AGILE: GIScience Series, 5, 28, 2024. https://doi.org/10.5194/agile-giss-5-28-2024
Proceedings of the 27th AGILE Conference on Geographic Information Science, 4–7 Sept. 2024.
Editors: Alison Heppenstall, Mingshu Wang, Urska Demsar, Rob Lemmens, and Jing Yao.
This contribution underwent peer review based on a full paper submission.

Multimodal Geo-Information Extraction from Social Media for Supporting Decision-Making in Disaster Management

David Hanny [1] and Bernd Resch [1,2]

[1]Department of Geoinformatics, University of Salzburg, Austria
[2]Center for Geographic Analysis, Harvard University, Cambridge, MA, USA

Correspondence: David Hanny (david.hanny@plus.ac.at)

Abstract. Effective decision-making in natural disaster management relies heavily on a comprehensive understanding of the situation in affected areas. Social media has been established as a tool to monitor human response and damage assessment. Given the vast amounts of data available, computational methods such as topic modelling are typically employed to reduce information complexity. However, these methods mostly neglect aspects such as geographic location and emotional response, which frequently results in sequential workflows of initial semantic filtering and subsequent spatial or spatio-temporal analysis. This study presents a novel approach for multimodal information extraction from geo-social media data for aiding decision support in disaster management. The method leverages a spatial, temporal, semantic, and sentiment-based clustering approach of social media posts to extract clusters that provide insights into disaster-related content. A case study in the Ahr Valley region in Germany demonstrates the method's effectiveness in providing actionable insights for disaster response and management. The approach offers a tool for the quick assessment of disaster-related information from social media, potentially aiding timely and informed decision-making.

Keywords. disaster management, social media, natural language processing, spatial machine learning

1 Introduction

With billions of posts available, the analysis of online social networks has proven increasingly useful for decision support, both during and after the occurrence of natural disasters (Phengsuwan et al., 2021; Wang and Ye, 2018; Xiao et al., 2015). As the number of active social media users is constantly rising (Newman et al., 2023; Ortiz-Ospina and Roser, 2023), disasters such as floods, wildfires, or earthquakes are one of many topics discussed on social media platforms. Potentially, these discussions can provide insights into disaster response actions and situations on-site. A fraction of social media posts are georeferenced (Sloan and Morgan, 2015; Serere et al., 2023), especially microblogging posts such as tweets. Such georeferenced information can be valuable for event detection, damage assessment, aid improvement, or exploratory data analyses in disaster situations (Chae et al., 2012; Crooks et al., 2013; Kar et al., 2018; Ragini et al., 2018; Resch et al., 2018).

A major research challenge has always been the extraction of relevant information from the myriad of posts these platforms provide. Previous efforts have mostly utilized topic modelling or text classification combined with spatio-temporal analysis to extract disaster-related information from large collections of social media posts (Kar et al., 2018; Chae et al., 2012; Crooks et al., 2013). The joint analysis of semantics, sentiments, space, and time – potentially providing insights into what is happening where, and when – has hardly been considered. This study addresses this research gap by introducing a novel approach for aiding decision support in disaster management by leveraging multimodal social media analysis. The approach computes multimodal clusters using a joint spatio-temporal topic-sentiment model, providing human-interpretable outputs. The utility of these outputs is demonstrated in a case study regarding the 2021 Ahr Valley flood in Germany.

2 Related Work

Social media data has been used in a multitude of ways to aid disaster management efforts in all four phases *mitigation*, *preparedness*, *response*, and *recovery* of the disaster management cycle (Albtoush et al., 2011). Kar et al. (2018) developed a pipeline to classify tweets by needs using a Support Vector Machine (SVM), geo-coding them using OpenStreetMap, and mapping the results with satellite imagery and other information for improved aid and medical supply. Resch et al. (2018) combined topic modelling with Latent Dirichlet Allocation (LDA) and spatio-temporal analysis for footprint and damage assessment. The detected hotspots were validated using data from the 2014 Napa Earthquake and official earthquake footprint reports. Chae et al. (2012) took a similar approach that also leveraged LDA and visual analytics for the exploratory analysis of abnormal topics on social media in various disaster situations. Crooks et al. (2013) argued that social media feeds represent a sensor system that allows for the identification and localization of earthquake impacts by analysing keyword-filtered, geo-referenced tweets from the Mineral Earthquake in Virginia in 2011. Aulov and Halem (2012) presented a comparable method for viewing social media data as human sensors for disaster monitoring. Huang et al. (2018) integrated remote sensing imagery with social media data to compute a near real-time flood probability map that could be used by emergency responders to identify areas in need of immediate attention. Similarly, Havas et al. (2017) combined information derived from social media, remote sensing technologies and crowdsourcing to improve disaster management systems. Chae et al. (2014) analysed Twitter data regarding Hurricane Sandy to extract spatial and temporal behavioral patterns. Huang and Xiao (2015) encoded social media posts into classes within the different disaster phases and trained a classifier using logistic regression. Lastly, Wang et al. (2016) examined the distribution of information during the 2012 Beijing rainstorm with the help of content classification and statistical models. Categorising these efforts into the initial framework, social media data has mostly been utilized to support the phases *preparedness*, *response* and *recovery* of disaster management (Phengsuwan et al., 2021).

