

Enhancing Crisis-Related Tweet Classification with Entity-Masked Language Modeling and Multi-Task Learning

Philipp Seeberger and Korbinian Riedhammer

Technische Hochschule Nürnberg Georg Simon Ohm

{philipp.seeberger, korbinian.riedhammer}@th-nuernberg.de

Abstract

Social media has become an important information source for crisis management and provides quick access to ongoing developments and critical information. However, classification models suffer from event-related biases and highly imbalanced label distributions which still poses a challenging task. To address these challenges, we propose a combination of entity-masked language modeling and hierarchical multi-label classification as a multi-task learning problem. We evaluate our method on tweets from the TREC-IS dataset and show an absolute performance gain w.r.t. F1-score of up to 10% for actionable information types. Moreover, we found that entity-masking reduces the effect of overfitting to in-domain events and enables improvements in cross-event generalization. Our source code is publicly available on GitHub.¹

1 Introduction

Messages on social media during disaster events have become an important information source in crisis management (Reuter et al., 2018). In contrast to traditional sources (e.g., official news), social media posts immediately provide details about developments, first-party observations, and affected people in an ongoing emergency situation (Sakaki et al., 2010). Having access to this information is crucial for developing situational awareness and supporting relief providers, government agencies, and other official institutions (Kruspe et al., 2021).

One key challenge poses the information refinement of high-volume social media streams which requires automatic methods for reliable detection of relevant content (Kaufhold, 2021). Most recent work has focused on binary, multi-class, and multi-label text classification techniques to classify posts into coarse (e.g., *Relevant*, *Irrelevant*) or fine-grained (e.g., *InfrastructureDamage*, *Missing-*

¹<https://github.com/th-nuernberg/crisis-tapt-hmc>

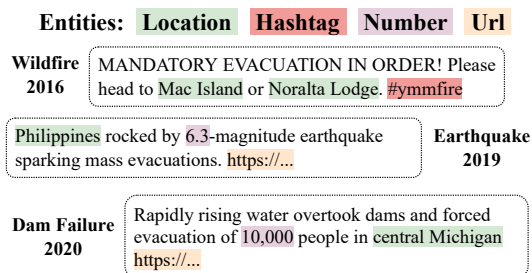


Figure 1: Example tweets of several disasters over time, annotated with entities. The short posts are mostly biased towards specific events.

People) categories composed of flattened or hierarchical structures (Alam et al., 2018b, 2021; Buntain et al., 2021).

Another challenge in Natural Language Processing (NLP) is the nature of data prevalent in social media and microblogging platforms. For example, most works in the crisis-related domain focus on Twitter data (Kruspe et al., 2021) which inherits properties such as short texts (280 characters limitation per tweet), less contextual information, hashtags, and noise (e.g., misspellings, emojis) (Wiegmann et al., 2020; Zahera et al., 2021). According to Sarmiento and Poblete (2021), different types of disasters (e.g., flood, wildfire) can be identified by only a few text-based features. However, event-related biases and entities as shown in Figure 1 prevent models from generalizing to unseen disaster events and therefore degrade w.r.t. detection performance.

To circumvent this problem, approaches such as adversarial training (Medina Maza et al., 2020), domain adaptation (Alam et al., 2018a), and hierarchical label embeddings (Miyazaki et al., 2019) have been proposed but suffer from mixed event types, assume unlabeled data or require semantic label descriptions. Contrary to this work, we aim to enhance the detection of rare actionable information for unseen events by masking out entities,

	Train	Test
Event Ids	1 - 52	53 - 75
# Events	51	21
# tweets	50,412	22,003
<i>Upper classes</i>		
# Report (14)	30,389	16,059
# Other (5)	32,105	10,709
# CallToAction (3)	1,458	389
# Request (3)	683	144

Table 1: Overview of the dataset split; the values within the brackets of the upper classes corresponds to the number of unique low-level information types.

applying adaptive pre-training, and incorporating the hierarchical structure of labels.

Contributions Our main contributions are as follows: (1) We introduce an adaptive pre-training strategy based on entity-masking. (2) We incorporate the hierarchical structure of labels as multi-task learning (MTL) problem. (3) We empirically show that our approach improves generalization to new events and increases detection performance for actionable information types.

