

Combining Deep Neural Reranking and Unsupervised Extraction for Multi-Query Focused Summarization

Philipp Seeberger and Korbinian Riedhammer

Technische Hochschule Nürnberg Georg Simon Ohm

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

Abstract

The CrisisFACTS Track aims to tackle challenges such as multi-stream fact-finding in the domain of event tracking; participants' systems extract important facts from several disaster-related events while incorporating the temporal order. We propose a combination of retrieval, reranking, and the well-known Integer Linear Programming (ILP) and Maximal Marginal Relevance (MMR) frameworks. In the former two modules, we explore various methods including an entity-based baseline, pre-trained and fine-tuned Question Answering systems, and ColBERT. We then use the latter module as an extractive summarization component by taking diversity and novelty criteria into account. The automatic scoring runs show strong results across the evaluation setups but also reveal shortcomings and challenges.

1 Introduction

Natural and human-made disasters can result in significant loss of life, property, and environment if situational awareness is insufficient due to a lack of critical information in an ongoing emergency event. Today's information ecosystem offers new opportunities and directions for emergency response by integrating various online information sources (Buntain et al., 2021; Kruspe et al., 2021). Additional information sources such as social media and microblogging platforms can immediately provide details about current developments (Sakaki et al., 2010; Reuter et al., 2018). This leads to a multi-stream setting in which traditional sources are complemented with a variety of recently emerged online sources. Previous research efforts acknowledged this setting as a promising venue, as shown by the evolving tasks over several decades (Allan et al., 1998; Aslam et al., 2015; Sequiera et al., 2018; Buntain et al., 2021).

However, the high-velocity nature of content generation and the inherent properties of different

information sources (Kaufhold, 2021) and events (Seeberger and Riedhammer, 2022) face present models with new challenges. Those provide relevant event-related results to the user but are still ill-suited to multi-stream fact-finding and summarization needs. The novel CrisisFACTS Track aims to tackle these issues and challenges the community to develop systems more suited for factoid extraction over time. Overall the task asks participants' systems to extract a query-focused list of facts from crisis-related datasets, including *Twitter*, *Reddit*, *Web News*, and *Facebook* as data sources. Each of these extracted lists of facts is based on an event-day pair and shipped with importance scores which serve as a basis for downstream summarization. In fact, this can be considered as an extension of previous tasks in the area of Information Retrieval (IR) and summarization.

The recent incorporation of pre-trained language models such as BERT (Devlin et al., 2019) has significantly improved ad-hoc ranking (Lin et al., 2022) and summarization (Ma et al., 2022) results. In particular, BERT-based cross-encoders achieve notable improvements over classical retrieval and dual-encoder approaches but have infeasible computational costs (Khattab and Zaharia, 2020). To mitigate this issue, deep neural ranking models are typically deployed as second stage rerankers, whereby the first stage often represents an efficient retriever to create a subset of candidate documents (MacAvaney et al., 2022). Similarly, modern summarization methods rely on cross-encoders to fetch relevant documents, paragraphs, or sentences for both extractive summarization (Xu and Lapata, 2020; Ahuja et al., 2022) or as preliminary selection for abstractive summarization (Xu and Lapata, 2021). Finally, the resulting pool of ranked documents can be further refined through approaches such as MMR (Carbonell and Goldstein, 1998), ILP (McDonald, 2007), and TextRank (Mihalcea and Tarau, 2004).




Figure 1: Overview of our proposed framework. All queries and documents for each event and time period are separately processed by the following three major components: (1) Retrieve, (2) Rerank, and (3) Summarize. The symbols E1 to E4 represent the concepts w.r.t. the ILP formulation. For final scoring, the selected \mathcal{S}_{sel} and past summary \mathcal{S}_{past} documents are used in terms of redundancy penalization.

In this work, we explore various information retrieval and reranking pipelines ranging from pre-trained to fine-tuned state-of-the-art models. Complementary, we propose to subsequently process the list of facts in an extractive summarization setup by leveraging a combination of the well-known ILP and MMR frameworks. In this way, we aim to overcome issues related to diversity and redundancy.

