Multiple View Summarization Framework for Social Media

Chih-Yuan Li[1], Soon Ae Chun[2], James Geller[3]
[1,3]New Jersey Institute of Technology, [2]City University of New York – College of Staten Island
{cl524, james.geller}@njit.edu, soon.chun@csi.cuny.edu

Abstract

Social Media provide voluminous posts about current topics and events. When a user desires to investigate a popular topic, it is not feasible as there are many posts. Besides, posts show different biases, viewpoints, perspectives, and emotions. Thus, providing summaries of large post sets with different viewpoints is necessary. We develop a multiple view summarization framework to generate different view-based summaries of Twitter posts. Users can apply different methods to generate summaries: 1) Entity-centered, 2) Social feature-based, 3) Event-based summarization, using all triple embeddings and 4) Sentiment-based summarization to generate summaries of positive or negative views of tweets. These summarization methods are compared with BertSum, SBert, T5, and Bart-Large-CNN with a gold standard dataset. Our results, based on Rouge scores, were better than these published extractive and abstractive summarization models.

Introduction

Microblogging sites, such as Twitter, Facebook, etc. have become important sources of information about socio-political events, natural disasters, pandemics, etc. Twitter has around 450 million monthly active users, posting 200 billion tweets in 2022 (Ruby 2022). It is infeasible for humans to review all the tweets on a certain topic, e.g., for identifying fake news to prevent their spread.

Most current summarization models generate one summary for one or multiple coherent documents, e.g., a news article. Single posts on social media are short and incoherent, but sets of posts contain voluminous content, and convey many aspects of the same event. Thus, generating one summary of microblogging post sets may be inadequate. It is better to provide users with a system to generate multiple summaries focusing on different perspectives. For example, a summary view can be *entity focused*. A COVID-related summary can be generated about Pfizer vaccines, while an election summary can be about a new candidate. Another summary may focus on significant *events*, e.g., what an election candidate has done or lied about. Alternatively, a summary view can focus on *influential posts* that are most often shared or liked. It would be helpful to derive a summary with combined views, e.g., a summary of posts on an election candidate expressing negative sentiments.

We developed a Multiple View Summarization Framework (MVSF), which generates summaries of a dataset reflecting different perspectives: 1) Entity-based summarization (E), 2) Triple Distance-based summarization (D), and 3) Social Feature-based Summarization (SF), which can be combined with 4) Sentiment-based summarization (S). These views can uncover different aspects in posts. Views can be combined (e.g., S+E, or S+D).

For each post, we extracted <subject, predicate, object> triples (Sintek and Decker 2002), applied the view-based summarization, identified top groups of actions and entities, and ranked triples. We achieved better Rouge F-1 scores (Lin 2004) than those generated by extractive and abstractive summarization models from the literature.

Related Work

Algorithms for microblog summarization have been proposed, e.g., (Modhe et al. 2021; Geng et al. 2020; Saini et al. 2019; Dutta et al. 2019). *Abstractive summarization* methods include bart-large-cnn (Lewis et al. 2019), T5 (Raffel et al. 2019), FactSumm (Heo 2021), XNLG (Chi et al. 2019), FAIRSEQ (Ott et al. 2019), etc. *Extractive summarization* methods that retain a subset of the sentences that best represent the original document include BertSum, SBert, HAHSum (Jia et al. 2020), NeRoBERTa (Kwon et al. 2021), DebateSum (Roush and Balaji 2020), MemSum (Gu et al. 2021), etc. *Entity summarization* methods include FACES and FACES-E for single and multi-entity-based summaries (Gunaratna et al. 2015, 2016).

Dutta et al. (2019) performed the first systematic analysis of microblog summaries posted during disasters. They compared performances of different algorithms that generated summaries with significant differences from the same input. Joshi et al. (2023) proposed a method based on LDA topic modeling and GloVe word embeddings for the extractive summarization of single documents.

