

supported by the Catalonia Trade and Investment Agency (ACCIÓ). P.C.M. was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 683253/GraphInt).

REFERENCES

- [1] Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. 1993. Mining association rules between sets of items in large databases. In *Acm sigmod record*, Vol. 22. ACM, 207–216.
- [2] L. M. Aiello, G. Petkos, C. Martin, D. Corney, S. Papadopoulos, R. Skrabar, A. Gker, I. Kompatsiaris, and A. Jaimes. 2013. Sensing Trending Topics in Twitter. *IEEE Transactions on Multimedia* 15, 6 (Oct 2013), 1268–1282. DOI: <http://dx.doi.org/10.1109/TMM.2013.2265080>
- [3] Mohammad Akbari, Xia Hu, Nie Liqiang, and Tat-Seng Chua. 2016. From Tweets to Wellness: Wellness Event Detection from Twitter Streams. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*. AAAI Press, 87–93.
- [4] Nasser Alsaedi, Pete Burnap, and Omer Rana. 2016. Sensing Real-World Events Using Arabic Twitter Posts. (2016).
- [5] Gaurav Baruah, Mark D. Smucker, and Charles L.A. Clarke. 2015. Evaluating Streams of Evolving News Events. In *Proc. of Conference on Research and Development in Information Retrieval (SIGIR) (SIGIR '15)*. ACM, New York, NY, USA, 675–684.
- [6] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching Word Vectors with Subword Information. CoRR abs/1607.04606 (2016). <http://arxiv.org/abs/1607.04606>
- [7] Xingyuan Chen, Yunqing Xia, Peng Jin, and John Carroll. 2015. Dataless Text Classification with Descriptive LDA. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15)*. AAAI Press, 2224–2231.
- [8] T. Cover and P. Hart. 2006. Nearest Neighbor Pattern Classification. *IEEE Trans. Inf. Theory* 13, 1 (Sept. 2006), 21–27.
- [9] Lishan Cui, Xiuzhen Zhang, Xiangmin Zhou, and Flora Salim. 2016. *Topical Event Detection on Twitter*. Springer International Publishing, Cham, 257–268. DOI: http://dx.doi.org/10.1007/978-3-319-46922-5_20
- [10] Qiming Diao, Jing Jiang, Feida Zhu, and Ee-Peng Lim. 2012. Finding Bursty Topics from Microblogs. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers - Volume 1 (ACL '12)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 536–544.
- [11] S. Girtelschmid, A. Salfinger, B. Prll, W. Retschitzegger, and W. Schwinger. 2016. Near real-time detection of crisis situations. In *2016 39th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. 247–252. DOI: <http://dx.doi.org/10.1109/MIPRO.2016.7522146>
- [12] Qi He, Kuiyu Chang, and Ee-Peng Lim. 2007. Analyzing Feature Trajectories for Event Detection. In *Proc. of Conference on Research and Development in Information Retrieval (SIGIR) (SIGIR '07)*. ACM, New York, NY, USA, 207–214.
- [13] Weijing Huang, Wei Chen, Lamei Zhang, and Tengjiao Wang. 2016. *An Efficient Online Event Detection Method for Microblogs via User Modeling*. Springer International Publishing, Cham, 329–341.
- [14] Shan Jiang, Yuening Hu, Changsung Kang, Tim Daly, Jr., Dawei Yin, Yi Chang, and Chengxiang Zhai. 2016. Learning Query and Document Relevance from a Web-scale Click Graph. In *Proc. of Conference on Research and Development in Information Retrieval (SIGIR) (SIGIR '16)*. ACM, New York, NY, USA, 185–194.
- [15] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of Tricks for Efficient Text Classification. *arXiv preprint arXiv:1607.01759* (2016).
- [16] S. Katragadda, S. Virani, R. Benton, and V. Raghavan. 2016. Detection of event onset using Twitter. In *2016 International Joint Conference on Neural Networks (IJCNN)*. 1539–1546. DOI: <http://dx.doi.org/10.1109/IJCNN.2016.7727381>
- [17] Tom Kenter and Maarten de Rijke. 2015. Short Text Similarity with Word Embeddings. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management (CIKM '15)*. ACM, New York, NY, USA, 1411–1420.
- [18] Matt J. Kusner, Yu Sun, Nicholas I. Kolkin, and Kilian Q. Weinberger. 2015. From Word Embeddings to Document Distances. In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37 (ICML '15)*. JMLR.org, 957–966.
- [19] Chenliang Li, Haoran Wang, Zhiqian Zhang, Aixin Sun, and Zongyang Ma. 2016. Topic Modeling for Short Texts with Auxiliary Word Embeddings. In *Proc. of Conference on Research and Development in Information Retrieval (SIGIR) (SIGIR '16)*. ACM, New York, NY, USA, 165–174.
- [20] Chenliang Li, Jian Xing, Aixin Sun, and Zongyang Ma. 2016. Effective Document Labeling with Very Few Seed Words: A Topic Model Approach. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management (CIKM '16)*. ACM, New York, NY, USA, 85–94.
- [21] Rui Li, Kin Hou Lei, Ravi Khadiwala, and Kevin Chen-Chuan Chang. 2012. TEDAS: A Twitter-based Event Detection and Analysis System. In *Proceedings of the 2012 IEEE 28th International Conference on Data Engineering (ICDE '12)*. IEEE Computer Society, Washington, DC, USA, 1273–1276.
