

Examining Bias in Opinion Summarisation Through the Perspective of Opinion Diversity

Nannan Huang Lin Tian Haytham Fayek Xiuzhen Zhang

RMIT University, Australia

{amber.huang, lin.tian2}@student.rmit.edu.au

haytham.fayek@ieee.org

xiuzhen.zhang@rmit.edu.au

Abstract

Opinion summarisation is a task that aims to condense the information presented in the source documents while retaining the core message and opinions. A summary that only represents the majority opinions will leave the minority opinions unrepresented in the summary. In this paper, we use the stance towards a certain target as an opinion. We study bias in opinion summarisation from the perspective of opinion diversity, which measures whether the model generated summary can cover a diverse set of opinions. In addition, we examine opinion similarity, a measure of how closely related two opinions are in terms of their stance on a given topic, and its relationship with opinion diversity. Through the lens of stances towards a topic, we examine opinion diversity and similarity using three debatable topics under COVID-19. Experimental results on these topics revealed that a higher degree of similarity of opinions did not indicate good diversity or fairly cover the various opinions originally presented in the source documents. We found that BART (Lewis et al., 2020) and ChatGPT can better capture diverse opinions presented in the source documents.

1 Introduction

The aim of opinion summarisation is to reduce the amount of information in the source text while maintaining the core message and opinions expressed therein. It can be in the form of a summary of a product review (Alam et al., 2016; Bražinskas et al., 2020; Chu and Liu, 2019), online discourse using platforms like Twitter, Reddit (Fabbri et al., 2021; Bilal et al., 2022) or other types of text with opinions. There are two major types of models, extractive (Meng et al., 2012; Ku et al., 2006; Erkan and Radev, 2004; Liu, 2019) and abstractive (Lewis et al., 2020; Raffel et al., 2020; Zhang et al., 2020). The extractive models extract the key information by selecting the most salient sentences in the source documents. Whereas the abstractive summarisation models generate new phrases that reflect the

key information in the source documents. Applications for this activity include tracking consumer opinions, assessing political speeches, or internet conversations, among many others.

A summarisation model’s output will reflect any biases present in the data used for training the model. Moreover, summarisation models are used in many applications where the fairness of the outputs is critical. For instance, opinion summarisation models can be used to summarise product reviews and social media posts. If these models produce biased summaries, they have the risk of being misused as a tool to influence and manipulate people’s opinions. Therefore, it is essential to look into the fairness of the models.

Earlier studies of bias in opinion summarisation have mainly evaluated biases in summarisation by comparing whether the information is selected from different social groups equally, using attributes including gender, dialect, and other societal properties (Dash et al., 2019; Blodgett et al., 2016; Keswani and Celis, 2021; Olabisi et al., 2022). Such representation is fair from the perspective of information input, leaving the other side of information consumption under studied. It is equally important to look at fairness from end-users’ perspective. Shandilya et al. (2020) pointed out that fair representation from the information input perspective does not always imply fair representation from end-users’ perspective. From the standpoint of end-users, it is more important for summaries to cover the various opinions presented in the source documents so that the various opinions in the source documents can be heard equally (Shandilya et al., 2020). In this study, we examine bias in opinion summarisation from end-users’ perspective by comparing the output of different modern summarisation models.

Blodgett et al. (2020) noted that many research work on bias in natural language processing lacks a precise description and definition of *bias* in their

study, making it challenging for readers to understand. Our working definitions of the key terms are as follows. According to the *stance triangle theory* (Du Bois, 2007), a stance is composed of three components: the object, the subject, and the attitude, which shape sociocultural value systems and opinion expression. While studies such as Misra et al. (2016) and Reimers et al. (2019) utilise arguments with similarity and clustering techniques to capture similar opinions, our study takes a different approach based on the stance triangle theory. In our study, an *opinion* is a personal stance about a particular target, and a stance is a position or attitude that an individual takes about a given target. Our definition of *bias in opinion summarisation* is when the summary focuses more on certain opinions than on the diversity of opinions presented in the source documents. Note that it is generally agreed that diverse opinions should be taken into account regardless of their quantitative value in order to include more diverse opinions when using sentiment information (Angelidis and Lapata, 2018; Siledar et al., 2023). Hence, our focus is on measuring the diversity of opinions rather than quantity. This is measured through *opinion diversity* which assesses opinion equivalence relative to the source documents. It measures the opinions in the source documents that the generated summary contains. In addition, we compare opinion similarity between the source and generated documents and further examine the relationship between opinion diversity and opinion similarity. *Opinion similarity* is a measure of how closely related two opinions are in terms of their stance on a given topic. We use BERT (Devlin et al., 2019) to compare the semantic closeness in the embedding space. We aim to understand whether models perform well in capturing overall opinions that are less biased by covering diverse opinions. We examine opinions on three COVID-19 topics using stances on these topics.

