

An Incremental Clustering Baseline for Event Detection on Twitter

Marjolaine Ray¹, Qi Wang¹, Frédérique Mélanie-Becquet¹, Thierry Poibeau¹,
Béatrice Mazoyer²

¹Lattice (CNRS & École normale supérieure-PSL & U. Sorbonne nouvelle), Paris, France

²médialab, Sciences Po, Paris, France

Correspondence: beatrice.mazoyer@sciencespo.fr

Abstract

Event detection in text streams is a crucial task for the analysis of online media and social networks. One of the current challenges in this field is establishing a performance standard while maintaining an acceptable level of computational complexity. In our study, we use an incremental clustering algorithm combined with recent advancements in sentence embeddings. Our objective is to compare our findings with previous studies, specifically those by Cao et al. (2024) and Mazoyer et al. (2020). Our results demonstrate significant improvements and could serve as a relevant baseline for future research in this area.

1 Introduction

With the development of social media, the ability to recognize events in streams of short texts—particularly tweets—has become increasingly important. This process, called event recognition, involves identifying significant occurrences within large volumes of data, posing various challenges. A key component of this task is defining a clear and operational concept of what qualifies as an event. In this paper, we will use a working definition of event, as proposed by McMinn et al. (2013). The authors propose in fact a double definition: “**Definition 1:** An event is a **significant** thing that happens at some specific time and place”. This needs to be completed by the definition of what ‘significant’ means, so they add: “**Definition 2:** Something is significant if it may be discussed in the media. For example you may read a news article or watch a news report about it”. Because this definition has been used to build other corpora, it can be considered functional. As a result, corpora created with this definition should be comparable, with different annotators likely producing similar outcomes.

One of the main challenges in event recognition is then being able to cluster different texts

that refer to the same event. This difficulty arises from the wide range of expressions used to describe similar events. Different sources and users may refer to the same event using different expressions, making it essential for recognition systems to account for synonymy, paraphrasing, and other linguistic variations. Moreover, the temporal dimension is also a critical parameter in event recognition. The timing of events and the sequence in which they are reported can significantly impact the interpretation and relevance of the information extracted. Another challenge in event detection on social networks is the sheer volume of messages posted on these platforms: an effective algorithm must be capable of processing millions of tweets within a reasonable time frame. Many studies propose computationally intensive methodologies that are impractical for many real-world applications. Therefore, research in this field needs to establish baselines on publicly accessible datasets that are both performant and time-efficient.

The primary objective of this paper is thus to establish a performance standard for event detection while maintaining an acceptable level of computational complexity. Our approach involves the use of an incremental clustering algorithm enhanced by recent advancements in sentence embeddings. Specifically, we build upon the incremental clustering algorithm introduced by Mazoyer et al. (2020) in their dataset publication. While effective at the time, their approach relied on lexical descriptions that may now be outdated due to the development of new word embedding techniques, particularly those stemming from recent large language models based on the transformer architecture (Vaswani et al., 2017). In our study, we utilize SentenceBERT (Reimers and Gurevych, 2019), a model that is especially noteworthy for its ability to encode entire sentences from individual word encodings.

The structure of the paper is as follows: First, we will review recent work in the domain. Next,

we will detail our method and experiments. Finally, we will present and discuss our results, concluding with a broader discussion. We conduct experiments on two large public Twitter datasets to demonstrate the state of the art performance, efficiency, and robustness of this method (note that our code is publicly accessible¹). We then aim to compare our results with previous studies in the field, specifically those by Cao et al. (2024) and Mazoyer et al. (2020). By leveraging these advanced sentence embeddings, we demonstrate that our implementation surpasses more recent and complex approaches in both time-efficiency and the quality of detected events. These short-text representations provide a sophisticated understanding of language and context, allowing for more accurate and nuanced event recognition.

2 Related Work

Hasan et al. (2018) conducted a comprehensive review of event detection techniques on Twitter. Like these authors, we identify three main categories of methods: ‘term-interestingness-based’ approaches, topic modeling, and incremental clustering. However, we expand upon their typology by adding a fourth category: graph-based approaches.

“Term-Interestingness-Based” Approaches. These methods involve monitoring terms that are probably associated with an event, often identified by a sudden increase in the frequency of certain terms. Typically, they return the top trending events on Twitter. These approaches generally do not allow the detection of low-bursty events.