Unlike most of the approaches in the literature that focus on a single summary, our work is able to generate variant summaries with different perspectives from tweet posts.

These summaries with multiple views can also be either utilized alone or combined to generate fine-grained summaries.

Multiple View Summarization

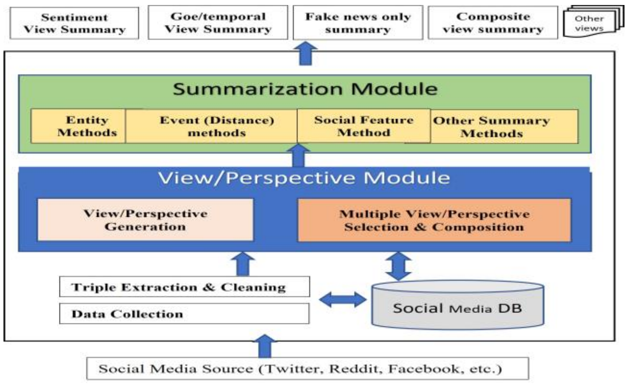

Figure 1: Multiple View Summarization Framework (MVSF).

Social media posts often cover divergent aspects on a topic (e.g., elections, pandemics). Our MVSF addresses these diverse aspects that should be reflected in the summaries. A dataset can be summarized according to different perspectives ("views"). The architecture of MVSF is in Figure 1.

Knowledge Triple Extraction Social media users are often more interested in tweets describing events or activities (Liu et al. 2012). Therefore, we focus on <SPO> triples describing them. We used the (Angeli et al. 2015) model to extract triples from post sentences. The algorithm splits a sentence into shorter clauses. It predicts at each step whether an edge should yield a new independent clause (Niklaus et al. 2018).

View Generation Methods

Multiple Views Analysis:

Social media data can be analyzed into various views or perspectives. *Negative sentiment tweets* express one perspective. *Fake news* can be another view one may want to extract and compare with *real news views*. These view generation/analysis components generate datasets using data analytics and data selection. For instance, the Sentiment-based view (S) requires sentiment analysis on the dataset to select a one polarity view (positive or negative). We applied a Sentiment Analyzer (Socher et al. 2013). We can choose to focus only on the triples with a specified sentiment. MVSF provides the flexibility to compose views to create more focused summaries, by combining S with the E, D, or SF. Similarly, it can generate a set of data of only fake news, or a conservative/liberal political view, etc.

Multiple View Composition: The power of the MVSF is its capability to choose views and compose them together. For instance, when a user is interested in a *fake news only view* with negative sentiments, the view generation could deliver this combined view. The multiple view composition module generates a composite view, $cv = compose\ (v_1, v_2, \ldots)$ where v_i is a single view.

Summarization Methods

Entity-based Summarization

In this view, the entities define the primary perspective for obtaining summaries. Unlike other research where the user provides the entities of interest (e.g., "Vaccine"), we start with posts mentioning a vast number of entities. The important entities can appear in different events, marked by different actions. Thus, we identify predicates expressing similar meanings with different words, such as "offer," "offered," "provide," and "provided," all of which express a meaning of "giving something to someone." Groups of similar verbs can be subsumed by one root verb, using a synset from WordNet (Fellbaum 2005). A synset is a set of one or more synonyms. We then identify the frequent events by selecting triples whose root verbs occur more often than a threshold θ. Given a triple set TS= {ts1, ts2 …} where tsi= <s, p, o>, find tsi= <s, root(p), o>. We select triples having frequent root verbs root(p). Then we identify important entities to focus on. Each triple includes a subject and an object. To identify the best summary, we experimented with different weights (α) for subjects and objects. We select the top-scoring $\mathbf{m_E}$ triples and their original sentences.

$$\textbf{Triple Score} = \alpha * (\text{Subject Score}) + (1-\alpha) * (\text{Object Score}) \quad (1)$$
$$\text{where } \alpha \in [0.0, 1.0]$$

