



This is a repository copy of *RumourEval 2019: Determining rumour veracity and support for rumours*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/143173/>

Version: Submitted Version

Article:

Gorrell, G., Bontcheva, K. orcid.org/0000-0001-6152-9600, Derczynski, L. et al. (3 more authors) (Submitted: 2018) *RumourEval 2019: Determining rumour veracity and support for rumours*. arXiv. (Submitted)

© 2018 The Author(s). For reuse permissions, please contact the Author(s).

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

RumourEval 2019: Determining Rumour Veracity and Support for Rumours

**Genevieve Gorrell¹, Kalina Bontcheva¹, Leon Derczynski²,
Elena Kochkina³, Maria Liakata³, and Arkaitz Zubiaga³**

¹University of Sheffield, UK
(g.gorrell,k.bontcheva)@sheffield.ac.uk

²IT University of Copenhagen, Denmark
ld@itu.dk

³University of Warwick, UK
(e.kochkina,m.liakata,a.zubiaga)@warwick.ac.uk

Abstract

This is the proposal for RumourEval-2019, which will run in early 2019 as part of that year’s SemEval event. Since the first RumourEval shared task in 2017, interest in automated claim validation has greatly increased, as the dangers of “fake news” have become a mainstream concern. Yet automated support for rumour checking remains in its infancy. For this reason, it is important that a shared task in this area continues to provide a focus for effort, which is likely to increase. We therefore propose a continuation in which the veracity of further rumours is determined, and as previously, supportive of this goal, tweets discussing them are classified according to the stance they take regarding the rumour. Scope is extended compared with the first RumourEval, in that the dataset is substantially expanded to include Reddit as well as Twitter data, and additional languages are also included.

Overview

Since the first RumourEval shared task in 2017 (Derczynski et al. 2017), interest in automated verification of rumours has only deepened, as research has demonstrated the potential impact of false claims on important political outcomes (Allcott and Gentzkow 2017). Living in a “post-truth world”, in which perceived truth can matter more than actual truth (Dale 2017), the dangers of unchecked market forces and cheap platforms, alongside often poor discernment on the part of the reader, are evident. For example, the need to educate young people about critical reading is increasingly recognised.¹ The European Commission’s High Level Expert Group on Fake News cite provision of tools to empower users and journalists to tackle disinformation as one

of the five pillars of their recommended approach.² Simultaneously, research in stance prediction and assembling systems to understand and assess rumours expressed in written text have made some progress over baselines, but a broader understanding of the relation between stance and veracity – and a more extensive dataset – are required.

In a world where click-bait headlines mean advertising revenue, incentivising stories that are more attractive than they are informative, we are experiencing a deluge of fake news. Automated approaches offer the potential to keep up with the increasing number of rumours in circulation. Initial work (Qazvinian et al. 2011) has been succeeded by more advanced systems and annotation schemas (Kumar and Geethakumari 2014; Zhang et al. 2015; Shao et al. 2016; Zubiaga et al. 2016). Full fact checking is a complex task that may challenge the resourcefulness of even a human expert. Statistical claims, such as “we send the EU 350 million a week”, may offer a more achievable starting point in full fact checking, and have inspired engagement from researchers (Vlachos and Riedel 2015) and a new shared task. FEVER;³ whilst this work has a different emphasis from rumour verification, it shows the extent of interest in this area of research. Other research has focused on stylistic tells of untrustworthiness in the source itself (Conroy, Rubin, and Chen 2015; Singhania, Fernandez, and Rao 2017). Stance detection is the task of classifying a text according to the position it takes with regards to a statement. Research supports the value of this subtask in moving toward veracity detection (Ferreira and Vlachos 2016; Enayet and El-Beltagy 2017).

UK fact-checking charity Full Fact provides a roadmap⁴

¹<http://www.bbc.co.uk/mediacentre/latestnews/2017/fake-news>

²http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=50271

³<https://sheffieldnlp.github.io/fever/>

⁴https://fullfact.org/media/uploads/full_fact-the_state_of_automated_factchecking_aug_2016.pdf

Veracity prediction. Example 1:**u1:** Hostage-taker in supermarket siege killed, reports say. #ParisAttacks LINK [true]**Veracity prediction. Example 2:****u1:** OMG. #Prince rumoured to be performing in Toronto today. Exciting! [false]

Table 1: Examples of source tweets with veracity value

for development of automated fact checking. They cite open and shared evaluation as one of their five principles for international collaboration, demonstrating the continuing relevance of shared tasks in this area. Shared datasets are a crucial part of the joint endeavour. Twitter continues to be a highly relevant platform, being popular with politicians. By including Reddit data in the 2019 RumourEval we also provide diversity in the types of users, more focussed discussions and longer texts.

