets. It extracts features based on a graph-guided multi-task learning model, and classify data based on a novel supervised model that takes task relatedness into account. **TAREVEN** [151] detects social media events that are related to news reports. It extracts features from both tweets and news reports to find out relevant tweets. Then, relevant tweets are splits into positive and negative examples through an EM-based refinement algorithm and final rele-

**Table 5** Summary of techniques that detect specific types of events from microblogs data.

	Event Type	Features	Classification	Event-related Info
TEDAS [203]	Crimes	Textual, Spatial, Temporal	Lexical matching	Keywords
TRAFEVER [211]	Traffic	Textual	Wavlet analysis	Traffic data
DYNKEYGEN [387]	User-defined	Textual	Expectation maximization	Dynamic keywords
EARTHQUAKEVEN [291]	Earthquakes	Textual, Statistical, Contextual	SVM	Labeled data
CONTRAEVER [268]	Controversial topics	Statistical, Sentiment, Linguistic	Regression	Labeled data
WELLEVEN [11]	Wellness	Textual	Novel supervised model	Labeled data
TAREVEN [151]	News-related	Hashtags, Mentions, Replies	Novel semi-supervised model	News articles, historical tweets
STED [150]	User-defined	Textual, Named entities	SVM	News articles
PERSONALIFE [362]	Personal life	Textual, POS tags	Multi-task LSTM model	Labeled data
CROWDEVEN [141]	Bus-related	Sentiment, Named entities	SVM	Labeled data

vance is computed based on textual, spatial, and temporal similarities. The data is then fed to a novel semi-supervised approach for detecting spatio-temporal events from tweets. **STED** [150] proposes a semi-supervised approach that enables automatic detection and visualization of user-defined specific events. The framework first applies transfer learning and label propagation to automatically generate labeled data, then learns an SVM text classifier based on tweet mini-clusters obtained by graph partitioning. Then it finally applies fast spatial scan statistics to estimate the locations of events. **PERSONALIFE** [362] detects personal life events from users' tweets using Multi-Task LSTM model with attention. The system detects whether the tweet is an explicit event, implicit event, or not an event and then detects category of the event from predefined life events categories. **CROWDEVEN** [141] treats each bus-related tweet as a micro-event which can be further analyzed for event type categorization, entity extraction, and sentiment mining. It uses CRF for entity extraction and one-against-one classification strategy with SVM as the classifier.

**Lexical techniques.** **TEDAS** [203] detects crime events based on crime-related keywords along with lexical matching to classify relevant data. **TRAFEVER** [211] detects traffic events using related keywords along with wavlet analysis to classify relevant tweets. **DYNKEYGEN** [387] proposes a semi-supervised solution based on expectation maximization mechanism that leverages word information to infer tweet labels. The candidate tweets are selected based on a set of keywords, which are generated and updated dynamically based on word importance score that changes over time.


#### 4.3.3 Event Analysis

Unlike event detection techniques, where new events are outputs, event analysis techniques take an event as an input and analyze its data in different ways. In specific, event analysis work focuses more on providing exploration tools for known predefined events rather than detecting new events that are not known beforehand. For example, the Syrian revolution is a long-term event that is known beforehand with a

set of features such as keywords and locations. So, an event analysis module is interested more in analyzing data of this well-known event rather than discovering a new event that is not known beforehand. Another example is King Tut festival in Hayward, California. This is a short-term event that is known beforehand with a set of keywords, locations, and a time period. Again, an event analysis module focuses more on analyzing data of this event without discovering any new events. Thus, most of existing event analysis work follows a simple framework that is depicted in Figure 11(c). The framework has two stages: (a) *filtering* and (b) *visualization & analysis*. The filtering stage employs simple filters on different attributes, e.g., keywords [28, 237, 238, 326], spatial [28, 326], and temporal [28], to extract relevant microblogs to a certain event, e.g., Hurricane Sandy. Then, extracted data is forwarded to a rich visualization module that enables end users to analyze event data based on multiple views, e.g., map view, aggregate views, frequent keywords, influential users, timeline view, sentiment view, or individual microblogs. The features of analysis and visualization views are highly variant and depend on the application and the analysis purpose. The rest of this section presents examples of event analysis applications in the literature.

**TWEETTRACKER** [326] provides an event analysis framework for long-term events, such as Arab Spring uprisings, Occupy Wall Street, and US presidential elections. Users can define new *jobs* to define new events to analyze. Events data are filtered based on keywords, locations, and usernames features. Newly incoming data is tracked based on the event features for a long term. Then, the collected and new data is visualized based on a timeseries view, geographic map view, trending keywords view, entities view, and individual tweets view. TweetTracker has currently collected 3.2 billion tweets, and it is adding new  $\sim 700,000$  tweet every day. **TWITTERPOLITICALINDEX** [335] is a social media index for US presidential elections co-developed by Twitter and Topsy Labs, a social search and analytics company that owns all Twitter data and is acquired by Apple Inc. [26]. The index visualizes





**Fig. 12** An overview of recommendations using microblogs literature.

tweets relevant to US elections based on political party, sentiment, locations such as states and counties, and timeline view. **TWITINFO** [237, 238] provides a timeline-based event analysis framework that allows users to define events based on relevant keywords. Then, the system collects relevant tweets, categorizes them based on sentiment, and organize them in timeline and map views in both aggregate and individual data records forms. The system addresses scalability problems that are associated with analyzing and visualizing such large number of data records. **STEVENT** [28] analyzes events based on three aspects. First, how topic initiators influence popularity of the topic. Second, the impact of geography on popularity by partitioning the Twitter network according to regional divisions and studying the behavior of popular and non-popular topics. Third, the effect of topology and the dynamics of topic spread on popularity.

#### 4.3.4 Events and Microblogs Aggregate Queries

Several aggregate querying techniques (Section 2.2.2) have been motivated by detecting events from large-scale microblogs data [50, 224, 304]. This includes detecting highly frequent [304] and highly trending [224] keywords that identify popular topics among users, and detecting highly correlated keywords with different locations [50] that identify localized topics of people interests. Such techniques can be used as scalable infrastructures to detect events from large amount of data. However, the core research methods focus on indexing and query processing on a large scale, which lies in a lower-level of the data analysis stack compared to techniques that are reviewed in this section.

### 4.4 Recommendations Using Microblogs

Microblogs represent a rich and up-to-date source for user-generated content. Therefore, they are appealing for several recommendation applications to extract up-to-date user preferences, which is essential to recommend relevant items. Although recommendation applications that exploit microblogs data are diverse, being an up-to-date source for user preferences is the common theme that links all of them. From a data management perspective, having such large and highly changing data as a source of preferences introduces significant challenges in updating recommendation models in practice. In fact, this has triggered deep research discussions in the data management community on the ability to

support recommendation models efficiently in data management systems [111, 181, 191, 192, 294, 296, 297, 298, 363]. This clearly makes a transformative shift towards a new generation of recommender systems that should be able to recommend relevant items accurately through updating models much more efficient than their ancestor generations of recommender systems. Microblogs data plays a major role as a source of preferences for this new generation of recommender systems and the data management research community is in the heart of addressing their challenges.

Figure 12 depicts a high-level overview about recommendation techniques using microblogs. The literature includes two major recommendation problems, recommending content and recommending friends, in addition to a set of diverse miscellaneous recommendation applications. Such applications are as diverse as recommending news items, products, question answers, events, and scholarly information. The rest of this section highlights each category.

**Recommending user-generated content.** One of the major recommendation problems that is widely studied in the literature is recommending user-generated content, such as recommending other microblogs to read, hashtags to search, and mentions to post. **NETREC** [13] recommends tweets that are not visible to the user, e.g., posted by friends of friends or further, by exploiting the social network, content, and retweet analysis. The importance of invisible tweets is initially estimated by the social distance. Then, both content and user analysis is performed to rank highly relevant users to recommend their tweets. Content analysis is based on textual analysis using bigrams, while user analysis is based on comparing timelines and mutual retweets. **BLGREC** [173] leverages and combines the user's location, social network feeds, and in-app actions to infer the user's interest and develop a personalized recommendation model. A user's feed is then made up of recommended content, including trending news, social network feeds, and social content, on either local or global scales based on the user spatial interests. **TWIMER** [235] performs tweet recommendations based on formulating a query based on the user's interest profile to probabilistic language models. Then, irrelevant and near-duplicate tweets are discarded using threshold-based filtering, locality sensitive hashing, and tweet freshness. **SIMGRAPH** [76] is a scalable recommendation model based on a similarity graph that induces the mutual interest among users by analyzing retweets. The probability of a certain user to like incoming microblogs are estimated based on a propagation model that aggregates top-k tweets and recommends them to the user. **CMPREC** [66] tackles a more fundamental functionality in microblogs recommendation through comparing two approaches to compute similarity among microblogs with brief content: a topic-based approach and WordNet corpus-based approach. The study

shows the superiority of WordNet corpus to catch similarity between brief textual content of microblogs.

Hashtag and mention recommendation is another content recommendation task that is popular in the literature, so users can easily search for their topics of interest. **EMTAGGER** [90] is a trained model for learning word embeddings and assigning hashtags with the trained embedding system. **COGREC** [186] proposes two cognitive-inspired hashtag recommendation techniques based on the Base-Level Learning (BLL) equation:  $BBL_{I,S}$  and  $BBL_{I,S,C}$ . BLL accounts for the time-dependent decay of item exposure in human memory, once with the current tweet content ( $BBL_{I,S,C}$ ) and once without ( $BBL_{I,S}$ ). **MRTM** [204] is a personalized hashtag recommendation model based on collaborative filtering and topic modeling. It integrates user adoption behaviors, user hashtag content, and contextual information into a joint probabilistic latent factor model to recommend hashtags to users. **MENREC** [221] addresses the problem of using both texts and images of microblogs for mention recommendation. A cross-attention memory network is proposed which considers the content of a tweet, interests of the user, and interests of the author to recommend a user to be mentioned for a certain tweet.

**Recommending friends.** Another recommendation problem that is widely studied in the literature by researchers from academia and industry (specifically Twitter) is recommending users to follow to expand and enhance the social graph connected components. **TWIT-TOMENDER** [134] started exploiting the real-time nature of microblogs by dynamically profiling the users through their recent microblogs. Then, collaborative filtering techniques are used to recommend users with similar interests. **FUREC** [353] tackles the problem from a different angle and recommends top-k users who will likely interact with microblog posts of a certain focal user. It uses the existing follower network and creates a new network based on retweets and mentions, then a single hybrid network is composed to recommend the new users. The problem is also studied and realized by Twitter Inc. [131, 132, 300], where substantial contributions in enriching connections between Twitter users are made. The *Who to Follow* (**WTF**) project [131, 300] started to recommend users to follow and enrich Twitter social graph. The core of the system is the *Cassovary* in-memory graph processing engine and a novel technique for performing user recommendation, called *Stochastic Approach for Link-Structure Analysis* (SALSA). SALSA constructs a bipartite graph that include the user's circle of trust on the left side, while the right side includes users who are followed by the users in the left side. Then, this bipartite graph is traversed and ranking scores are assigned, on which users are recommended accordingly. Approximation algorithms are also provided in the second generation of WTF to reduce the complexity of processing

hundreds of millions of users. To exploit the time aspect of Twitter data, they added MagicRecs [132] that recommends users who are followed by friends within certain temporal constraints. To expand Twitter's recommendation services, they added content recommendation through GraphJet [300] that is based on a bipartite graph similar to the one maintained in WTF system, except the right side models actual user tweets. A random walk on this graph with a fixed probability of reset outputs a ranked list of vertices that represents the tweets to be recommended to the user.

**Miscellaneous recommendation applications.** A significant portion of the literature is recommending miscellaneous items/users, where the common theme is using microblogs as an up-to-date source for user preferences. **NEWSREC** [267] recommends news items re-ranking based on user preferences extracted from tweets. The user tweets and RSS news feeds are both processed by a preference extraction module that finds out common keywords in both. Then, these keywords are used to promote relevant news in the news feeds timeline, so important news appear early to users. **METIS** [385] recommends products based on detecting purchase intent from microblogs data in near real time fashion, combining their model with the offline traditional models that are similar to e-commerce website recommendations, e.g., Amazon. Such exploitation of real-time user-generated data has enhanced the effectiveness of product recommendation models. Another recommendation model that handles cold-start problem for product recommendation exploiting user-generated microblogs is **CSPR** [386]. CSPR uses data from microblogging users with no historical purchase records to map users attributes extracted from microblogs into feature representations learned from e-commerce websites. Thus, given a microblogging user, a personalized ranking of recommended products can be generated to overcome the cold-start problem. **EVENREC** [231] exploits geotagged microblogs to recommend events from Eventbrite, a popular event organization website. The extracted events depend on microblogs locations that are fed to item-user models. This work is orthogonal from event recommendation in event-based social networks [86, 125, 347], e.g., Meetup.com, which has different nature compared to microblogging platforms, and thus it is beyond the scope of this paper. **CRAQ** [311] recommends potential answers to a posted question through selecting a group of potential authority users who are selected based on their topically relevant microblogs. Then, the candidate group is iteratively filtered by discarding non-informative users, and top-k relevant microblogs are determined as potential answers. **JURY** [53] recommends potential authority users who are able to answer a given question. It adapts a probabilistic model that selects a set of users so that the probability of having wrong answer is minimized. **SCHREC** [365] recommends scholarly information through microblogs posted





Fig. 13 An overview of microblogs automatic geotagging literature.

by researchers who post about their latest findings or research resources. Two neural embedding methods are proposed to learn the vector representations for both users and microblogs. Recommendation is made by measuring the cosine distance of a given microblog and user.


#### 4.5 Automatic Geotagging

Geo-locations are heavily exploited in several microblogs applications, such as localized event detection [2], geo-targeted advertising [256], local news extraction [293], user interest inference [128], and finding local active users [162]. With all such importance of geo-location data in microblogs applications, still the majority of microblogs are not associated with precise location information. In fact, a small percentage ( $< 4\%$ ) of popular microblogging data, e.g., Twitter, is associated with locations sourced from user devices. This triggered a need to associate location information with more microblogs data automatically to exploit as much microblogs as possible in location-aware applications. However, traditional geotagging techniques are limited for enriching microblogs location data due to the brevity of microblogs textual content. Such brief text contains a lot of abbreviations and noisy words that make it hard for named entity recognizers to extract accurate places and locations. In this section, we give an overview about new techniques in the literature that are designed to extract locations from microblogs data. Although traditional geotagging techniques purely depend on linguistic analysis to extract locations, recent geotagging techniques on microblogs go beyond this to identify top-k locations for both users and data records, as elaborated later in this section. This recent paradigm overlaps and makes use of certain indexing and query processing techniques from the data management literature. Thus, automatic geotagging on microblogs is leaning toward making more use of data management infrastructures in addition to the linguistic techniques.