2 Approach

As illustrated in Figure 1, our proposed framework first retrieves and reranks a set of documents $\mathcal{D} = \{d_1, \dots, d_N\}$ based on an information request $\mathcal{Q} = \{q_1, \dots, q_M\}$ where each query q_i typically consists of a short text or list of indicative terms. The number of documents¹ and queries are given as N and M , respectively. Let $\mathcal{C} = \{\mathcal{C}^{(q_1)}, \dots, \mathcal{C}^{(q_M)}\}$ denote the resulting set of query-related clusters with $\mathcal{C}^{(q_i)} = \{d_1^{(q_i)}, \dots, d_k^{(q_i)}\}$ consisting of the top k candidates ranked by relevance. Then, a summarization component further selects L candidates from the cluster pool $\bigcup_{i=1}^M \mathcal{C}^{(q_i)}$ to create a summary \mathcal{S} . Following the track design, each summary \mathcal{S}_t is created w.r.t. a time period $p_t \in \{p_1, \dots, p_T\}$ with T as the number of time periods. Within the scope of CrisisFACTS, the time period p_i corresponds to one day. In the following, we detail each individual component used for our submissions.

¹Throughout this work, we use the term document interchangeably with the CrisisFACTS stream items. These stream items are rather sentences or short posts than long documents.

2.1 Stage 1: Retrieve

In the first stage, we employ lexical retrieval approaches to mitigate the infeasible computational costs of deep neural models such as BERT-based cross-encoders. Hence, we first retrieve the top $k^{(1)}$ candidates for each query $q \in \mathcal{Q}$ from a set of documents \mathcal{D} using a list of indicative terms. For each query q , the reduced set of $k^{(1)}$ candidates is then subsequently processed by the given reranking stage. For retrieval, we adopt the well-known BM25 model (Robertson and Zaragoza, 2009) but one can easily replace it with more sophisticated methods. We empirically found at preliminary experiments that the number of candidates is relatively low, limiting the subsequent reranking components. To address this problem, one can implement query expansion (Amati and Van Rijsbergen, 2002), document expansion (Nogueira et al., 2019), or adaptive reranking (MacAvaney et al., 2022) methods. In this work, we integrate query expansion in order to grow the candidates pool for overcoming the recall limitation.

2.2 Stage 2: Rerank

Classical retrieval approaches may not be sufficient to capture the semantics expressed in the query and its relation to the text contents. Therefore, each query-related cluster $\mathcal{C}^{(q)}$ of the previous retrieval stage is reranked and the top $k^{(2)}$ candidates are selected, resulting in a new set $\mathcal{C}'^{(q)}$. Here, we exploit supervision signals from other existing datasets by using pre-trained deep neural models. Furthermore,

we include an entity and keyword-based baseline in order to assess the performance gain due to the models pre-trained on large text corpora. In the following, we detail the considered reranking models:

BoE As baseline we adopt the Bag-of-Entities representation for document ranking introduced by Xiong et al. (2016) which relies on the semantic information achieved by entity linking systems. We implement a simplified version by constructing Bag-of-Entities vectors for each document based on entity-types and extend them with Bag-of-Keywords vectors. The final model scores a document by summing over the frequency of expected query entity-types and keywords present.

QA Similar to Xu and Lapata (2020), we employ Question Answering (QA) systems to leverage distant supervision signals related to best answer selection. While QA approaches support both sentence and span selection, we rely only on sentence level selection which suits more to the queries and stream items provided by the organizers. That is, we concatenate the query q and a candidate document d into a sequence [CLS] q [SEP] d and predict the relevance score with a BERT-based cross-encoder. In this work, we consider a pre-trained and fine-tuned QA version whereby the fine-tuned system is adapted to the crisis domain.

ColBERT This model follows a contextualized late interaction approach which makes use of both a first-stage approximate nearest neighbor search and a reranking stage to calculate ranking scores (Khattab and Zaharia, 2020). In particular, ColBERT supports the reranking mechanism to produce more precise scores based on a candidate pool but can also be used for end-to-end retrieval. We follow the end-to-end retrieval approach. In this way, we consider an approach without the limitations related to classical retrieval models such as BM25. Note that this alternative skips the first retrieval stage by directly retrieving the set of documents \mathcal{D} to obtain the top $k^{(2)}$ candidates for each query q .