Topic Modelling. Topic models are widely used techniques derived from Latent Dirichlet Allocation (LDA) (Blei et al., 2003) to uncover the thematic structure within a collection of textual documents. Several works have been interested in adapting this method to make topics evolve over time, and to adapt to the short format of tweets by restricting the number of topics associated with a document. Likhitha et al. (2019) propose a survey of topic modeling methods adapted to short texts.

Incremental Clustering. This family of methods derives from the Topic Detection and Tracking (TDT) initiative (Allan et al., 1998), aimed at identifying and following events in a stream of broadcast news stories. The task of detecting new events

(First Story Detection) involves representing documents as vectors in a semantic space. Each new document is compared to existing ones (or to a set of past documents within a time-window) and if its similarity to the closest document (or centroid) falls below a defined threshold, it is identified as a new story. This methodology was then adapted to event detection on Twitter (Petrovic et al., 2010; McMinn and Jose, 2015) with tf-idf (Sparck Jones, 1972) as a vector representation of tweets. More recent works (Mazoyer et al., 2020; Qiu et al., 2021; Pradhan et al., 2024) use BERT (Devlin et al., 2019) or Sentence Transformers (Reimers and Gurevych, 2019) to produce a vector representation of tweets.

Graph-Based Approaches. These methods (Peng et al., 2022; Ren et al., 2022; Cao et al., 2024) leverage the semantic structure of social media, using anchors such as hashtags, user mentions, hyperlinks and named entities. They construct message graphs that include all candidate messages, linking those that share common attributes. The event detection task is then framed as a graph-partitioning problem.

3 Methodology

When working with social media data, one needs to consider both the textual similarity of the documents and their temporal proximity to avoid grouping together tweets posted at significantly different times. Since the number of events is not known in advance, the chosen algorithm does not require the number of events given a priori. Following the method by Mazoyer et al. (2020), we use an incremental clustering algorithm derived from the Topic Detection and Tracking (Allan et al., 1998) initiative.

Algorithm. This mini-batch First Story Detection (FSD) algorithm works as follows: documents are vectorized (we develop embedding methods in the subsequent section), sorted chronologically, and processed in batches of b documents. Each new batch is compared to a window of w previous documents in terms of cosine distance. For each batch document, if the distance to its nearest neighbor is smaller than a threshold t , it joins the same cluster as its nearest neighbor. Otherwise, the document joins a new cluster. The procedure is detailed in Algorithm 1, where δ denotes the cosine distance.

Short-Text Embeddings. In the work published by Mazoyer et al. (2020), the best performing

¹<https://github.com/medialab/twitter-incremental-clustering/>

Algorithm 1: "mini-batch" FSD

input: threshold t , window size w , batch size b , corpus $C = \{d_0 \dots d_{n-1}\}$ of n documents in chronological order

output: a list T of cluster ids for each document

```
1  $T \leftarrow []$ ;  $i \leftarrow 0$ ;  $j \leftarrow 0$ ;  
2 while  $i < n - b$  do  
3    $batch = \{d_i, \dots, d_{i+b-1}\}$ ;  
4   do in parallel  
5     for document  $d$  in batch do  
6       if  $T$  is empty then  
7          $cluster\_id(d) \leftarrow j$ ;  
8          $j \leftarrow j + 1$ ;  
9       else  
10         $d_{nearest} \leftarrow$  nearest neighbor of  $d$  in  $T$ ;  
11        if  $\delta(d, d_{nearest}) < t$  then  
12           $cluster\_id(d) \leftarrow cluster\_id(d_{nearest})$ ;  
13        else  
14           $cluster\_id(d) \leftarrow j$ ;  
15           $j \leftarrow j + 1$ ;  
16        if  $|T| \geq w$ ;  
17        then  
18          remove first document from  $T$   
19          add  $d$  to  $T$ ;  
20    $i \leftarrow i + b$ ;
```

embedding method is a tf-idf score where the df (document-frequency) is computed over the entire tweet dataset (millions of tweets). Over the past five years, numerous models have emerged, particularly large language models (LLMs), which are especially suited for this task as they encode both linguistic and world knowledge, making them highly effective in capturing the nuances and complexities of event detection. We use Sentence Transformers, also known as SBERT (Reimers and Gurevych, 2019), a BERT/RoBERTa (Devlin et al., 2019; Liu et al., 2019) fine-tuning architecture using Siamese networks. This model ensures that the resulting sentence embeddings are both semantically meaningful and comparable, using cosine distance.