These previous approaches all require the training of classifiers or leverage topic modelling techniques that have been outperformed by newer techniques (Egger and Yu, 2022). Furthermore, they neglect the multimodal nature of textual data which contains not only semantic but also emotional information that might be of use for efforts such as improving disaster response (Ragini et al., 2018). Neppalli et al. (2017) highlighted the connection between sentiments expressed in social media posts and the distance to the disaster during Hurricane Sandy in the United States of America (USA).

Parimala et al. (2021) partly addressed this issue and presented a multimodal algorithm for risk assessment after and during disasters using social media data from Twitter. The method combines semantic text classification, sentiment analysis, and spatio-temporal analysis to predict the severity of a disaster. The algorithm is limited by its supervised nature, meaning that it requires the training of a semantic classification model for the respective disaster. Additionally, the utilized sentiment analysis technique can only handle English tweets and has been outperformed by newer transformer-based methods (e.g. Camacho-Collados et al., 2022).

The model presented in this study addresses the issues of previous research efforts by introducing a fully unsupervised approach for the multimodal assessment of social media posts. The required input is merely a collection of social media posts and no specific classification model must be trained. The approach not only categorizes posts into interpretable clusters but also extracts relevant semantic information for emergency response entities which has not been achieved so far.

3 Methods

The proposed methodology is based on a clustering approach that leverages multimodal feature vectors. Given textual social media posts as input data, multimodal clusters are computed using a multi-step procedure that consists of three phases: (1) feature engineering, (2) clustering, and (3) information extraction. Each output cluster is associated with a semantic topic, a cluster sentiment, location, time and emergency-relevant information. A visual overview of the methodology is depicted in Fig. 1, while each step is explained separately in the following sub-sections.

3.1 Feature Engineering

To obtain a unified representation for each post consisting of multiple modalities, numeric feature vectors are computed. Each feature vector consists of semantic components $e_1, ..., e_5$, sentiment components $s_1, ..., s_3$, planar coordinates x, y and a timestamp t. The semantic part of the feature vector is computed using the multilingual *distiluse-base-multilingual-cased-v1* transformer model, a knowledge-distilled version of the multilingual Universal Sentence Encoder (USE) by Yang et al. (2020) available in the Sentence-BERT (SBERT) implementation of Reimers and Gurevych (2020). The model yields high-dimensional embedding vectors in a semantic space. These high-dimensional representations are subsequently reduced to five dimensions using Uniform Manifold Approximation and Projection (UMAP) and scaled to unit length. This configuration was chosen since SBERT outperformed similar approaches such as doc2vec (Le and Mikolov, 2014) and has successfully been employed in combination with UMAP for previous topic modelling efforts (e.g. Grootendorst, 2022). Allaoui et al. (2020) further showed that clustering results can be improved significantly when UMAP is applied as a pre-processing step. The sentiment part of the feature vector consists of sentiment probabilities for the classes *negative*, *neutral*, and *positive*. It is computed using the multilingual RoBERTa model pre-trained on tweets and fine-tuned for sentiment classification by Camacho-Collados et al. (2022). The coordinates of each post are latitude/longitude coordinates projected onto a plane using an equidistant map projection and normalised to a range of $[0, 1]$. Analogously, the time is converted to Unix time and normalised to a range of