- [22] Bing Liu, Xiaoli Li, Wee Sun Lee, and Philip S. Yu. 2004. Text Classification by Labeling Words. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI'04)*. AAAI Press, 425–430.
- [23] Xiaomo Liu, Quanzhi Li, Armineh Nourbakhsh, Rui Fang, Merine Thomas, Kajsa Anderson, Russ Kociuba, Mark Vedder, Steven Pomerville, Ramdev Wudali, Robert Martin, John Duprey, Arun Vachher, William Keenan, and Sameena Shah. 2016. Reuters Tracer: A Large Scale System of Detecting and Verifying Real-Time News Events from Twitter. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management (CIKM '16)*. ACM, New York, NY, USA, 207–216.
- [24] Jonathan Matusitz. 2012. *Terrorism and communication: A critical introduction*. Sage Publications.
- [25] R. McCreadie, C. Macdonald, I. Ounis, M. Osborne, and S. Petrovic. 2013. Scalable distributed event detection for Twitter. In *2013 IEEE International Conference on Big Data*. 543–549. DOI: <http://dx.doi.org/10.1109/BigData.2013.6691620>
- [26] Duc T. Nguyen and Jason J. Jung. 2015. Real-time Event Detection on Social Data Stream. *Mob. Netw. Appl.* 20, 4 (Aug. 2015), 475–486.
- [27] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2014. CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises (ICWSM '14).
- [28] Nikolaos Panagiotou, Ioannis Katakis, and Dimitrios Gunopoulos. 2016. *Detecting Events in Online Social Networks: Definitions, Trends and Challenges*. Springer International Publishing, Cham, 42–84. DOI: http://dx.doi.org/10.1007/978-3-319-41706-6_2
- [29] Min Peng, Jiahui Zhu, Xuhui Li, Jiajia Huang, Hua Wang, and Yanchun Zhang. 2015. Central Topic Model for Event-oriented Topics Mining in Microblog Stream. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management (CIKM '15)*. ACM, New York, NY, USA, 1611–1620.
- [30] Saša Petrović, 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)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 181–189.
- [31] Alan Ritter, Mausam, Oren Etzioni, and Sam Clark. 2012. Open Domain Event Extraction from Twitter. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '12)*. ACM, New York, NY, USA, 1104–1112.
- [32] Bella Robinson, Robert Power, and Mark Cameron. 2013. A Sensitive Twitter Earthquake Detector. In *Proceedings of the 22Nd International Conference on World Wide Web (WWW '13 Companion)*. ACM, New York, NY, USA, 999–1002.
- [33] Yu Rong, Qiankun Zhu, and Hong Cheng. 2016. A Model-Free Approach to Infer the Diffusion Network from Event Cascade. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management (CIKM '16)*. ACM, New York, NY, USA, 1653–1662.
- [34] Mehdi Sadri, Sharad Mehrotra, and Yaming Yu. 2016. Online Adaptive Topic Focused Tweet Acquisition. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management (CIKM '16)*. ACM, New York, NY, USA, 2353–2358.
- [35] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors. In *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*. ACM, New York, NY, USA, 851–860.
- [36] Roy Schwartz, Roi Reichart, and Ari Rappoport. 2015. Symmetric Pattern Based Word Embeddings for Improved Word Similarity Prediction. In *Proceedings of CoNLL 2015*.
- [37] Yangqiu Song and Dan Roth. 2014. On Dataless Hierarchical Text Classification. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI'14)*. AAAI Press, 1579–1585.
- [38] Alberto Tonon, Philippe Cudré-Mauroux, Albert Blarer, Vincent Lenders, and Boris Motik. 2017. ArmaTweet: Detecting Events by Semantic Tweet Analysis. In *ESWC*. 138–153. DOI: http://dx.doi.org/10.1007/978-3-319-58451-5_10
- [39] Zeynep Tufekci. 2017. *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. Yale University Press.
- [40] Maximilian Walther and Michael Kaiser. 2013. Geo-spatial Event Detection in the Twitter Stream. In *Proceedings of the 35th European Conference on Advances in Information Retrieval (ECIR'13)*. Springer-Verlag, Berlin, Heidelberg, 356–367.
- [41] Chao Wang, Xue Zhao, Ying Zhang, and Xiaojie Yuan. 2016. *Online Hot Topic Detection from Web News Based on Bursty Term Identification*. Springer International Publishing, Cham, 393–397.
- [42] Mohammed Javeed Zaki, Srinivasan Parthasarathy, and Wei Li. 1997. A localized algorithm for parallel association mining. In *Proceedings of the ninth annual ACM symposium on Parallel algorithms and architectures*. ACM, 321–330.
- [43] Hamed Zamani and W. Bruce Croft. 2016. Estimating Embedding Vectors for Queries. In *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval (ICTIR '16)*. ACM, New York, NY, USA, 123–132.
- [44] Shanshan Zhang and Slobodan Vucetic. 2016. Semi-supervised Discovery of Informative Tweets During the Emerging Disasters. CoRR abs/1610.03750 (2016).