Triple Distance-based Summarization

In this view, we represent the event expressed in a triple sentence. To create sentence representations of triples, we use BERT sentence embeddings. BERT reads the entire input sequence at one time, which allows it to learn the context of each word based on the neighbors (left and right). To achieve a lower dimensional representation, we used an autoencoder to learn a 32-dimensional vector representation of each triple. Instead of focusing on specific entities, this view summary reflects top contextual information conveyed among triples. We identified the centroid, which is the avg of triple vectors in the cleaned triple set (see Cleaning Process below). Each triple t is represented as a vector of dimension d:

$$\mathbf{v_t} = [\mathrm{v_{t1}}, \mathrm{v_{t2}}, \ldots, \mathrm{v_{td}}] \quad (2)$$

The centroid vector $\mathbf{v_c}$ of all triple vectors is:

$$\boldsymbol{v_c} = \frac{1}{n}\left(\sum_{t=1}^{n} \mathrm{v_t}\right) = [\mathrm{v_{c1}}, \mathrm{v_{c2}}, \ldots, \mathrm{v_{cd}}] \quad (3)$$

where n is the triple count.
The Euclidean distance d from a triple t to the centroid c is:

$$\mathbf{d}_{tc} = sqrt\left(\sum_{i=1}^{32} (\mathrm{v_{ti}} - \mathrm{v_{ci}})^2\right) \quad (4)$$

We selected $\mathbf{m_D}$ triples with the minimum $\mathbf{d}_{tc}$. The triples' corresponding original sentences were recovered to form the summary.

Social Feature-based Summarization

In this view, we exploit the user's followers and retweets to identify the saliences of tweets. By (Liu et al. 2012), a tweet is more important if it has been retweeted many times, and posted by a user with many followers. They defined the salience score of a tweet as the product of retweet #, user follower # and readability. Our goal is to identify summaries with high impact. Thus, we redefined the salience score:

$$\textbf{Salience Score} = \begin{cases} \text{follower} + \text{retweet} * \text{Avg(f)}, & \text{follower} > 0 \\ 0, & \text{follower} = 0 \end{cases} \quad (5)$$

where Avg(f) is the avg # of followers of a user. When a post is retweeted, there will be on avg 707 people who will see it (Petrov 2023). We rank triples based on the scores of their original tweets, and select the $\mathbf{m_{SF}}$ top-scoring triples. The triples' original sentences are selected as the summary.

Summarization Methods Evaluation

We evaluated our MVSF by comparing with extractive and abstractive summarization models from the literature. Extractive models are BertSum and SBert, which are variants of BERT we used in one of the views. Abstractive summarization models are bart-large-cnn (Lewis et al. 2019) and T5 (Raffel et al. 2019), which according to (Li et al. 2022b), generate better summaries compared with TextRank (Barrios et al. 2016) and GPT-2 (Radford et al. 2019).

To our knowledge, there is no social media dataset with gold standard summaries for evaluation. To work around this, we used a dataset of BBC news items (Greene and Cunningham 2006) with their human-generated summaries. The BBC dataset contains 2,225 news documents, covering business, entertainment, sport, politics, and tech. Each document is paired with an extractive summary. We selected 20 news documents from the business category. As this dataset doesn't have social signals, we couldn't apply the SF summarization results. Each of the generated summaries is approximately 300 words long, which applies to our methods and the published models. We used Rouge scores as the evaluation metrics, which compare a generated summary against a gold standard summary, in 1-gram, 2-gram, and longest common subsequence (LCS) units.

All of our view summaries perform better than those by the models shown in Table 1. The view (S+D) performed best with the highest score for 1-gram and LCS, while the view (D) performed best for 2-gram. Both views (E) and (S+E) achieved the best performances when α = 0.7. This result shows that the summaries generated by our summary methods with views are more consistent with the gold standard than existing models.