Figure 13 classifies the literature at high-level into techniques that use a single microblog record at a time for geotagging and techniques that use collections of microblogs. Figure 14 shows frameworks for the two types of techniques. In fact, most of microblogs geotagging techniques in the literature depend on classification models to assign location(s) to one microblog at a time. Figure 14(a) shows a geotagging framework that is induced based on existing work on microblogs. The framework consists of two stages. The first



(a) Geotagging for Single Microblogs



(b) Geotagging for Collections of Microblogs

Fig. 14 Frameworks for microblogs automatic geotagging.

stage is a feature extraction stage that extracts keywords and named entities places from the brief textual content of training microblogs. The extracted keywords and places are used to train the classification model. For each incoming microblog, the classifier assigns a location based on its textual content features. The location classification is performed through different models, such as probabilistic models [199, 272, 290], multinomial naive bayes [140], lexical matching [155], ensemble of statistical and heuristic classifiers [234], pure place entity recognition [210], gazetteer verification [12, 97], and matrix factorization [96].

A common problem in these techniques is the trade-off between error distance and classification precision. The precision is significantly dropped down for the practical margins of error distance, which represents the distance between actual location and predicted location. For example, with error within 100 meters, the precision ranges from 10-20% for different techniques. On increasing the error distance to 30 KM, the precision is raised to 60-80%. With 100+ KM error distance, the precision reaches 80-90%. Therefore, accurate location prediction provides very low precision where over 90% of data is mistakenly geo-located. On another hand, the significant increase of error distance makes predicted locations not useful for practical applications.

To overcome this problem, a state-of-the-art technique [199] proposed to process microblogs as collections instead of individual records as depicted in Figure 14(b). The technique is collecting all microblogs of each user as one collection and perform exact and fuzzy location extraction on them to identify all possible locations for this user. Then, top-k locations for each user are predicted and identified as the most likely locations where the user is posting microblogs. When a new microblog arrives, a set of top-k locations are extracted from microblogs content and meta-

data. Then, the  $k$  microblog locations and the  $k$  user locations are fed into a location refinement module that predicts the final top- $k$  microblog locations. This technique has shown tremendous enhancement in prediction precision and recall (95+%) within 100 meters error distance, which is the threshold for accurate location prediction.

#### 4.6 Other Analysis Tasks

The reviewed analysis tasks in previous sections represent the major high-level analysis tasks on microblogs that are of interest to the community of data management and analysis researchers. However, the microblogs literature and applications are so rich to enumerate all possible analysis types or techniques. In fact, other analysis are sporadically addressed on microblogs in both (1) academic community, such as news extraction [267, 293], topic extraction [142, 201, 275], summarization [46, 98, 121], situational awareness [288, 303], and resource needs matching [40, 41], and (2) industrial community, such as geo-targeted advertising [256] and generic social media analysis [323, 383]. Yet, the reviewed literature represents the main performed high-level analysis tasks that span a wide variety of interests, applications, and novel research challenges as well as future research opportunities.

### 5 Conclusions and Future Directions

This paper has provided a comprehensive review for major research work and systems for microblogs data management and the corresponding analysis tasks. The paper categorized the literature into three parts: data indexing and querying, data management systems, and data analysis, where each part is further divided into sub-topics. The data indexing and querying part has reviewed microblogs query languages, individual indexing and query processing techniques, and main-memory management techniques. The systems part has reviewed characteristics of different genres of big data systems, e.g., batch processing systems, big data indexing systems, and key-value stores, in terms of their adequacy to handle microblogs query workloads. It has also discussed challenges and solutions that are provided through these systems for fast data, highlighting their potential limitations to handle certain microblogs applications. The data analysis part provided a detailed roadmap for the major analysis tasks that are directly or indirectly make use of the data management literature: visual analysis, user analysis, event detection and analysis, recommendations, and automatic geo-tagging. For each task, we presented a generic framework, when applicable, that is induced from major techniques in the literature and drives main research innovations for this task. In addition, we classified the literature based on the

major component of this framework to provide better understanding for different techniques and highlight existing challenges and future opportunities in this research direction.

The rich literature of research on microblogs data faces several big challenges and is still rich with opportunities on different fronts. In terms of data management, there are several research opportunities in real-time indexing, query optimization, and system-level integration. For real-time indexing, the microblogs literature does not provide a comprehensive study for supporting spatial-keyword queries on real-time data. This has not been studied before either in existing spatial-keyword querying techniques [54, 55, 56, 63, 64, 65, 71, 74, 75, 129, 196, 200, 206, 218, 222, 232, 233, 343, 372, 374] that focus on traditional static datasets, e.g., restaurants, or in existing microblog indexing that considers the spatial-keyword combination only in aggregate queries that retrieve frequent or trending keywords [50, 224, 304]. Existing specialized systems for microblogs supports two separate indexes, a keyword index and a spatial index, as a generic option that allows supporting various queries with few system assets. However, it is not clear how much performance is lost compared to hybrid indexing strategies. Quantifying such performance losses will enable better understanding for parameters that control querying performance on different indexes, which in turn will allow optimizing each index. Such understanding contributes to developing query optimization models for real-time data management as elaborated below. In addition to spatial-keyword queries, social information is still underutilized in supporting scalable personalized queries on real-time microblogs data. Although there exist few techniques that exploit this information [205, 352], these queries still suffer from inherent scalability limitations due to the overhead of supporting hundreds of millions of users while sustaining efficient data digestion, indexing, and querying in real time.

Despite the richness of exploring real-time indexing on microblogs, there is almost no work on studying the implications of these novel indexing techniques on query optimization models. For example, the traditional selectivity estimation models assume relatively stable index content that is dominated by read operations and encounter much less write operations. This assumption does not hold on microblogs real-time indexes that have highly dynamic content. In addition, microblogs indexes are segmented based on temporal and spatial ranges, which gives a room for estimation model compression to serve such excessive amount of data with limited storage requirements. In general, the implications of new real-time indexing techniques on traditional query optimization models need to be revisited on microblogs.

Integrating all existing and future techniques of microblogs data management in end-to-end systems is a must to widen the impact of existing data management technology in microblogs applications. Recently, extensive efforts

started to develop end-to-end systems to support microblogs data as elaborated in Section 3. However, there is still a gap between the available research techniques and their applicability for system-level integration. For example, existing aggregate queries techniques face challenges to be integrated with microblogs systems as they cannot be supported efficiently using existing indexes and require separate indexes. This is not favorable from system point of view to maintain additional indexes. So, new ways need to be innovated to integrate aggregate data structures within index cells of non-aggregate queries at a system level. Another example is flushing policies that are way developed in separate indexes than the ones supported at system level. This due to a lack of integration techniques that allow flexible flushing policies while maintaining the real-time performance.

In terms of data analysis, there are several untackled challenges on two levels: enhancing the analysis modules and integrating them with microblogs systems to extend their functionality for enriching and facilitating microblogs applications. There are many examples that can be induced from the reviewed literature. We will highlight few of them in different analysis areas. First, developing a unified event detection framework that allows users to express different types of event-based queries. Such framework will exploit the rich literature of event detection and analysis on microblogs to provide common utilities that allow effective and efficient event queries. Second, real-time geotagging of microblogs data. Although recent work started to tackle this problem [96, 97], there are still challenges in reducing the geotagging time due to the high computational cost of this task. Achieving the goal of attaching locations to microblogs as they come will widely impact a plethora of location-aware applications that are built on top of microblogs. Third, integrating the rich literature of user analysis techniques with the scalable data management infrastructures, e.g., indexes and query processors, in microblogs systems. Such integration will allow a variety of user-centric applications to be supported at scale. Fourth, developing a unified recommendation framework that exploits microblogs data and allow users to express a variety of recommendation queries flexibly. Such framework will serve a diverse set of applications that are reviewed in Section 4.4. The envisioned unified framework could exploit existing work on supporting generic recommendation queries in data management systems [198, 295, 297, 298].

In addition to enhancing different analysis modules, there is a dire need to integrate such rich literature of analysis techniques with microblogs data systems to widen the impact of microblogs research in a practical sense. Such integration will have tremendous impact of a plethora of applications that benefit the society, the research community, and business applications, including public health, disaster response, public safety, and education. The feasible

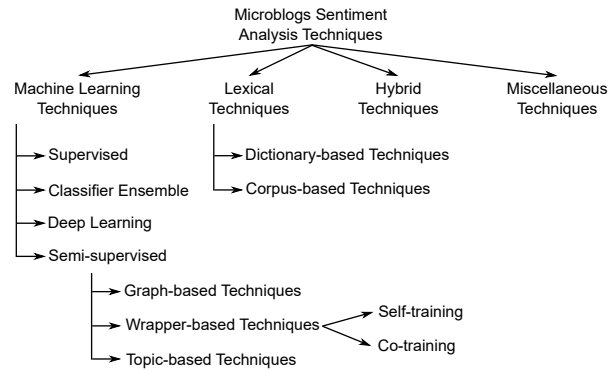


Fig. 15 An overview of microblogs sentiment analysis literature.

way to achieve such goal is abstracting different analysis tasks on microblogs into basic building blocks that can be supported in microblogs systems, inspired by SELECT-PROJECT-JOIN building blocks in SQL database management systems. Such task is huge and shall be started with developing generic frameworks for different analysis tasks, as discussed earlier for event detection and recommendations as well as provided throughout Section 4.

## Appendix

### A Orthogonal Research Directions

This appendix gives an overview about sentiment and semantic analysis in microblogs as an example of an orthogonal research direction, from the natural language processing literature, that does not exploit much of the data management infrastructures. The appendix highlights the differences of new techniques on micro-length data with the corresponding techniques on traditional long data. For detailed surveys about these topics, the reader can refer to [82, 119].

#### A.1 Sentiment Analysis

Sentiment analysis automatically discovers the polarity of feelings expressed in a chunk of text, e.g., a citizen posts positive or negative opinions about certain election candidate. Traditional sentiment analysis techniques make use of the microblogs brevity to enhance the classification accuracy of user sentiment. As reported in [45], using a traditional sentiment classification technique on microblogs boosts the accuracy up to 10% higher for binary sentiment. This boost is an absolute advantage of the content brevity that makes it less confusing and more decisive to catch positive and negative feelings in user-generated content. However, microblogs brevity introduces both challenges and differences compared with traditional data. For example, feature extraction is more challenging due to lots of abbreviations and noise, e.g., extracting meaningful keywords is harder. In addition, compared with traditional data where sentiment is analyzed on three different level, document level, sentence level, and entity level, microblogs short content mostly limits the sentiment scope to a single sentence or a single entity that represents the whole micro document. Moreover, microblogs come with additional advantageous features that were not available in traditional data, such as links, user information, and their interactions with different topics. Thus, the sentiment analysis research on microblogs has addressed a wider variety of challenges compared

with traditional sentiment analysis. In this section, we give an overview about this rich literature.

Figure 15 depicts an overview of the microblogs sentiment analysis literature. The major techniques can be categorized into four main categories, namely, *machine learning techniques*, *lexical techniques*, *hybrid techniques*, and *miscellaneous techniques*. The machine learning techniques represent the majority of techniques in the literature. It could be further categorized into four sub-categories as depicted in Figure 15, namely, *supervised*, *classifier ensemble*, *deep learning*, and *semi-supervised*. The first sub-category of techniques use supervised machine learning, i.e., traditional classifiers [4, 35, 67, 89, 202, 219, 254, 255, 274, 279, 317, 351, 394]. The differences among these techniques are the classifier type, stages, and features used to distinguish sentiment. The major used classifiers are support vector machines (SVM) [4, 5, 30, 34, 38, 85, 89, 122, 133, 159, 163, 171, 180, 248, 261, 273, 305], (multinomial) naive bayes (MNB and NB) [30, 34, 122, 133, 247, 261], k-nearest neighbor (kNN) [30, 84], MaxEnt [94, 122, 180], random forest (RF) [91, 318], logistic regression [35, 202, 394], and AdaBoost [185]. The used features include different types of language-based features such as unigrams [5, 34, 122, 163, 247, 261], bigrams [34, 122, 261], trigrams [261], n-grams [84, 180, 185, 248], and POS tags [5, 38, 122, 180, 185, 247, 248, 261], microblog-specific features [185, 248] such as retweets [38], hashtags [34, 38, 163], emotions [34, 38, 163, 180], links [34, 38], and other features such as punctuation-based [5, 84, 180, 248], pattern-based [5, 34, 84, 163, 180, 248] and semantic-based [180, 247].

To enhance the classification accuracy, techniques of the second sub-category ensemble multiple classifiers [61, 72, 77, 81, 88, 136, 171, 185, 194, 208, 244, 265, 316, 364]. The set of used features is almost identical to the single classifier techniques, while the used classification algorithms are overlapping but not identical. In specific, SVM [81, 136, 265], NB [72, 136], MNB [81], logistic regression [81, 136, 208], and AdaBoost [185] are still used, while new classifiers are also introduced such as neural models [88, 136, 364] and bayes network [136]. A third sub-category is deep learning, which is an emerging field in machine learning. In the past few years, deep learning is getting increasing popularity and many learning problems migrated to deep learning frameworks. Deep learning offers a black box of neural networks that are trained with huge amounts of data that offer better accuracy over traditional classifiers. In the case of microblog platforms, huge amount of data is generated daily, which has motivated the use of deep learning techniques for sentiment analysis on microblogs. Existing deep learning techniques is exploited in short textual contexts in two-step fashion [36, 62, 73, 88, 92, 135, 144, 145, 164, 264, 278, 287, 320, 321, 337, 342, 360]. It first learns word embeddings then it applies them to produce representations for the text sentiment.

The main limitation of all supervised techniques, either with single classifier, multiple classifiers, or deep learning models, is the sensitivity to dataset size. For increasing their performance, there is a high reliance on the manually annotated labels which is extremely expensive. To alleviate this problem, distance supervision has been employed where the labels are generated based on the emoticons and hashtags [257, 309]. However, this approach did not perform well. This encouraged a fourth sub-category of semi-supervised techniques to rise. The semi-supervised techniques rely on both a small set of manually annotated data as well as unlabeled data to train the model. They can be further divided into three main types as depicted in Figure 15: *graph-based*, *wrapper-based*, and *topic-based* techniques. The graph-based techniques [79, 312, 319, 344] use label propagation to label the unlabeled training data based on the similarity metric between two nodes in the graph. Then, a classifier is trained and used as previous techniques. The wrapper-based techniques [43, 44, 213, 214, 381] rely either on self-training [43, 44, 381] or co-training [213, 214]. In both types, the classification process is an iterative process, starting with the initial labelled data, classify the other unlabeled data, and use the high confident ones in the next iteration of the classification till all data are labelled or




Fig. 16 An overview of microblogs semantic analysis literature.

it hits the maximum number of iterations. The difference between the self-training and the co-training is that in self-training only one classifier is used whereas in the co-training two classifiers with different feature sets are used to provide two different views for the data. The more confident classification within the two classifiers is chosen to be within the labelled data in the next iteration. The last semi-supervised types is topic-based techniques [10, 126, 139, 165, 241, 301, 356], where topic information is extracted with sentiment analysis simultaneously under the observation that the context of the content affects the sentiment.