2 Related Work

Crisis Tweet Classification Besides conventional detection approaches such as filtering (Kumar et al., 2011) or crowdsourcing (Poblet et al., 2014), machine learning has received much attention in this area. Researchers experimented with several methods such as Naive Bayes, Support Vector Machines, and Decision Trees either with term-frequency features (Habdank et al., 2017) or static embeddings (Kejriwal and Zhou, 2019). More recently, the combination of Word2Vec (Mikolov et al., 2013) with Convolutional and Recurrent Neural Networks achieved remarkable improvement in this field (Kersten et al., 2019; Snyder et al., 2019). Due to the success of Transformers (Vaswani et al., 2017) and the follow-up language models (Devlin et al., 2019), most works have been built upon this and outperformed previous approaches (Alam et al., 2021; Wang et al., 2021).

Adaptive Pre-Training Transfer learning with language models essentially contributes to state-of-the-art results in a variety of NLP tasks (Devlin et al., 2019; Liu et al., 2019; Clark et al., 2020). Typically, such language models follow the three training steps (Howard and Ruder, 2018; Ben-

David et al., 2020): (1) Pre-training on massive corpora; (2) Optional pre-training on task-specific data; (3) Supervised fine-tuning on target tasks. However, the second step is often neglected due to computational constraints whereby adaptive pre-training has shown to be effective (Howard and Ruder, 2018). Hence, Gururangan et al. (2020) introduced domain-adaptive pre-training (DAPT) and task-adaptive pre-training (TAPT) which cover continual pre-training on corpora tailored for a specific task. Moreover, strategies such as adding special tokens for tweets (Nguyen et al., 2020; Wiegmann et al., 2020) or additional masked language modeling (MLM) approaches (Ben-David et al., 2020) have been proven beneficial.

Hierarchical Multi-Label Classification Hierarchical multi-label classification (HMC) covers local and global approaches and the combination of both worlds (Wehrmann et al., 2018). A popular categorization of local methods is the subdivision into local classifier per parent node (LCPN) (Dumais and Chen, 2000), local classifier per node (LCN) (Banerjee et al., 2019), and local classifier per level (LCL) (Wehrmann et al., 2018). Hybrid approaches integrate the global part as a particular constraint such as hierarchical softmax (Brinkmann and Bizer, 2021) or combine multiple local and global prediction heads (Wehrmann et al., 2018). Recent work in information type classification introduced label embeddings which utilize the hierarchical structure (Miyazaki et al., 2019). Finally, the classification can also be viewed as MTL by combining certain loss functions (Yu et al., 2021; Wang et al., 2021).

3 TREC-IS

In this work, we mainly focus on the dataset of the shared-task TREC-IS, which represents a collection of annotated crisis-related tweets (Buntain et al., 2021). Each tweet belongs to a disaster event and is annotated with high-level information types which are derived from an ontology composed of hierarchical stages. However, information type labels are only shipped as a two-level hierarchy with four upper classes L_T and 25 lower classes L_B . Thus, both hierarchy levels represent a multi-label classification task. Following the TREC-IS track design, we split the dataset into train and test events which corresponds to the TREC-IS 2020B task. This split poses a challenging setup due to the requirement of cross-event generalization (Wiegmann et al., 2020).