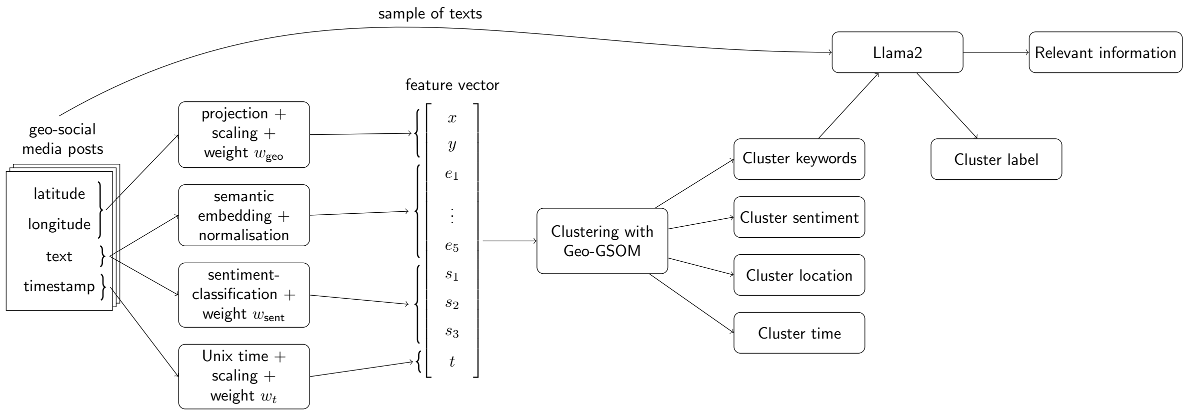

Figure 1. Overview of the methodology used to obtain multimodal clusters from social media data.

$[0, 1]$. All features, including semantic components, sentiment probabilities, planar coordinates, and timestamps are normalised to prevent skew from large values and can be weighted to emphasize certain characteristics.

3.2 Clustering

The numeric vectors resulting from the feature engineering process are clustered such that coherent groups concerning their semantic content, sentiment, location, and time are achieved. For the clustering procedure, a modified Growing Self-Organizing Map (GSOM) (Alahakoon et al., 2000) that follows the idea of the Geographic Self-Organizing Map (Geo-SOM) of Bação et al. (2005) is employed. In this study, it is called the Geographic Growing Self-Organizing Map (Geo-GSOM) and describes an unsupervised neural network that grows neurons on a flexible grid based on the accumulated error of data points assigned to neurons. Each input data point is mapped to the closest neuron called the Best Matching Unit (BMU) which is updated together with its neighbours during training such that the quantisation error is minimized. Geographic coordinates are redistributed evenly over the entire study area during each training iteration ensuring that the neuron grid stays spatially coherent at all times. As a result, neurons that are nearby in the grid represent clusters that are geographically close. After training, the identity of the BMU represents the cluster identity of each respective post. The Geo-GSOM was chosen over the classic Geo-SOM to allow the network to capture topological information with its structure.

3.3 Information Extraction

To provide a semantic representation of each cluster, the top k keywords are extracted using a modified term frequency and inverse document frequency (tf-idf) procedure (Sparck Jones, 1972; Grootendorst, 2022) where the document frequencies are computed using the original input data, and the term frequencies are calculated using the contamination of all documents in the respective cluster. By definition, the Geo-GSOM learns a representative vector in the feature space for each cluster which, in turn, is associated with exactly one neuron in the network. These representative vectors can be used to extract a learned approximate geospatial location for each cluster by re-scaling and re-projecting the learned coordinate dimensions of the Geo-GSOM units. Additionally, the most common sentiment and an approximate time can be obtained by calculating the respective mode and mean. For non-spatial properties, cluster statistics were consistently more reliable during manual experiments because the Geo-GSOM tries to learn many dimensions at once. As the keyword representation of topics is not easily intuitively interpretable, summarising cluster labels are computed using the *Llama-2-70b-chat* generative language model (Touvron et al., 2023). The model is prompted with the cluster keywords and a sample of the associated posts, asking it to compute a short, informative cluster label that provides an overview of the content for emergency responders. In addition, the language model is presented with a second prompt consisting of the posts and a request to extract the most important information for emergency responders in one or two sentences. As a result, each cluster has semantic keywords, a cluster label, a sentiment, a location, a time, and cluster-specific relevant information for emergency response entities.

3.4 Data and Software Availability

The methodological workflow including the Geo-GSOM was implemented in Python (van Rossum, 1995). The code is available by request. The dataset collected for the case study was obtained using the former Twitter (now X) Application Programming Interface (API) and can only be redistributed as a list of post IDs due to X's terms of service[1]. It is also available by request.