Table 1: Rouge F-1 Scores of Summaries.

Rouge Score	1-gram	2-gram	LCS
View (E), α = 0.7	0.465	0.310	0.444
View (D)	0.557	**0.427**	0.551
View (S+D)	**0.58**	0.416	**0.572**
View (S+E), α = 0.7	0.459	0.299	0.433
BertSum	0.406	0.174	0.394
SBert	0.444	0.152	0.428
bart-large-cnn	0.402	0.15	0.394
T5	0.304	0.104	0.304

We observed the summaries are notably different, especially between view E and view D. Either with or without view S, there is NO overlapping sentence between view E and view D. There are two sentences common in D and S+D, and three sentences common in view E and view S+E.

COVID-19 Vaccine Tweet Summarization

We applied our multi-view summarization framework to COVID-19 Vaccine tweets as a case study.

Data: We used a dataset of 18,047 tweets about the COVID vaccines Pfizer/BioNTech, Sinopharm, Sinovac, Moderna, Oxford/AstraZeneca (AZ), Covaxin, and Sputnik V (Preda 2021). We removed the non-ASCII codes, URLs, and redundant spaces/punctuations. We identified the vaccine types by the hashtags in the posts. We assigned each post an index, starting from 0. We applied Stanza (Qi et al. 2020) to split posts into sentences, as it is beneficial to work with post fragments (sentences) rather than entire posts (Rudra et al. 2015). We assigned each sentence of a post an index. We obtained 28,242 sentences. We extracted 101,432 triples, and linked them to their corresponding original sentence and post indices. Table 2 shows the two triples from the two sentences (0 and 1) from the same post with ID 484. By cleaning triples with auxiliary verbs (that carry little meaning) and removing duplicates generated from the same sentence, we reduced the triple sets from 67,049 to 16,270, which is the input for each summarization view.

Table 2: Extracted Triples and Index Labels.

Post Ind.	Sent. Index.	Triple Index.	Subject	Verb	Obj.
484	0	0	shipment	arrived in	vietnam
484	1	0	handover ceremony	took place at	noi bai airport

Entity Summarization Process

We used Entity summarization as follows. For triple verbs, we lemmatized them and used *WordNet* to find their root verb synsets. A verb word might have multiple meanings, and each synset expresses a meaning. For example, the meaning of the word "die" is different when applied to a human, a computer, or a star in the Milky Way. To disambiguate the meaning of a predicate, and find the closest *synsets*, we compared two methods and used a human evaluator to determine the better one. The first method is the Lesk algorithm (Lesk 1986). Given a verb and the triple where it occurs, Lesk returns a synset representing the meaning in context. However, Lesk often failed at finding the correct synset. Among a set of 80 randomly picked triples, only 37 verb synsets were correctly identified.

The second method was that we selected the "v.01" (primary meaning) as the synset for each verb. We used the same 80 triples for this experiment. Based on human evaluation, we found 75 verb synsets were correct. Even though this is a simple approach, it produced a better result. Therefore, we used the second method to find the root synset. For each verb (vx) synset *"vx.v.01,"* we identified its hypernym chain until reaching the root. If a root verb occurs fewer times than a threshold, all its corresponding triples were deleted. This reduced the triple set from 16,270 to 16,163.

Entity-based (E) summarization is based on the entities (subject/object) occurring most often in triple sets. To identify the frequencies of meaningful words, we first removed numbers, stop words, and words with fewer than three characters. Next, we lemmatized each word. If a pre-processed subject or object became an empty string, we removed the whole triple. This step reduced the triple set size to 13,616. If a pre-processed subject consisted of more than two words, we removed all except for the last two words.

