Who? What? Event Tracking Needs Event Understanding

Nicholas Mamo[a], Joel Azzopardi and Colin Layfield[b]
Faculty of ICT, University of Malta, Malta

Keywords: Twitter, Topic Detection and Tracking, Information Retrieval, Event Modelling and Mining.

Abstract: Humans first learn, then think and finally perform a task. Machines neither learn nor think, but we still expect them to perform tasks as well as humans. In this position paper, we address the lack of understanding in Topic Detection and Tracking (TDT), an area that builds timelines of events, but which hardly understands events at all. Without understanding events, TDT has progressed slowly as the community struggles to solve the challenges of modern data sources, like Twitter. We explore understanding from different perspectives: what it means for machines to understand events, why TDT needs understanding, and how algorithms can generate knowledge automatically. To generate understanding, we settle on a structured definition of events based on the four Ws: the Who, What, Where and When. Of the four Ws, we focus especially on the Who and the What, aligning them with other research areas that can help TDT generate event knowledge automatically. In time, understanding can lead to machines that not only track events better, but also model and mine them.

1 INTRODUCTION

There is nothing revolutionary about the idea that understanding could improve machine performance—certainly not in Topic Detection and Tracking (TDT). In 1996, TDT started as a way to discover news topics from a stream of documents, and just two years later, Allan et al. (1998) had already mooted understanding as a way to improve event tracking. Later, TDT embraced Twitter as its main data source, and the research community again evoked knowledge as an intuitive way to make sense of events (Bontcheva and Rout, 2014). Even today, understanding is too intuitive to be considered revolutionary, which makes it all the more perplexing why the TDT community never explored event understanding in depth.

While TDT generates understanding about events, such as by detecting when something happens, knowledge rarely drives event tracking (De Boom et al., 2015). However, event knowledge does not have to be complex: it can be as simple as Kubo et al. (2013)'s 33-term football lexicon, with words like *goal* or *foul*. In this position paper, we argue that TDT no longer affords to ignore understanding, and examine how machines can generate knowledge automatically. We make the following contributions:

- TDT's ventures into understanding have created several interpretations of event knowledge. In this paper, we explore and contrast these perspectives in the context of TDT and related areas.
- TDT's challenges have increased since Allan et al. (1998) first proposed understanding, but so have the opportunities. In this paper, we discuss how event understanding can give TDT a new relevance to describe and model events.
- It is difficult for TDT to generate event knowledge without understanding events. In this paper, we argue that the TDT community needs to interpret events in a structured manner and we settle on the four Ws as a solution: the Who, What, Where and When. The Who and the What have historically been challenging for TDT to define, but we align them with other research areas that TDT can use to generate understanding.

The rest of this position paper is structured as follows. In Section 2 we explore how the TDT community has interpreted events and understanding, and what event knowledge can look like. Next, in Section 3 we explore how TDT can apply understanding to improve performance, describe events and eventually model and mine events. Then, in Section 4 we propose a structured definition of events based on the four Ws and discuss how TDT can exploit research in adjacent areas. Section 5 summarizes our position.

[a] https://orcid.org/0000-0001-5452-5548
[b] https://orcid.org/0000-0002-1868-4258

Mamo, N., Azzopardi, J. and Layfield, C.
Who? What? Event Tracking Needs Event Understanding.
DOI: 10.5220/0010650500003064
In *Proceedings of the 13th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2021) - Volume 1: KDIR*, pages 139-146
ISBN: 978-989-758-533-3; ISSN: 2184-3228

Figure 1: Liverpool F.C.'s terse, but expressive tweet after conceding a goal against Leeds United.