In our study, we aim to answer the following questions:

1. How well can summarisation models present various opinions in a document from the perspective of stance towards a topic?
2. Does a greater degree of opinion similarity with the source documents suggest a lack of bias in summarisation models?

2 Related Work

2.1 Opinion summarisation

Opinion summarisation has received significant attention in recent years, with extractive models such as Hybrid-TFIDF (Inouye and Kalita, 2011) a frequency-based summarisation method, having great performance for summarising social media data like tweets. Recent studies have introduced the concept of key point analysis (Bar-Haim et al., 2020a,b, 2021), which uses extractive models to identify key arguments from the source documents and match them to the main opinions' associated key points. Abstractive opinion summarisation models, such as Copycat (Bražinskas et al., 2020) and MeanSum (Chu and Liu, 2019), are designed to address the problem of summarising product or business reviews. MeanSum (Chu and Liu, 2019) is based on LSTM, while Copycat (Bražinskas et al., 2020) uses a variational autoencoder (VAE) to generate latent vectors of given reviews.

2.2 Biases in opinion summarisation

Existing studies of bias in opinion summarisation have focused on the perspective of using sensitive attributes of social media users and categorising them under different social groups (Dash et al., 2019). These attributes include social identities like gender, race and political leaning information. Other studies focused on the perspective of dialect used in text and whether the generated summaries cover such dialects (Blodgett et al., 2016; Keswani and Celis, 2021; Olabisi et al., 2022).

One limitation of these approaches is that they require obtaining sensitive hidden attributes from the text producers. Due to other factors including privacy, this is not always practicable. Another limitation, as pointed out in Shandilya et al. (2020) is the assumption of fairness in using information from different social groups equally and this is from the input information perspective. Fairness should also be evaluated from the perspective of end-users. Shandilya et al. (2020) discovered in their study that, from the perspective of end-users, it is more important to present diverse information in the source documents. We argue that considering fairness from the perspective of end-users is equally important since they are the users of the final product and make informed decisions based on the generated summaries.

Therefore, in our work, we focus on the fair representation of information generated by summarisation models.




Figure 1: The process of calculating opinion diversity: COVID-Stance-BERT is applied to each source tweet and aggregated into a single set of opinions. The same process is applied to the generated summaries at the sentence level. The F1 score is applied among the two sets of opinions, where the true positive (TP) measures opinions captured by the model generated summary that is also presented in the source documents; false positive (FP) measures opinions captured by the model generated summary that is not presented in the source documents; false negative (FN) measures opinions not captured by the model generated summary that is presented in the source documents.

sation models from end-users’ perspective. We study bias in opinion summarisation by examining whether the summary focuses more on certain opinions than on the diverse opinions presented in the source documents.

3 Methodology

As previously stated, we capture opinion in our study by utilising stance and its target. When the generated summaries reflect opinions that diverge from those of the source texts, the summary is considered biased. This could take the form of concentrating on a narrow range of opinions or going beyond what was expressed in the source documents. It should be noted that the biased information in the source text is not the main focus of our study; rather, we are interested in how opinions are presented and whether the opinions in the generated summaries match those in the source documents.

We formulate our problems in three steps. We first use a pre-trained language model to capture opinions from the input sequences. The opinion diversity is then calculated using the F1 score between the set of stances in the source tweets and the generated summary for each cluster (discussed below) under each topic. Finally, we compare the cosine similarity between the source tweet cluster and the summary at the sentence level to measure opinion similarity using model representation.