We labelled each subject with its frequency in the triple set. If there are two words in a subject, we obtained the subject score from the word with the higher frequency, e.g., if there are 77 occurrences of "covid" and 100 of "vaccine," then the subject score of "covid" is 77 and of "vaccine" is 100, and for "covid vaccine" it is also 100. We apply the same method to objects for object scores. We used $\alpha = 0.7$ to calculate the triple scores, which resulted in the best performance in the evaluation in the previous section. As described in the previous section, we also generated summaries by views (D) and (SF), besides by the view (E).

Results

From the 13,616 triples (after subject/object preprocessing), we selected the top-scoring $\mathbf{m_E}$ triples with their corresponding sentences, such that the total number of words in the sentences was approximately 300. The summary triples of this view are in Table 3. For views (D) and (SF), we generated summaries from 16,270 triples (after the cleaning process of extracted triples). The goal was again to retain the $\mathbf{m_D}$ and $\mathbf{m_{SF}}$ top-ranking triples to generate ~300-word readable summaries. Tables of summary triples of view (D) and view (SF) are omitted due to space limitations.

Different views yielded completely different triple sets. There is **no** common sentence among our views. The sentiments are represented by integers, 1 for *negative*, 2 for *neutral*, and 3 for *positive*. In view (D), more than half of the summary triples are *neutral*, while in views (E) and (SF), more than half are *negative*. The AstraZeneca (AZ) vaccine is not mentioned in any summary. In view (D), each vaccine except for AZ is evenly mentioned 2 to 3 times. Among 24 triples in view (SF), Covaxin and Sputnik V are mentioned most frequently, 14 and 9 times respectively. In view (E), Covaxin, Moderna and Sputnik V are mentioned most frequently, 5, 4, and 3 times among 16 triples. In general, Covaxin and Sputnik V are the most frequently mentioned vaccines in our summaries. The average length of summary sentences of view (D) is the longest, almost twice as long as for view (SF) (Table 4).

SF-based summaries are more likely to contain summary triples from the same sentences, as this view is based on the user's followers and retweets. In this view, the average length of summary sentences is the shortest among the three views. This might indicate that popular tweeters often generate short posts, tending to appeal to social media users.

Table 3: Entity-based Summary triples with vaccines.

tripleS	sentiment	tripleScore	vaccine
j&j vaccine is essentially first shot of sputnik v vaccine	1	1646	sputnik
new recombinant covid19 vaccine developed by national vaccine	3	1646	sinopharm
covid19 vaccine of batch is international vaccine campaign co-led by world health organization	1	1646	sinopharm
vaccine be listed on whos list of approved vaccines	2	1646	covaxin
covaxin shows vaccine efficacy of 81 % in phase 3 trial	1	1153.2	covaxin
covaxin will from june 21 will most expensive of three vaccines	1	1153.2	covaxin
moderna seeks regulatory approval for its covid vaccine in india	1	1129.4	moderna
moderna will fourth vaccine used for vaccination drive in india	1	1129.4	moderna
sputnikv joins indias vaccine utsav	2	1028.6	sputnik
india will have its third vaccine	2	1002.7	sputnik
china has approved emergency use of sinovac biotechs covid-19 vaccine in people aged	1	999.2	sinovac
other countries are also relying heavily on pfizer vaccine	2	994.3	pfizer
bharat biotech had submitted eoi for its vaccine andwho informed in document	2	978.2	covaxin
people develop reaction about week after moderna vaccine	2	974	moderna
bharatbiotech support worldwide amid covid19 surge coronavirus vaccine india bharatbiotech	1	962.8	covaxin
fda warning to moderna covid-19 vaccine shots	1	939.7	moderna pfizer

Table 4: Statistics about each view.

	Summary Triple #	Summary Sentence #	Word # in Sentences	Avg word # per sentence
SF	24	21	290	13.81
E	16	16	314	19.63
D	11	11	295	26.82

Discussion

Our framework can include additional views, e.g., fake news analysis (Li et al. 2022a), and other summarization methods (e.g., BertSum, etc.) to generate new results. A disadvantage of transformers is that when the input sequence is long, the performance of the attention mechanism is undermined (Rohde et al. 2021). When breaking long input text into chunks, and combining summarized chunks as output, this loses context. Some abstractive models that generate new sentences not in the original text may also suffer when using Rouge scores with a gold standard extractive summary.