2 WHAT DO WE UNDERSTAND BY UNDERSTANDING?

The TDT community does not really understand understanding, although interpretations abound. Some researchers approached understanding through linguistics, like synonymy (Madani et al., 2014) or topic modelling (De Boom et al., 2015). Others adopted a more structured approach, splitting events into the Who, What, Where and When (Allan et al., 1998; Panagiotou et al., 2016). Such interpretations appear throughout TDT research, but no consensus exists on how to understand events, nor even on how to define events (Saeed et al., 2019). The most common definition (Farzindar and Khreich, 2015) is also the earliest, but it is no more expressive than it was in 1998: an event is "something that happens at a particular time and place" (Allan et al., 1998).

It is difficult to define event understanding without a comprehensive definition of events. In the absence of an expressive definition, in this position paper we ask a different question: what would the ideal understanding look like? We focus only on specified, or planned, events (Farzindar and Khreich, 2015), like football matches and elections, because most have a fixed structure, which facilitates understanding. Before a football match starts, for example, we know what can happen and how many teams and players will participate. We draw inspiration from how humans understand specified events.

Humans need very little information to understand events. When Liverpool F.C. conceded a late equalizer against Leeds United on 19 April, 2021, their community manager had the unpleasant task of tweeting about the goal. The community manager tweeted just two, one-word sentences, shown in Figure 1: "Goal. Leeds."[1] Not only do these two sentences, short even for a tweet, retain the message's significance, but their brevity makes the tweet more incisive. Liverpool F.C.'s tweet shows just how little information is required for a human to understand what happened. A handful of well-chosen keywords, in no particular order, suffice for humans to understand topics (Hsieh et al., 2012). A few more keywords can describe entire domains; with just 33 phrases, Kubo et al. (2013) describe football matches, from goals to fouls.

More complex knowledge can improve understanding further. Löchtefeld et al. (2015)'s football knowledge base combined two elements: manually-crafted patterns to extract topics, such as goals and yellow cards, and a manually-constructed database of the German Bundesliga's teams and players. These structures represent the ideal form of understanding: machine-readable information that allows machines to infer new and more complex machine-readable information, such as which team scored.

Nevertheless, expecting a human to transfer their knowledge to a machine, like Löchtefeld et al. (2015) did, is unreasonable, infeasible and not scalable (Buntain et al., 2016; Syed et al., 2016). Existing semantic structures, like WordNet, are not comprehensive enough either (Syed et al., 2016). WordNet lists 16 senses of the term *cross*, which appears in Kubo et al. (2013)'s football lexicon, but none relate to football. Therefore TDT needs automatic ways to generate understanding.

Unfortunately, the existing automatic interpretations of understanding appear insufficient next to manually-defined knowledge. Rudra et al. (2015)'s algorithm, for example, automatically extracts content words, like *killed* or *stranded*, to describe disaster events. Rudra et al. (2015) interpreted content words from a linguistic perspective based on Part of Speech (POS) tagging: numerals, nouns and verbs. Outwardly, these linguistic choices make sense; most of Kubo et al. (2013)'s terms are nouns, like *corner* or *feint*. However, hundreds of other nouns and verbs fit Rudra et al. (2015)'s broad interpretation of content words without being content-bearing in most domains, like *seat* in football matches. The contrast is stark: Rudra et al. (2015)'s automatic understanding is ambiguous and inaccurate, while Kubo et al. (2013)'s and Löchtefeld et al. (2015)'s manually-defined knowledge is unambiguous and accurate. Bridging the gap between the raw understanding of machines and the refined understanding of humans remains a challenge.

[1]twitter.com/LFC/status/1384246547257327625, last accessed on July 25, 2021

3 WHY DO WE NEED UNDERSTANDING?

Although the event tracking community has not studied understanding in depth, applications for event knowledge are plentiful literature. It is, at least, clear why we need understanding—and why would understanding *not* be beneficial?

TDT faced challenges from the start, which is what prompted Allan et al. (1998) to propose understanding. Traditionally, the research community approached the TDT problem through clustering, grouping together formal news articles to form topics (Fung et al., 2005). Clustering-based techniques, better known as document-pivot approaches, exploited a well-researched area, but they also inherited many of clustering's problems. For example documents, and by extension topics, often end up fragmented in different clusters, and clustering algorithms are generally highly-parametric (Fung et al., 2005; Aiello et al., 2013; Ifrim et al., 2014).