4 Case Study

To demonstrate the method's use, the 2021 Ahr Valley flood in Germany served as a real-world scenario. From 14 to 15 July 2021, heavy rainfall hit the region causing a major flood, severe structural and economic damage as well as human casualties (Bundesministerium des Innern und für Heimat, 2022; Koks et al., 2021). For a detailed analysis of the response to the catastrophe via social media, 11 177 geo-referenced tweets posted within the Ahr Valley region in Germany and surroundings in July 2021 were obtained via the former Twitter v1.1 API using the filtered stream and recent search endpoints. The tweets were filtered using a bounding box and using the month of July as the timeframe. The resulting data set consists of tweets in multiple languages, mainly German and English. While this would have been a challenge for many previous topic modelling approaches such as LDA (Blei et al., 2003), the proposed multimodal model can handle mixtures of languages natively due to the multilingual embedding model that is employed.

During feature engineering, the individual feature weights were set to $w_{\text{geo}} = \frac{3}{2}$, $w_{\text{sent}} = \frac{1}{3}$, and $w_t = \frac{1}{4}$ based on the results of exploratory analyses. Specifically, the geographic features of the input vectors were weighted highly such that the approach yielded geographically well-formed clusters. Sentiments and time were weighted slightly lower. The Geo-GSOM was parametrised to grow a large number of neurons to obtain granular local clusters of social media posts. Specifically, after experimenting with different configurations, the number of iterations where new neurons are grown based on the accumulated error was fixed to 10 full passes over the input data. This setup was chosen because it produced clusters that were both (1) balanced in size and (2) still differentiated between different topics and sentiments in areas with high tweet density.

Overall, the Geo-GSOM algorithm yielded 92 clusters. To identify clusters that might be concerned with the flood, the clusters were filtered, allowing only those with keywords "hochwasser", "flut", "flood", "hilfe", "notfall" to remain. These keywords were chosen based on the event and with emergency responders as a target group in mind. Using this process, 13 clusters consisting of a total of 1 592 tweets were identified as flood-related. Subsequently, cluster labels were computed with Llama-2.

As discussed in Sec. 3.3, each output cluster is associated with a neuron of the Geo-GSOM which, in turn, has a representative feature space vector. The feature space vectors were used to obtain a learned geospatial location for the respective clusters. Moreover, the top 25 semantic keywords, the most common sentiment and the average posting time were extracted for all clusters. Subsequently, summarising cluster labels were computed using the keywords and a subset of the underlying tweets.

In Fig. 2, the computed cluster labels are mapped at the respective cluster locations with additional information about the cluster sentiment. The convex hull of the underlying posts conveys a sense of the spatial spread. The map was further enriched with a heatmap of the posts within each cluster. Exemplary mappings of cluster keywords and tweets to labels are depicted in Table 1.

The cluster labels alone do not necessarily capture all disaster-relevant information present in the posts. Some clusters were only partly disaster-related (e.g. 8, 38 or 55) but still contained posts with explicit information about weather updates and local events. Therefore, additional disaster-relevant information was extracted from each cluster based on the posts using Llama-2. The results are visible in Table 2 together with the cluster sentiment and the average posting time of the tweets in the cluster. Notably, clusters 8, 14, 39, 54, 55, 76, 82 and 87 all had their temporal mean around July 15 and mostly contained information about warnings, evacuation efforts and in-situ reports. In contrast, 13, 27, 38, 71, 74 had their temporal mean a few days after the flood event and were more concerned with reports about damage, displacements and impact. The computed clusters are therefore coherent with the disaster's timeline.

All clusters that contained relevant information were negatively or neutrally connoted aligning with the findings of Neppalli et al. (2017) who showed that the sentiments of social media posts were linked to the distance to the disaster. The cities and municipalities impacted most heavily by the flood were located along the Ahr and Rhine rivers (Bundesministerium des Innern und für Heimat, 2022; Koks et al., 2021) which - as visible in Figure 2 - is also where a majority of the flood-related posts occurred. Based on the results presented in this case, the methodology can be used to obtain an overview of the overall situation on social media while simultaneously extracting the most important information for emergency response entities. It could particularly provide insights for the phases *prepardness*, *response* and *recovery* of the disaster management cycle, allowing for monitoring before, during and after the event.