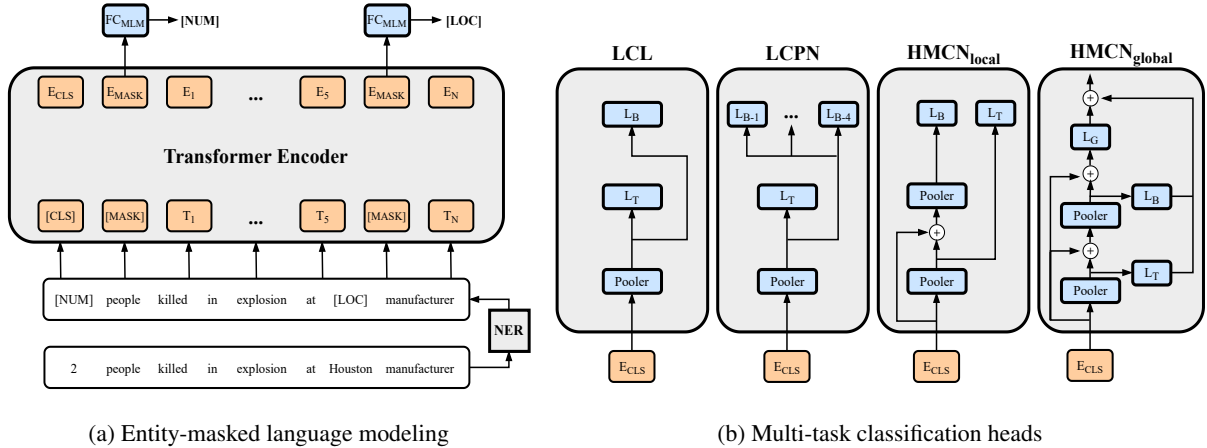


Figure 2: Illustration of the concepts E-MLM with named entity recognition (NER) and MTL. FC_{MLM} represents the prediction head for MLM. The classification heads will be placed on top of the pre-trained encoder. The building blocks Pooler, L_T , L_B and L_G are fully connected layers and use the CLS token as sentence embedding.

Table 1 gives an overview of each split; obviously, the information type distribution is highly imbalanced. For example, information types with low criticality such as *MultimediaShare* (31.7%) and *News* (25.4%) are prevalent. In contrast, the highly critical information types *MovePeople* (0.9%) and *SearchAndRescue* (0.4%) occur only rarely (McCreadie et al., 2019).²

4 Method

As depicted in Figure 2 our approach combines the two concepts entity-masked language modeling (E-MLM) and MTL. In the following, we briefly describe our method as a combination of those two.

4.1 Entity-Masked Language Modeling

Based on adaptive pre-training, we extend on masked language modeling of a transformer encoder pre-trained on a large corpus such as BERT (Devlin et al., 2019). Here, the mitigation of event-related biases is facilitated by replacing entities – which are prone to be event-specific – with special tokens (see Figure 2a). This way we intend to capture disaster-related language patterns independently of the concrete entities. Following Ben-David et al. (2020), we further introduce a masking probability α tailored to entities in addition to the standard word masking with probability β . That is, with a typically higher probability α we select random entity-tokens such as locations and lower probability β random standard subword-tokens. Finally, these selected tokens will be replaced by

²We provide an overview of the labels with some example posts in Appendix A.

$[MASK]$, random tokens or the unchanged tokens in order to learn the linguistic patterns related to those entities. For the rest of this paper, we rely on the pre-trained $BERT_{BASE}$ as the encoder model and the corresponding default MLM setup for pre-training ($[MASK]$ with 80%, random tokens with 10%, and unchanged tokens with 10%).

4.2 Multi-Task Learning

The next step represents the fine-tuning of a classification head. We implement four basic hierarchical multi-label classification approaches as shown in Figure 2b. The LCL classification head jointly trains a flattened classification layer for each of the two hierarchy levels. In contrast, the LCPN model consists of a classification layer for each parent node. The hierarchical multi-label classification network (HMCN) is adapted from Wehrmann et al. (2018) and introduces a pooling layer on top of the preceding pooling layer. We experiment with a local and a global variant, whereas the global one additionally consists of a global classification layer. All pooling and classification layers are composed of a single feed-forward layer with *tanh* and *sigmoid* as activation functions, respectively. Finally, we minimize the binary cross-entropy $\mathcal{L}_{MTL} = \lambda \mathcal{L}_{L_T} + (1 - \lambda) \mathcal{L}_{L_B}$ as a weighted loss function whereby \mathcal{L}_{L_T} represents the upper classes and \mathcal{L}_{L_B} the lower classes loss.

5 Experiments

5.1 Evaluation Metric

We follow the TREC-IS evaluation scheme: macro-averaged F1-score across information types for the

Model	L_T	L_B	AIT
<i>Single-Task</i>			
TF-IDF+LR	0.657	0.499	0.462
BERT _{BASE}	0.717	0.531	0.513
BERT _{MLM}	0.714	0.551	0.546
BERT _{E-MLM}	0.701	0.481	0.444

Table 2: Overall results on the development set.

two hierarchy levels in addition to the actionable information types (AIT) (McCreadie et al., 2019). The latter include rare information types with high priority consisting of: *MovePeople*, *EmergingThreats*, *NewSubEvent*, *ServiceAvailable*, *GoodsServices*, and *SearchAndRescue*.