Let $\mathcal{C} = \{c_0, c_1, c_2, \dots, c_l\}$ be a set of clusters. These clusters are derived from the three main topics (“CDC”, “Stay at Home Orders” and “Wearing a Face Mask”). For each c , we have a set of source tweets \mathcal{T} , defined as $\mathcal{T} = \{t_0, t_1, t_2, \dots, t_n\}$. For each c , we have a set of generated summaries $\mathcal{S} = \{s_0, s_1, s_2, \dots, s_l\}$, where each summary s consists of a lists of sentences, defined as $s = \{e_0, e_1, e_2, \dots, e_q\}$, where e refers to the textual content of each input sentence.

3.1 Capturing Opinions

We train COVID stance classification models related to several COVID-19 topics or targets using a publicly available dataset from Glandt et al. (2021) on COVID-19 related stance detection. The dataset consists of four different COVID-19 related topics and targets. See Appendix A.1 for further detail on the summary of the data distribution. Similar to Glandt et al. (2021), using the further pre-trained BERT model (bert-large-uncased)¹ with the COVID-19 tweet corpus (Müller et al., 2020), we fine-tune the model with stance labeled data from Glandt et al. (2021). Thus, our COVID-Stance-BERT is a further pre-trained BERT fine-tuned with standard cross entropy loss to do a three-class classification of stance labels (support,

¹<https://huggingface.co/bert-large-uncased>

against, and neutral). Each tweet t_i is associated with a ground-truth label $d \in \mathcal{D}$, where \mathcal{D} represents the label set (support, against or neutral; 3 classes).

$$v_i = \text{BERT}([\text{CLS}] \oplus x_i) \quad (1)$$

$$\hat{d}_i = \text{softmax}(Wv_i + b) \quad (2)$$

We call these models COVID-Stance-BERT in the remainder of the paper. The average accuracy and macro F1 scores across three targets are 0.8208 and 0.8026 respectively. Similar levels of accuracy and F1 scores were obtained across these topics compared to Glandt et al. (2021). The detailed result is reported in Appendix A.2.

3.2 Opinion Diversity

The overall process is presented in Fig 1. In order to determine whether different opinions stated in the source documents can be captured using summarisation models, we apply COVID-Stance-BERT to the source documents to produce a collection of opinions represented in the input tweets and the generated summaries. The majority of the source tweets have only a single sentence. We therefore treat each sentence in the generated summaries as a tweet and apply the same COVID-Stance-BERT to retrieve its opinions. We apply the in-domain stance detection model on both the source documents and the summarised sentences across the tweet clusters under different topics. That is, a stance detection model fine-tuned to the target of "Stay at Home Orders" is applied to the collection of tweets and the generated summaries towards the "Stay at Home Orders" topic.

Once the prediction is done on both the source documents and the generated summaries, we compare the sets of opinions and examine if the generated summaries cover the various opinions presented in the source document. We compute the F1 score by comparing the two sets of stances under each discussed topic across all clusters and use them to represent how well the summarisation model captures the stances in relation to the input documents.

The detailed calculation can be found as follows: we apply COVID-Stance-BERT stance prediction to calculate the opinion diversity on both the set of source tweets and all generated summaries. For each generated summary $s_j \in \mathcal{S}$,

we adopt sentence splitting function² as $s_j = \{ss_0, ss_1, \dots, ss_m\}$. For each sentence e in generated summary s_j , we take the associated stance label d_e , formally,

$$\begin{aligned} d_t &= \text{BERT}(\text{emb}([\text{CLS}], t_p)), \\ d_{ss} &= \text{BERT}(\text{emb}([\text{CLS}], ss_q)), \end{aligned} \quad (3)$$

where t_p represents the text of the source tweet and s_q is the summary, $\text{emb}()$ the embedding function and d_t and d_{ss} are the stance predictions based on the $[\text{CLS}]$ token produced by our COVID-Stance-BERT.