Understanding news topics better might have improved clustering (Allan et al., 1998), but the predominant solution became feature-pivot approaches, the second broad family of TDT techniques (Fung et al., 2005). Feature-pivot algorithms rely on the changes of certain features in the document stream instead of on the documents themselves. These features can be an unexpected increase in volume when a disaster occurs, or an associated keyword, like *crash*, suddenly becoming more popular.

Feature-pivot techniques eliminated some of clustering's challenges, but they also contributed new issues. An isolated keyword extracted to describe an emergent topic, like *president*, does not express the narrative adequately, unlike a cluster, which tells a comprehensive story through its news articles. Groups of terms are not cohesive either, especially since term correlations can be deceptive (Aiello et al., 2013; Hasan et al., 2019).

Again, understanding might have solved some of feature-pivot techniques' issues. With better understanding, algorithms might have been able to improve accuracy by focusing on a selection of descriptive keywords, like the ones in Liverpool F.C.'s tweet. However, instead of exploring understanding, TDT looked for a new relevance on Twitter.

Twitter launched in 2006, a year after Fung et al. (2005) proposed feature-pivot methods. Twitter's launch gave TDT a new utility. Before, TDT reproduced topics after the media had reported the news; now, TDT could discover the news for itself from Twitter. The research community exploited TDT's new-found relevance, but Twitter also introduced new challenges.

On Twitter, TDT algorithms must handle large volumes of tweets arriving too fast for heavy processing in real-time systems (Farzindar and Khreich, 2015; Panagiotou et al., 2016; Saeed et al., 2019). Twitter's volume and velocity impact document-pivot approaches the worst (Panagiotou et al., 2016). Clustering requires heavy processing and some researchers consider document-pivot approaches to be infeasible on Twitter (Panagiotou et al., 2016). Document-pivot techniques did survive Twitter, after all, but inferior on-line clustering methods have to suffice (McMinn and Jose, 2015). Even here, understanding events could have pushed clustering algorithms to work smarter, if not faster. For example, McMinn and Jose (2015) assume that named entities drive events and remove any tweet without named entities, filtering 90% of tweets.

The brevity of tweets represents a bigger problem than volume or velocity, however, and not just for TDT. Most Information Retrieval (IR) approaches are designed for longer and more formal content than tweets. Brevity leads to sparsity, which harms even well-established IR methods, like term-weighting schemes (Samant et al., 2019). We cannot assume that the traditional IR methods that worked on formal documents work just as well on tweets (Panagiotou et al., 2016; Saeed et al., 2019). Mishra and Diesner (2016) discuss Named Entity Recognition (NER)'s difficulties on Twitter at length, and propose their own NER algorithm, tailored specifically to Twitter's unruly orthography. Tellingly, Mishra and Diesner (2016)'s algorithm uses gazetteers, themselves a form of understanding.

Above all, Twitter is noisy. While news articles become newsworthy as soon as the media publishes them (Hua et al., 2016), users talk about their daily lives, react to news and share opinions (Hua et al., 2016; Panagiotou et al., 2016; Hasan et al., 2019; Saeed et al., 2019)—sometimes in the same tweet. Dealing with Twitter's noise is a momentous task, but what is noise if not a failure to understand what is relevant to an event and what is irrelevant? Very few algorithms have explored noise filtering by understanding which keywords are relevant to events (Hua et al., 2016; Zhou et al., 2017; Hossny and Mitchell, 2018).