5.2 Named Entity Recognition

As event-specific entities, we use the special tokens *hashtag*, *url*, *person*, *location*, *organization*, *event*, *address*, *phone number*, *date*, and *number*. All entities except the tokens *hashtag* and *url* are extracted with the Natural Language API of the Google Cloud Platform.³ We manually annotated 300 tweets and calculated a strict F1-score (Segura-Bedmar et al., 2013) of 0.692 which represents a reasonable good result for tweets.

5.3 Baseline and Hyper-Parameters

As baseline, we use TF-IDF with Logistic Regression (TF-IDF+LR) and BERT_{BASE} with a single-task classification head. Furthermore, we apply the standard MLM of BERT in contrast to E-MLM in order to validate the effect of masking entities. Lastly, we train the MTL model (MTL_{prio}) from Wang et al. (2021) which combines lower classes as classification and priority scores as regression task. We choose the best hyper-parameters for each model based on a stratified split with a ratio of 90% for train and 10% for development data, respectively. In terms of hyper-parameters, we set $\alpha = 0.5$ and $\beta = 0.1$ for E-MLM; other parameters were set according to other work, including learning rate of $5e-5$, batch size of 32, and $\lambda = 0.1$ for fine-tuning. The detailed hyper-parameter selection process is shown in Appendix B.

5.4 Results

In the following, we report the performance for the upper classes L_T , lower classes L_B , and AIT. However, for our evaluation we do not focus on

³We extracted the entities on 29 March 2022.

Model	L_T	L_B	AIT
MTL _{prio} *	-	0.278	0.279
<i>Single-Task</i>			
TF-IDF+LR	0.460	0.201	0.168
BERT _{BASE}	<u>0.553</u>	0.269	0.236
BERT _{MLM}	0.524	0.245	0.229
BERT _{E-MLM}	<u>0.553</u>	0.307	0.306
<i>Multi-Task</i>			
LCL	0.548	0.314	0.309
LCPN	0.548	0.305	0.307
HMCN _{global}	0.546	0.310	<u>0.320</u>
HMCN _{local}	0.558	<u>0.312</u>	0.335

Table 3: Overall results of information type classification; bold and underlined values indicate the best and second-best results, respectively. *We fine-tuned the approach of Wang et al. (2021) with BERT_{BASE} and without ensembling.

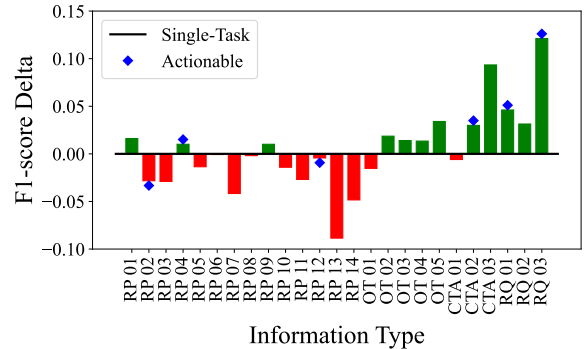


Figure 3: Absolute performance differences w.r.t. F1-score between the single-task and HMCN_{local} model.

L_T since the experiments did not show large differences across all BERT models. The MTL models are only reported with BERT_{E-MLM}.

E-MLM Table 3 displays the results of all single-task and MTL runs. For E-MLM, we observe an absolute performance gain w.r.t. F1-score for both L_B and AIT by up to 4% and 7%, respectively. To validate the event-generalization effect, we additionally analyzed the development set, as a proxy to estimate the in-domain event performance as shown in Table 2. Contrary to the test set, standard MLM increases the absolute L_B performance by 2% whereas the E-MLM approach drops by 5% which is a confirmation of our assumption about event-related overfitting.