Conclusions and Future Work

We presented a multiple view-based summarization framework that can generate different summary views that can be composed with other views. We introduced the Entity-based, Distance-based, and Social-Feature-based summarization methods for social media data, which were compared with published extractive models (SBert, BertSum) and abstractive models (bart-large-cnn, T5). We found that summaries generated by different views have few sentences in common. We used a COVID-19 vaccine tweet set with different sentiment views. The ratios of mentioned vaccine types and expressed sentiments, and the average lengths of summary sentences are different for different views. Furthermore, there are no common sentences in different vaccine tweet views. This supports the claim that our different summaries capture different perspectives. We are currently building a real-time summarization app of our framework.

Acknowledgment

Research reported in this publication was supported by the National Center for Advancing Translational Sciences (NCATS), a component of the National Institute of Health (NIH) under award number UL1TR003017. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

Angeli, G., Premkumar, M. J. J., and Manning, C. D. (2015). Leveraging Linguistic Structure for Open Domain Information Extraction. In Proc. of the 53rd ACL: 344–354.

Barrios, F., López, F., Argerich, L., & Wachenchauzer, R. (2016). Variations of the Similarity Function of TextRank for Automated Summarization.

Chi, Z., Dong, L., Wei, F., Wang, W., Mao, X. L., and Huang, H. (2019). Cross-Lingual Natural Language Generation via Pre-Training. arXiv.

Dutta, S., Chandra, V., Mehra, K., Ghatak, S., Das, A. & Ghosh, S. (2019). "Summarizing Microblogs during Emergency Events: A Comparison of Extractive Summarization Algorithms." IEMIS, 859-872.

Fellbaum, C. (2005). WordNet and wordnets. In: Brown, Keith et al. (eds.), Encyclopedia of Language and Linguistics, Second Edition, Oxford: Elsevier, 665-670.

Geng, F., Liu, Q., Zhang, P. (2020). A time-aware query-focused summarization of an evolving microblogging stream via sentence extraction, Digital Communications and Networks, Volume 6, Issue 3, 2020, 389-397, ISSN 2352-8648,

Greene, D., and Cunningham, P. (2006). Practical Solutions to the Problem of Diagonal Dominance in Kernel Document Clustering. In Proc. of the 23rd ICML, ACM., 377–384.

Gu, N., Ash, E., and Hahnloser, R. H. R. (2021). MemSum: Extractive Summarization of Long Documents Using Multi-Step Episodic Markov Decision Processes. arXiv.

Gunaratna, K., Thirunarayan, K., and Sheth, A. (2015). FACES: Diversity-Aware Entity Summarization Using Incremental Hierarchical Conceptual Clustering. Proc. of the AAAI Conference on Artificial Intelligence. 29, 1.

Gunaratna, K., Thirunarayan, K., Sheth, A., Cheng, G. (2016). Gleaning Types for Literals in RDF Triples with Application to Entity Summarization. In: Sack, H., Blomqvist, E., d'Aquin, M., Ghidini, C., Ponzetto, S., Lange, C. (eds) The Semantic Web. Latest Advances and New Domains. ESWC 2016. vol 9678. Springer, Cham.

Heo, H. (2021). FactSumm: Factual Consistency Scorer for Abstractive Summarization. Retrieved from https://github.com/Huffon/factsumm

Jia, R., Cao, Y., Tang, H., Fang, F., Cao, C and Wang, S. (2020). Neural Extractive Summarization with Hierarchical Attentive Heterogeneous Graph Network. In Proc. of the 2020 Conference on EMNLP. ACL. Online, 3622–3631.