5 Conclusions

This study introduced a novel method for the multimodal analysis of social media data and its use in disaster management. It leverages state-of-the-art language processing models and an unsupervised neural network called the Geo-GSOM to compute clusters of posts that share a semantic topic and sentiment, and that are geographically and temporally close. The outputs demonstrated in a case study are human-interpretable and the topic labels com-

[1]https://developer.twitter.com/en/developer-terms/policy

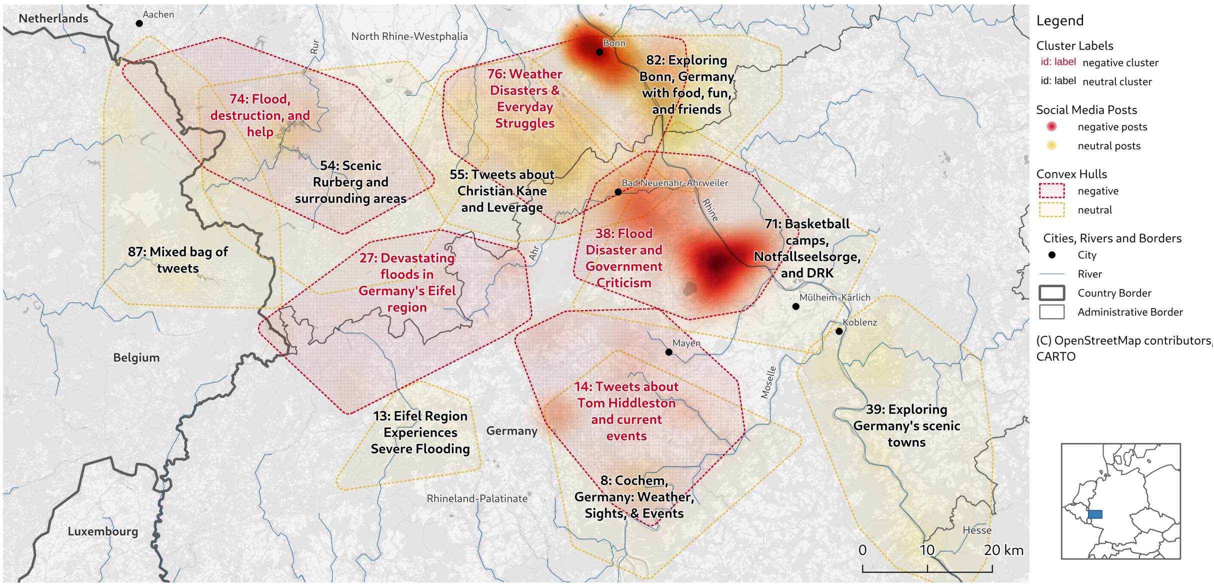

Figure 2. Visualisation of the multimodal output cluster labels together with the associated sentiments, post concentration and the convex cluster hulls.

Table 1. Exemplary labels computed from the cluster keywords and a sample of the underlying tweets.

ID	top 15 keywords	exemplary tweets	cluster label
38	ahrweiler, menschen, regierung, nena, erkrankung, hochwasser, ahrtal, jahren, personen, home, sinzig, macht, people, schützen, impfung	"Earlier German chancellor promised these victims would not be forgotten. One question we keep hearing from people here, when will official aid arrive?", "Es gibt z.B. Hochwasserwarnstufen 1-4 aber wenn die zuständigen öffentlichen Stellen nichts tun (schlafen, versagen) wie immer bei solchen Katastrophen ist man hilflos.", "Es ist wirklich schrecklich was das Hochwasser hinterlassen hat. Eine reine Katastrophe. Wir sind sprachlos von der riesigen Spendenaktion und der großen Hilfsbereitschaft. Vielen Dank an Alle. Folgt uns gerne mal auf Twitter."	Flood Disaster and Government Criticism
76	unwetter, halt, laschet, grade, hochwasser, hunger, since, bonn, möchte, wetter, stadt, wasser, richtig, leider, tierheim	"Got called at 3 am in the morning. Luckily the water is getting pumped out by firefighters! Unfortunately, the whole street was flooded - which is why they had to work at night too. A visiting old lady also lost her car in the underground garage. The flood came way too quick.", "Dass die Nachrichten nach dem #Unwetter immer noch schlimmer werden, hätte ich nicht gedacht... Soviele Tote und Zerstörung ....", "RIP washing machines. One car was trapped inside the underground garage when it started flooding....I couldn't find my gummi boots and it was very dark in the basement. It was a great mistake..."	Weather Disasters & Everyday Struggles

puted by a generative Large Language Model (LLM) are significantly more readable than the topic-word-based output of standard topic models (e.g. Blei et al., 2003; Grootendorst, 2022). In addition, potentially relevant information for emergency responders was explicitly extracted from the disaster-related clusters. This way, an overview of the discussion about a disaster on social media and the relevant information for emergency responders could be obtained quickly. The clusters are regionally constrained, allowing the user to draw local conclusions about specific areas. Since the output clusters are inherently multimodal, additional dimensions could be provided as input and leveraged for improved filtering or decision-making.