Multi-Task Learning In terms of MTL, the HMCN_{local} model achieved the best results for AIT. Overall the MTL classification outperforms

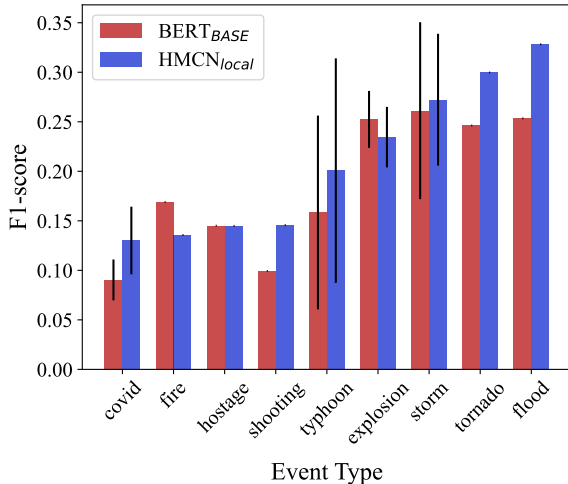


Figure 4: Comparison across event types w.r.t. F1-score between the BERT_{BASE} and HMCN_{local} model. We plot the mean and standard deviation for multiple events within a event type.

the single-task models for actionable categories and in addition the L_B classes except for LCPN. We assume that the L_T classification objective implicitly clusters the internal representation w.r.t. the high-level information types and therefore mitigates overfitting towards the major classes. As depicted in Figure 3, the HMCN_{local} model improves the detection of rare actionable information types over the single-task model while at the same time decreasing the performance on the category with the most information types. This can be caused by the ambiguous label definitions and semantic similarities with other information types (Mehrotra et al., 2022).

5.5 Analysis of Events

In Figure 4 we illustrate the model performance for L_B across different event types. For multiple events, we report the mean and standard deviation, respectively. We observe an increase in performance for the event types *covid*, *shooting*, *typhoon*, *storm*, *tornado*, and *flood* and a small decrease for the event types *fire*, *hostage*, and *explosion*. As shown by the variance for multiple events, the performance highly differs across specific events. Surprisingly, the event type *covid* achieved the worst performance for both models despite the existence of three *covid* events within the train data. These results indicate that even regional differences about the same global event predominantly affect the generalization performance across events.

Method	L_T	L_B	AIT
HMCN _{local}	0.558	0.312	0.335
- Hierarchy	0.548	0.314	0.309
- Multi-Task	0.553	0.307	0.306
- MLM	0.529	0.276	0.242
- Entities	0.553	0.269	0.236

Table 4: Overall results of the ablation study.

5.6 Ablation Study

As ablation study we removed several proposed components to assess the performance impact of our model. Thereby, the component entities represents the additional special tokens and replacement within the input text. As shown in Table 4, we started with the HMCN_{local} model and demonstrate that entities, MLM and MTL contribute to an increase w.r.t. F1-score for both L_B and AIT. The results indicate that the variant which removes the hierarchical component only degrades the performance for the low-resource actionable information types. Removing the E-MLM mechanism degrades the model’s performance most in our experiments.

6 Conclusion and Future Work

In this work, we identified shortcomings in the field of crisis tweet classification for unseen events. For the TREC-IS data, we found contrasting effects in terms of pre-training and observed an absolute improvement of up to 3% w.r.t. F1-score for actionable information types by incorporating the hierarchical structure. Furthermore, we confirmed the effectiveness of our method based on the shared-task TREC-IS. Future work includes pre-training on a larger corpus, the mitigation of the trade-off between major and minor classes performances, and to analyse the influence of label semantics.

Ethical and Societal Implications

Open Source Intelligence (OSINT) has become a significant role for various authorities and NGOs for advancing struggles in global health, human rights, and crisis management (Bernard et al., 2018; Evangelista et al., 2021; Kaufhold, 2021). Following the view of OSINT as a tool, our work pursues the goal to support relief providers, government agencies, and other disaster-response stakeholders during ongoing and evolving crisis events.

We argue that NLP for disaster response can have a positive impact on comprehensive situational awareness and in decision-making processes

such as coordination of particular services or physical goods. In the context of this work, positive impact means to supplement traditional information sources with social media streams that enable faster access to ongoing developments, first-party observations, and more fine-grained information content. For example, NLP for social media can enrich the information with the public as co-producers which may reveal critical subevents like missed or trapped people (Li et al., 2018). Retrieving this kind of information could positively affect disaster management strategies and relief efforts during natural and human-made disasters.