Understanding has rarely been the solution to TDT's challenges. Instead, TDT's solution has been to pummel Twitter's best virtue, its large volume, by aggressively filtering all retweets (McMinn and Jose, 2015; Huang et al., 2018) because they are redundant or introduce bias (McMinn and Jose, 2015; Saeed et al., 2019). And TDT's solution has been to retain only the largest clusters because they are more likely to be newsworthy (McMinn and Jose, 2015; Hasan et al., 2019). The solution has rarely been understanding, and TDT keeps suffering the consequences (Bontcheva and Rout, 2014; Madani et al., 2014; De Boom et al., 2015; Panagiotou et al., 2016).

Today, TDT's lack of understanding shows through its many difficulties. Few approaches support unpopular events because aggressive filtering is infeasible in small datasets. Excessive filtering also penalizes algorithms, such as when McMinn and Jose (2015), and Hasan et al. (2019) reject any cluster with fewer than 10 tweets—a steep threshold in the small datasets of unpopular events. Even in massively-popular events, an overabundance of caution leads to TDT algorithms missing most non-key topics, like yellow cards in football matches, although they reliably capture key topics, like goals (Löchtefeld et al., 2015).

Understanding can address TDT's performance problems on Twitter, but it can also allow machines to describe events, not just detect them (Kubo et al., 2013; Panagiotou et al., 2016). Eventually, understanding can also be a way to make sense of events (Bontcheva and Rout, 2014) and automatically create machine-readable knowledge through event modelling and mining (Chen and Li, 2020). First, however, TDT needs to understand understanding.

4 THE WHO, WHAT, WHERE AND WHEN

Identifying what knowledge best characterizes events is challenging (Mohd, 2007; Madani et al., 2014), and as a result, the TDT community has struggled to explore understanding. Ironically, the earliest research in TDT had the clearest idea of what elements make up events. Allan et al. (1998) and others (Makkonen et al., 2004; Mohd, 2007; Zhou et al., 2017) dissect events into the four Ws: the Who, What, Where and When. Although the four Ws were never widely-adopted in TDT, today they are resurging as part of more intensive research into event modelling and mining (Rudnik et al., 2019; Chen and Li, 2020).

The four Ws make sense because they are not new to events, even beyond TDT. For a long time, these four elements, along with the Why and How, have been a journalistic best practice to describe events (Rudnik et al., 2019). BBC's tweet, shown in Figure 2, hinges on the Who, What and Where to tell a story in one sentence: "Indonesian Navy [Who] hunting for submarine that has gone missing [What] in waters north of island of Bali [Where]"[2]. The tweet's publication time implies the When.

Figure 2: Journalists rely on the Who, What, Where and When to describe events.

In this paper, we too advocate for the four Ws, described in Table 1, to characterize events. We focus specifically on the Who and the What since our primary aim is to understand specified events. The Where is more useful in unspecified TDT to distinguish between events happening in different locations. The When is implicit in TDT since algorithms detect *when* something happens. Therefore we describe the Who and the What in detail next.

4.1 The Who

The Who is understood well throughout IR thanks to years of research into NER. NER gave TDT a rare relief from having to define the Who because we understand what a named entity looks like: usually an organization, a person or a place. Moreover, barring NER's difficulties on Twitter, TDT exploited existing tools to identify the Who.

In the beginning, the TDT community assigned named entities a simple, distinguishing role. An election candidate runs for office in one place at a time, the reasoning went, so named entities can separate similar events. Therefore Makkonen et al. (2004) and Zhou et al. (2017) represent events as four vectors, one of which stores named entities. Others boost the importance of named entities (Aiello et al., 2013; Ifrim et al., 2014), or prioritize tweets containing

[2]twitter.com/BBCBreaking/status/1384817851219988480, last accessed on July 25, 2021

Table 1: The four basic elements of an event's structure: the Who, What, Where and When.

Element	Description
Who	The participants who affect or are affected by the specified event while the event is ongoing (Mamo et al., 2021)
What	The concepts that are related to or describe the specified event and its topics
Where	The place where the specified event is taking place
When	The date and time when a topic happens during the specified event

names when summarizing (Kubo et al., 2013).