Joshi, A., Fidalgo, E., Alegre, E., & Fernández-Robles, L. (2023). DeepSumm: Exploiting topic models and sequence to sequence networks for extractive text summarization. Expert Systems with Applications, 211, 118442.

Kwon, J., Kobayashi, N., Kamigaito, H., and Okumura, M. (2021). Considering Nested Tree Structure in Sentence Extractive Summarization with Pre-trained Transformer. In Proc. of the 2021 Conference on EMNLP, ACL. Online, 4039–4044.

Lesk, M. (1986). Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. In Proc. of the 5th annual international conference on Systems documentation (SIGDOC '86). ACM, 24–26.

Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V. and Zettlemoyer, L. (2019). BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension.

Li, C. Y, Chun, S., & Geller, J. (2022) a. Stemming the Tide of Fake News about the COVID-19 Pandemic. The International FLAIRS Conference Proc., 35.

Li, C. Y., Renda, M., Yusuf, F., Geller, J., Chun, S.A. (2022) b. Public Health Policy Monitoring through Public Perceptions: A Case of COVID-19 Tweet Analysis. Information 2022, *13*, 543.

Lin, C. Y. (2004). ROUGE: A Package for Automatic Evaluation of Summaries. In Text Summarization Branches Out, ACL. 74–81.

Liu, X. & Li, Y. & Wei, F. & Zhou, M. (2012). Graph-Based Multi-Tweet Summarization using Social Signals. Proc. of COLING 2012: Technical Papers. 1699-1714.

Modhe, S. V., Panse, N. S., Rathi, M. U., Chavan, S. (2021). Extractive Based Approach For Microblogs Summarization. International Research Journal of Modernization in Engineering Technology and Science Volume:03/Issue:03/March-2021.

Niklaus, C., Cetto, M., Freitas, A., & Handschuh, S. (2018). A Survey on Open Information Extraction. arXiv.

Ott, M, Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. (2019). fairseq: A Fast, Extensible Toolkit for Sequence Modeling. In Proc. of the 2019 Conference of the North American Chapter of the ACL. (Demonstrations), 48–53.

Petrov, C. (2023). 50+ Stunning Twitter Statistics You Need to Know in 2023. https://techjury.net/blog/twitter-statistics/

Preda, G. (2021). COVID-19 All Vaccines Tweets. Kaggle.

Qi, P., Zhang, Y., Zhang, Y., Bolton, J., and Manning, C. D. (2020). Stanza: A Python Natural Language Processing Toolkit for Many Human Languages. In Proc. of the 58th Annual Meeting of the ACL: System Demonstrations, ACL., 101–108.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I. Language Models are Unsupervised Multitask Learners. 2019. https://openai.com/blog/better-language-models/

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2019). Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. arXiv.

Rohde, T., Wu, X., and Liu, Y. (2021). Hierarchical Learning for Generation with Long Source Sequences. arXiv.

Roush, A., and Balaji, A. (2020). DebateSum: A large-scale argument mining and summarization dataset. arXiv.

Ruby, D. (2022). Twitter Statistics: New Trends, Figures & Data. https://www.demandsage.com/Twitter-statistics/

Rudra, K., Ghosh, S., Ganguly, N., Goyal, P., and Ghosh, S. (2015). “Extracting situational information from microblogs during disaster events: a classification-summarization approach,” in Proc. of the 24th ACM International on CIKM, pp. 583–592.

Saini, N., Saha, S., and Bhattacharyya, P. (2019). “Multiobjective-based approach for microblog summarization,” IEEE Transactions on Computational Social Systems, vol. 6, no. 6, pp. 1219–1231.

Sintek, M., and Decker, S. (2002). TRIPLE—A Query, Inference, and Transformation Language for the Semantic Web.

Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., and Potts, C. (2013). Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In Proc. of the 2013 Conference on EMNLP, ACL, 1631–1642.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention Is All You Need. arXiv.