The methodology was particularly tailored to providing emergency response entities with rapid, interpretable information about the discussion on social media during and after the disaster. In future studies, its applicability for different types of disasters in different geographical regions might be explored to examine its generalisability. Furthermore, a sensitivity analysis of the Geo-GSOM and feature

Table 2. Relevant information for emergency responders extracted from the posts using Llama-2 by cluster ID.

ID	sentiment	size	relevant information	temporal mean
8	neutral	37	Flood warning in Germany, especially in the Eifel region. Rescue services are on high alert.	2021-07-15 20:50
13	neutral	17	Updates on current flood situation in Eifel region, Germany.	2021-07-19 06:04
14	negative	74	Severe weather warning in Germany, with heavy rain and thunderstorms expected.	2021-07-16 14:42
27	negative	27	Damage and displacement caused by floods in Schuld and Bad Münstereifel, Germany.	2021-07-17 02:56
38	negative	287	Request for help from people affected by floods in Germany.	2021-07-20 10:22
39	neutral	118	Reports of flooding, landslides, and road closures in Germany.	2021-07-13 22:57
54	neutral	78	Updates on rescue operations and aid distribution in flooded regions of Germany.	2021-07-15 15:43
55	neutral	254	Flood warning and evacuation alert in the Voreifel region due to heavy rainfall and rising river levels.	2021-07-15 16:06
71	neutral	31	Deployments by volunteer fire department and rescue service in Mayen-Koblenz area, Germany.	2021-07-21 20:35
74	negative	34	Impact of floods on people's lives in Germany.	2021-07-20 11:01
76	negative	152	Power outages and evacuation requests due to severe weather in Rhine Valley, Germany.	2021-07-15 10:59
82	neutral	397	Evacuation alerts due to overfilled dams in Wuppertal, Rheinbach, and Radevormwald, Germany.	2021-07-15 20:20
87	neutral	86	Updates on flood situation in Belgium and Netherlands, including rising water levels and evacuations.	2021-07-15 11:42

engineering parameters should be conducted. For some situations, it might be more valuable to lay more emphasis on time to obtain a timeline of topics associated with sentiments in geographic space. In a real-world application, posts can be collected and analysed in all stages of the disaster. Yet, a critical aspect of disaster information is the abstraction of complexity and the simplest possible visualisation, which is intuitively and quickly understandable and convertible into concrete disaster management actions. This is subject to further usability research together with disaster managers.

Author contributions. Conceptualisation, D.H.; methodology, D.H.; software, D.H.; validation, D.H.; formal analysis, D.H.; investigation, D.H.; resources, B.R.; data curation, D.H.; writing—original draft preparation, D.H.; writing—review and editing, D.H. and B.R.; visualisation, D.H.; supervision, B.R.; project administration, B.R.; funding acquisition, B.R. All authors have read and agreed to the published version of the manuscript.

Competing interests. The authors declare that they have no known competing interests or personal relationships that could have influenced the work presented in this paper. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Acknowledgements. This project has received funding from the European Commission - European Union under HORIZON EUROPE (HORIZON Research and Innovation Actions) under grant agreement 101093003 (HORIZON-CL4-2022-DATA-01-01). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union - European Commission. Neither the European Commission nor the European Union can be held responsible for them. This work has also received funding from the Austrian Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology (BMK) project GeoSHARING (Grant Number 878652).

References

Alahakoon, D., Halgamuge, S., and Srinivasan, B.: Dynamic Self-Organizing Maps with Controlled Growth for Knowledge Discovery, IEEE Transactions on Neural Networks, 11, 601–614, https://doi.org/10.1109/72.846732, 2000.

Albtoush, R., Dobrescu, R., and Ionescu, F.: A Hierarchical Model for Emergency Management Systems, 73, 2011.

Allaoui, M., Kherfi, M. L., and Cheriet, A.: Considerably Improving Clustering Algorithms Using UMAP Dimensionality Reduction Technique: A Comparative Study, in: Image and Signal Processing: 9th International Conference, ICISP 2020, Marrakesh, Morocco, June 4–6, 2020, Proceedings, pp. 317–325, Springer-Verlag, Berlin, Heidelberg, https://doi.org/10.1007/978-3-030-51935-3_34, 2020.

Aulov, O. and Halem, M.: Human Sensor Networks for Improved Modeling of Natural Disasters, Proceedings of the IEEE, 100, 2812–2823, https://doi.org/10.1109/JPROC.2012.2195629, 2012.