Later iterations gave named entities more prominent roles. Löchtefeld et al. (2015)'s knowledge base allows their pattern-matching algorithm to create machine-readable information about teams and players in football matches. McMinn and Jose (2015), and Huang et al. (2018) use NER to build separate timelines for each named entity, making it possible to explore topics related to individual entities.

It is tempting to think of NER as providing understanding, but these applications do not stand up to scrutiny. Even ignoring NER's difficulties with Twitter's noise, TDT uses named entities too recklessly. Most applications of the Who in TDT literature conveniently, but mistakenly, assume that named entities are equivalent to understanding. There is a stark contrast between Löchtefeld et al. (2015)'s knowledge base and NER. On the one hand, Löchtefeld et al. (2015)'s knowledge base contains every single team and player in the Bundesliga, and not a single extra named entity. On the other hand, NER captures many incorrect named entities and misses many correct ones (Mamo et al., 2021).

Working towards event understanding, in our previous work (Mamo et al., 2021) we distinguished between named entities and participants. Participants may be named entities, but more importantly, they play an active role in the event. Therefore to reconcile named entities with event participants, we proposed Automatic Participant Detection (APD). APD first confirms which named entities qualify as participants, and then looks for other participants missed by NER. With APD, we demonstrated how machines can achieve broad coverage of the participants even before the event starts.

4.2 The What

Understanding the What of events is more difficult than the Who. We recognize Kubo et al. (2013)'s 33 words and phrases, like *cross* and *foul*, as football-related terms, but what makes these words terms? Unlike NER, which is "relatively well-assessed", the notion of a term remains "underspecified" (Velardi et al., 2001).

Early TDT research adopted a simplified view of the What based on linguistics. Makkonen et al. (2004) represent event profiles using four vectors corresponding to the four Ws. The event profile accepts anything that is not already in the Who, Where and When as part of the What, with only minimal filtering based on POS tagging. Similar interpretations of the What are common, with a particular focus on nouns and verbs because they describe events best (Liu et al., 2013). Later, synonymy (Madani et al., 2015) and word embedding (Farnaghi et al., 2020) also became a way of understanding the event domain more broadly.

Like NER substituting for the Who, linguistics purporting to understand the What failed TDT (Mohd, 2007). TDT does not need any understanding; it needs good understanding applied right. Linguistic choices may be convenient, but they are not good understanding. Terms are supposed to be meaningful—"the subject, occasion, body or activity [...] involved in the event" (Mohd, 2007)—but what meaning do linguistics represent? POS tagging understands the syntax of languages, not events, and synonymy and word embedding understand the semantics of languages. Unsurprisingly, Makkonen et al. (2004)'s simple interpretation of the What worsened results.

To explore real understanding, TDT needs to understand the What of events better. Already, the TDT community recognizes that most events are part of a broader domain; all football matches, for example, share a similar vocabulary (Yang et al., 2002; Hua et al., 2016). Hua et al. (2016) distinguish between the general domain terms and particular event terms, which change from one event to the other. Adopting Hua et al. (2016)'s interpretation, Kubo et al. (2013)'s 33 words and phrases would qualify as domain terms because they are relevant to all football matches, whereas players and teams belong to particular events, so they are event terms.

We identify domain terms as the most promising avenue to form a basic understanding of the What in events. Extracting domain terms aligns with the research area of Automatic Term Extraction (ATE) (Astrakhantsev et al., 2015). In fact, ATE commonly shares many of the same linguistic constraints as POS tagging in TDT, normally nouns and verbs (Astrakhantsev et al., 2015). Unlike TDT, however, ATE complements linguistics with a statistical measure, termhood, to analyse the fit of a word or a phrase as a domain term (Maldonado and Lewis, 2016).