Summary of RumourEval 2017

RumourEval 2017 comprised two subtasks:

- In subtask A, given a source claim, tweets in a conversation thread discussing the claim are classified into support, deny, query and comment categories.
- In subtask B, the source tweet that spawned the discussion is classified as true, false or unverified.
 - In the open variant, this was done on the basis of the source tweet itself, the discussion and additional background information.
 - In the closed variant, only the source tweet and the ensuing discussion were used.

Eight teams entered subtask A, achieving accuracies ranging from 0.635 to 0.784. In the open variant of subtask B, only one team participated, gaining an accuracy of 0.393 and demonstrating that the addition of a feature for the presence of the rumour in the supplied additional materials does improve their score. Five teams entered the closed variant of task B, scoring between 0.286 and 0.536. Only one of these made use of the discussion material, specifically the percentage of responses querying, denying and supporting the rumour, and that team scored joint highest on accuracy and achieved the lowest RMSE. A variety of machine learning algorithms were employed. Among traditional approaches, a gradient boosting classifier achieved the second best score in task A, and a support vector machine achieved a fair score in task A and first place in task B. However, deep learning approaches also fared well; an LSTM -based approach took first place in task A and an approach using CNN took second place in task B, though performing less well in task A. Other teams used different kinds of ensembles and cascades of traditional and deep learning supervised approaches. In summary, the task attracted significant interest and a variety of approaches. However, for 2019 it is worth considering

how participants might be encouraged to be more innovative in the information they make use of, particularly exploiting the output of task A in their task B approaches.

How RumourEval 2019 will be different

For RumourEval 2019 we plan to extend the competition through the addition of new data, including Reddit data, and through extending the dataset to include new languages, namely Russian (Lozhnikov, Derczynski, and Mazzara 2018) and Danish.

In order to encourage more information-rich approaches, we will combine variants of subtask B into a single task, in which participants may use the additional materials (selected to provide a range of options whilst being temporally appropriate to the rumours in order to mimic the conditions of a real world rumour checking scenario) whilst not being obliged to do so. In this way, we prioritise stimulation of innovation in pragmatic approaches to automated rumour verification, shifting the focus toward success at the task rather than comparing machine learning approaches. At the same time, closed world entries are not excluded, and the task still provides a forum via which such approaches might be compared among themselves.

Subtask A - SDQC support classification

Rumour checking is challenging, and the number of data-points is relatively low, making it hard to train a system and to demonstrate success convincingly. Therefore, as a first step toward this, in task A participants track how replies to an initiating post orientate themselves to the accuracy of the rumour presented in it. Success on this task supports success on task B by providing information/features; for example, where the discussion ends in a number of agreements, it could be inferred that human respondents have verified the rumour. In this way, task A provides an intermediate challenge on which a greater number of participants may be able to gain traction, and in which a much larger number of data-points can be provided. Table 2 gives examples of the material.

Subtask B - Veracity prediction

As previously, the goal of subtask B is to predict the veracity of a given rumour, presented in the form of a post reporting an update associated with a newsworthy event. Given such a claim, plus additional data such as stance data classified

SDQC support classification. Example 1:

u1: We understand that there are two gunmen and up to a dozen hostages inside the cafe under siege at Sydney.. ISIS flags remain on display #7News [**support**]

u2: @u1 not ISIS flags [**deny**]

u3: @u1 sorry - how do you know its an ISIS flag? Can you actually confirm that? [**query**]

u4: @u3 no she cant cos its actually not [**deny**]

u5: @u1 More on situation at Martin Place in Sydney, AU LINK [**comment**]

u6: @u1 Have you actually confirmed its an ISIS flag or are you talking shit [**query**]

SDQC support classification. Example 2:

u1: These are not timid colours; soldiers back guarding Tomb of Unknown Soldier after today's shooting #StandforCanada PICTURE [**support**]

u2: @u1 Apparently a hoax. Best to take Tweet down. [**deny**]

u3: @u1 This photo was taken this morning, before the shooting. [**deny**]

u4: @u1 I dont believe there are soldiers guarding this area right now. [**deny**]

u5: @u4 wondered as well. Ive reached out to someone who would know just to confirm that. Hopefully get response soon. [**comment**]

u4: @u5 ok, thanks. [**comment**]

Table 2: Examples of tree-structured threads discussing the veracity of a rumour, where the label associated with each tweet is the target of the SDQC support classification task.

in task A and any other information teams choose to use from the selection provided, systems should return a label describing the anticipated veracity of the rumour. Examples are given in table 1. In addition to returning a classification of true or false, a confidence score should also be returned, allowing for a finer grained evaluation. A confidence score of 0 should be returned if the rumour is unverified.