Bação, F., Lobo, V., and Painho, M.: GEO-SOM and ITS Integration with Geographic Information Systems, WSOM 2005 - 5th Workshop on Self-Organizing Maps, 2005.

Blei, D. M., Ng, A. Y., and Jordan, M. I.: Latent Dirichlet Allocation, The Journal of Machine Learning Research, 3, 993–1022, 2003.

Bundesministerium des Innern und für Heimat: Bericht zur Hochwasserkatastrophe 2021, Tech. rep., Bundesministerium des Innern und für Heimat, 2022.

Camacho-Collados, J., Rezaee, K., Riahi, T., Ushio, A., Loureiro, D., Antypas, D., Boisson, J., Espinosa Anke, L., Liu, F., and Martínez Cámara, E.: TweetNLP: Cutting-Edge Natural Language Processing for Social Media, in: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pp. 38–49, Association for Computational Linguistics, Abu Dhabi, UAE, 2022.

Chae, J., Thom, D., Bosch, H., Jang, Y., Maciejewski, R., Ebert, D. S., and Ertl, T.: Spatiotemporal Social Media Analytics for Abnormal Event Detection and Examination Using Seasonal-Trend Decomposition, in: 2012 IEEE Conference on Visual Analytics Science and Technology (VAST), pp. 143–152, https://doi.org/10.1109/VAST.2012.6400557, 2012.

Chae, J., Thom, D., Jang, Y., Kim, S., Ertl, T., and Ebert, D. S.: Public Behavior Response Analysis in Disaster Events Utilizing Visual Analytics of Microblog Data, Computers & Graphics, 38, 51–60, https://doi.org/10.1016/j.cag.2013.10.008, 2014.

Crooks, A., Croitoru, A., Stefanidis, A., and Radzikowski, J.: #Earthquake: Twitter as a Distributed Sensor System, Transactions in GIS, 17, 124–147, https://doi.org/10.1111/j.1467-9671.2012.01359.x, 2013.

Egger, R. and Yu, J.: A Topic Modeling Comparison Between LDA, NMF, Top2Vec, and BERTopic to Demystify Twitter Posts, Frontiers in Sociology, 7, 886 498, https://doi.org/10.3389/fsoc.2022.886498, 2022.

Grootendorst, M.: BERTopic: Neural Topic Modeling with a Class-Based TF-IDF Procedure, 2022.

Havas, C., Resch, B., Francalanci, C., Pernici, B., Scalia, G., Fernandez-Marquez, J. L., Van Achte, T., Zeug, G., Mondardini, M. R. R., Grandoni, D., Kirsch, B., Kalas, M., Lorini, V., and Rüping, S.: E2mC: Improving Emergency Management Service Practice through Social Media and Crowdsourcing Analysis in Near Real Time, Sensors, 17, 2766, https://doi.org/10.3390/s17122766, 2017.

Huang, Q. and Xiao, Y.: Geographic Situational Awareness: Mining Tweets for Disaster Preparedness, Emergency Response, Impact, and Recovery, ISPRS International Journal of Geo-Information, 4, 1549–1568, https://doi.org/10.3390/ijgi4031549, 2015.

Huang, X., Wang, C., and Li, Z.: A near Real-Time Flood-Mapping Approach by Integrating Social Media and Post-Event Satellite Imagery, Annals of GIS, 24, 113–123, https://doi.org/10.1080/19475683.2018.1450787, 2018.

Kar, S., Al-Olimat, H. S., Thirunarayan, K., Shalin, V. L., Sheth, A., and Parthasarathy, S.: D-Record: Disaster Response and Relief Coordination Pipeline, in: Proceedings of the 1st ACM SIGSPATIAL Workshop on Advances on Resilient and Intelligent Cities, ARIC'18, pp. 13–16, Association for Computing Machinery, New York, NY, USA, https://doi.org/10.1145/3284566.3284572, 2018.

Koks, E., Van Ginkel, K., Van Marle, M., and Lemnitzer, A.: Brief Communication: Critical Infrastructure Impacts of the 2021 Mid-July Western European Flood Event, Preprint, Risk Assessment, Mitigation and Adaptation Strategies, Socioeconomic and Management Aspects, https://doi.org/10.5194/nhess-2021-394, 2021.

Le, Q. and Mikolov, T.: Distributed Representations of Sentences and Documents, in: Proceedings of the 31st International Conference on Machine Learning, pp. 1188–1196, PMLR, 2014.