Once the above is completed, we get the set of non-repeated opinion(s) in both the source documents and the generated summaries. We use the F1 score, which measures the harmonic mean of opinion precision and opinion recall, to evaluate the performance of opinion diversity, where the opinion precision measures the proportion of important opinions in the generated summary. The opinion recall measures the degree of salient opinion in the source documents that the generated summary contains. A higher F1 score indicates the model generated summary can better cover the various opinions presented in the source documents. The true positive (TP), false positive (FP), and false negative (FN) are measured as follows:

- TP = opinions captured by the model generated summary that is also presented in the source documents.
- FP = opinions captured by the model generated summary that is not presented in the source documents.
- FN = opinions not captured by the model generated summary that is presented in the source documents.

More detail on illustrations of various scenarios of opinion precision and recall and their associated F1 scores can be found in Appendix A.3. We report the average across all clusters under each topic as the overall opinion diversity for each model.

3.3 Opinion Similarity

The overall process of evaluating opinion similarity is visualised in Fig 2. To aggregate all the source tweets in a cluster to get a source representation

²<https://www.nltk.org/api/nltk.tokenize.html>

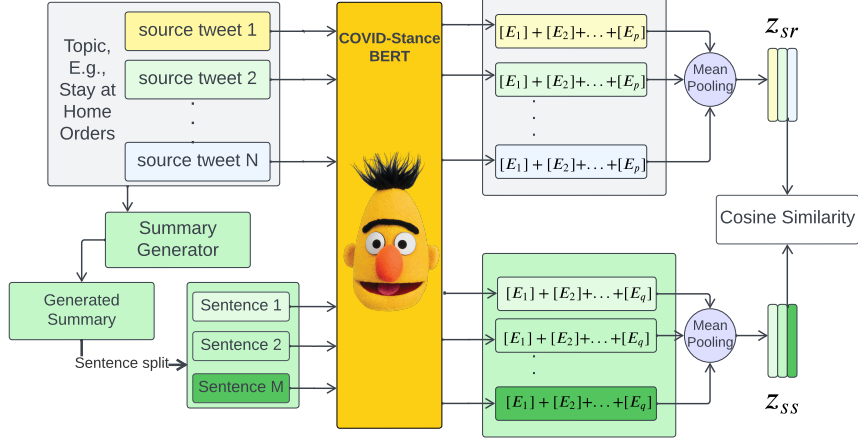


Figure 2: The process of calculating opinion similarity. For the source tweets, a single representation is obtained using mean-pooling approach based on the mean representations of all input tokens. For the generated summaries, the same process is applied at the sentence level. The cosine similarity is then applied to these representations to measure opinion similarity.

(z_{sr}), we adopt the mean-pooling approach based on the mean representations of all input tokens:

$$z_{sr} = \frac{1}{|n|} \sum_{i \in n} h_i^c, \quad (4)$$

where h_i^c is the mean stance representation derived from source tweet t_i .

To extract the representation for each generated summary (z_{ss}), we applied the mean-pooling approach over the average of all token representations from the last layer of our language model encoder as shown in Eq. 5:

$$z_{ss} = \frac{1}{|m|} \sum_{i \in m} v_i^c, \quad (5)$$

where v_i^c is the mean stance representation based on each token representation from a given summary s_i .

For each generated summary, we compute the cosine similarity to the average source representation in the cluster, with the following equation:

$$\cos(z_{sr}, z_{ss}) = \frac{z_{sr} z_{ss}}{\|z_{sr}\| \|z_{ss}\|}. \quad (6)$$

A higher cosine similarity score between the source and the summary indicates a model is better at capturing the overall opinions and semantic information presented in the source documents.

In summary, we investigate bias in opinion summarisation models from the standpoint of opinion

diversity. The idea is to look at whether summarisation models can capture the various opinions presented in the source documents. In addition, from the perspective of opinion similarity, we evaluate how closely related the generated summary and the source documents are in terms of their stance and semantic information on a given topic. Combining the results of opinion diversity and opinion similarity, we aim to understand whether summaries that express opinions that are overall similar to the source documents also indicate they are less biased.