Impact

RumourEval 2019 will aid progress on stance detection and rumour extraction, both still unbested NLP tasks. They are currently moderately well performed for English short texts (tweets), with data existing in a few other languages (notably as part of IberEval). We will broaden this with a multi-lingual task, having the largest dataset to date, and providing a new baseline system for stance analysis.

Rumour verification and automated fact checking is a complex challenge. Work in credibility assessment has been around since 2011 (Castillo, Mendoza, and Poblete 2011), making use initially of local features. Vosoughi (2015) demonstrated the value of propagation information, i.e. the ensuing discussion, in verification. Crowd response and propagation continue to feature in successful approaches, for example Chen et al (2016) and the most successful system in RumourEval 2017 (Enayet and El-Beltagy 2017), which might be considered a contender for the state of the art (Zubiaga et al. 2018). It is clear that the two part task formulation proposed here has continued relevance.

Platforms are increasingly motivated to engage with the problem of damaging content that appears on them, as society moves toward a consensus regarding their level of responsibility. Independent fact checking efforts, such as

Snopes⁵, Full Fact⁶, Chequeado⁷ and many more, are also becoming valued resources. Zubiaga et al (2018) present an extensive list of projects. Effort so far is often manual, and struggles to keep up with the large volumes of online material. It is therefore likely that the field will continue to grow for the foreseeable future.

Datasets are still relatively few, and likely to be in increasing demand. In addition to the data from RumourEval 2017, another dataset suitable for veracity classification is that released by Kwon et al (2017), which includes 51 true rumours and 60 false rumours. Each rumour includes a stream of tweets associated with it. A Sina Weibo corpus is also available (Wu, Yang, and Zhu 2015), in which 5000 posts are classified for veracity, but associated posts are not available. Partially generated statistical claim checking data is now becoming available in the context of the FEVER shared task, mentioned above, but isn't suitable for this type of work. A further RumourEval would provide additional data for system development as well as encouraging researchers to compare systems on a shared data resource.

Data and Resources

The data are structured as follows. Source texts assert a rumour, and may be true or false. These are joined by an ensuing discussion (tree-shaped) in which further users support, deny, comment or query (SDCQ) the source text. This is illustrated in figure 1 with a Putin example.

⁵<https://www.snopes.com/>

⁶<https://fullfact.org/>

⁷<http://chequeado.com/>

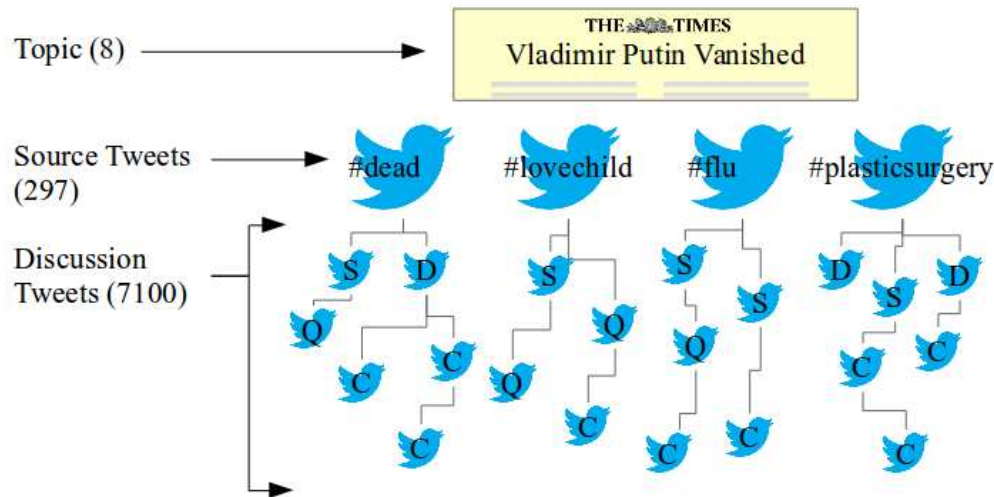


Figure 1: Structure of the first rumours corpus

The RumourEval 2017 corpus contains 297 source tweets grouped into eight overall topics, and a total of 7100 discussion tweets. This will become training data in 2019. We propose to augment this with at least the following:

- New English Twitter test data
- Reddit data in English
- Twitter data in Russian
- Twitter data in Danish

Topics will be identified using Snopes and similar debunking projects. Potential source posts within these topics are identified on the basis of the amount of attention they attract; for tweets, number of retweets has been used successfully as an indicator of a good source tweet. Source texts are then manually selected from among these and labeled for veracity by an expert.