The second major category in Figure 15 is lexical techniques [29, 143, 149, 207, 259, 276, 299, 322, 340], where a predefined list of positive and negative words is employed to classify the sentiment of the new microblog. There are two main sub-categories in lexical techniques, namely, *dictionary-based* and *corpus-based*. The dictionary-based techniques [29, 143, 149, 207, 299, 340] use dictionaries as lexical resources and approximate lexical matching techniques are used to account for microblogs noise and abbreviations. The corpus-based technique [276] uses statistical or semantic methods to match incoming data with existing lexical resources. The third major category in Figure 15 is hybrid techniques [109, 117, 175, 177, 182, 189, 377] that combine both machine learning and lexical methods to detect microblogs sentiment. These techniques use lexical terms either to train a machine learning model or to filter data in a first stage that is fed to a classifier for further processing on a second stage.

Other miscellaneous techniques are proposed for microblogs sentiment analysis. **CONSENT** [183] uses concept analysis to determine sentiment based on associated topic. **APPSSENT** [184] uses appraisal terms to outperform supervised techniques. **SOCIOSENT** [147] uses sociological information in the supervised learning process to improve the performance. **CHINESENTIMENT** [366] proposes a rule-based model for analyzing sentiment features of different linguistic components, and a corresponding methodology for calculating sentiment using emoticon elements as auxiliary affective factors.

## A.2 Semantic Analysis

Semantic analysis is a popular analysis task that is widely used in microblogs literature for different applications, such as topic modeling [339], knowledge extraction [308], community detection [190, 310], stance detection [391], sentiment analysis [182], event analysis [47, 260], effective microblog retrieval and ranking [380], and user recommendations [108]. This task automates discovering the meanings of a chunk of text by discovering semantic relationships that relate to real-world entities, such as places, persons, and organizations. For example, a text like *Trump to campaign for Cindy Hyde-Smith in Mississippi* can be related to two persons, *Trump* and *Cindy Hyde-Smith*, and one place, *Mississippi*. This relatedness connects the input text to a predefined set of semantic concepts or categories that are commonly extracted from human-contributed content, such as Wikipedia, or professionally-maintained ontologies, such as FOAF and DBpedia ontologies. Such type of analysis used to be performed on long chunks of text, e.g., news articles, blog posts, or web documents. However, in microblogs, the textual content is very short and contains a plenty of abbreviations, informality, and noisy terms. Such brevity hurts the performance of traditional semantic analysis techniques, as shown in [242], that depend on lexical matching and search-based retrieval in, for example, Wikipedia concepts.

To overcome the brevity problem, a general theme of semantic analysis research on microblogs is exploring different ways to enrich the microblogs short textual content to enable accurate semantic relations discovery. Existing techniques in the literature can be categorized into four categories, as depicted in Figure 16, based on the source of enrichment through: *external documents-based techniques*, *machine learning-based techniques*, *hashtag-based techniques*, and *lexical techniques*. Techniques of the first category [57, 115, 339, 350] depend on linking the microblog short document to external long documents, e.g., news articles or web documents, which allow traditional semantic analysis techniques to be applied with high precision. **TOSEM** [339] performs semantic enrichment based on explicit web links that are included in the microblog to associate the linked web document. Then, it extracts both named entities and top-k terms from the web document to be appended to the microblog as auxiliary terms. **NWSEM** [57] identifies online news articles that are related to the microblog post in order to extract named entities and include them in the user profile as semantic tags. **USRSEM** [350] explores semantics of user interactions, specifically retweets and links that are embedded in tweets, and their role in inferring notions such as quality of user relationships, trust, and other attributes of user relationships. This could be applied to re-ranking microblogs based on importance, user interest, quality, etc. **DISEM** [115] maps microblog posts to Wikipedia articles, then use the Wikipedia ontology for semantic categorization.

The second category is machine learning-based techniques [112, 146, 215, 216, 242, 285, 310, 371] that use either: (1) clustering to group different related microblogs and use their collective content to semantically label the whole cluster, or (2) classification that exploits annotated training data as an external source of information to learn different semantic classes of new microblogs. **TRSEM** [371] introduces a novel transfer learning approach, namely transfer latent semantic learning, that utilizes a large number of tagged documents with rich information from domain-specific sources to discover latent semantics of the abbreviated text. **ACCSEM** [146] clusters related microblogs and use the collective content of each cluster to automatically assign semantically meaningful labels. The semantic labels are solicited from external knowledge sources, such as Wikipedia and WordNet, based on informative fragments parsed from microblogs contents. **ADSEM** [242] uses SVM and naive bayes classifiers to enhance the precision of mapping tweets to Wikipedia-based concepts. For this, it obtain an initial ranked list of candidate concepts through lexical matching, language modeling, and traditional techniques. Then, annotated training data is used to train classifiers that further classify microblogs based on different feature vectors to the correct semantic category, which significantly boosts both precision and recall. **NOMSEM** [215] uses SVM classifiers to identify nominal predicates in tweets. Then, a factor graph for each nominal predicate is constructed and joined with graphs of other predicates so their semantic arguments are jointly resolved. **COMSEM** [310] clusters related microblogs to detect user groups within sub-communities. Then, a probabilistic model is employed to measure the semantic, or topical, coherence of the user group and filter out non-coherent groups. **GEOSEM** [313] clusters microblogs based on spatial, temporal, and semantic features, including LDA topics, to evaluate the performance of combining different features in retrieving insights from microblogs data. **ST-SRL** [216] proposes a semi-supervised self-training approach that utilizes a small training dataset to label unlabeled tweets in an iterative way to increase the training dataset size. Labeled data records with highest confidence from two different labelers are used to enhance the classification accuracy in the following iterations. **VECSEM** [112] has performed a unique study that explores the effect of changing microblog-specific semantic representation features on the performance of semantic prediction. It studies a set of 13 microblog-specific prediction tasks to understand both textual and social aspects of different representations.

The third category is hashtag-based techniques [37, 263, 345]. Hashtags are user-defined tags included in microblog posts, which in-

dicates the discussed topics and enable posts related to the same topics to be searched easily. These hashtags are used in different ways to discover latent semantic content in microblogs. **SMOB** [263] uses hashtags as seeds to generate potential related links to web documents and ontology entries from both FOAF and DBpedia ontologies. Then, relevant semantic relations to the discovered entities are appended to the microblog. **ENTSEM** [37] enriches semantics through retrieving a ranked list of the top-k hashtags that are relevant to a user's query and segments them into relevant individual words. Then, it retrieves a set of Wikipedia articles that are related to tweet text, hashtags, and segmented hashtags. **HGTM** [345] introduces a new topic model through using hastags to determine semantic relatedness to each other through a graph structure. A graph of hashtag relatedness is constructed using probabilistic models, then related hashtags are grouped in coherent topics.

The fourth category is lexical techniques [8, 9, 47, 83, 108, 179, 260, 271, 308, 393] that improve traditional techniques that are used for long text to be effective for short textual microblog content. **INDUCSEM** [271] induces semantic entities using lexical pattern-based approach that match microblog text with seed keywords of each semantic category, e.g., food, sports, or vehicles. **KNOSEM** [308] uses lexical resources that include corpus and POS-tagged terms to label tweets with semantic frames for knowledge extraction purposes. **EVESEM** [260] analyzes word co-occurrences to discover relationships among word pairs. Then, such features are used to calculate the pairwise similarity of tweets for event detection purposes. **HIVSEM** [9] uses lexical matching techniques to analysis the presence of an HIV prevention drug on Twitter. **PLCSEM** [179] extracts place semantics through LDA topic modeling from a collection of microblogs to abstract their content through probabilistic models into a set of coherent topics. Then, the extracted place semantics is analyzed for temporal changes, e.g., a sports arena could evolve over time to be a place for concerts and exhibitions. **MONSEM** [47] uses lexical matching to match microblog content with semantic knowledge bases to monitor unexpected events on social media. **LIKSEM** [8] uses semantic user attributes to enhance link prediction among social media users. **RETSEM** [393] uses lexical semantic features to enhance microblogs retrieval performance. **TRISEM** [83] uses semantic relevance to filter tweets based on Wikipedia concepts and trigrams. **RECSEM** [108] uses semantic relatedness to recommend users to follow. It links users to Wikipedia through lexical and disambiguation algorithms, then similar users are recommended.

## References

1. Hamed Abdelhaq, Michael Gertz, and Ayser Armiti. Efficient Online Extraction of Keywords for Localized Events in Twitter. *GeoInformatica*, 21(2):365–388, 2017.
2. Hamed Abdelhaq, Christian Sengstock, and Michael Gertz. EventTweet: Online Localized Event Detection from Twitter. In *VLDB*, 2013.
3. Youcef Abdelsadek, Kamel Chelghoum, Francine Herrmann, Imed Kacem, and Benoît Otjacques. Community Extraction and Visualization in Social Networks Applied to Twitter. *Information Sciences*, 424:204–223, 2018.
4. José Abreu, Iván Castro, Claudia Martínez, Sebastián Oliva, and Yoan Gutiérrez. UCSC-NLP at SemEval-2017 Task 4: Sense n-grams for Sentiment Analysis in Twitter. In *SemEval-2017*, 2017.
5. Apoorv Agarwal, Boyi Xie, Iliia Vovsha, Owen Rambow, and Rebecca Passonneau. Sentiment Analysis of Twitter Data. In *LSM@ACL*, 2011.

6. Manoj K Agarwal, Divyam Bansal, Mridul Garg, and Krithi Ramamritham. Keyword Search on Microblog Data Streams: Finding Contextual Messages in Real Time. In *EDBT*, 2016.
7. Manoj K. Agarwal, Krithi Ramamritham, and Manish Bhide. Real Time Discovery of Dense Clusters in Highly Dynamic Graphs: Identifying Real World Events in Highly Dynamic Environments. *PVLDB*, 5(10):980–991, 2012.
8. Cherry Ahmed and Abeer ElKorany. Enhancing Link Prediction in Twitter using Semantic User Attributes. In *ASONAM*, 2015.
9. Zheng Ahn, Margaret McLaughlin, Jinghui Hou, Yujung Nam, Chih-Wei Hu, Mina Park, and Jingbo Meng. Social Network Representation and Dissemination of Pre-Exposure Prophylaxis (PrEP): A Semantic Network Analysis of HIV Prevention Drug on Twitter. In *Springer SCSM*, 2014.
10. Aman Ahuja, Wei Wei, and Kathleen M Carley. Microblog Sentiment Topic Model. In *ICDM Workshops*, 2016.
11. Mohammad Akbari, Xia Hu, Liqiang Nie, and Tat-Seng Chua. From Tweets to Wellness: Wellness Event Detection from Twitter Streams. In *AAAIz*, 2016.
12. Hussein Al-Olimat, Krishnaprasad Thirunarayan, Valerie L. Shalin, and Amit P. Sheth. Location Name Extraction from Targeted Text Streams using Gazetteer-based Statistical Language Models. In *COLING*, 2018.
13. Noor Aldeen Alawad, Anagnostopoulos Aris, Leonardi Stefano, Mele Ida, and Silvestri Fabrizio. Network-aware Recommendations of Novel Tweets. In *SIGIR*, 2016.
14. Zeynep Zengin Alp and Sule Gündüz Ögüdücü. Influential User Detection on Twitter: Analyzing Effect of Focus Rate. In *ASONAM*, 2016.
15. Nasser Alsaedi, Pete Burnap, and Omer Rana. Can We Predict A Riot? Disruptive Event Detection using Twitter. *ACM TOIT*, 17(2):18, 2017.
16. Nasser Alsaedi, Pete Burnap, and Omer F. Rana. Automatic Summarization of Real World Events Using Twitter. In *ICWSM*, 2016.
17. Sattam Alsubaiee, Yasser Altowim, Hotham Altwaijry, Alexander Behm, Vinayak R. Borkar, Yingyi Bu, Michael J. Carey, Inci Cetindil, Madhusudan Cheelangi, Khurram Faraaz, Eugenia Gabrielova, Raman Grover, Zachary Heilbron, Young-Seok Kim, Chen Li, Ji Mahn Ok, Nicola Onose, Pouria Pirzadeh, Vasilis Tsotras, Rares Vernica, Jian Wen, and Till Westmann. AsterixDB: A Scalable, Open Source BDMS. *PVLDB*, 7(14), 2014.
18. Apache AsterixDB. <http://asterixdb.apache.org/>, 2018.
19. Apache Cassandra. <http://cassandra.apache.org/>, 2018.
20. Apache Flink. <https://flink.apache.org/>, 2018.
21. Apache Ignite. <https://ignite.apache.org/>, 2018.
22. Apache Impala. <https://impala.apache.org/>, 2018.
23. Apache Spark. <https://spark.apache.org/>, 2014.
24. Apache Spark Streaming. <https://spark.apache.org/streaming/>, 2018.
25. Apache Storm. <https://storm.apache.org/>, 2014.
26. Apple buys social media analytics firm Topsy Labs. [www.bbc.co.uk/news/business-25195534](http://www.bbc.co.uk/news/business-25195534), 2013.
27. New Study Quantifies Use of Social Media in Arab Spring. [www.washington.edu/news/2011/09/12/new-study-quantifies-use-of-social-media-in-arab-spring/](http://www.washington.edu/news/2011/09/12/new-study-quantifies-use-of-social-media-in-arab-spring/), 2011.
28. Sebastien Ardon, Amitabha Bagchi, Anirban Mahanti, Amit Ruhela, Aaditeshwar Seth, Rudra Mohan Tripathy, and Sipat Triukose. Spatio-temporal and Events based Analysis of Topic Popularity in Twitter. In *CIKM*, 2013.
29. Yusuf Arslan, Aysenur Birturk, Bekjan Djumabaev, and Dilek Küçük. Real-time Lexicon-based Sentiment Analysis Experiments on Twitter with a Mild (More Information, Less Data) Approach. In *IEEE Big Data*, 2017.
30. Amir Asiaee, Mariano Tepper, Arindam Banerjee, and Guillermo Sapiro. If You Are Happy and You Know It... Tweet. In *CIKM*, 2012.
31. Neela Avudaiappan, Alexander Herzog, Sneha Kadam, Yuheng Du, Jason Thatche, and Ilya Safro. Detecting and Summarizing Emergent Events in Microblogs and Social Media Streams by Dynamic Centralities. In *IEEE Big Data*, 2017.
32. Brian Babcock, Mayur Datar, and Rajeev Motwani. Load Shedding for Aggregation Queries Over Data Streams. In *ICDE*, 2004.
33. Shuotian Bai, Bibo Hao, Ang Li, Sha Yuan, Rui Gao, and Ting-shao Zhu. Predicting Big Five Personality Traits of Microblog Users. In *WI*, 2013.
34. Akshat Bakliwal, Piyush Arora, Senthil Madhappan, Nikhil Kapre, Mukesh Singh, and Vasudeva Varma. Mining Sentiments from Tweets. In *WASSA@ACL*, 2012.
35. Georgios Balikas. TwiSe at SemEval-2017 Task 4: Five-point Twitter Sentiment Classification and Quantification. In *SemEval-2017*, 2017.
36. Georgios Balikas, Simon Moura, and Massih-Reza Amini. Multitask Learning for Fine-Grained Twitter Sentiment Analysis. In *SIGIR*, 2017.
37. Piyush Bansal, Somay Jain, and Vasudeva Varma. Towards Semantic Retrieval of Hashtags in Microblogs. In *WWW Companion*, 2015.
38. Luciano Barbosa and Junlan Feng. Robust Sentiment Detection on Twitter from Biased and Noisy Data. In *COLING*, 2010.
39. Massimo Bartoletti, Stefano Lande, and Alessandro Massa. Faderank: An Incremental Algorithm for Ranking Twitter Users. In *WISE*, 2016.
40. Moumita Basu, Kripabandhu Ghosh, Somenath Das, Ratnadeep Dey, Somprakash Bandyopadhyay, and Saptarshi Ghosh. Identifying Post-Disaster Resource Needs and Availabilities from Microblogs. In *ASONAM*, 2017.
41. Moumita Basu, Anurag Shandilya, Kripabandhu Ghosh, and Saptarshi Ghosh. Automatic Matching of Resource Needs and Availabilities in Microblogs for Post-Disaster Relief. In *WWW Companion*, 2018.
42. Leilani Battle, Remco Chang, and Michael Stonebraker. Dynamic Prefetching of Data Tiles for Interactive Visualization. In *SIGMOD*, 2016.
43. Wesley Baugh. Bwbaugh: Hierarchical Sentiment Analysis with Partial Self-training. In *SemEval*, volume 2, 2013.
44. Lee Becker, George Erhart, David Skiba, and Valentine Matula. Avaya: Sentiment Analysis on Twitter with Self-training and Polarity Lexicon Expansion. In *SemEval*, volume 2, 2013.
45. Adam Birmingham and Alan F. Smeaton. Classifying Sentiment in Microblogs: Is Brevity an Advantage? In *CIKM*, 2010.