4 Experiments and Results

4.1 Data

In this study, we use the COVID-19 tweet clusters dataset provided by Bilal et al. (2022). The dataset contained tweet clusters that are coherently opinionated, coherently non-opinionated, and incoherent subsets. Each topic contains a different number of clusters. A cluster of tweets contains a collection of tweets obtained during a particular timestamp on a specific topic. In Bilal et al. (2022) each cluster was used for generating the final summary. Hence, in our work, we focus on the cluster level when generating summaries for each topic. For example, under the topic "Wearing a Face Mask" there can be multiple clusters obtained at different times; each cluster of tweets could be discussing whether wearing a face mask is a good idea, obtained at a different time of the year.

In our experiment, we are only considering tweet

Table 1: The statistic of the number of tweet clusters and the average number of tweets in each cluster under different topics.

Topic	No. Clusters	Ave No. Tweets
CDC	78	21.77
Stay at Home Orders	48	20.54
Wearing a Face Mask	52	22.42

clusters that are coherently opinionated, with a similar discussion of targets as the ones mentioned in [Glandt et al. \(2021\)](#). This is to ensure that all tweets contain opinions and, at the same time, to utilise the stance detection model in an in-domain setting to evaluate the opinions expressed in these clusters. We obtained coherently opinionated clusters including "Stay at Home Orders", "Wearing a Face Mask" and a highly related topic, "CDC/Centres for Disease Control and Prevention", where the clusters of tweets centred around the aforementioned topics are mainly focusing on the discussion and the expression of opinions towards them during the COVID-19 pandemic. The in-domain stance detection models were then applied to the clusters apart from "CDC/Centres for Disease Control and Prevention" where we applied the stance detection model that was trained on the "Fauci" topic since they are both public figures.

The overall data distribution after our selection can be found in Table 1. There are 78, 48, and 52 clusters of tweets; and on average, 21.77, 20.54, and 22.42 tweets under each cluster for the "CDC", "Stay at Home Orders", and "Wearing a Face Mask" topics, respectively.

4.2 Baseline Models

Several summarisation models are used to generate summaries in our experiments, including extractive summarisation models TextRank ([Mihalcea and Tarau, 2004](#)), LexRank ([Erkan and Radev, 2004](#)) and Hybrid-TFIDF ([Inouye and Kalita, 2011](#)); and abstractive summarisation models BART ([Lewis et al., 2020](#)), Pegasus ([Zhang et al., 2020](#)) and T5 ([Raffel et al., 2020](#)), a summarisation model for review Copycat ([Bražinskis et al., 2020](#)) and a recently released Large Language Model (LLM) - ChatGPT³. Following [Bilal et al. \(2022\)](#) we limit the abstract summarisation models word limit to the generated summary within [90%, 110%] of

³<https://chat.openai.com/>

the gold standard length; and the average token length of the gold standard length for ChatGPT. For extractive models only allows selecting sentences we limit to the average number of sentences of the gold standard length. We apply all the models mentioned in a zero-shot setting. More in-depth discussion on each of the models is below:

- **BART** ([Lewis et al., 2020](#)) is an encoder-decoder model with a bidirectional encoder and a left-to-right decoder. Pretrained using a novel in-filling technique by replacing a span of text with a single mask token. Making it useful for language generation tasks We use the BART large model, pre-trained on CNN/Daily Mail⁴.
- **Pegasus** ([Zhang et al., 2020](#)) is a model that employs the Transformer encoder and decoder, self-supervised learning and pre-training on predicting the removed sentences, and tokens similar to the masked language model. We use the Pegasus model pre-trained on CNN/Daily Mail⁵.
- **T5** ([Raffel et al., 2020](#)) is an encoder-decoder model pretrained on a multi-task setting using both supervised and unsupervised settings where the tasks are converted into a set of input-output text pairs. This allows it to understand a large variety of relationships between texts. We use the T5 base model pre-trained on CNN/Daily Mail⁶.
- **ChatGPT** OpenAI's ChatGPT, a recently released Large Language Model (LLM), was developed by employing reinforcement learning from human feedback (RLHF) ([Christiano et al., 2017](#)) to train a GPT-3.5 series model. We employ OpenAI's ChatGPT API (gpt-3.5-turbo-0301) for our experiments. We adjusted the maximum tokens to the average gold standard token length. We prompt ChatGPT as below "Summarise the following tweets: You should write it in tweet style. You should use no more than 4 sentences. Tweets: [Source tweets]".