An existing methodology, used successfully in RumourEval 2017, allows us to harvest the ensuing discussions. The stances of discussion texts will be crowdsourced. Multiple annotators will be used, as well as testing, to ensure quality. Previous experience with annotating for this task shows that it can be achieved with a high interannotator agreement.

Twitter’s developer terms, (Developer Policy, 1F.2a ⁸) state that up to 50,000 tweets may be shared in full via non-automated means such as download. This limit is sufficient for the tweets we envisage sharing. Twitter also requires that reasonable effort be made to ensure that tweets deleted by the author are also deleted by us (Developer Policy, 1C.3). To this end, the corpus will be checked for deleted tweets before release. In the event that Twitter requests a tweet be removed from the dataset, a new version of the data will be released to participants. It is unlikely that this would have a

⁸<https://developer.twitter.com/en/developer-terms/agreement-and-policy>

major impact on outcomes. Reddit places no restrictions on data redistribution.

Evaluation

In task A, stance classification, care must be taken to accommodate the skew towards the “comment” class, which dominates, as well as being the least helpful type of data in establishing rumour veracity. Therefore we aim to reward systems that perform well in classifying support, denial and query datapoints. To achieve this, we use macroaveraged F1, aggregated for each of these three types, and disregard the comment type entirely. Individual scores for the three main types will be provided separately in the final results.

In task B participants supply a true/false classification for each rumour, as well as a confidence score. Microaveraged accuracy will be used to evaluate the overall classification. For the confidence score, a root mean squared error (RMSE, a popular metric that differs only from the Brier score in being its square root) will be calculated relative to a reference confidence of 1. By providing these two scores, we give firstly a measure of system performance in the case of a real world scenario where the system must choose, and secondly a more fine-grained indicator of how well the system performed, that might be more relevant in the case that rumours are being automatically triaged for manual review. Note that it is possible for a system to score lower on accuracy but higher on RMSE compared with another system.

Baseline

For task A, we will provide code for a state-of-the-art baseline from RumourEval 2017 Task A (Kochkina, Liakata, and Augenstein 2017) together with later higher-performing entry published at RANLP that year (Aker, Derczynski, and Bontcheva 2017). The latest state of the art system for stance classification on RumourEval 2017 Task A dataset (Veysch

et al. 2017) may be provided in case of successful implementation.

For task B, we will provide our implementation of state-of-the-art baseline from RumourEval 2017 Task B (Enayet and El-Beltagy 2017) incorporating the best performing stance classification system.

Task Organizers

Kalina Bontcheva, University of Sheffield, UK. Research interests in social media analysis, information extraction, rumour analysis, semantic annotation, and NLP applications. Experience organising and running both workshops (including RDSM2015) and conferences (the RANLP series, UMAP2015). Email k.bontcheva@sheffield.ac.uk.

Leon Derczynski, IT University of Copenhagen, Denmark. Research interests in social media processing, semantic annotation, and information extraction. Experience running prior SemEval tasks (TempEval-3, Clinical TempEval, RumourEval) and major conferences (COLING). Email leod@itu.dk, telephone +45 51574948.

Genevieve Gorrell, University of Sheffield, UK. Research interests in social media analysis, text mining, text mining for health, natural language processing. Chaired the 2010 GATE summer school. Email g.gorrell@sheffield.ac.uk.

Elena Kochkina, University of Warwick, UK. Research interests in natural language processing, automated rumour classification, machine learning, deep learning for NLP. Experience participating in shared tasks, including RumourEval 2017. Email e.kochkina@warwick.ac.uk.

Maria Liakata, University of Warwick, UK. Research interests in text mining, natural language processing (NLP), biomedical text mining, sentiment analysis, NLP for social media, machine learning for NLP and biomedical applications, computational semantics, scientific discourse analysis. Co-chair of workshops in rumours in social media (RDSM) and discourse in linguistic annotations (LAW 2013). Co-organized RumourEval 2017. Email m.liakata@warwick.ac.uk.

Arkaitz Zubiaga, University of Warwick, UK. Research interests in social media mining, natural language processing, computational social science and human-computer interaction. Experience running 4 prior social media mining shared tasks including RumourEval 2017, and co-chair of workshops on social media mining at ICWSM and WWW. Email a.zubiaga@warwick.ac.uk.