In contrast, relying on social media as an information source runs the risk of introducing mis- and disinformation. This can cause adverse effects on relief efforts and requires tailored strategies and particular care before the deployment of such models. Furthermore, data privacy issues may arise due to the inherited properties of social media data. Various anonymization processes should be taken into account for identifying and neutralizing sensitive references (Medlock, 2006). In this work, the use of entity tokens as categorization can be seen as one kind of anonymization procedure. However, model training with such entities could be task-specific and prone to error propagation by named entity recognition systems.

Acknowledgments

The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the project ISAKI (project number 13N15572). We also would like to thank the anonymous reviewers for their constructive feedback.

References

- Firoj Alam, Shafiq Joty, and Muhammad Imran. 2018a. [Domain adaptation with adversarial training and graph embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1077–1087, Melbourne, Australia. Association for Computational Linguistics.
- Firoj Alam, Ferda Ofli, and Muhammad Imran. 2018b. [Crisismmd: Multimodal twitter datasets from natural disasters](#). In *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM)*.
- Firoj Alam, Hassan Sajjad, Muhammad Imran, and Ferda Ofli. 2021. [Crisisbench: Benchmarking crisis-related social media datasets for humanitarian information processing](#). In *15th International Conference on Web and Social Media (ICWSM)*.
- Siddhartha Banerjee, Cem Akkaya, Francisco Perez-Sorrosal, and Kostas Tsioutsoulouklis. 2019. [Hierarchical transfer learning for multi-label text classification](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6295–6300, Florence, Italy. Association for Computational Linguistics.
- Eyal Ben-David, Carmel Rabinovitz, and Roi Reichart. 2020. [PERL: Pivot-based domain adaptation for pre-trained deep contextualized embedding models](#). *Transactions of the Association for Computational Linguistics*, 8:504–521.
- Rose Bernard, G. Bowsher, C. Milner, P. Boyle, P. Patel, and R. Sullivan. 2018. [Intelligence and global health: assessing the role of open source and social media intelligence analysis in infectious disease outbreaks](#). *Journal of Public Health*, 26(5):509–514.
- Alexander Brinkmann and Christian Bizer. 2021. [Improving hierarchical product classification using domain-specific language modelling](#). *Bulletin of the Technical Committee on Data Engineering / IEEE Computer Society*, 44(2):14–25.
- Cody L. Buntain, Richard McCreadie, and Ian Soboroff. 2021. [Incident Streams 2020: TREC-IS in the Time of COVID-19](#). In *ISCRAM 2021: 18th International Conference on Information Systems for Crisis Response and Management*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: Pre-training text encoders as discriminators rather than generators](#). In *ICLR*.
- 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)*, Minneapolis, Minnesota. Association for Computational Linguistics.
- Susan Dumais and Hao Chen. 2000. [Hierarchical classification of web content](#). In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '00*, page 256–263, New York, NY, USA. Association for Computing Machinery.
- João Rafael Gonçalves Evangelista, Renato José Sassi, Márcio Romero, and Domingos Napolitano. 2021. [Systematic Literature Review to Investigate the Application of Open Source Intelligence \(OSINT\) with Artificial Intelligence](#). *Journal of Applied Security Research*, 16(3):345–369. Publisher: Routledge_eprint: <https://doi.org/10.1080/19361610.2020.1761737>.

- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don't stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Matthias Habscheid, Nikolai Rodehutsors, and Rainer Koch. 2017. [Relevancy assessment of tweets using supervised learning techniques: Mining emergency related tweets for automated relevancy classification](#). In *2017 4th International Conference on Information and Communication Technologies for Disaster Management (ICT-DM)*, pages 1–8.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Marc-André Kaufhold. 2021. *Information Refinement Technologies for Crisis Informatics: User Expectations and Design Principles for Social Media and Mobile Apps*. Springer Fachmedien Wiesbaden, Wiesbaden.
- M. Kejriwal and P. Zhou. 2019. [Low-supervision urgency detection and transfer in short crisis messages](#). In *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 353–356, Los Alamitos, CA, USA. IEEE Computer Society.
- Jens Kersten, Anna Kruspe, Matti Wiegmann, and Friederike Klan. 2019. [Robust filtering of crisis-related tweets](#). In *ISCRAM 2019: 16th International Conference on Information Systems for Crisis Response and Management*.
- Anna Kruspe, Jens Kersten, and Friederike Klan. 2021. [Review article: Detection of actionable tweets in crisis events](#). *Natural Hazards and Earth System Sciences*, 21(6):1825–1845.
- Shamant Kumar, Geoffrey Barbier, Mohammad Abbasi, and Huan Liu. 2011. [Tweetracker: An analysis tool for humanitarian and disaster relief](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 5(1):661–662.
- Lifang Li, Qingpeng Zhang, Jun Tian, and Haolin Wang. 2018. [Characterizing information propagation patterns in emergencies: A case study with Yiliang Earthquake](#). *International Journal of Information Management*, 38(1):34–41.
- 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](#).
- Richard McCreddie, Cody L. Buntain, and Ian Soboroff. 2019. [TREC Incident Streams: Finding Actionable Information on Social Media](#). In *ISCRAM 2019: 16th International Conference on Information Systems for Crisis Response and Management*.
- Salvador Medina Maza, Evangelia Spiliopoulou, Edward Hovy, and Alexander Hauptmann. 2020. [Event-related bias removal for real-time disaster events](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3858–3868, Online. Association for Computational Linguistics.
- Ben Medlock. 2006. [An introduction to NLP-based textual anonymisation](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Harshit Mehrotra, Akanksha Mishra, and Sukomal Pal. 2022. [A Multi-stage Classification Framework for Disaster-Specific Tweets](#). *SN Computer Science*, 3(1):24.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Taro Miyazaki, Kiminobu Makino, Yuka Takei, Hiroki Okamoto, and Jun Goto. 2019. [Label embedding using hierarchical structure of labels for Twitter classification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6317–6322, Hong Kong, China. Association for Computational Linguistics.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.
- Marta Poblet, Esteban García-Cuesta, and Pompeu Casanovas. 2014. [Crowdsourcing tools for disaster management: A review of platforms and methods](#). In *AI Approaches to the Complexity of Legal Systems*, pages 261–274, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Christian Reuter, Amanda Lee Hughes, and Marc-André Kaufhold. 2018. [Social Media in Crisis Management: An Evaluation and Analysis of Crisis Informatics Research](#). *International Journal of Human-Computer Interaction*, 34(4):280–294.
- 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*, page 851–860, New York, NY, USA. Association for Computing Machinery.

- Hernan Sarmiento and Barbara Poblete. 2021. [Crisis communication: A comparative study of communication patterns across crisis events in social media](#). In *Proceedings of the 36th Annual ACM Symposium on Applied Computing, SAC '21*, page 1711–1720, New York, NY, USA. Association for Computing Machinery.
- Isabel Segura-Bedmar, Paloma Martínez, and María Herrero-Zazo. 2013. [SemEval-2013 Task 9 : Extraction of Drug-Drug Interactions from Biomedical Texts \(DDIExtraction 2013\)](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 341–350, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Luke S. Snyder, Yi-Shan Lin, Morteza Karimzadeh, Dan Goldwasser, and David S. Ebert. 2019. [Interactive Learning for Identifying Relevant Tweets to Support Real-time Situational Awareness](#). *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Congcong Wang, Paul Nulty, and David Lillis. 2021. [Transformer-based Multi-task Learning for Disaster Tweet Categorisation](#). In *ISCRAM 2021: 18th International Conference on Information Systems for Crisis Response and Management*.
- Jonatas Wehrmann, Ricardo Cerri, and Rodrigo Barros. 2018. [Hierarchical multi-label classification networks](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5075–5084. PMLR.
- Matti Wiegmann, Jens Kersten, Friederike Klan, Martin Potthast, and Benno Stein. 2020. [Analysis of Detection Models for Disaster-Related Tweets](#). In *ISCRAM 2020: 17th International Conference on Information Systems for Crisis Response and Management*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yipeng Yu, Zixun Sun, Chi Sun, and Wenqiang Liu. 2021. [Hierarchical multilabel text classification via multitask learning](#). In *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (IC-TAI)*, pages 1138–1143.
- Hamada M. Zahera, Rricha Jalota, Mohamed Ahmed Sherif, and Axel-Cyrille Ngonga Ngomo. 2021. [I-aid: Identifying actionable information from disaster-related tweets](#). *IEEE Access*, 9:118861–118870.