46. Jingwen Bian, Yang Yang, and Tat-Seng Chua. Multimedia summarization for trending topics in microblogs. In *CIKM*, 2013.
47. Federica Bisio, Claudia Meda, Rodolfo Zunino, Roberto Surlinelli, Eugenio Scillia, and Augusto Ottaviano. Real-time Monitoring of Twitter Traffic by using Semantic Networks. In *ASONAM*, 2015.
48. Imen Bizid, Nibal Nayef, Patrice Boursier, Sami Faïz, and Antoine Doucet. Identification of Microblogs Prominent Users during Events by Learning Temporal Sequences of Features. In *CIKM*, 2015.
49. After Boston Explosions, People Rush to Twitter for Breaking News. <http://www.latimes.com/business/technology/la-fi-tn-after-boston-explosions-people-rush-to-twitter-for-breaking-news-20130415,0,3729783.story>, 2013.
50. Ceren Budak, Theodore Georgiou, Divyakant Agrawal, and Amr El Abbadi. GeoScope: Online Detection of Geo-Related Information Trends in Social Networks. In *VLDB*, 2014.
51. Michael Busch, Krishna Gade, Brian Larson, Patrick Lok, Samuel Luckenbill, and Jimmy Lin. Earlybird: Real-Time Search at Twitter. In *ICDE*, 2012.
52. Hongyun Cai, Zi Huang, Divesh Srivastava, and Qing Zhang. Indexing Evolving Events from Tweet Streams. *TKDE*, 27(11):3001–3015, 2015.
53. Caleb Chen Cao, Jieying She, Yongxin Tong, and Lei Chen. Whom to Ask? Jury Selection for Decision Making Tasks on Micro-blog Services. *PVLDB*, 5(11), 2012.
54. Xin Cao, Gao Cong, Tao Guo, Christian S Jensen, and Beng Chin Ooi. Efficient Processing of Spatial Group Keyword Queries. *TODS*, 40(2), 2015.
55. Xin Cao, Gao Cong, Christian S Jensen, and Beng Chin Ooi. Collective Spatial Keyword Querying. In *SIGMOD*, 2011.
56. Ariel Cary, Ouri Wolfson, and Naphtali Rishe. Efficient and Scalable Method for Processing Top-k Spatial Boolean Queries. In *SSDBM*, 2010.
57. Ilknur Celik, Fabian Abel, and Geert-Jan Houben. Learning Semantic Relationships between Entities in Twitter. In *ICWE*, 2011.
58. Sirish Chandrasekaran, Owen Cooper, Amol Deshpande, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong, Sailesh Krishnamurthy, Samuel Madden, Frederick Reiss, and Mehul A. Shah. TelegraphCQ: Continuous Dataflow Processing. In *SIGMOD*, 2003.
59. Harshada Chavan and Mohamed F. Mokbel. Scout: A GPU-Aware System for Interactive Spatio-temporal Data Visualization. In *SIGMOD*, 2017.
60. Chun Chen, Feng Li, Beng Chin Ooi, and Sai Wu. TI: An Efficient Indexing Mechanism for Real-Time Search on Tweets. In *SIGMOD*, 2011.
61. Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. NLG301 at SemEval-2017 Task 5: Fine-Grained Sentiment Analysis on Financial Microblogs and News. In *SemEval*, 2017.
62. Fuhai Chen, Rongrong Ji, Jinsong Su, Donglin Cao, and Yue Gao. Predicting Microblog Sentiments via Weakly Supervised Multimodal Deep Learning. *IEEE Transactions on Multimedia*, 20(4):997–1007, 2018.
63. Lisi Chen, Gao Cong, and Xin Cao. An Efficient Query Indexing Mechanism for Filtering Geo-textual Data. In *SIGMOD*, 2013.
64. Lisi Chen, Gao Cong, Christian S Jensen, and Dingming Wu. Spatial Keyword Query Processing: an Experimental Evaluation. In *VLDB*, 2013.
65. Lisi Chen, Yan Cui, Gao Cong, and Xin Cao. SOPS: A System for Efficient Processing of Spatial-keyword Publish/Subscribe. *PVLDB*, 7(13), 2014.
66. Xing Chen, Lin Li, Guandong Xu, Zhenglu Yang, and Masaru Kitsuregawa. Recommending Related Microblogs: A Comparison Between Topic and WordNet based Approaches. In *AAAI*, 2012.
67. Xuetong Chen, Martin D Sykora, Thomas W Jackson, and Suzanne Elayan. What About Mood Swings: Identifying Depression on Twitter with Temporal Measures of Emotions. In *WWW Companion*, 2018.
68. Daniel Cheng, Peter Schretlen, Nathan Kronenfeld, Neil Bozowsky, and William Wright. Tile Based Visual Analytics for Twitter Big Data Exploratory Analysis. In *IEEE Big Data*, 2013.
69. How Twitter, Facebook, WhatsApp And Other Social Networks Are Saving Lives During Disasters. [http://www.huffingtonpost.in/2017/01/31/how-twitter-facebook-whatsapp-and-other-social-networks-are-sa\\_a\\_21703026/](http://www.huffingtonpost.in/2017/01/31/how-twitter-facebook-whatsapp-and-other-social-networks-are-sa_a_21703026/), 2017.
70. Sina Weibo, China Twitter, comes to rescue amid flooding in Beijing. <http://thenextweb.com/asia/2012/07/23/sina-weibo-chinas-twitter-comes-to-rescue-amid-flooding-in-beijing/>, 2012.
71. Maria Christoforaki, Jinru He, Constantinos Dimopoulos, Alexander Markowetz, and Torsten Suel. Text vs. Space: Efficient Geo-search Query Processing. In *CIKM*, 2011.
72. Sam Clark and Rich Wicentwoski. SwatCS: Combining Simple Classifiers with Estimated Accuracy. In *SemEval@NAACL-HLT*, 2013.
73. Mathieu Cliche. BB.twtr at SemEval-2017 Task 4: Twitter Sentiment Analysis with CNNs and LSTMs. *arXiv:1704.06125*, 2017.
74. Gao Cong and Christian S Jensen. Querying Geo-textual Data: Spatial Keyword Queries and Beyond. In *SIGMOD*, 2016.
75. Gao Cong, Christian S. Jensen, and Dingming Wu. Efficient Retrieval of the Top-k Most Relevant Spatial Web Objects. *PVLDB*, 2(1):337–348, 2009.
76. Camelia Constantin, Quentin Grossetti, Cédric Du Mouza, and Nicolas Travers. An Homophily-based Approach for Fast Post Recommendation in Microblogging Systems. In *EDBT*, 2018.
77. Edilson A Corrêa Jr, Vanessa Queiroz Marinho, and Leandro Borges dos Santos. Nilc-usp at SemEval-2017 Task 4: A Multi-view Ensemble for Twitter Sentiment Analysis. *arXiv:1704.02263*, 2017.
78. Scott Counts and Kristie Fisher. Taking It All In? Visual Attention in Microblog Consumption. In *ICWSM*, 2011.
79. Anqi Cui, Min Zhang, Yiqun Liu, and Shaoping Ma. Emotion Tokens: Bridging The Gap Among Multilingual Twitter Sentiment Analysis. In *Asia Information Retrieval Symposium*, 2011.
80. Anqi Cui, Min Zhang, Yiqun Liu, Shaoping Ma, and Kuo Zhang. Discover Breaking Events with Popular Hashtags in Twitter. In *CIKM*, 2012.
81. Nádia Félix F. da Silva, Eduardo R. Hruschka, and Estevam R. Hruschka Jr. Tweet Sentiment Analysis with Classifier Ensembles. *DSS Journal*, 66:170–179, 2014.
82. Nádia Félix Felipe da Silva, Luiz F. S. Coletta, and Eduardo R. Hruschka. A Survey and Comparative Study of Tweet Sentiment Analysis via Semi-Supervised Learning. *ACM Computing Surveys*, 49(1):15:1–15:26, 2016.
83. Anh Dang, Raheleh Makki, Abidalrahman Moh'd, Aminul Islam, Vlado Keselj, and Evangelos E. Milios. Real Time Filtering of Tweets Using Wikipedia Concepts and Google Tri-gram Se-

- semantic Relatedness. In *TREC*, 2015.
84. Dmitry Davidov, Oren Tsur, and Ari Rappoport. Enhanced Sentiment Learning Using Twitter Hashtags and Smileys. In *COLING*, 2010.
  85. Dayan de França Costa and Nadia Felix Felipe da Silva. INF-UFG at FiQA 2018 Task 1: Predicting Sentiments and Aspects on Financial Tweets and News Headlines. In *WWW Companion*, 2018.
  86. Augusto Q. de Macedo, Leandro Balby Marinho, and Rodrygo L. T. Santos. Context-Aware Event Recommendation in Event-based Social Networks. In *RecSys*, 2015.
  87. Justin DeBrabant, Andrew Pavlo, Stephen Tu, Michael Stonebraker, and Stanley B. Zdonik. Anti-Caching: A New Approach to Database Management System Architecture. In *VLDB*, 2013.
  88. Amit Ajit Deshmane and Jasper Friedrichs. TSA-INF at SemEval-2017 Task 4: An Ensemble of Deep Learning Architectures Including Lexicon Features for Twitter Sentiment Analysis. In *SemEval-2017*, 2017.
  89. Kuntal Dey, Ritvik Shrivastava, and Saroj Kaushik. Twitter Stance Detection: A Subjectivity and Sentiment Polarity Inspired Two-Phase Approach. In *ICDM Workshops*, 2017.
  90. Kuntal Dey, Ritvik Shrivastava, Saroj Kaushik, and L. Venkata Subramaniam. EmTagger: A Word Embedding Based Novel Method for Hashtag Recommendation on Twitter. In *ICDM Workshops*, 2017.
  91. Jiayuan Ding, Yongquan Dong, Tongfei Gao, Zichen Zhang, and Yali Liu. Sentiment Analysis of Chinese Micro-Blog Based on Classification and Rich Features. In *Web Information Systems and Applications Conference*, 2016.
  92. Li Dong, Furu Wei, Chuanqi Tan, Duyu Tang, Ming Zhou, and Ke Xu. Adaptive Recursive Neural Network for Target-dependent Twitter Sentiment Classification. In *ACL*, 2014.
  93. Nikolaos D. Doulamis, Anastasios D. Doulamis, Panagiotis C. Kokkinos, and Emmanouel Manos Varvarigos. Event Detection in Twitter Microblogging. *IEEE Transactions Cybernetics*, 46(12):2810–2824, 2016.
  94. Enkhzhol Dovdon and José Saias. ej-sa-2017 at SemEval-2017 Task 4: Experiments for Target oriented Sentiment Analysis in Twitter. In *SemEval@ACL*, 2017.
  95. Christian Drescher, Guenter Wallner, Simone Kriglstein, Rafet Sifa, Anders Drachen, and Margit Pohl. What Moves Players?: Visual Data Exploration of Twitter and Gameplay Data. In *CHI*, 2018.
  96. Nghia Duong-Trung, Nicolas Schilling, and Lars Schmidt-Thieme. Near Real-time Geolocation Prediction in Twitter Streams via Matrix Factorization based Regression. In *CIKM*, 2016.
  97. Ritam Dutt, Kaustubh Hiware, Avijit Ghosh, and Rameshwar Bhaskaran. SAVITR: A System for Real-time Location Extraction from Microblogs during Emergencies. *CoRR*, abs/1801.07757, 2018.
  98. Soumi Dutta, Vibhash Chandra, Kanav Mehra, Asit Kumar Das, Tanmoy Chakraborty, and Saptarshi Ghosh. Ensemble Algorithms for Microblog Summarization. *IEEE Intelligent Systems*, 33(3):4–14, 2018.
  99. Wolfgang Effelsberg and Theo Härder. Principles of Database Buffer Management. *TODS*, 9(4):560–595, 1984.
  100. Christodoulos Efstathiades, Helias Antoniou, Dimitrios Skoutas, and Yannis Vassiliou. TwitterViz: Visualizing and Exploring the Twittersphere. In *SSTD*, 2015.
  101. Humaira Ehsan, Mohamed A. Sharaf, and Panos K. Chrysanthis. MuVE: Efficient Multi-Objective View Recommendation for Visual Data Exploration. In *ICDE*, 2016.
  102. Ahmed Eldawy, Mohamed F. Mokbel, and Christopher Jonathan. HadoopViz: A MapReduce Framework for Extensible Visualization of Big Spatial Data. In *ICDE*, 2016.
  103. Miki Enoki, Yohei Ikawa, and Rudy Raymond. User Community Reconstruction Using Sampled Microblogging Data. In *WWW Companion*, 2012.
  104. Ahmet Enis Erdoğan, Tolga Yilmaz, Onur Can Sert, Mirun Akyüz, Tansel Özyer, and Reda Alhajj. From Social Media Analysis to Ubiquitous Event Monitoring: The Case of Turkish Tweets. In *ASONAM*, 2017.
  105. How Facebook Is Transforming Disaster Response. <https://www.wired.com/2016/11/facebook-disaster-response/>, 2016.
  106. Facebook Statistics. <http://newsroom.fb.com/company-info/>, 2018.
  107. Ronald Fagin, Amnon Lotem, and Moni Naor. Optimal Aggregation Algorithms for Middleware. In *PODS*, 2001.
  108. Stefano Faralli, Giorgia Di Tommaso, and Paola Velardi. Semantic Enabled Recommender System for Micro-Blog Users. In *ICDM*, 2016.
  109. Shi Feng, Kaisong Song, Daling Wang, and Ge Yu. A Word-emoticon Mutual Reinforcement Ranking Model for Building Sentiment Lexicon from Massive Collection of Microblogs. *WWW Journal*, 18(4):949–967, 2015.
  110. Wei Feng, Chao Zhang, Wei Zhang, Jiawei Han, Jianyong Wang, Charu Aggarwal, and Jianbin Huang. STREAMCUBE: Hierarchical Spatio-temporal Hashtag Clustering for Event Exploration Over the Twitter Stream. In *ICDE*, 2015.
  111. Rana Forsati, Mehrdad Mahdavi, Mehrnosh Shamsfard, and Mohamed Sarwat. Matrix factorization with explicit trust and distrust side information for improved social recommendation. *ACM Trans. Inf. Syst.*, 32(4):17:1–17:38, 2014.
  112. J Ganesh, Manish Gupta, and Vasudeva Varma. Interpretation of Semantic Tweet Representations. In *ASONAM*, 2017.
  113. Lianli Gao, Yuan Wang, Dongsheng Li, Junming Shao, and Jingkuan Song. Real-time social media retrieval with spatial, temporal and social constraints. *Neurocomputing*, 253:77–88, 2017.
  114. Bugra Gedik, Kun-Lung Wu, Philip S. Yu, and Ling Liu. A Load Shedding Framework and Optimizations for M-way Windowed Stream Joins. In *ICDE*, 2007.
  115. Yegin Genc, Yasuaki Sakamoto, and Jeffrey V. Nickerson. Discovering Context: Classifying Tweets through a Semantic Transform Based on Wikipedia. In *Springer FAC*, 2011.
  116. Thanaa Ghanem, Amr Magdy, Mashaal Musleh, Sohaib Ghani, and Mohamed Mokbel. VisCAT: Spatio-Temporal Visualization and Aggregation of Categorical Attributes in Twitter Data. In *SIGSPATIAL*, 2014.
  117. Manoochehr Ghiassi, James Skinner, and David Zimbra. Twitter Brand Sentiment Analysis: A hybrid System using n-gram Analysis and Dynamic Artificial Neural Network. *Expert Systems with applications*, 40(16):6266–6282, 2013.
  118. Saptarshi Ghosh, Naveen Kumar Sharma, Fabrício Benevenuto, Niloy Ganguly, and P. Krishna Gummadi. Cognos: Crowdsourcing Search for Topic Experts in Microblogs. In *SIGIR*, 2012.