2.3 Stage 3: Summarize

Selection Finding a diverse set of facts with less redundancy is crucial for summarization tasks. However, without any post-processing, reranked candidates still suffer in terms of diversity and redundancy. To tackle this problem, we use an additional selection step formalized as ILP. We follow the concept-based model (Gillick and Favre, 2009;

Riedhammer et al., 2010) where concepts can be facts, events, or information units. In this problem setup, the objective function is maximized over the weighted sum of the concepts present in the selection, subject to a length constraint. Finally, we obtain an extractive summary $\mathcal{S}_t = \{d_1, \dots, d_l\}$ where $|\mathcal{S}_t|$ is limited by $l \leq L$ with L as the maximum number of documents.

Scoring The well-known MMR algorithm greedily selects documents by trading off query-based relevancy and redundancy to the previously selected documents, until a summary length constraint is met. However, this constraint can be relaxed to rerank a summary \mathcal{S}_t in order to increase the diversity in the top documents. Formally, we define the final score of a document i as

$$\lambda \cdot \text{Rel}_i - (1 - \lambda) \cdot \max_{j \in \mathcal{S}_{sel} \cup \mathcal{S}_{past}} \text{Red}_{ij} \quad (1)$$

where Rel_i is the relevance score of document i and Red_{ij} is the redundancy penalty for having both documents i and j in the summary \mathcal{S}_{sel} as well as past summaries $\mathcal{S}_{past} = \bigcup_{i=1}^{t-1} \mathcal{S}_i$. However, a single retrieved document might contain multiple scores due to multiple matched queries. We argue that a document that covers multiple queries expresses more relevant information content for the summary. Formally, we denote the relevance score as $\text{Rel}_i = |Q^{(i)}| \cdot \text{score}_i$ where score_i is the mean score of document i weighted by the number of matched queries $Q^{(i)} \subseteq Q$.

3 Experiments

In this section, we detail the experimental setup and discuss the results for our submitted runs. Throughout all experiments, we mainly consider the sources *Twitter*, *Reddit*, *Web News* and ignore *Facebook* due to the limited access to the post contents.

3.1 Preprocessing

We normalize all tweets in order to represent the text content similar to the other online sources. Specifically, all retweet-indicating prefixes, user mentions, emoticons, emojis, and URLs are removed. Furthermore, we remove any hashtag symbols and split the text into their corresponding words using *WordSegment*.²

²<https://grantjenks.com/docs/wordsegment>

	Summarization			Matching	
	ICS	NIST	Wiki	Comprehensiveness	Redundancy
ColBERT	.050/. 450	.139/.546	.031/.542	.189	.201
BM25 → BoE	.047/.436	.142/.560	.030/.533	.185	.176
BM25 → QA _{ASNQ}	.051 /.448	.147 /.563	.036 /.565	.213	.226
BM25 → QA _{Crisis}	.046/.443	.147 /.564	.034/.545	.210	.226
TREC best	.058/.459	.147/.564	.036/.565	.217	.125

Table 1: Overall results of our automatic submission runs. We report the Rouge-2/BERTScore for summarization and comprehensiveness and redundancy ratio for matching as defined in Appendix A. The top results across our proposed systems are in bold.

Event	ICS	NIST	Wiki
001	.116*/.522	.273 /.560*	.013*/.540
002	.066 /.561	.050*/.563	.043/. 579
003	.053/.516	.238/.611*	.021/.593*
004	.061/.480	.171*/ .585	.060 /. 582
005	-	.136/.544*	.032/.526*
006	.057/.506	.048/.533	.019*/.580
007	.040/.494*	.104*/.524*	.057 /. 554
008	.012/.501	.154 /.583	.044 /. 562

Table 2: Rouge-2/BERTScore results for each event w.r.t. the QA_{ASNQ} system. The TREC best results across all submissions are in bold. Results below the TREC median are in grey, while results below our baseline BoE are marked with *.