TDT has delved into ATE only briefly (Yang et al., 2002; Hua et al., 2016; Zhou et al., 2017; Hossny and Mitchell, 2018). Moreover, the existing work comes across as an afterthought, as if its sole purpose is to answer the question: can domain terms help TDT? TDT never evaluates the terms themselves, and rarely proposes its own sophisticated termhood measures. Instead, TDT relies on ATE's rudimentary baselines, like the chi-square (Yang et al., 2002). Our solace is in the fact that even in its raw form, TDT's experimentation with ATE improves results, if only by eliminating off-topic tweets (Yang et al., 2002; Hua et al., 2016; Zhou et al., 2017; Hossny and Mitchell, 2018).

Naturally, resting on the domain terms means missing the exceptional event terms. No one would think to include *parachute* in Kubo et al. (2013)'s list of football terms, so the lexicon would miss a parachutist landing on the pitch during a football match[3]. Automatic understanding is a trade-off for the exceptional, but the exceptional remains extraordinary. While Buntain et al. (2016) question this trade-off, we question whether TDT affords to continue ignoring understanding just to subserve the extraordinary. Even Buntain et al. (2016) revise their position in the end.

Just like APD adapted NER to fit the Who, TDT also needs to adapt ATE to fit the What, but it is not a straightforward endeavour. To the best of our knowledge, so far the ATE community has never studied event domains. Neither could we find any ATE research on Twitter' noisy and informal content. Moreover, aside from Hossny and Mitchell (2018)'s work, the TDT research that uses domain terms on Twitter normally extracts the vocabulary from formal documents (Hua et al., 2016; Zhou et al., 2017). ATE might not be ready for events or tweets, but until it is, TDT will not be ready to understand events either.

5 CONCLUSION

Since its early days, the TDT community has perceived event tracking as little more than a sensor, as if TDT's only purpose is to detect when something happens, not explain what happens and who is involved. As a result, TDT has ended up isolated, barely able to drive summarization, or event modelling and mining. Worse still, Twitter only exacerbated TDT's performance challenges signalled by Allan et al. (1998).

In this position paper, we argued that TDT no longer affords to ignore event understanding. We proposed the four Ws as a solution, and linked TDT with APD and ATE to automatically understand the Who and the What. The four Ws also align event tracking with event modelling and mining, eventually allowing machines to describe not just what happened, but also why and how it happened (Chen and Li, 2020).

However, the road to event understanding is a long one. For too long, TDT has relied on simple definitions and techniques to acquire knowledge about events, like NER and POS tagging. TDT needs to move beyond NER and linguistics and study event understanding more seriously, which means reconsidering NER's role in understanding the Who, and adapting ATE to understand the What. As much as it is clear that TDT needs understanding, it is also clear that TDT cannot progress alone.

[3] france24.com/en/20191020-parachutist-gatecrashes-inter-milan-s-win-at-sassuolo, last accessed on July 25, 2021

ACKNOWLEDGEMENTS

The research work disclosed in this publication is funded by the Tertiary Education Scholarships Scheme.

REFERENCES

Aiello, L. M., Petkos, G., Martin, C., Corney, D., Papadopoulos, S., Skraba, R., Goker, A., Kompatsiaris, I., and Jaimes, A. (2013). Sensing Trending Topics in Twitter. *IEEE Transactions on Multimedia*, 15(6):1268–1282.

Allan, J., Papka, R., and Lavrenko, V. (1998). On-Line New Event Detection and Tracking. In *SIGIR '98: Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 37–45, Melbourne, Australia. ACM.

Astrakhantsev, N., Fedorenko, D., and Turdakov, D. (2015). Methods for Automatic Term Recognition in Domain-Specific Text Collections: A Survey. *Programming and Computer Software*, 41(6):336–349.

Bontcheva, K. and Rout, D. (2014). Making Sense of Social Media Streams through Semantics: a Survey. *Semantic Web*, 5(5):373–403.

Buntain, C., Lin, J., and Golbeck, J. (2016). Discovering Key Moments in Social Media Streams. In *2016 13th IEEE Annual Consumer Communications & Networking Conference (CCNC)*, page 366–374, Las Vegas, NV, USA. IEEE.