Time Complexity. The time complexity of the FSD algorithm is $O(nw)$ (with n the number of documents in the collection and w the number of documents in the time window), since each document in the corpus is compared only with the last w documents in chronological order. In practice, when using the "mini-batch" FSD, computation time is inversely proportional to batch size, as illustrated in Figure 1.




Figure 1: Evolution of execution time and adjusted mutual information (AMI) of the "mini-batch" FSD algorithm depending on batch size b on the entire Event2012 corpus (68,841 documents).

4 Experiments

Baselines. We compare our results (FSD-SBERT) with HISEvent² (HE), the most recent paper on event detection: Cao et al.’s (2024) work on the partition of a graphical neural network representation of tweets using structural entropy minimization. We also evaluate the performance improvement achieved by using Sentence Transformers in

²<https://github.com/SELGroup/HISEvent>

comparison to the tf-idf vectors used in (TW)³ by Mazoyer et al. (2020).

Datasets. We conducted experiments using two extensive, publicly accessible tweets datasets: Event2012 (McMinn et al., 2013) and Event2018 (Mazoyer et al., 2020). The Event2012 dataset contains 150,000 English tweet IDs related to 506 distinct events over a four-week period. In contrast, Event2018 comprises 96,000 French tweet IDs corresponding to 257 unique events, all posted within a span of 23 days. For a fair comparison with baseline methods, we limit our analysis to the subset of the dataset used by Cao et al. (2024). Indeed, these authors downloaded the tweets recently after many were deleted. Their dataset, therefore, contains 68,841 tweets related to 503 events for Event2012 and 64,516 tweets related to 257 events for Event2018. We do not use the distinction adopted by Cao et al. (2024) between open-set (day-by-day detection) and closed-set (detection across the entire corpus), as we argue that events should be allowed to span multiple consecutive days. Therefore, we only evaluate our method on the complete corpus.

Short-Text Embeddings. We use Sentence Transformers (SBERT) models pre-trained on English and French corpora to compute vectors from tweets. Specifically, we use **all-mpnet-base-v2**⁴ for the English dataset and **Sentence-CamemBERT-Large**⁵ (Martin et al., 2020) for the French dataset.

Parameters. The mini-batch FSD algorithm takes three input parameters: the cosine distance threshold (t), the time-window size (w) and the batch-size (b). Consistently with Mazoyer et al. (2020), we set w to the average number of documents per day in each dataset, and the batch size to 8 documents. The threshold t depends on the type of text-embedding. It was optimized using grid-search and set to 0.5 for English and 0.55 for French.

Evaluation Metrics. We use the scikit-learn (Pedregosa et al., 2011) implementation of adjusted mutual information (AMI) (Vinh et al., 2009) and adjusted rand index (ARI) (Rand, 1971), which

³<https://github.com/ina-foss/twembeddings>

⁴<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

⁵<https://huggingface.co/dangvantuan/sentence-camembert-large>

		FSD-SBERT	HE	TW
dataset				
2012	ARI	0.63	0.50	0.39
	AMI	0.86	0.81	0.82
2018	ARI	0.55	0.44	0.25
	AMI	0.81	0.66	0.72

Table 1: ARI and AMI scores on two datasets: Event2012 (in English) and Event2018 (in French).

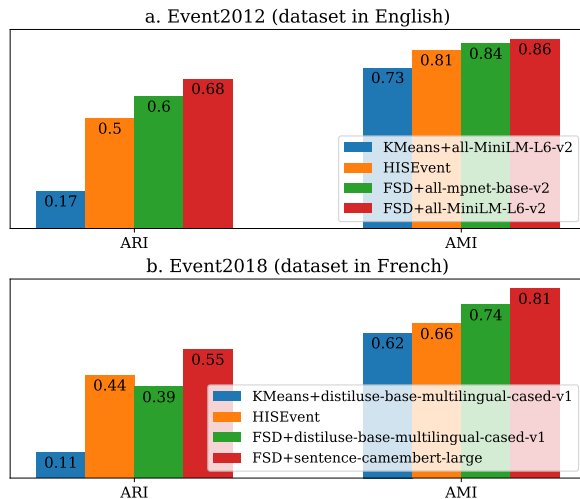


Figure 2: ARI and AMI scores with different SBERT models and different clustering algorithms. All FSD tests ran with $b = 8$ and $t = 0.55$.

are widely employed in event detection evaluation (Cao et al., 2024).