3.2 Crisis-QA

Since the first CrisisFACTS Track does not provide any annotations w.r.t. the task, we decided to create a synthetic version that reflects the query-focused sentence selection. We leverage the DocEE dataset (Tong et al., 2022), a recently published benchmark for document-level event extraction. We extract a subset of 6818 documents which only covers crisis-related events and their corresponding event arguments.³ First, we manually create coarse-grained questions for each event argument. Second, the dataset is augmented with a T5_{BASE} question generation model⁴ for obtaining fine-grained questions. Last, we synthesize question-sentence pairs based on the argument position and label this pairs as binary relevance classification task. For model validation, we use the published dataset splits.

³We checked for an overlap between the DocEE and CrisisFACTS events. In fact, some of the events are part of the DocEE dataset and thus we removed the corresponding documents prior to our experiments.

⁴<https://huggingface.co/mrm8488/t5-base-finetuned-question-generation-ap>

3.3 Experimental Setup

Retrieve For the first stage, we use the BM25 model with default settings of the PyTerrier library (Macdonald and Tonello, 2020) and extend it with the Bo1 query expansion (Amati and Van Rijsbergen, 2002) component. For each query, we concatenate the query text and indicative terms, retrieve the top $k^{(1)} = 100$ candidates, and drop exact duplicates. The majority of duplicates appear in the tweet documents which is mostly related to retweets.

Rerank The BoE model is based on a manually curated set of entity-types that mostly fits the expected information needs w.r.t. each query. For example, queries about missing peoples typically cover numbers and locations, respectively. The indicative terms provided by the organizers are used for the keywords. The QA_{ASNQ} system is based on RoBERTa_{BASE} pre-trained on the ASNQ dataset (Garg et al., 2020) without any further adjustments. Similarly, we employ the ColBERT_{v2} version (Santhanam et al., 2022) which is trained on the MS MARCO Passage Ranking task. In terms of QA_{Crisis}, we follow the adaptation step of Garg et al. (2020) by fine-tuning the QA_{ASNQ} model on the domain-specific Crisis-QA dataset. This results in an adapted version of the QA system. Although the synthesized dataset relies on a broad range of labeled event arguments, we still observe a significant proportion of false negatives within the question-sentence pairs. Hence, we use the model QA_{Crisis-0} in a first step to denoise the dataset with an upper threshold of 0.1 and then train a new model QA_{Crisis-1} in a second step, which is in line with previous work such as RocketQA (Qu et al., 2021). We use the Transformers library (Wolf et al., 2020) for the QA models, the official implemen-

tation of ColBERT, and select the top $k^{(2)} = 25$ candidates for each query.

Summarize To enable a fair comparison among the different retrieval and reranking components, we re-use the selection and scoring procedure for each run. Specifically, inspired by information extraction (Martinez-Rodriguez et al., 2020), we extract entities⁵ as concepts, entity-frequency as weights, and set $L = 150$ for the ILP formulation. For MMR, we select $\lambda = 0.8$ and calculate the redundancy Red_{ij} based on TF-IDF features and cosine similarity.

3.4 Results

In Table 1, we present the overall performance of our pipeline setups. Since this is the first installment of the CrisisFACTS Track, we mainly limit the analysis across our submission runs. However, we provide the reader a comparison of our models to the medians and top results for the summarization task (Table 2).

Overall The QA models outperform the baseline BoE and ColBERT in almost all evaluation settings. These results reflect the findings of previous text retrieval work which report higher performance for cross-encoder architectures. Interestingly, the fine-tuned QA model decreases the performance in two summarization setups and in terms of comprehensiveness. We assume that the adapted QA is biased towards the entities of the Crisis-QA dataset. This might result into higher scores for only a subset of facts. Furthermore, we are aware of concerns about potential data overlap due to the time intersection between the CrisisFACTS and Crisis-QA events. However, the performance increase appears only for the NIST reference summaries and we therefore leave the analysis for future work.

Summarization In depth analysis in Table 2 found that the pre-trained QA model achieves top results for a variety of events and reference summaries. When compared to the BoE baseline, the performance increase differs among the events, metrics, and reference summaries. However, only three performance measures are below the TREC medians which suggests strong results for the overall pipeline. Nevertheless, in contrast to automatic summarization evaluation, manual matching re-

⁵<https://stanfordnlp.github.io/stanza/ner.html>




Figure 2: Comprehensiveness trend for all events. The QA_{ASNO} system is displayed in bold, while the min-max region of all models is highlighted.

veals high variance for different days within the same event.