119. Anastasia Giachanou and Fabio Crestani. Like It or Not: A Survey of Twitter Sentiment Analysis Methods. *ACM Computing Surveys*, 49(2):28:1–28:41, 2016.
120. Zafar Gilani, Ekaterina Kochmar, and Jon Crowcroft. Classification of Twitter Accounts into Automated Agents and Human Users. In *ASONAM*, 2017.
121. Mehreen Gillani, Muhammad U. Ilyas, Saad Saleh, Jalal S. Alowibdi, Naif R. Aljohani, and Fahad S. Alotaibi. Post Summarization of Microblogs of Sporting Events. In *WWW Companion*, 2017.
122. Alec Go, Richa Bhayani, and Lei Huang. Twitter Sentiment Classification Using Distant Supervision. Technical report, Stanford University, 2009.
123. Raman Grover and Michael Carey. Data Ingestion in AsterixDB. In *EDBT*, 2015.
124. Hansu Gu, Mike Gartrell, Liang Zhang, Qin Lv, and Dirk Grunwald. AnchorMF: Towards Effective Event Context Identification. In *CIKM*, 2013.
125. Yulong Gu, Jiaying Song, Weidong Liu, Lixin Zou, and Yuan Yao. Context Aware Matrix Factorization for Event Recommendation in Event-Based Social Networks. In *WI*, 2016.
126. Satarupa Guha, Tanmoy Chakraborty, Samik Datta, Mohit Kumar, and Vasudeva Varma. TweetGrep: Weakly Supervised Joint Retrieval and Sentiment Analysis of Topical Tweets. In *ICWSM*, 2016.
127. Long Guo, Dongxiang Zhang, Guoliang Li, Kian-Lee Tan, and Zhifeng Bao. Location-Aware Pub/Sub System: When Continuous Moving Queries Meet Dynamic Event Streams. In *SIGMOD*, 2015.
128. Long Guo, Dongxiang Zhang, Yuan Wang, Huayu Wu, Bin Cui, and Kian-Lee Tan. Co 2: Inferring personal interests from raw footprints by connecting the offline world with the online world. *ACM Transactions on Information Systems (TOIS)*, 36(3):31, 2018.
129. Tao Guo, Xin Cao, and Gao Cong. Efficient Algorithms for Answering the M-closest Keywords Query. In *SIGMOD*, 2015.
130. Tao Guo, Kaiyu Feng, Gao Cong, and Zhifeng Bao. Efficient Selection of Geospatial Data on Maps for Interactive and Visualized Exploration. In *SIGMOD*, 2018.
131. Pankaj Gupta, Ashish Goel, Jimmy J. Lin, Aneesh Sharma, Dong Wang, and Reza Zadeh. WTF: The Who to Follow Service at Twitter. In *WWW*, 2013.
132. Pankaj Gupta, Venu Satuluri, Ajeet Grewal, Siva Gurumurthy, Volodymyr Zhabiuk, Quannan Li, and Jimmy J. Lin. Real-Time Twitter Recommendation: Online Motif Detection in Large Dynamic Graphs. *PVLDB*, 7(13):1379–1380, 2014.
133. Hussam Hamdan, Frédéric Béchet, and Patrice Bellot. Experiments with DBpedia, WordNet and SentiWordNet as Resources for Sentiment Analysis in Micro-blogging. In *SemEval@NAACL-HLT*, 2013.
134. John Hannon, Mike Bennett, and Barry Smyth. Recommending Twitter Users to Follow Using Content and Collaborative Filtering Approaches. In *RecSys*, 2010.
135. Yazhou Hao, Yangyang Lan, Yufei Li, and Chen Li. XJSA at SemEval-2017 Task 4: A Deep System for Sentiment Classification in Twitter. In *SemEval-2017*, 2017.
136. Ammar Hassan, Ahmed Abbasi, and Daniel Zeng. Twitter Sentiment Analysis: A Bootstrap Ensemble Framework. In *SocialCom*, 2013.
137. Ling He and Jiebo Luo. What Makes a Pro Eating Disorder Hashtag: Using Hashtags to Identify Pro Eating Disorder Tumblr Posts and Twitter Users. In *IEEE Big Data*, 2016.
138. Yeye He, Siddharth Barman, and Jeffrey F. Naughton. On Load Shedding in Complex Event Processing. In *ICDT*, 2014.
139. Yulan He, Chenghua Lin, Wei Gao, and Kam-Fai Wong. Tracking Sentiment and Topic Dynamics from Social Media. In *ICWSM*, 2012.
140. Brent J. Hecht, Lichan Hong, Bongwon Suh, and Ed H. Chi. Tweets from Justin Bieber’s Heart: The Dynamics of the Location Field in User Profiles. In *CHI*, 2011.
141. Thong Hoang, Pei Hua Cher, Philips Kokoh Prasetyo, and Ee-Peng Lim. Crowdsensing and Analyzing Micro-event Tweets for Public Transportation Insights. In *IEEE Big Data*, 2016.
142. Liangjie Hong, Amr Ahmed, Siva Gurumurthy, Alexander J. Smola, and Kostas Tsioutsoulouklis. Discovering Geographical Topics In The Twitter Stream. In *WWW*, 2012.
143. Amal Htait, Sébastien Fournier, and Patrice Bellot. LSIS at SemEval-2017 Task 4: Using Adapted Sentiment Similarity Seed Words For English and Arabic Tweet Polarity Classification. In *SemEval*, 2017.
144. Guoning Hu, Preeti Bhargava, Saul Fuhrmann, Sarah Ellinger, and Nemanja Spasojevic. Analyzing Users’ Sentiment Towards Popular Consumer Industries and Brands on Twitter. *arXiv:1709.07434*, 2017.
145. Qinmin Hu, Yijun Pei, Qin Chen, and Liang He. SG++: Word Representation with Sentiment and Negation for Twitter Sentiment Classification. In *SIGIR*, 2016.
146. Xia Hu, Lei Tang, and Huan Liu. Enhancing Accessibility of Microblogging Messages using Semantic Knowledge. In *CIKM*, 2011.
147. Xia Hu, Lei Tang, Jiliang Tang, and Huan Liu. Exploiting Social Relations for Sentiment Analysis in Microblogging. In *WSDM*, 2013.
148. Yuheng Hu, Ajita John, Fei Wang, and Subbarao Kambhampati. ET-LDA: Joint Topic Modeling for Aligning Events and their Twitter Feedback. In *AAAI*, volume 12, 2012.
149. Yuheng Hu, Tingting Nian, and Cheng Chen. Mood Congruence or Mood Consistency? Examining Aggregated Twitter Sentiment Towards Ads in 2016 Super Bowl. In *ICWSM*, 2017.
150. Ting Hua, Feng Chen, Liang Zhao, Chang-Tien Lu, and Naren Ramakrishnan. STED: Semi-supervised Targeted-interest Event Detection in Twitter. In *SIGKDD*, 2013.
151. Ting Hua, Feng Chen, Liang Zhao, Chang-Tien Lu, and Naren Ramakrishnan. Automatic Targeted-domain Spatio-temporal Event Detection in Twitter. *Geoinformatica*, 20(4):765–795, 2016.
152. Rocío B. Hubert, Elsa Estevez, Ana Gabriela Maguitman, and Tomasz Janowski. Examining Government-Citizen Interactions on Twitter Using Visual and Sentiment Analysis. In *DG.O.*, 2018.
153. Hurricane Harvey Victims Turn to Twitter and Facebook. <http://time.com/4921961/hurricane-harvey-twitter-facebook-social-media/>, 2017.
154. In Irma, Emergency Responders New Tools: Twitter and Facebook. <https://www.wsj.com/articles/for-hurricane-irma-information-officials-post-on-social-media-1505149661>, 2017.
155. Yohei Ikawa, Miki Enoki, and Michiaki Tsubori. Location Inference Using Microblog Messages. In *WWW*, 2012.

156. Prediction, Optimization and Control for Information Propagation on Networks: A Differential Equation and Mass Transportation Based Approach. [https://www.nsf.gov/awardsearch/showAward?AWD\\_ID=1620342](https://www.nsf.gov/awardsearch/showAward?AWD_ID=1620342), 2017.
157. A Nobel Peace Prize for Twitter? [www.csmonitor.com/Commentary/Opinion/2009/0706/p09s02-coop.html](http://www.csmonitor.com/Commentary/Opinion/2009/0706/p09s02-coop.html), 2009.
158. Masahiko Itoh, Daisaku Yokoyama, Masashi Toyoda, Yoshimitsu Tomita, Satoshi Kawamura, and Masaru Kitsuregawa. Visual Exploration of Changes in Passenger Flows and Tweets on Mega-City Metro Network. *IEEE Transactions on Big Data*, 2(1):85–99, 2016.
159. Mohammed Jabreel and Antonio Moreno. SiTAKA at SemEval-2017 Task 4: Sentiment Analysis in Twitter Based on a Rich Set of Features. In *SemEval*, 2017.
160. Japan earthquake: how Twitter and Facebook helped. <http://www.telegraph.co.uk/technology/twitter/8379101/Japan-earthquake-how-Twitter-and-Facebook-helped.html>, 2011.
161. Jianfeng Jia, Chen Li, Xi Zhang, Chen Li, Michael J. Carey, and Simon Su. Towards Interactive Analytics and Visualization on One Billion Tweets. In *SIGSPATIAL*, 2016.
162. Jinling Jiang, Hua Lu, Bin Yang, and Bin Cui. Finding Top-k Local Users in Geo-Tagged Social Media Data. In *ICDE*, 2015.
163. Long Jiang, Mo Yu, Ming Zhou, Xiaohua Liu, and Tiejun Zhao. Target-dependent Twitter Sentiment Classification. In *ACL*, 2011.
164. Zhao Jianqiang, Gui Xiaolin, and Zhang Xuejun. Deep Convolution Neural Networks for Twitter Sentiment Analysis. *IEEE Access*, 6:23253–23260, 2018.
165. Yohan Jo and Alice H Oh. Aspect and Sentiment Unification Model for Online Review Analysis. In *WSDM*, 2011.
166. Christopher Jonathan, Amr Magdy, Mohamed F Mokbel, and Albert Jonathan. GARNET: A Holistic System Approach For Trending Queries in Microblogs. In *ICDE*, 2016.
167. Alexander J Jones and Eric Carlson. TwitterViz: A Robotics System for Remote Data Visualization. In *ICWSM*, 2013.
168. Robert Kallman, Hideaki Kimura, Jonathan Natkins, Andrew Pavlo, Alex Rasin, Stanley B. Zdonik, Evan P. C. Jones, Samuel Madden, Michael Stonebraker, Yang Zhang, John Hugg, and Daniel J. Abadi. H-Store: A High-performance, Distributed Main Memory Transaction Processing System. *PVLDB*, 1(2):1496–1499, 2008.
169. Janani Kalyanam, Sumithra Velupillai, Mike Conway, and Gert Lanckriet. From Event Detection to Storytelling on Microblogs. In *ASONAM*, 2016.
170. Takamu Kaneko and Keiji Yanai. Visual Event Mining from the Twitter Stream. In *WWW Companion*, 2016.
171. Maria Karanasou, Annetta Ampla, Christos Doukeridis, and Maria Halkidi. Scalable and Real-time Sentiment Analysis of Twitter Data. In *ICDM Workshops*, 2016.
172. Embrace of Social Media Aids Flood Victims in Kashmir. <https://www.nytimes.com/2014/09/13/world/asia/embrace-of-social-media-aids-flood-victims-in-kashmir.html>, 2014.
173. Gabriella Kazai, Yusof Iskander, and Clarke Daoud. Personalised News and Blog Recommendations based on User Location, Facebook and Twitter User Profiling. In *SIGIR*, 2016.
174. Renato Kempter, Valentina Sintsova, Claudiu Cristian Musat, and Pearl Pu. EmotionWatch: Visualizing Fine-Grained Emotions in Event-Related Tweets. In *ICWSM*, 2014.
175. Farhan Hassan Khan, Saba Bashir, and Usman Qamar. TOM: Twitter Opinion Mining Framework using Hybrid Classification Scheme. *DSS Journal*, 57:245–257, 2014.
176. Apalak Khatua and Aparup Khatua. Cricket World Cup 2015: Predicting User’s Orientation through Mix Tweets on Twitter Platform. In *ASONAM*, 2017.
177. Vinh Ngoc Khuc, Chaitanya Shivade, Rajiv Ramnath, and Jay Ramanathan. Towards Building Large-scale Distributed Systems for Twitter Sentiment Analysis. In *ACM SAC*, 2012.
178. Albert Kim, Eric Blais, Aditya G. Parameswaran, Piotr Indyk, Samuel Madden, and Ronitt Rubinfeld. Rapid Sampling for Visualizations with Ordering Guarantees. *PVLDB*, 8(5):521–532, 2015.
179. Eunyoung Kim, Hwon Ihm, and Sung-Hyon Myaeng. Topic-based Place Semantics Discovered from Microblogging Text Messages. In *WWW Companion*, 2014.
180. Svetlana Kiritchenko, Xiaodan Zhu, and Saif M. Mohammad. Sentiment Analysis of Short Informal Texts. *JAIR*, 50:723–762, 2014.
181. Takuya Kitazawa and Makoto Yui. Query-based Simple and Scalable Recommender Systems with Apache Hivemall. In *RecSys*, 2018.
182. Athanasia Kolovou, Filippos Kokkinos, Aris Fergadis, Pnelopi Papalampidi, Elias Iosif, Nikolaos Malandrakis, Elisavet Palogiannidi, Haris Papageorgiou, Shrikanth Narayanan, and Alexandros Potamianos. Tweester at SemEval-2017 Task 4: Fusion of Semantic-Affective and Pairwise Classification Models for Sentiment Analysis in Twitter. In *SemEval@ACL*, 2017.
183. Efstratios Kontopoulos, Christos Berberidis, Theologos Dergiades, and Nick Bassiliades. Ontology-based Sentiment Analysis of Twitter Posts. *Expert Systems with Applications*, 40(10):4065–4074, 2013.
184. Peter Korenek and Marián Simko. Sentiment Analysis on Microblog Utilizing Appraisal Theory. *WWW Journal*, 17(4):847–867, 2014.
185. Efthymios Kouloumpis, Theresa Wilson, and Johanna D. Moore. Twitter Sentiment Analysis: The Good the Bad and the OMG! In *ICWSM*, 2011.
186. Dominik Kowald, Subhash Chandra Pujari, and Elisabeth Lex. Temporal Effects on Hashtag Reuse in Twitter: A Cognitive-inspired Hashtag Recommendation Approach. In *WWW*, 2017.
187. John Krumm and Eric Horvitz. Eyewitness: Identifying Local Events via Space-time Signals in Twitter Feeds. In *SIGSPATIAL*, 2015.
188. Tadahiko Kumamoto, Tomoya Suzuki, and Hitomi Wada. Visualizing Impression-Based Preferences of Twitter Users. In *SCSM-HCI*, 2014.
189. Akshi Kumar and Teeja Mary Sebastian. Sentiment Analysis on Twitter. *IJCSI*, 9(4):372, 2012.
190. Toshiya Kuramochi, Naoki Okada, Kyohei Tanikawa, Yoshinori Hijikata, and Shogo Nishida. Applying to Twitter Networks of a Community Extraction Method Using Intersection Graph and Semantic Analysis. In *Springer HCI*, 2013.
191. Emanuel Lacic. Real-Time Recommendations in a Multi-Domain Environment. In *ACM HT*, 2016.
192. Emanuel Lacic, Dominik Kowald, Denis Parra, Martin Kahr, and Christoph Trattner. Towards A Scalable Social Recommender Engine for Online Marketplaces: The Case of Apache Solr. In *WWW Companion*, 2014.