Acknowledgements

This work is supported by the European Commissions Horizon 2020 research and innovation programme under grant agreement No. 654024, SoBigData.

References

Aker, A.; Derczynski, L.; and Bontcheva, K. 2017. Simple open stance classification for rumour analysis. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, 31–39.

Allcott, H., and Gentzkow, M. 2017. Social media and fake news in the 2016 election. *Journal of Economic Perspectives* 31(2):211–36.

Castillo, C.; Mendoza, M.; and Poblete, B. 2011. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, 675–684. ACM.

Chen, W.; Yeo, C. K.; Lau, C. T.; and Lee, B. S. 2016. Behavior deviation: An anomaly detection view of rumor pre-emption. In *Information Technology, Electronics and Mobile Communication Conference (IEMCON), 2016 IEEE 7th Annual*, 1–7. IEEE.

Conroy, N. J.; Rubin, V. L.; and Chen, Y. 2015. Automatic deception detection: Methods for finding fake news. *Proceedings of the Association for Information Science and Technology* 52(1):1–4.

Dale, R. 2017. Nlp in a post-truth world. *Natural Language Engineering* 23(2):319–324.

Derczynski, L.; Bontcheva, K.; Liakata, M.; Procter, R.; Hoi, G. W. S.; and Zubiaga, A. 2017. Semeval-2017 task 8: Rumoureval: Determining rumour veracity and support for rumours. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, 69–76.

Enayet, O., and El-Beltagy, S. R. 2017. Niletmg at semeval-2017 task 8: Determining rumour and veracity support for rumours on twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, 470–474.

Ferreira, W., and Vlachos, A. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: Human language technologies*, 1163–1168.

Kochkina, E.; Liakata, M.; and Augenstein, I. 2017. Turing at semeval-2017 task 8: Sequential approach to rumour stance classification with branch-lstm. In *Proceedings of SemEval.ACL*.

Kumar, K. K., and Geethakumari, G. 2014. Detecting misinformation in online social networks using cognitive psychology. *Human-centric Computing and Information Sciences* 4(1):14.

Kwon, S.; Cha, M.; and Jung, K. 2017. Rumor detection over varying time windows. *PLoS one* 12(1):e0168344.

Lozhnikov, N.; Derczynski, L.; and Mazzara, M. 2018. Stance Prediction for Russian: Data and Analysis. *arXiv preprint arXiv:1809.01574*.

Qazvinian, V.; Rosengren, E.; Radev, D. R.; and Mei, Q. 2011. Rumor has it: Identifying misinformation in microblogs. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 1589–1599. Association for Computational Linguistics.

Shao, C.; Ciampaglia, G. L.; Flammioni, A.; and Menczer, F. 2016. Hoaxy: A platform for tracking online misinformation. In *Proceedings of the 25th international conference companion on world wide web*, 745–750. International World Wide Web Conferences Steering Committee.

- Singhania, S.; Fernandez, N.; and Rao, S. 2017. 3han: A deep neural network for fake news detection. In *International Conference on Neural Information Processing*, 572–581. Springer.
- Veysseh, A. P. B.; Ebrahimi, J.; Dou, D.; and Lowd, D. 2017. A temporal attentional model for rumor stance classification. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 2335–2338. ACM.
- Vlachos, A., and Riedel, S. 2015. Identification and verification of simple claims about statistical properties. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2596–2601. Association for Computational Linguistics.
- Vosoughi, S. 2015. *Automatic detection and verification of rumors on Twitter*. Ph.D. Dissertation, Massachusetts Institute of Technology.
- Wu, K.; Yang, S.; and Zhu, K. Q. 2015. False rumors detection on sina weibo by propagation structures. In *Data Engineering (ICDE), 2015 IEEE 31st International Conference on*, 651–662. IEEE.
- Zhang, Q.; Zhang, S.; Dong, J.; Xiong, J.; and Cheng, X. 2015. Automatic detection of rumor on social network. In *Natural Language Processing and Chinese Computing*. Springer. 113–122.
- Zubiaga, A.; Liakata, M.; Procter, R.; Hoi, G. W. S.; and Tolmie, P. 2016. Analysing how people orient to and spread rumours in social media by looking at conversational threads. *PloS one* 11(3):e0150989.
- Zubiaga, A.; Aker, A.; Bontcheva, K.; Liakata, M.; and Procter, R. 2018. Detection and resolution of rumours in social media: A survey. *ACM Computing Surveys (CSUR)* 51(2):32.