Matching If we plot the comprehensiveness evolution along the number of days (Figure 2), we see that the performance decreases by large extent across a variety of events. Since this trend holds for all models, we hypothesize that this is due to at least two factors. First, the retrieval and reranking stages of the pipeline setup does not consider diversity for each query and might cut off rare facts in favor of facts with higher relevance, spread along the timeline of the event. Second, the diversification in the selection stage w.r.t. past summaries still displays a challenging task. For example, specific sentences only differ by a single number (e.g. burned acres) and might unintentionally penalize new facts by unsophisticated similarity measures.

4 Conclusion

In this work, we have investigated the combination of deep neural reranking and global unsupervised extraction for a multi-query focused summarization task. Our experiments demonstrated the strength of cross-encoders with QA based on distant supervision. However, we identified shortcomings and challenges in the face of temporal aspects which underlines the downstream summarization as a critical component. We believe there is much room for improvement, especially by integrating more sophisticated extractive approaches, abstractive summarization techniques, or even joint optimization.

Acknowledgments

The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the project ISAKI (project number 13N15572).

References

- Ojas Ahuja, Jiacheng Xu, Akshay Gupta, Kevin Horecka, and Greg Durrett. 2022. [ASPECTNEWS: Aspect-Oriented Summarization of News Documents](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6494–6506, Dublin, Ireland. Association for Computational Linguistics.
- James Allan, Jaime G. Carbonell, George Doddington, Jonathan Yamron, and Yiming Yang. 1998. [Topic Detection and Tracking Pilot Study Final Report](#). Publisher: Carnegie Mellon University.
- Gianni Amati and Cornelis Joost Van Rijsbergen. 2002. [Probabilistic models of information retrieval based on measuring the divergence from randomness](#). *ACM Transactions on Information Systems*, 20(4):357–389.
- Javed A. Aslam, Fernando Diaz, Matthew Ekstrand-Abueg, Richard McCreadie, Virgil Pavlu, and Tetuya Sakai. 2015. [TREC 2015 Temporal Summarization Track Overview](#). In *Proceedings of The Twenty-Fourth Text REtrieval Conference, TREC 2015, Gaithersburg, Maryland, USA, November 17-20, 2015*, volume 500-319 of *NIST Special Publication*. National Institute of Standards and Technology (NIST).
- 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*.
- Jaime Carbonell and Jade Goldstein. 1998. [The use of MMR, diversity-based reranking for reordering documents and producing summaries](#). In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval - SIGIR '98*, pages 335–336, Melbourne, Australia. ACM.
- 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*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Siddhant Garg, Thuy Vu, and Alessandro Moschitti. 2020. [TANDA: Transfer and Adapt Pre-Trained Transformer Models for Answer Sentence Selection](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7780–7788.
- Dan Gillick and Benoit Favre. 2009. [A Scalable Global Model for Summarization](#). In *Proceedings of the Workshop on Integer Linear Programming for Natural Language Processing, ILP '09*, pages 10–18, USA. Association for Computational Linguistics. Event-place: Boulder, Colorado.
- 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.
- Omar Khattab and Matei Zaharia. 2020. [CoBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT](#). In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 39–48, Virtual Event China. ACM.
- A. Kruspe, J. Kersten, and F. Klan. 2021. [Review article: Detection of actionable tweets in crisis events](#). *Natural Hazards and Earth System Sciences*, 21(6):1825–1845.
- Jimmy Lin, Rodrigo Nogueira, and Andrew Yates. 2022. [Pretrained Transformers for Text Ranking: BERT and Beyond](#). Synthesis Lectures on Human Language Technologies. Springer International Publishing, Cham.
- Congbo Ma, Wei Emma Zhang, Mingyu Guo, Hu Wang, and Quan Z. Sheng. 2022. [Multi-document summarization via deep learning techniques: A survey](#). *ACM Comput. Surv.*, 55(5).
- Sean MacAvaney, Nicola Tonello, and Craig Macdonald. 2022. [Adaptive re-ranking with a corpus graph](#). In *Proceedings of the 31st ACM International Conference on Information Knowledge Management, CIKM '22*, page 1491–1500, New York, NY, USA. ACM.
- Craig Macdonald and Nicola Tonello. 2020. [Declarative Experimentation in Information Retrieval Using PyTerrier](#). In *Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval, ICTIR '20*, pages 161–168, New York, NY, USA. ACM. Event-place: Virtual Event, Norway.
- Jose L. Martinez-Rodriguez, Aidan Hogan, and Ivan Lopez-Arevalo. 2020. [Information extraction meets the Semantic Web: A survey](#). *Semantic Web*, 11(2):255–335. Publisher: IOS Press.
- Ryan McDonald. 2007. [A Study of Global Inference Algorithms in Multi-Document Summarization](#). In *Proceedings of the 29th European Conference on IR Research, ECIR'07*, pages 557–564, Berlin, Heidelberg. Springer-Verlag. Event-place: Rome, Italy.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing Order into Text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.

- Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. 2019. [Document Expansion by Query Prediction](#). Publisher: arXiv Version Number: 2.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. [RocketQA: An Optimized Training Approach to Dense Passage Retrieval for Open-Domain Question Answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5835–5847, Online. Association for Computational Linguistics.
- 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.
- Korbinian Riedhammer, Benoit Favre, and Dilek Hakkani-Tür. 2010. [Long story short – Global unsupervised models for keyphrase based meeting summarization](#). *Speech Communication*, 52(10):801–815.
- Stephen Robertson and Hugo Zaragoza. 2009. [The Probabilistic Relevance Framework: BM25 and Beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389. Place: Hanover, MA, USA Publisher: Now Publishers Inc.
- 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, Raleigh, North Carolina, USA. ACM.
- Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. [ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction](#). pages 3715–3734, Seattle, United States. Association for Computational Linguistics.
- Philipp Seeberger and Korbinian Riedhammer. 2022. [Enhancing crisis-related tweet classification with entity-masked language modeling and multi-task learning](#). Publisher: arXiv Version Number: 1.
- Royal Sequiera, Luchen Tan, and Jimmy Lin. 2018. [Overview of the TREC 2018 Real-Time Summarization Track](#). In *Proceedings of the Twenty-Seventh Text REtrieval Conference, TREC 2018, Gaithersburg, Maryland, USA, November 14-16, 2018*, volume 500-331 of *NIST Special Publication*. National Institute of Standards and Technology (NIST).
- MeiHan Tong, Bin Xu, Shuai Wang, Meihuan Han, Yixin Cao, Jiangqi Zhu, Siyu Chen, Lei Hou, and Juanzi Li. 2022. [DocEE: A Large-Scale and Fine-grained Benchmark for Document-level Event Extraction](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3970–3982, Seattle, United States. Association for Computational Linguistics.
- 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.
- Chenyan Xiong, Jamie Callan, and Tie-Yan Liu. 2016. [Bag-of-Entities Representation for Ranking](#). In *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval*, pages 181–184, Newark Delaware USA. ACM.
- Yumo Xu and Mirella Lapata. 2020. [Coarse-to-Fine Query Focused Multi-Document Summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3632–3645, Online. Association for Computational Linguistics.
- Yumo Xu and Mirella Lapata. 2021. [Generating Query Focused Summaries from Query-Free Resources](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6096–6109, Online. Association for Computational Linguistics.

A Matching Metrics

The submitted stream items for a system and event-day pair are ordered by the importance scores and formed to a summary S by a rank cut-off k . Based on a manual matching of the stream items against a fact list F , the comprehensiveness is calculated as

$$\frac{1}{\sum_{f \in F} R(f)} \sum_{\{f \in F: M(f, S) \neq \emptyset\}} R(f) \quad (2)$$

where $f \in F$ is a unique fact, $M(f, S)$ is the set of stream items matching fact f , and $R(f)$ is the gain assigned to the fact f . Similarly, the redundancy ratio is measured for a system and event-day pair as

$$\frac{\sum_{\{f \in F: M(f, S) \neq \emptyset\}} R(f)}{\sum_{\{f \in F\}} R(f) \cdot |M(f, S)|} \quad (3)$$

All runs are macro-averaged across days within an event, and then across the eight events.