193. Preethi Lahoti, Gianmarco De Francisci Morales, and Aristides Gionis. Finding Topical Experts in Twitter via Query-dependent Personalized PageRank. In *ASONAM*, 2017.
194. Naveen Kumar Laskari and Suresh Kumar Sanampudi. TWINA at SemEval-2017 Task 4: Twitter Sentiment Analysis with Ensemble Gradient Boost Tree Classifier. In *SemEval-2017*, 2017.
195. George Lee, Jimmy Lin, Chuang Liu, Andrew Lorek, and Dmitriy V. Ryaboy. The Unified Logging Infrastructure for Data Analytics at Twitter. *PVLDB*, 5(12), 2012.
196. Taesung Lee, Jin-woo Park, Sanghoon Lee, Seung-won Hwang, Sameh Elnikety, and Yuxiong He. Processing and Optimizing Main Memory Spatial-Keyword Queries. *PVLDB*, 9(3), 2015.
197. Justin J. Levandoski, Per-Åke Larson, and Radu Stoica. Identifying Hot and Cold Data in Main-memory Databases. In *ICDE*, 2013.
198. Justin J. Levandoski, Mohamed Sarwat, Mohamed F. Mokbel, and Michael D. Ekstrand. RecStore: An Extensible and Adaptive Framework for Online Recommender Queries Inside the Database Engine. In *EDBT*, 2012.
199. Guoliang Li, Jun Hu, Jianhua Feng, and Kian-Lee Tan. Effective Location Identification from Microblogs. In *ICDE*, 2014.
200. Guoliang Li, Yang Wang, Ting Wang, and Jianhua Feng. Location-aware Publish/Subscribe. In *KDD*, 2013.
201. Jing Li, Ming Liao, Wei Gao, Yulan He, and Kam-Fai Wong. Topic Extraction from Microblog Posts Using Conversation Structures. In *ACL*, 2016.
202. Quanzhi Li, Sameena Shah, Armineh Nourbakhsh, Rui Fang, and Xiaomo Liu. funSentiment at SemEval-2017 Task 5: Fine-Grained Sentiment Analysis on Financial Microblogs Using Word Vectors Built from StockTwits and Twitter. In *SemEval*, 2017.
203. Rui Li, Kin Hou Lei, Ravi Khadiwala, and Kevin Chen-Chuan Chang. TEDAS: A Twitter-based Event Detection and Analysis System. In *ICDE*, 2012.
204. Yang Li, Jing Jiang, Ting Liu, Minghui Qiu, and Xiaofei Sun. Personalized Microtopic Recommendation on Microblogs. *ACM TIST*, 8(6):77, 2017.
205. Yuchen Li, Zhifeng Bao, Guoliang Li, and Kian-Lee Tan. Real Time Personalized Search on Social Networks. In *ICDE*, 2015.
206. Zhisheng Li, Ken C. K. Lee, Baihua Zheng, Wang-Chien Lee, Dik Lun Lee, and Xufa Wang. IR-Tree: An Efficient Index for Geographic Document Search. *TKDE*, 23(4):585–599, 2011.
207. Kwan Hui Lim, Kate E Lee, Dave Kendal, Lida Rashidi, Elham Naghizade, Stephan Winter, and Maria Vasardani. The Grass Is Greener On The Other Side: Understanding the Effects of Green Spaces on Twitter User Sentiments. In *WWW Companion*, 2018.
208. Jimmy Lin and Alek Kolcz. Large-scale machine learning at twitter. In *SIGMOD*, 2012.
209. Jimmy Lin and Gilad Mishne. A Study of "Churn" in Tweets and Real-Time Search Queries. In *ICWSM*, 2012.
210. John Lingad, Sarvnaz Karimi, and Jie Yin. Location Extraction From Disaster-related Microblogs. In *WWW*, 2013.
211. Meiling Liu, Kaiqun Fu, Chang-Tien Lu, Guangsheng Chen, and Huiqiang Wang. A Search and Summary Application for Traffic Events Detection Based on Twitter Data. In *SIGSPATIAL*, 2014.
212. Nian Liu, Lin Li, Guandong Xu, and Zhenglu Yang. Identifying Domain-Dependent Influential Microblog Users: A Post-Feature Based Approach. In *AAAI*, 2014.
213. Shenghua Liu, Fuxin Li, Fangtao Li, Xueqi Cheng, and Huawei Shen. Adaptive Co-training SVM for Sentiment Classification on Tweets. In *CIKM*, 2013.
214. Shenghua Liu, Wenjun Zhu, Ning Xu, Fangtao Li, Xue-qi Cheng, Yue Liu, and Yuanzhuo Wang. Co-training and Visualizing Sentiment Evolvement for Tweet Events. In *WWW*, 2013.
215. Xiaohua Liu, Zhongyang Fu, Furu Wei, and Ming Zhou. Collective Nominal Semantic Role Labeling for Tweets. In *AAAI*, 2012.
216. Xiaohua Liu, Kuan Li, Ming Zhou, and Zhongyang Xiong. Enhancing Semantic Role Labeling for Tweets Using Self-Training. In *AAAI*, 2011.
217. Xiaomo Liu, Quanzhi Li, Armineh Nourbakhsh, Rui Fang, Merline Thomas, Kajsa Anderson, Russ Kociuba, Mark Vedder, Steven Pomerville, Ramdev Wudali, et al. Reuters Tracer: A Large Scale System of Detecting & Verifying Real-time News Events from Twitter. In *CIKM*, 2016.
218. Cheng Long, Raymond Chi-Wing Wong, Ke Wang, and Ada Wai-Chee Fu. Collective Spatial Keyword Queries: a Distance Owner-driven Approach. In *SIGMOD*, 2013.
219. David Lozić, Doria Šarić, Ivan Tokić, Zoran Medić, and Jan Šnajder. TakeLab at SemEval-2017 Task 4: Recent Deaths and the Power of Nostalgia in Sentiment Analysis in Twitter. In *SemEval-2017*, 2017.
220. Xiao Lu, Peng Li, Hongyuan Ma, Shuxin Wang, Anying Xu, and Bin Wang. Computing and Applying Topic-level User Interactions in Microblog Recommendation. In *SIGIR*, 2014.
221. Renfeng Ma, Qi Zhang, Jiawen Wang, Lizhen Cui, and Xuanjing Huang. Mention Recommendation for Multimodal Microblog with Cross-attention Memory Network. In *SIGIR*, 2018.
222. Amr Magdy, Louai Alarabi, Saif Al-Harathi, Mashaal Musleh, Thanaa Ghanem, Sohaib Ghani, and Mohamed Mokbel. Taghreed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs. In *SIGSPATIAL*, 2014.
223. Amr Magdy, Rami Alghamdi, and Mohamed F. Mokbel. On Main-memory Flushing in Microblogs Data Management Systems. In *ICDE*, 2016.
224. Amr Magdy, Ahmed M Aly, Mohamed F Mokbel, Sameh Elnikety, Yuxiong He, Suman Nath, and Walid G Aref. GeoTrend: Spatial Trending Queries on Real-time Microblogs. In *SIGSPATIAL*, 2016.
225. Amr Magdy and Mohamed Mokbel. Towards a Microblogs Data Management System. In *MDM*, 2015.
226. Amr Magdy and Mohamed Mokbel. Microblogs Data Management and Analysis (Tutorial). In *ICDE*, 2016.
227. Amr Magdy and Mohamed Mokbel. Demonstration of Kite: A Scalable System for Microblogs Data Management. In *ICDE*, 2017.
228. Amr Magdy, Mohamed F. Mokbel, Sameh Elnikety, Suman Nath, and Yuxiong He. Mercury: A Memory-Constrained Spatio-temporal Real-time Search on Microblogs. In *ICDE*, 2014.
229. Amr Magdy, Mohamed F. Mokbel, Sameh Elnikety, Suman Nath, and Yuxiong He. Venus: Scalable Real-time Spatial Queries on Microblogs with Adaptive Load Shedding. *TKDE*, 28(2), 2016.
230. Amr Magdy, Mashaal Musleh, Kareem Tarek, Louai Alarabi, Saif Al-Harathi, Hicham G. Elmongui, Thanaa M. Ghanem, Sohaib Ghani, and Mohamed F. Mokbel. Taqreer: A System for Spatio-temporal Analysis on Microblogs. *IEEE Data Engineering Bulletin*, 38(2), 2015.

231. Axel Magnuson, Vijay Dialani, and Deepa Mallela. Event Recommendation using Twitter Activity. In *RecSys*, 2015.
232. Ahmed R. Mahmood, Walid G. Aref, and Ahmed M. Aly. FAST: Frequency-Aware Indexing for Spatio-Textual Data Streams. In *ICDE*, 2018.
233. Ahmed R Mahmood, Walid G Aref, Ahmed M Aly, and Mingjie Tang. Atlas: On the Expression of Spatial-keyword Group Queries Using Extended Relational Constructs. In *SIGSPATIAL*, 2016.
234. Jalal Mahmud, Jeffrey Nichols, and Clemens Drews. Where Is This Tweet From? Inferring Home Locations of Twitter Users. In *ICWSM*, 2012.
235. Axel J. Soto Stephen Brooks Makki, Raheleh and Evangelos E. Milios. Twitter Message Recommendation based on User Interest Profiles. In *ASONAM*, 2016.
236. Raheleh Makki, Eder J. de Carvalho, Axel J. Soto, Stephen Brooks, Maria Cristina Ferreira de Oliveira, Evangelos E. Milios, and Rosane Minghim. ATR-Vis: Visual and Interactive Information Retrieval for Parliamentary Discussions in Twitter. *TKDD*, 12(1):3:1–3:33, 2018.
237. Adam Marcus, Michael S. Bernstein, Osama Badar, David R. Karger, Samuel Madden, and Robert C. Miller. Tweets as Data: Demonstration of TweepQL and TwitInfo. In *SIGMOD*, 2011.
238. Adam Marcus, Michael S. Bernstein, Osama Badar, David R. Karger, Samuel Madden, and Robert C. Miller. Twitinfo: Aggregating and Visualizing Microblogs for Event Exploration. In *CHI*, 2011.
239. Dean McCullough, Jimmy Lin, Craig Macdonald, Iadh Ounis, and Richard M. C. McCreddie. Evaluating Real-Time Search over Tweets. In *ICWSM*, 2012.
240. Andrew J McMinn, Daniel Tsvetkov, Tsvetan Yordanov, Andrew Patterson, Rrobi Szk, Jesus A Rodriguez Perez, and Joemon M Jose. An Interactive Interface for Visualizing Events on Twitter. In *SIGIR*, 2014.
241. Qiaozhu Mei, Xu Ling, Matthew Wondra, Hang Su, and ChengXiang Zhai. Topic Sentiment Mixture: Modeling Facets and Opinions in Weblogs. In *WWW*, 2007.
242. Edgar Meij, Wouter Weerkamp, and Maarten de Rijke. Adding Semantics to Microblog Posts. In *WSDM*, 2012.
243. Ahmed Metwally, Divyakant Agrawal, and Amr El Abbadi. Efficient Computation of Frequent and Top-k Elements in Data Streams. In *ICDT*, 2005.
244. Sabino Miranda-Jiménez, Mario Graff, Eric Sadit Tellez, and Daniela Moctezuma. INGEOTEC at SemEval 2017 Task 4: A B4MSA Ensemble based on Genetic Programming for Twitter Sentiment Analysis. In *SemEval*, 2017.
245. Gilad Mishne, Jeff Dalton, Zhenghua Li, Aneesh Sharma, and Jimmy Lin. Fast Data in the Era of Big Data: Twitter's Real-time Related Query Suggestion Architecture. In *SIGMOD*, 2013.
246. Gilad Mishne and Jimmy Lin. Twanchor Text: A Preliminary Study of the Value of Tweets as Anchor Text. In *SIGIR*, 2012.
247. Saif Mohammad. #Emotional Tweets. In *\*SEM@NAACL-HLT*, 2012.
248. Saif Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. NRC-Canada: Building the State-of-the-Art in Sentiment Analysis of Tweets. In *SemEval@NAACL-HLT*, 2013.
249. Mohamed Mokbel and Amr Magdy. Microblogs Data Management Systems: Querying, Analysis, and Visualization (Tutorial). In *SIGMOD*, 2016.
250. Mohamed F. Mokbel and Walid G. Aref. SOLE: Scalable On-Line Execution of Continuous Queries on Spatio-temporal Data Streams. *VLDB Journal*, 17(5), 2008.
251. Mohamed F. H. Mokbel and Amr M. M. Ahmed. System and Method for Microblogs Data Management, Provisionally filed in U.S. Patent and Trademark Office on August 31, 2015, Application number: 14/841299. <http://appft1.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PG01&p=1&u=/netahtml/PTO/srchnum.html&r=1&f=G&l=50&s1=20160070754.PGNR>.
252. MongoDB. <https://www.mongodb.com/>, 2018.
253. Lin Mu, Peiquan Jin, Lizhou Zheng, En-Hong Chen, and Lihua Yue. Lifecycle-Based Event Detection from Microblogs. In *WWW Companion*, 2018.
254. Hala Mulki, Hatem Haddad, Mourad Gridach, and Ismail Babaoğlu. Tw-StAR at SemEval-2017 Task 4: Sentiment Classification of Arabic Tweets. In *SemEval-2017*, 2017.
255. Zarmeen Nasim. IBA-Sys at SemEval-2017 Task 5: Fine-Grained Sentiment Analysis on Financial Microblogs and News. In *SemEval*, 2017.
256. New Enhanced Geo-targeting for Marketers. <https://blog.twitter.com/2012/new-enhanced-geo-targeting-for-marketers>, 2012.
257. Nikolaos Nodarakis, Spyros Sioutas, Athanasios K Tsakalidis, and Giannis Tzimas. Large Scale Sentiment Analysis on Twitter with Spark. In *EDBT Workshops*, 2016.
258. One Million Tweet Map. <http://onemilliontweetmap.com/>, 2016.
259. Reynier Ortega, Adrian Fonseca, and Andres Montoyo. SSA-UO: Unsupervised Twitter Sentiment Analysis. In *Joint Conference on Lexical and Computational Semantics (\*SEM)*, volume 2, 2013.
260. Ozer Ozdakis, Pinar Senkul, and Halit Oguztüzün. Semantic Expansion of Tweet Contents for Enhanced Event Detection in Twitter. In *ASONAM*, 2012.
261. Alexander Pak and Patrick Paroubek. Twitter as a Corpus for Sentiment Analysis and Opinion Mining. In *LREC*, 2010.
262. Yongjoo Park, Michael J. Cafarella, and Barzan Mozafari. Visualization-aware Sampling for Very Large Databases. In *ICDE*, 2016.
263. Alexandre Passant, Uldis Bojars, John G. Breslin, Tuukka Hasstrup, Milan Stankovic, and Philippe Laublet. An Overview of SMOB 2: Open, Semantic and Distributed Microblogging. In *ICWSM*, 2010.
264. Debjyoti Paul, Feifei Li, Murali Krishna Teja, Xin Yu, and Richie Frost. Compass: Spatio Temporal Sentiment Analysis of US Election What Twitter Says! In *SIGKDD*, 2017.
265. Carlos Rodríguez Penagos, Jordi Atserias Batalla, Joan Codina-Filbà, David García Narbona, Jens Grivolla, Patrik Lambert, and Roser Saurí. FBM: Combining Lexicon-based ML and Heuristics for Social Media Polarities. In *SemEval@NAACL-HLT*, 2013.
266. Min Peng, Jiahui Zhu, Hua Wang, Xuhui Li, Yanchun Zhang, Xizhen Zhang, and Gang Tian. Mining Event-Oriented Topics in Microblog Stream with Unsupervised Multi-View Hierarchical Embedding. *TKDD*, 12(3):38, 2018.
267. Owen Phelan, Kevin McCarthy, and Barry Smyth. Using Twitter to Recommend Real-Time Topical News. In *RecSys*, 2009.