Chen, X. and Li, Q. (2020). Event Modeling and Mining: a

Long Journey Toward Explainable Events. *The VLDB journal*, 29(1):459–482.

De Boom, C., Van Canneyt, S., and Dhoedt, B. (2015). Semantics-driven event clustering in Twitter feeds. In *Proceedings of the 5th Workshop on Making Sense of Microposts*, page 2–9, Florence, Italy. CEUR.

Farnaghi, M., Ghaemi, Z., and Mansourian, A. (2020). Dynamic Spatio-Temporal Tweet Mining for Event Detection: A Case Study of Hurricane Florence. *International Journal of Disaster Risk Science*, 11(3):378–393.

Farzindar, A. and Khreich, W. (2015). A Survey of Techniques for Event Detection in Twitter. *Computational Intelligence*, 31(1):132–164.

Fung, G. P. C., Yu, J. X., Yu, P. S., and Lu, H. (2005). Parameter Free Bursty Events Detection in Text Streams. In *VLDB '05: 31st International Conference on Very Large Data Bases*, page 181–192, Trondheim, Norway. VLDB Endowment.

Hasan, M., Orgun, M. A., and Schwitter, R. (2019). Real-Time Event Detection from the Twitter Data Stream Using the TwitterNews+ Framework. *Information Processing & Management*, 56(3):1146–1165.

Hossny, A. H. and Mitchell, L. (2018). Event Detection in Twitter: A Keyword Volume Approach. In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*, page 1200–1208, Singapore. IEEE.

Hsieh, L.-C., Lee, C.-W., Chiu, T.-H., and Hsu, W. (2012). Live Semantic Sport Highlight Detection Based on Analyzing Tweets of Twitter. In *2012 IEEE International Conference on Multimedia and Expo*, page 949–954, Melbourne, VIC, Australia. IEEE.

Hua, T., Chen, F., Zhao, L., Lu, C.-T., and Ramakrishnan, N. (2016). Automatic Targeted-Domain Spatiotemporal Event Detection in Twitter. *GeoInformatica*, 20(4):765–795.

Huang, Y., Shen, C., and Li, T. (2018). Event Summarization for Sports Games using Twitter Streams. *World Wide Web*, 21(3):609–627.

Ifrim, G., Shi, B., and Brigadir, I. (2014). Event Detection in Twitter using Aggressive Filtering and Hierarchical Tweet Clustering. In *Proceedings of the SNOW 2014 Data Challenge*, page 33–40, Seoul, Korea. CEUR.

Kubo, M., Sasano, R., Takamura, H., and Okumura, M. (2013). Generating Live Sports Updates from Twitter by Finding Good Reporters. In *Proceedings of the 2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, volume 1, pages 527–534, Atlanta, Georgia, USA. IEEE Computer Society.

Liu, C., Xu, R., and Gui, L. (2013). Burst Events Detection on Microblogging. In *Proceedings of the 2013 International Conference on Machine Learning and Cybernetics*, page 1921–1924, Tianjin, China. IEEE.

Löchtefeld, M., Jäckel, C., and Krüger, A. (2015). TwitSoccer: Knowledge-Based Crowd-Sourcing of Live Soccer Events. In *MUM '15: Proceedings of the 14th International Conference on Mobile and Ubiquitous Multimedia*, pages 148–151, Linz, Austria. ACM.

Madani, A., Boussaid, O., and Zegour, D. (2015). Real-Time Trending Topics Detection and Description from Twitter Content. *Social Network Analysis and Mining*, 5(1):1–13.

Madani, A., Boussaid, O., and Zegour, D. E. (2014). What's Happening: A Survey of Tweets Event Detection . In *INNOV 2014 : Proceedings of the Third International Conference on Communications, Computation, Networks and Technologies*, page 16–22, Nice, France. IARIA.

Makkonen, J., Ahonen-Myka, H., and Salmenkivi, M. (2004). Simple Semantics in Topic Detection and Tracking. *Information Retrieval*, 7(3):347–368.