Neppalli, V. K., Caragea, C., Squicciarini, A., Tapia, A., and Stehle, S.: Sentiment Analysis during Hurricane Sandy in Emergency Response, International Journal of Disaster Risk Reduction, 21, 213–222, https://doi.org/10.1016/j.ijdrr.2016.12.011, 2017.

Newman, N., Fletcher, R., Eddy, K., Robertson, C. T., and Nielsen, R. K.: Reuters Institute Digital News Report 2023, 2023.

Ortiz-Ospina, E. and Roser, M.: The Rise of Social Media, Our World in Data, 2023.

Parimala, M., Swarna Priya, R. M., Praveen Kumar Reddy, M., Lal Chowdhary, C., Kumar Poluru, R., and Khan, S.: Spatiotemporal-Based Sentiment Analysis on Tweets for Risk Assessment of Event Using Deep Learning Approach, Software: Practice and Experience, 51, 550–570, https://doi.org/10.1002/spe.2851, 2021.

Phengsuwan, J., Shah, T., Thekkummal, N. B., Wen, Z., Sun, R., Pullarkatt, D., Thirugnanam, H., Ramesh, M. V., Morgan, G., James, P., and Ranjan, R.: Use of Social Media Data in Disaster Management: A Survey, Future Internet, 13, 46, https://doi.org/10.3390/fi13020046, 2021.

Ragini, J. R., Anand, P. M. R., and Bhaskar, V.: Big Data Analytics for Disaster Response and Recovery through Sentiment Analysis, International Journal of Information Management, 42, 13–24, https://doi.org/10.1016/j.ijinfomgt.2018.05.004, 2018.

Reimers, N. and Gurevych, I.: Making Monolingual Sentence Embeddings Multilingual Using Knowledge Distillation, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 4512–4525, Association for Computational Linguistics, Online, https://doi.org/10.18653/v1/2020.emnlp-main.365, 2020.

Resch, B., Usländer, F., and Havas, C.: Combining Machine-Learning Topic Models and Spatiotemporal Analysis of Social Media Data for Disaster Footprint and Damage Assessment, Cartography and Geographic Information Science, 45, 362–376, https://doi.org/10.1080/15230406.2017.1356242, 2018.

Serere, H. N., Resch, B., and Havas, C. R.: Enhanced Geocoding Precision for Location Inference of Tweet Text Using spaCy, Nominatim and Google Maps. A Comparative Analysis of the Influence of Data Selection, PLOS ONE, 18, e0282 942, https://doi.org/10.1371/journal.pone.0282942, 2023.

Sloan, L. and Morgan, J.: Who Tweets with Their Location? Understanding the Relationship between Demographic Characteristics and the Use of Geoservices and Geotagging on Twitter, PLOS ONE, 10, e0142 209, https://doi.org/10.1371/journal.pone.0142209, 2015.

Sparck Jones, K.: A STATISTICAL INTERPRETATION OF TERM SPECIFICITY AND ITS APPLICATION IN RETRIEVAL, Journal of Documentation, 28, 11–21, https://doi.org/10.1108/eb026526, 1972.

Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P. S., Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subramanian, R., Tan, X. E., Tang, B., Taylor, R., Williams, A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., and Scialom, T.: Llama 2: Open Foundation and Fine-Tuned Chat Models, https://doi.org/10.48550/arXiv.2307.09288, 2023.

van Rossum, G.: Python Tutorial, Tech. Rep. R 9526, CWI (National Research Institute for Mathematics and Computer Science), NLD, 1995.

Wang, Y., Wang, T., Ye, X., Zhu, J., and Lee, J.: Using Social Media for Emergency Response and Urban Sustainability: A Case Study of the 2012 Beijing Rainstorm, Sustainability, 8, 25, https://doi.org/10.3390/su8010025, 2016.

Wang, Z. and Ye, X.: Social Media Analytics for Natural Disaster Management, International Journal of Geographical Information Science, 32, 49–72, https://doi.org/10.1080/13658816.2017.1367003, 2018.

Xiao, Y., Huang, Q., and Wu, K.: Understanding Social Media Data for Disaster Management, Natural Hazards, 79, 1663–1679, https://doi.org/10.1007/s11069-015-1918-0, 2015.

Yang, Y., Cer, D., Ahmad, A., Guo, M., Law, J., Constant, N., Hernandez Abrego, G., Yuan, S., Tar, C., Sung, Y.-h., Strope, B., and Kurzweil, R.: Multilingual Universal Sentence Encoder for Semantic Retrieval, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 87–94, Association for Computational Linguistics, Online, https://doi.org/10.18653/v1/2020.acl-demos.12, 2020.