268. Ana-Maria Popescu and Marco Pennacchiotti. Detecting Controversial Events from Twitter. In *CIKM*, 2010.
269. Presto. <http://prestodb.io/>, 2018.
270. Public Health Emergency, Department of Health and Human Services. <http://nowtrending.hhs.gov/>, 2015.
271. Ashequl Qadir, Pablo N. Mendes, Daniel Gruhl, and Neal Lewis. Semantic Lexicon Induction from Twitter with Pattern Relatedness and Flexible Term Length. In *AAAI*, 2015.
272. Yujie Qian, Jie Tang, Zhilin Yang, Binxuan Huang, Wei Wei, and Kathleen M. Carley. A Probabilistic Framework for Location Inference from Social Media. *CoRR*, abs/1702.07281, 2017.
273. Lirong Qiu, Qiao Lei, and Zhen Zhang. Advanced Sentiment Classification of Tibetan Microblogs on Smart Campuses Based on Multi-Feature Fusion. *IEEE Access*, 6:17896–17904, 2018.
274. S Milton Rajendram, TT Mirnalinee, et al. SSN\_MLRG1 at SemEval-2017 Task 4: Sentiment Analysis in Twitter Using Multi-Kernel Gaussian Process Classifier. In *SemEval*, 2017.
275. Daniel Ramage, Susan T. Dumais, and Daniel J. Liebling. Characterizing Microblogs with Topic Models. In *ICWSM*, 2010.
276. Jaishree Ranganathan, Allen S Irudayaraj, and Angelina A Tzacheva. Action Rules for Sentiment Analysis on Twitter Data Using Spark. In *ICDM Workshops*, 2017.
277. Redis. <https://redis.io/>, 2018.
278. Yafeng Ren, Yue Zhang, Meishan Zhang, and Donghong Ji. Context-Sensitive Twitter Sentiment Classification Using Neural Network. In *AAAI*, 2016.
279. Yafeng Ren, Yue Zhang, Meishan Zhang, and Donghong Ji. Improving Twitter Sentiment Classification Using Topic-Enriched Multi-Prototype Word Embeddings. In *AAAI*, 2016.
280. Manoel Horta Ribeiro, Pedro H. Calais, Yuri A. Santos, Virgílio A. F. Almeida, and Wagner Meira Jr. Characterizing and Detecting Hateful Users on Twitter. *CoRR*, abs/1803.08977, 2018.
281. Miguel Rios and Jimmy J Lin. Visualizing the "Pulse" of World Cities on Twitter. In *ICWSM*. Citeseer, 2013.
282. Ricardo Araújo Rios, Paulo A. Pagliosa, Renato Porfirio Ishii, and Rodrigo Fernandes de Mello. TSViz: A Data Stream Architecture to Online Collect, Analyze, and Visualize Tweets. In *SAC*, 2017.
283. Alan Ritter, Oren Etzioni, Sam Clark, et al. Open Domain Event Extraction from Twitter. In *SIGKDD*, 2012.
284. RocksDB. <https://rocksdb.org/>, 2018.
285. Simone Aparecida Pinto Romero and Karin Becker. A Framework for Event Classification in Tweets Based on Hybrid Semantic Enrichment. *Expert Systems with Applications*, 118, 2019.
286. Guilherme Coletto Rotta, Vinícius Silva de Lemos, Felipe Lammel, Isabel Harb Manssour, Milene Selbach Silveira, and André Fagundes Pase. Visualization Techniques for the Analysis of Twitter Users' Behavior. In *ICWSM*, 2013.
287. Alon Rozental and Daniel Fleischer. Amobee at SemEval-2017 Task 4: Deep Learning System for Sentiment Detection on Twitter. *arXiv:1705.01306*, 2017.
288. Koustav Rudra, Subham Ghosh, Niloy Ganguly, Pawan Goyal, and Saptarshi Ghosh. Extracting Situational Information from Microblogs during Disaster Events: a Classification-Summarization Approach. In *CIKM*, 2015.
289. Koustav Rudra, Pawan Goyal, Niloy Ganguly, Prasenjit Mitra, and Muhammad Imran. Identifying Sub-events and Summarizing Disaster-Related Information from Microblogs. In *SIGIR*, 2018.
290. KyoungMin Ryoo and Sue Moon. Inferring Twitter User Locations with 10 km Accuracy. In *WWW Companion*, 2014.
291. Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. Earthquake Shakes Twitter Users: Real-Time Event Detection by Social Sensors. In *WWW*, 2010.
292. Jitao Sang, Dongyuan Lu, and Changsheng Xu. A Probabilistic Framework for Temporal User Modeling on Microblogs. In *CIKM*, 2015.
293. Jagan Sankaranarayanan, Hanan Samet, Benjamin E. Teitler, Michael D. Lieberman, and Jon Sperling. TwitterStand: News in Tweets. In *SIGSPATIAL*, 2009.
294. Mohamed Sarwat. Recdb: towards DBMS support for online recommender systems. In *Proceedings of the ACM SIGMOD/PODS PhD Symposium 2012, Scottsdale, AZ, USA, May 20, 2012*, pages 33–38, 2012.
295. Mohamed Sarwat, James L. Avery, and Mohamed F. Mokbel. A RecDB in Action: Recommendation Made Easy in Relational Databases. *PVLDB*, 6(12):1242–1245, 2013.
296. Mohamed Sarwat, James L. Avery, and Mohamed F. Mokbel. A recdb in action: Recommendation made easy in relational databases. *PVLDB*, 6(12):1242–1245, 2013.
297. Mohamed Sarwat, James L. Avery, and Mohamed F. Mokbel. RECATHON: A Middleware for Context-Aware Recommendation in Database Systems. In *MDM*, 2015.
298. Mohamed Sarwat, Raha Moraffah, Mohamed F. Mokbel, and James L. Avery. Database System Support for Personalized Recommendation Applications. In *ICDE*, 2017.
299. Ranjan Satapathy, Claudia Guerreiro, Iti Chaturvedi, and Erik Cambria. Phonetic-Based Microtext Normalization for Twitter Sentiment Analysis. In *ICDM Workshops*, 2017.
300. Aneesh Sharma, Jiang Jerry, Bommannavar Praveen, Larson Brian, and Lin. Jimmy. GraphJet: Real-Time Content Recommendations at Twitter. In *VLDB*, pages 1281–1292, 2016.
301. Jianfeng Si, Arjun Mukherjee, Bing Liu, Qing Li, Huayi Li, and Xiaotie Deng. Exploiting Topic Based Twitter Sentiment for Stock Prediction. In *ACL*, volume 2, 2013.
302. Bas Sijtsma, Pernilla Qvarfordt, and Francine Chen. Tweetviz: Visualizing Tweets for Business Intelligence. In *SIGIR*, 2016.
303. Vivek K. Singh, Mingyan Gao, and Ramesh Jain. Situation Detection and Control using Spatio-temporal Analysis of Microblogs. In *WWW*, 2010.
304. Anders Skovsgaard, Darius Sidlauskas, and Christian S. Jensen. Scalable Top-k Spatio-temporal Term Querying. In *ICDE*, 2014.
305. Karin Sim Smith, Richard McCreddie, Craig Macdonald, and Iadh Ounis. Analyzing Disproportionate Reaction via Comparative Multilingual Targeted Sentiment in Twitter. In *ASONAM*, 2017.
306. Understanding Social and Geographical Disparities in Disaster Resilience Through the Use of Social Media. [https://www.nsf.gov/awardsearch/showAward?AWD\\_ID=1620451](https://www.nsf.gov/awardsearch/showAward?AWD_ID=1620451), 2017.
307. Sparsity Models for Forecasting Spatio-Temporal Human Dynamics. [https://www.nsf.gov/awardsearch/showAward?AWD\\_ID=1737770](https://www.nsf.gov/awardsearch/showAward?AWD_ID=1737770), 2017.
308. Anders Sogaard, Barbara Plank, and Héctor Martínez Alonso. Using Frame Semantics for Knowledge Extraction from Twitter.