5 Results and Discussion

Performance. Table 1 compares the performance of our method (FSD-SBERT) with HE and TW. We observe that Mazoyer et al.’s (2020) mini-batch FSD algorithm combined with Sentence Transformers pre-trained on large text corpora consistently outperforms the baselines on both datasets. The comparison between HISEvent (HE) and twembeddings (TW) seems to indicate that the mini-batch First Story Detection algorithm, even used with a simple tf-idf representation of tweets, is still a strong baseline, since its performance is comparable (and even superior on the French dataset) to HISEvent when using AMI as the indicator, though it is inferior when evaluated with ARI.

It is important to note that Cao et al. (2024) also use Sentence Transformers as a baseline in their article, with a different clustering algorithm (K-means). Their results are represented as the first

column in Figure 2, with the exact SBERT models they have used ("all-MiniLM-L6-v2"⁶ for English and "distiluse-base-multilingual-cased-v1"⁷ for French). Our experiments show that the type of SBERT model has an effect on performance: as shown on Figure 2 b, the "multilingual" model is less efficient for French than the language-specific "CamemBERT" model. Nevertheless, regardless of the model used, the FSD algorithm (see the last two columns) is much more efficient than the K-means for both datasets. This gap is explained by the fact that the FSD algorithm is able to take into account the temporality of tweets (by applying a sliding time window when searching for nearest neighbors) unlike the K-means. Moreover, FSD seems to be robust to changes in SBERT models without the need to adapt the parameters: on Figure 2, when using FSD, the same threshold $t = 0.55$ is used for all SBERT models. This common threshold explains the small difference between the values in Table 1 and Figure 2 for the all-mpnet-base-v2 model, since the threshold is set to 0.5 in Table 1 and to 0.55 in Figure 2.

Time efficiency. Increasing the batch size is a way to increase the computation speed with minimal loss in clustering performance: as shown in Figure 1, doubling the batch size only decreases the performance (measured by AMI) by 0.5%. This is why our experiments were all run with a batch size (b) set to 8 documents. With these parameters, our algorithm processes the Event2012 corpus, consisting of 68,841 documents (with a window size w of 2,368 documents), in 72 seconds. In contrast, HISEvent requires 1 hour and 45 minutes to process a block of 8,722 documents, and over 5 days to handle the entire corpus.

The experiments shown on Figure 1 were run on a notebook PC with 32GB of RAM and 8 2.4GHz CPUs. Note that these tests do not take into account the encoding of the tweets using Sentence Transformers, since we computed the embeddings only once on a GPU server and then stored them to be re-used for further experiments on a notebook computer without GPU. It took 65 seconds using a NVIDIA RTX A4500 GPU to encode the Event2012 corpus, and 240 seconds to encode the Event2018 corpus.

⁶<https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

⁷<https://huggingface.co/sentence-transformers/distiluse-base-multilingual-cased-v1>

Resources. We observed that executing HISEvent on the entire Event2012 dataset required substantial memory resources, exceeding 62 GB of RAM. In contrast, FSD operates with significantly lower memory requirements (less than 32GB of RAM).

Limitations. Twitter has been an invaluable resource for research on social media and real-time data streams. However, this is no longer possible due to the platform’s API restrictions. Nevertheless, we believe this study remains relevant, as other data streams and social networks continue to produce valuable data, and event recognition continues to be a crucial task.

Another limitation related to the mini-batch FSD algorithm is the need to pre-determine the hyperparameter t . However, the consistency of the results with the same t value across several SBERT models (see Figure 2) suggests that this threshold ($t = 0.55$) could be applied to other Sentence Transformers models pre-trained on corpora in different languages.