⁴<https://huggingface.co/facebook/bart-large-cnn>

⁵https://huggingface.co/google/pegasus-cnn_dailymail

⁶<https://huggingface.co/flax-community/t5-base-cnn-dm>

Table 2: Results of opinion diversity (Opi Div) and opinion similarity (Opi Sim) for various models under different discussed topics. The best results are bolded, and the ranking of the models is provided inside the brackets.

Events	CDC		Stay at Home Orders		Wearing a Face Mask	
Models	Opi Div	Opi Sim	Opi Div	Opi Sim	Opi Div	Opi Sim
BART	0.7449 (1)	0.8503 (4)	0.7681 (2)	0.8373 (7)	0.8147 (1)	0.8412 (6)
Pegasus	0.5265 (7)	0.8745 (3)	0.7576 (3)	0.8775 (3)	0.3692 (5)	0.8768 (3)
T5	0.6346 (3)	0.8451 (5)	0.7417 (5)	0.8407 (6)	0.4692 (4)	0.8327 (7)
ChatGPT	0.7282 (2)	0.8818 (2)	0.8014 (1)	0.8515 (5)	0.6006 (3)	0.8498 (5)
Copycat	0.5265 (7)	0.6725 (8)	0.7014 (8)	0.7288 (8)	0.6737 (2)	0.7177 (8)
TextRank	0.5338 (6)	0.8370 (6)	0.7417 (5)	0.8519 (4)	0.2615 (8)	0.8828 (2)
LexRank	0.5530 (5)	0.8208 (7)	0.7569 (4)	0.8817 (2)	0.3590 (7)	0.8607 (4)
Hybrid TFIDF	0.5697 (4)	0.8914 (1)	0.7063 (7)	0.8923 (1)	0.3667 (6)	0.8965 (1)

- **Copycat** (Bražinskas et al., 2020) is a Variational Autoencoder model overcomes issues with limited training data in the context of review summarisation using self-supervision and latent representation that represents the general opinions expressed in the source reviews. We use the model provided by the authors ⁷.
- **TextRank** (Mihalcea and Tarau, 2004) is a graph-based model that extracts the most important sentence from the input document based on the weight determined by the edges connected to the words or phrases.
- **LexRank** (Erkan and Radev, 2004) is a graph-based model represents each sentence in the document as a node. Edges between vertices are calculated using cosine similarity, and the importance of a sentence is determined by the number of connected edges. The model extracts the most important sentence from the document based on the connectivity matrix.
- **Hybrid-TFIDF** (Inouye and Kalita, 2011) is a graph-based model similar to TextRank (Mihalcea and Tarau, 2004) and LexRank (Erkan and Radev, 2004). Each word in the sentence is represented using the TF-IDF score, and similarity between sentences is computed to build edges between sentence vertices. Similar to other graph-based models, sentences with the most connected edges are deemed most important.

⁷<https://github.com/abrazinskas/Copycat-abstractive-opinion-summarizer>.

4.3 Opinion Diversity by Models

To answer the first question, we first gather opinions from the source input tweets by running COVID-Stance-BERT on each tweet in each cluster under different topics. Similarly, the same model is used to collect opinions captured by different summarisation models as discussed in Section 3.2. Once the sets of opinions from the source document and the generated summaries are obtained, we compare whether the opinions presented in the source document are also the opinions captured by the models. This is to answer the question of whether models can fairly preserve the various opinions being presented in the source documents.

We compare the three different discussed topics listed in Table 1 using the clusters within each topic. Summaries were first obtained using the models mentioned in Section 4.2 for each cluster. Then, using the opinion sets obtained for both the source documents and the generated summaries, we compare whether the summaries have high opinion diversity by comparing the set of stance(s) captured by the summarisation models with the set of stance(s) presented in the source documents. See Appendix A.4 for details of the stance detection example and an example calculation of opinion diversity.