Maldonado, A. and Lewis, D. (2016). Self-Tuning Ongoing Terminology Extraction Retrained on Terminology Validation Decisions. In *Proceedings of the 12th International Conference on Terminology and Knowledge Engineering*, page 91–100, Copenhagen, Denmark. Copenhagen Business School.

Mamo, N., Azzopardi, J., and Layfield, C. (2021). An Automatic Participant Detection Framework for Event Tracking on Twitter. *Algorithms*, 14(3):92.

McMinn, A. J. and Jose, J. M. (2015). Real-Time Entity-Based Event Detection for Twitter. In Mothe, J., Savoy, J., Kamps, J., Pinel-Sa, Pinel-Sauvagnat, K., Jones, G., San Juan, E., Capellato, L., and Nicola, F., editors, *CLEF 2015: Experimental IR Meets Multilinguality, Multimodality, and Interaction*, pages 65–77, Toulouse, France. Springer International Publishing.

Mishra, S. and Diesner, J. (2016). Semi-Supervised Named Entity Recognition in Noisy-Text. In *Proceedings of the 2nd Workshop on Noisy User-generated Text*, page 203–212, Osaka, Japan. The COLING 2016 Organizing Committee.

Mohd, M. (2007). Named Entity Patterns Across News Domains. In *Proceedings of the BCS IRSG Symposium: Future Directions in Information Access 2007*, pages 1–6, Glasgow, Scotland. BCS, The Chartered Institute for IT.

Panagiotou, N., Katakis, I., and Gunopulos, D. (2016). *Detecting Events in Online Social Networks: Definitions, Trends and Challenges*, volume 9580 of *Lecture Notes in Computer Science*. Springer International Publishing, Cham.

Rudnik, C., Ehrhart, T., Ferret, O., Teyssou, D., Troncy, R., and Tannier, X. (2019). Searching News Articles Using an Event Knowledge Graph Leveraged by Wikidata . In *WWW '19: Companion Proceedings of The 2019 World Wide Web Conference*, page 1232–1239, San Francisco, CA, USA. Association for Computing Machinery.

Rudra, K., Ghosh, S., Ganguly, N., Goyal, P., and Ghosh, S. (2015). Extracting Situational Information from Microblogs during Disaster Events. In *Proceedings of the 24th ACM International on conference on information and knowledge management*, pages 583–592, Melbourne, Australia. ACM.

Saeed, Z., Abbasi, R. A., Maqbool, O., Sadaf, A., Razzak, I., Daud, A., Aljohani, N. R., and Xu, G. (2019). What's Happening Around the World? A Survey and Framework on Event Detection Techniques on Twitter. *Journal of Grid Computing*, 17(2):279–312.

Samant, S. S., Bhanu Murthy, N. L., and Malapati, A. (2019). Improving Term Weighting Schemes for Short Text Classification in Vector Space Model. *IEEE Access*, 7:166578–166592.

Syed, S., Spruit, M., and Borit, M. (2016). Bootstrapping a Semantic Lexicon on Verb Similarities. In *Proceedings of the 8th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2016) - Volume 1: KDIR,*, volume 1, page 189–196, Porto, Portugal. SciTePress.

Velardi, P., Missikoff, M., and Basili, R. (2001). Identification of Relevant Terms to Support the Construction of Domain Ontologies. In *Proceedings of the ACL 2001 Workshop on Human Language Technology and Knowledge Management*, pages 1–8, Toulouse, France. Association for Computational Linguistics.

Yang, Y., Zhang, J., Carbonell, J., and Jin, C. (2002). Topic-Conditioned Novelty Detection. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 688–693, Alberta, Canada. ACM.

Zhou, D., Chen, L., Zhang, X., and He, Y. (2017). Unsupervised Event Exploration from Social Text Streams. *Intelligent Data Analysis*, 21(4):849–866.