Finally, a potential improvement for this method would be to better account for the nested nature of events in public discourse: for instance, a major political event might consist of numerous smaller sub-events, such as speeches, protests, and negotiations (for example the Yellow Vest protest in France lasted several months, with protests every week, discussions with the government, thousands of declarations, actors and reactions (Wagner-Egger et al., 2022)). Each of these sub-events can be reported separately (or not) in different messages. This layered structure would ideally necessitate more sophisticated models capable of capturing and integrating these various components to provide a coherent and comprehensive understanding of the overall event.

6 Conclusion

In this study, we aimed to investigate the performance of incremental clustering combined with Sentence Transformers models for automatically detecting events in a stream of tweets. Our results demonstrated that applying the mini-batch FSD algorithm to SBERT representations significantly improves event detection performance on Twitter. We suggest that future research in this area should adopt this straightforward approach as a baseline for deploying more complex algorithms.

Acknowledgements

This work was funded in part by the Agence Nationale de la Recherche, as part of the ANR Medialex Project (AAPG 2021). Thierry Poibeau is also funded by the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute).

References

- James Allan, Jaime Carbonell, George Doddington, Jonathan Yamron, and Yiming Yang. 1998. *Topic detection and tracking pilot study final report*. Carnegie Mellon University.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Yuwei Cao, Hao Peng, Zhengtao Yu, and Philip S. Yu. 2024. [Hierarchical and incremental structural entropy minimization for unsupervised social event detection](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(8):8255–8264. Number: 8.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Mahmud Hasan, Mehmet A Orgun, and Rolf Schwitter. 2018. [A survey on real-time event detection from the twitter data stream](#). *Journal of Information Science*, 44(4):443–463.
- S Likhitha, BS Harish, and HM Keerthi Kumar. 2019. A detailed survey on topic modeling for document and short text data. *International Journal of Computer Applications*, 178(39):1–9.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villamonte de la Clergerie, Djamel Seddah, and Benoît Sagot. 2020. Camembert: a tasty french language model. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Béatrice Mazoyer, Julia Cagé, Nicolas Hervé, and Céline Hudelot. 2020. [A french corpus for event detection on twitter](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, page 6220. European Language Resources Association (ELRA).
- Andrew J McMinn and Joemon M Jose. 2015. Real-time entity-based event detection for twitter. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 6th International Conference of the CLEF Association, CLEF'15, Toulouse, France, September 8-11, 2015, Proceedings 6*, pages 65–77. Springer.
- Andrew J. McMinn, Yashar Moshfeghi, and Joemon M. Jose. 2013. [Building a large-scale corpus for evaluating event detection on twitter](#). In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 409–418. ACM Press.
- Fabian Pedregosa, Gael Varoquaux, Alexandre Gramfort, Vincent Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Hao Peng, Ruitong Zhang, Shaoning Li, Yuwei Cao, Shirui Pan, and S Yu Philip. 2022. Reinforced, incremental and cross-lingual event detection from social messages. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):980–998.
- Sasa Petrovic, Miles Osborne, and Victor Lavrenko. 2010. Streaming first story detection with application to twitter. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics (HLT'10)*, pages 181–189. Association for Computational Linguistics.
- Abhaya Kumar Pradhan, Hrushikesh Mohanty, and Rajendra Prasad Lal. 2024. [EDTBERT: Event detection and tracking in twitter using graph clustering and pre-trained language model](#). *Procedia Computer Science*, 233:481–491.
- Xingfa Qiu, Qiaosha Zou, and C. J. Richard Shi. 2021. [Single-pass on-line event detection in twitter streams](#). In *2021 13th International Conference on Machine Learning and Computing*, pages 522–529. ACM.
- William M Rand. 1971. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336):846–850.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.
- Jiaqian Ren, Lei Jiang, Hao Peng, Yuwei Cao, Jia Wu, Philip S Yu, and Lifang He. 2022. From known to unknown: Quality-aware self-improving graph neural network for open set social event detection. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 1696–1705.

- Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Nguyen Xuan Vinh, Julien Epps, and James Bailey. 2009. Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *Proceedings of the 26th annual international conference on machine learning*, pages 1073–1080.
- Pascal Wagner-Egger, Pascal Adam-Troian, Laurent Cordonier, Florian Cafiero, and Gérald Bronner. 2022. [The Yellow Vests in France: Psychosocial Determinants and Consequences of the Adherence to a Social Movement in a Representative Sample of the Population](#). *International review of social psychology*, 35(1).