- In *AAAI*, 2015.
309. Kaisong Song, Ling Chen, Wei Gao, Shi Feng, Daling Wang, and Chengqi Zhang. Persentiment: A Personalized Sentiment Classification System for Microblog Users. In *WWW Companion*, 2016.
  310. Dionisios N. Sotiropoulos, Chris D. Kounavis, and George M. Giaglis. Semantically Meaningful Group Detection Within Subcommunities of Twitter Blogosphere: A Topic Oriented Multi-objective Clustering Approach. In *ASONAM*, 2013.
  311. Laure Soulier, Tamine Lynda, and Nguyen. Gia-Hung. Answering Twitter Questions: A Model For Recommending Answerers Through Social Collaboration. In *CIKM*, 2016.
  312. Michael Speriosu, Nikita Sudan, Sid Upadhyay, and Jason Baldrige. Twitter Polarity Classification with Label Propagation over Lexical Links and The Follower Graph. In *Workshop on Unsupervised Learning in NLP*, 2011.
  313. Enrico Steiger, Bernd Resch, and Alexander Zipf. Exploration of Spatiotemporal and Semantic Clusters of Twitter Data using Unsupervised Neural Networks. *IJGIS*, 30(9):1694–1716, 2016.
  314. Michael Stonebraker and Ariel Weisberg. The VoltDB Main Memory DBMS. *IEEE Data Engineering Bulletin*, 36(2):21–27, 2013.
  315. Dhanasekar Sundararaman and Sudharshan Srinivasan. Twigraph: Discovering and Visualizing Influential Words Between Twitter Profiles. In *Social Informatics*, 2017.
  316. Symeon Symeonidis, Dimitrios Effrosynidis, John Kordonis, and Avi Arampatzis. DUTH at SemEval-2017 Task 4: A Voting Classification Approach for Twitter Sentiment Analysis. In *SemEval*, 2017.
  317. Symeon Symeonidis, John Kordonis, Dimitrios Effrosynidis, and Avi Arampatzis. DUTH at SemEval-2017 Task 5: Sentiment Predictability in Financial Microblogging and News Articles. In *SemEval*, 2017.
  318. Narges Tabari, Armin Seyedtabari, and Wlodek Zadrozny. Sentiheros at SemEval-2017 Task 5: An application of Sentiment Analysis on Financial Tweets. In *SemEval*, 2017.
  319. Chenhao Tan, Lillian Lee, Jie Tang, Long Jiang, Ming Zhou, and Ping Li. User-level Sentiment Analysis Incorporating Social Networks. In *SIGKDD*, 2011.
  320. Duyu Tang, Furu Wei, Bing Qin, Ting Liu, and Ming Zhou. Coooolll: A Deep Learning System for Twitter Sentiment Classification. In *SemEval@COLING*, 2014.
  321. Duyu Tang, Furu Wei, Nan Yang, Ming Zhou, Ting Liu, and Bing Qin. Learning Sentiment-Specific Word Embedding for Twitter Sentiment Classification. In *ACL*, 2014.
  322. Mike Thelwall, Kevan Buckley, and Georgios Paltoglou. Sentiment Strength Detection For the Social Web. *JASIST*, 63(1):163–173, 2012.
  323. Topsy Analytics: Find the insights that matter. [www.topsy.com](http://www.topsy.com), 2014.
  324. Jean Gomes Turet and Ana Paula Cabral Seixas Costa. Big Data Analytics to Improve the Decision-Making Process in Public Safety: A Case Study in Northeast Brazil. In *Springer ICDSST*, 2018.
  325. Tweet Complete Index. [https://blog.twitter.com/engineering/en\\_us/a/2014/building-a-complete-tweet-index.html](https://blog.twitter.com/engineering/en_us/a/2014/building-a-complete-tweet-index.html).
  326. TweetTracker: track, analyze, and understand activity on Twitter. [tweettracker.fulton.asu.edu/](http://tweettracker.fulton.asu.edu/), 2014.
  327. Health Department Use of Social Media to Identify Foodborne Illness Chicago, Illinois, 2013-2014. <https://www.cdc.gov/mmwr/preview/mmwrhtml/mm6332a1.htm>, 2014.
  328. Harvard Medical School Researchers Awarded Twitter Data Grant. <https://hms.harvard.edu/news/harvard-medical-school-researchers-awarded-twitter-data-grant>, 2014.
  329. Twitter and Informal Science Learning and Engagement. [https://www.nsf.gov/awardsearch/showAward?AWD\\_ID=1438898](https://www.nsf.gov/awardsearch/showAward?AWD_ID=1438898), 2017.
  330. The Power of Images: A Computational Investigation of Political Mobilization via Social Media. [https://www.nsf.gov/awardsearch/showAward?AWD\\_ID=1727459](https://www.nsf.gov/awardsearch/showAward?AWD_ID=1727459), 2017.
  331. Twitter Data Changing Future of Population Research. <http://news.psu.edu/story/474782/2017/07/17/research/twitter-data-changing-future-population-research>, 2017.
  332. Twitter Statistics. <https://about.twitter.com/company>, 2018.
  333. The Twitter War: Social Media's Role in Ukraine Unrest. [news.nationalgeographic.com/news/2014/05/140510-ukraine-odessa-russia-kiev-twitter-world/](http://news.nationalgeographic.com/news/2014/05/140510-ukraine-odessa-russia-kiev-twitter-world/), 2014.
  334. Twitter a Big Winner in 2012 Presidential Election. [www.computerworld.com/article/2493332/social-media/twitter-a-big-winner-in-2012-presidential-election.html](http://www.computerworld.com/article/2493332/social-media/twitter-a-big-winner-in-2012-presidential-election.html), 2012.
  335. Topsy Analytics for Twitter Political Index. [https://blog.twitter.com/official/en\\_us/a/2012/a-new-barometer-for-the-election.html](https://blog.twitter.com/official/en_us/a/2012/a-new-barometer-for-the-election.html).
  336. Norases Vesdapunt and Hector Garcia-Molina. Identifying Users in Social Networks with Limited Information. In *ICDE*, 2015.
  337. Duy-Tin Vo and Yue Zhang. Target-Dependent Twitter Sentiment Classification with Rich Automatic Features. In *IJCAI*, 2015.
  338. VoltDB. <https://www.voltdb.com/>, 2018.
  339. Jan Vosecky, Di Jiang, Kenneth Wai-Ting Leung, Kai Xing, and Wilfred Ng. Integrating Social and Auxiliary Semantics for Multifaceted Topic Modeling in Twitter. *ACM TOIT*, 14(4):27:1–27:24, 2014.
  340. VG Vinod Vydiswaran, Daniel M Romero, Xinyan Zhao, Deahhan Yu, Iris N Gomez-Lopez, Jin Xiu Lu, Bradley Iott, Ana Baylin, Philippa Clarke, Veronica J Berrocal, et al. "Bacon Bacon Bacon": Food-Related Tweets and Sentiment in Metro Detroit. In *ICWSM*, 2018.
  341. Shoko Wakamiya, Adam Jatowt, Yukiko Kawai, and Toyokazu Akiyama. Analyzing Global and Pairwise Collective Spatial Attention for Geo-social Event Detection in Microblogs. In *WWW Companion*, 2016.
  342. Ming Wang, Biao Chu, Qingxun Liu, and Xiaobing Zhou. YNUDLG at SemEval-2017 Task 4: A GRU-SVM Model for Sentiment Classification and Quantification in Twitter. In *SemEval-2017*, 2017.
  343. Xiang Wang, Ying Zhang, Wenjie Zhang, Xuemin Lin, and Wei Wang. AP-Tree: Efficiently Support Continuous Spatial-keyword Queries Over Stream. In *ICDE*, 2015.
  344. Xiaolong Wang, Furu Wei, Xiaohua Liu, Ming Zhou, and Ming Zhang. Topic Sentiment Analysis in Twitter: A Graph-based Hashtag Sentiment Classification Approach. In *CIKM*, 2011.
  345. Yuan Wang, Jie Liu, Yalou Huang, and Xia Feng. Using Hashtag Graph-based Topic Model to Connect Semantically-related Words without Co-occurrence in Microblogs. *TKDE*, 28(7):1919–1933, 2016.



346. Yuanyuan Wang, Panote Siriaraya, Yusuke Nakaoka, Haruka Sakata, Yukiko Kawai, and Toyokazu Akiyama. A Twitter-Based Culture Visualization System by Analyzing Multilingual Geo-Tagged Tweets. In *ICADL*, 2018.
347. Zhibo Wang, Yongquan Zhang, Yijie Li, Qian Wang, and Feng Xia. Exploiting Social Influence for Context-aware Event Recommendation in Event-based Social Networks. In *INFOCOM*, 2017.
348. Kazufumi Watanabe, Masanao Ochi, Makoto Okabe, and Rikio Onai. Jasmine: A Real-time Local-event Detection System based on Geolocation Information Propagated to Microblogs. In *CIKM*, 2011.
349. Ingmar Weber and Venkata Rama Kiran Garimella. Visualizing User-Defined, Discriminative Geo-Temporal Twitter Activity. In *ICWSM*, 2014.
350. Michael J. Welch, Uri Schonfeld, Dan He, and Junghoo Cho. Topical Semantics of Twitter Links. In *WSDM*, 2011.
351. Fangzhao Wu and Yongfeng Huang. Personalized Microblog Sentiment Classification via Multi-Task Learning. In *AAAI*, 2016.
352. Lingkun Wu, Wenqing Lin, Xiaokui Xiao, and Yabo Xu. LSII: An Indexing Structure for Exact Real-Time Search on Microblogs. In *ICDE*, 2013.
353. Shanchan Wu, Leanna Gong, William Rand, and Louiqa Raschid. Making Recommendations in a Microblog to Improve the Impact of a Focal User. In *RecSys*, 2012.
354. Xiaodong Wu, Lyn Bartram, and Chris Shaw. Plexus: An Interactive Visualization Tool for Analyzing Public Emotions from Twitter Data. *CoRR*, abs/1701.06270, 2017.
355. Yanhui Wu. Language E-learning based on Learning Analytics in Big Data Era. In *International Conference on Big Data and Education*, 2018.
356. Bing Xiang and Liang Zhou. Improving Twitter Sentiment Analysis with Topic-Based Mixture Modeling and Semi-supervised Training. In *ACL*, volume 2, 2014.
357. Qing Xie, Xiangliang Zhang, Zhixu Li, and Xiaofang Zhou. Optimizing Cost of Continuous Overlapping Queries over Data Streams by Filter Adaption. *TKDE*, 28(5):1258–1271, 2016.
358. Chen Xing, Yuan Wang, Jie Liu, Yalou Huang, and Wei-Ying Ma. Hashtag-Based Sub-Event Discovery Using Mutually Generative LDA in Twitter. In *AAAI*, pages 2666–2672, 2016.
359. Xiaopeng Xiong, Mohamed F. Mokbel, and Walid G. Aref. SEA-CNN: Scalable Processing of Continuous K-Nearest Neighbor Queries in Spatio-temporal Databases. In *ICDE*, 2005.
360. Tzu-Hsuan Yang, Tzu-Hsuan Tseng, and Chia-Ping Chen. deepSA at SemEval-2017 Task 4: Interpolated Deep Neural Networks for Sentiment Analysis in Twitter. In *SemEval*, 2017.
361. Junjie Yao, Bin Cui, Zijun Xue, and Qingyun Liu. Provenance-based Indexing Support in Micro-blog Platforms. In *ICDE*, 2012.
362. An-Zi Yen, Hen-Hsen Huang, and Hsin-Hsi Chen. Detecting Personal Life Events from Twitter by Multi-Task LSTM. In *WWW Companion*, 2018.
363. Hongzhi Yin, Bin Cui, Ling Chen, Zhiting Hu, and Chengqi Zhang. Modeling Location-Based User Rating Profiles for Personalized Recommendation. *TKDD*, 9(3):19:1–19:41, 2015.
364. Yichun Yin, Yangqiu Song, and Ming Zhang. NNEMBs at SemEval-2017 Task 4: Neural Twitter Sentiment Classification: a Simple Ensemble Method with Different Embeddings. In *SemEval*, 2017.
365. Xiaojun Wan Yu, Yang and Xinjie Zhou. User Embedding for Scholarly Microblog Recommendation. In *ACL*, volume 2, 2016.
366. Zhiwen Yu, Zhitao Wang, Liming Chen, Bin Guo, and Wenjie Li. Featuring, Detecting, and Visualizing Human Sentiment in Chinese Micro-Blog. *TKDD*, 10(4):48, 2016.
367. Majed Al Zayer and Mehmet Hade Gunes. Analyzing the Use of Twitter to Disseminate Visual Impairments Awareness Information. In *ASONAM*, 2017.
368. Chao Zhang, Dongming Lei, Quan Yuan, Honglei Zhuang, Lance Kaplan, Shaowen Wang, and Jiawei Han. GeoBurst+: Effective and Real-Time Local Event Detection in Geo-Tagged Tweet Streams. *ACM TIST*, 9(3):34, 2018.
369. Chao Zhang, Liyuan Liu, Dongming Lei, Quan Yuan, Honglei Zhuang, Tim Hanratty, and Jiawei Han. Trioveevent: Embedding-based Online Local Event Detection in Geo-tagged Tweet Streams. In *SIGKDD*, 2017.
370. Chao Zhang, Guangyu Zhou, Quan Yuan, Honglei Zhuang, Yu Zheng, Lance Kaplan, Shaowen Wang, and Jiawei Han. Geoburst: Real-time Local Event Detection in Geo-tagged Tweet Streams. In *SIGIR*, 2016.
371. Dan Zhang, Yan Liu, Richard D. Lawrence, and Vijil Chen-thamarakshan. Transfer Latent Semantic Learning: Microblog Mining with Less Supervision. In *AAAI*, 2011.
372. Dongxiang Zhang, Chee-Yong Chan, and Kian-Lee Tan. Processing Spatial Keyword Query As a Top-k Aggregation Query. In *SIGIR*, 2014.
373. Dongxiang Zhang, Liqiang Nie, Huanbo Luan, Kian-Lee Tan, Tat-Seng Chua, and Heng Tao Shen. Compact Indexing and Judicious Searching for Billion-Scale Microblog Retrieval. *ACM TOIS*, 35(3):27, 2017.
374. Dongxiang Zhang, Kian-Lee Tan, and Anthony K. H. Tung. Scalable Top-k Spatial Keyword Search. In *EDBT*, 2013.
375. Hao Zhang, Gang Chen, Beng Chin Ooi, Weng-Fai Wong, Shensen Wu, and Yubin Xia. "Anti-Caching"-based Elastic Memory Management for Big Data. In *ICDE*, 2015.
376. Jinxue Zhang, Rui Zhang, Jingchao Sun, Yanchao Zhang, and Chi Zhang. TrueTop: A Sybil-Resilient System for User Influence Measurement on Twitter. *IEEE/ACM TON*, 24(5):2834–2846, 2016.
377. Lei Zhang, Riddhiman Ghosh, Mohamed Dekhil, Meichun Hsu, and Bing Liu. Combining Lexicon-based and Learning-based Methods for Twitter Sentiment Analysis. *HP Laboratories, Technical Report HPL-2011*, 89, 2011.
378. Yihong Zhang, Claudia Szabo, Quan Z Sheng, and Xiu Susie Fang. SNAF: Observation Filtering and Location Inference for Event Monitoring on Twitter. *WWW Journal*, 21(2):311–343, 2018.
379. Yiming Zhang, Yujie Fan, Yanfang Ye, Xin Li, and Erin Winstanley. Utilizing Social Media to Combat Opioid Addiction Epidemic: Automatic Detection of Opioid Users from Twitter. In *AAAI Workshops*, 2018.
380. Zihua Zhang and Man Lan. Estimating Semantic Similarity between Expanded Query and Tweet Content for Microblog Retrieval. In *TREC*, 2014.
381. Jiang Zhao, Man Lan, and Tiantian Zhu. ECNU: Expression- and Message-level Sentiment Orientation Classification in Twitter using Multiple Effective Features. In *SemEval*, 2014.
382. Jianqiang Zhao, Xiaolin Gui, and Feng Tian. A New Method of Identifying Influential Users in the Micro-Blog Networks. *IEEE*

- Access*, 5:3008–3015, 2017.
383. Junzhou Zhao, John C.S. Lui, Don Towsley, Pinghui Wang, and Xiaohong Guan. Sampling Design on Hybrid Social-Affiliation Networks. In *ICDE*, 2015.
  384. Liang Zhao, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. Online Spatial Event Forecasting in Microblogs. *ACM TSAS*, 2(4):15, 2016.
  385. Wayne Xin Zhao, Yanwei Guo, Yulan He, Han Jiang, Yuexin Wu, and Xiaoming Li. We Know What You Want to Buy: A Demographic-based System for Product Recommendation on Microblogs. In *KDD*, 2014.
  386. Wayne Xin Zhao, Li Sui, He Yulan, Edward Y. Chang, Wen Ji-Rong, and Xiaoming Li. Connecting Social Media to E-commerce: Cold-start Product Recommendation using Microblogging Information. *TKDE*, 28(5):1147–1159, 2016.
  387. Xin Zheng, Aixin Sun, Sibao Wang, and Jialong Han. Semi-Supervised Event-related Tweet Identification with Dynamic Keyword Generation. In *CIKM*, 2017.
  388. Deyu Zhou, Liangyu Chen, and Yulan He. An Unsupervised Framework of Exploring Events on Twitter: Filtering, Extraction and Categorization. In *AAAI*, 2015.
  389. Deyu Zhou, Tianmeng Gao, and Yulan He. Jointly Event Extraction and Visualization on Twitter via Probabilistic Modelling. In *ACL*, volume 1, 2016.
  390. Xiangmin Zhou and Lei Chen. Event Detection over Twitter Social Media Streams. *PVLDB*, 23(3):381–400, 2014.
  391. Yiwei Zhou, Alexandra I. Cristea, and Lei Shi. Connecting Targets to Tweets: Semantic Attention-Based Model for Target-Specific Stance Detection. In *WISE*, 2017.
  392. Rui Zhu, Bin Wang, Xiaochun Yang, Baihua Zheng, and Guoren Wang. SAP: Improving Continuous Top-K Queries over Streaming Data. In *ICDE*, 2018.
  393. Xiang Zhu, Jiuming Huang, Sheng Zhu, Ming Chen, Chenlu Zhang, Zhenzhen Li, Huang Dongchuan, Zhao Chengliang, Aiping Li, and Yan Jia. NUDTSNA at TREC 2015 Microblog Track: A Live Retrieval System Framework for Social Network based on Semantic Expansion and Quality Model. In *TREC*, 2015.
  394. Tiago Zini, Karin Becker, and Marcelo Dias. INF-UFRGS at SemEval-2017 Task 5: A Supervised Identification of Sentiment Score in Tweets and Headlines. In *SemEval*, 2017.