A Overview of Information Types

We list all information types of the TREC-IS dataset in Table 5. The value in the last column indicates the number of Twitter posts to which the corresponding labels were assigned. Table 6 displays example tweets for various events with the corresponding labels from the TREC-IS dataset.

B Hyper-Parameters

The search space for TF-IDF+LR included ngram-range, max features and regularization strength. In terms of BERT fine-tuning, we manually experimented with the same parameters as in Wang et al. (2021) and selected in line with this work the learning rate $5e - 5$ and batch size 32. Due to computational constraints, we used for BERT pre-training the TAPT parameters of Gururangan et al. (2020). Similar to Ben-David et al. (2020), we experimented with the MLM probabilities $\alpha \in \{0.1, 0.3, 0.5, 0.8\}$ and $\beta \in \{0.1, 0.3, 0.5, 0.8\}$ and found the setup $\alpha = 0.5$ and $\beta = 0.1$ to perform best. This is in line with Ben-David et al. (2020) which empirically show good results. For MTL we tuned $\lambda \in \{0.1, 0.5, 0.9\}$ and finally set $\lambda = 0.1$. We trained all transformer models with the *Transformers* library (Wolf et al., 2020) and *AdamW* for up to 50 (pre-training) and 15 (fine-tuning) epochs, evaluated the performance each 1000 steps on the development set and selected the best performing checkpoint. If not other mentioned, we used for the rest of the hyper-parameters the default setup of BERT_{BASE} from the *Transformers* library.

Id	Upper Class (L_T)	Lower Class (L_B)	Actionable (AIT)	# tweets
RQ 01	Request	GoodsServices	✓	194
RQ 02	Request	InformationWanted		395
RQ 03	Request	SearchAndRescue	✓	274
CTA 01	CallToAction	Donations		986
CTA 02	CallToAction	MovePeople	✓	679
CTA 03	CallToAction	Volunteer		242
O 01	Other	Advice		3,277
O 02	Other	ContextualInformation		4,583
O 03	Other	Discussion		5,303
O 04	Other	Irrelevant		23,053
O 05	Other	Sentiment		11,101
RP 01	Report	CleanUp		493
RP 02	Report	EmergingThreats	✓	6,930
RP 03	Report	Factoid		10,224
RP 04	Report	NewSubEvent	✓	2,806
RP 05	Report	FirstPartyObservation		5,290
RP 06	Report	Hashtags		15,787
RP 07	Report	Location		23,676
RP 08	Report	MultimediaShare		22,976
RP 09	Report	News		18,374
RP 10	Report	Official		2,836
RP 11	Report	OriginalEvent		4,148
RP 12	Report	ServiceAvailable	✓	2,184
RP 13	Report	ThirdPartyObservation		17,223
RP 14	Report	Weather		7,655

Table 5: Information types and hierarchical structure of labels.

Event	Labels	Tweet
Wildfire Colorado 2012	Irrelevant	From the train, showing the smoke filled sky from the #Lithgow #nswfires
Bushfire Australia 2013	ThirdPartyObservation, Factoid, Advice	FIRE UPDATE: Families told to be ready to run as a massive 300km wall of fire sweeps through Blue Mtns. #nswfires
Earthquake Chile 2014	News	New this morning: At least 6 people are dead after the massive M8.2 quake in #Chile
Explosion Beirut 2020	Location, Factoid, OriginalEvent, ContextualInformation	At least 25 dead and more than 2,500 injured as a result of the Beirut Port explosion according to the Lebanese Health Ministry
Flood Colorado 2013	Factoid	5 people confirmed dead in Colorado flooding, and 1,254 people unaccounted for statewide, official says
Hurricane Florence 2018	Weather, Location, Hash-tags	We have 2.5 inches here 2.6 miles northwest of Downtown awake Forest. #FlorenceHurricane2018

Table 6: Example tweets and labels for different events.