The results of the opinion diversity of different models under different discussion topics are listed in Table 2. We are reporting the F1 score for opinion diversity. The results of opinion precision and recall can be found in Appendix A.5. BART (Lewis et al., 2020) has the best opinion diversity for both the "CDC" and "Wearing a Face Mask" topics, while it has a competitive result for the "Stay at Home Orders". ChatGPT has the best opinion diversity for "Stay at Home Orders" and is




Figure 3: Stance distribution across various discussion topics for source documents and model generated summaries.

relatively competitive for the other topics.

Overall speaking, BART has the best performance in capturing diverse opinions and followed by ChatGPT. It is worth noticing, however, that previous studies found that ChatGPT tends to generate lengthy summaries when using the default parameters (Wang et al., 2023; Yang et al., 2023) and it is important to provide guidance in the prompt. The fact that we limit its maximum tokens and ask it to generate a fixed number of sentences may hurt its performance. For example, in the generated summary presented in Appendix A.5, ChatGPT generated an unfinished sentence.

4.4 Opinion Similarity by Models

The second question we aimed to answer is whether a model that generates summaries with high overall opinion similarity to the source documents would also be less biased by covering various opinions. To answer this question, we used the average input token representation obtained using the COVID-Stance-BERT model, for both the source documents and also the generated summaries.

For the source documents, the average input token representation for each tweet was obtained using COVID-Stance-BERT, followed by applying mean-pooling to aggregate them into a single representation. For the summary, a similar approach was applied to the generated summary at the sentence level. We obtained average input token representations for each sentence in the generated summary, and then a mean pooling was applied over the sentence representations to obtain a single representation for the generated summary. Once the representations are obtained for both the source documents and the model generated summaries. We compare the stance similarity between source documents

against different model generated summaries using the cosine similarity between the source representation and the summary representation.

Table 2 shows the cosine similarity for measuring how similar the overall stance presented in the generated summaries is to the source documents. From the result, we can see that the model with the overall highest similarity of opinions across different topics is Hybrid-TFIDF (Inouye and Kalita, 2011). We suspect this can be due to the superior performance of Hybrid-TFIDF (Inouye and Kalita, 2011) in summarising microblogging text. This is most likely because Twitter posts are not like typical documents and are unstructured, disconnected, and brief (Inouye and Kalita, 2011). As a result, when compared to the source documents, it maintains a similar overall stance.

4.5 Opinion Distribution

We computed the stance distribution across the discussion topics using the proportion of stances in the source documents and different summarisation models. The outcome is depicted in Fig. 3. From the distribution result, we can see that BART (Lewis et al., 2020) can better preserve the stance distribution for both the "CDC" and "Wearing a Face Mask" topics, whereas ChatGPT can better preserve the stance distribution for the "Stay at Home Orders" topic. While most models could pick up various stances and present a similar stance distribution compared to the source documents for the "Stay at Home Orders" topic. We suspect this is due to the fact that this topic has no obvious minority opinion, making it less challenging for models to cover diverse opinions. The above aligns with what we found in Section 4.3 where BART (Lewis et al., 2020) and ChatGPT outperformed the other

models by covering more diverse opinions. We believe one of the possible reasons could be due to BART’s impressive multi-document summarisation capabilities (Chen and Yang, 2020; Johner et al., 2021) and ChatGPT’s good performance in multiple downstream NLP tasks.

In conclusion, based on the results in Section 4.3 and 4.4 we found that when the generated summary has a higher degree of similarity in terms of overall opinion, that does not indicate it is fair in terms of covering more diverse opinions. In combination with the result from Section 4.5 we observed that the model’s ability to capture various opinions on different topics is case-dependent. We found that when no obvious minority stance is presented in the source documents, most models could capture various opinions in the source document.

5 Conclusion

In this work, we proposed a new way to examine bias in opinion summarisation from the perspective of presenting various opinions in the summary. We investigated various summarisation models for the COVID-19 event under three topics, using stance towards a target as a representation of opinions. In addition, we also examined overall stance similarity using model representation. We found that BART (Lewis et al., 2020) and ChatGPT are better at capturing diverse opinions when generating a summary. Whereas Hybrid-TFIDF has the highest similarity across the three discussed topics for overall opinion similarity. Based on the result, we found that higher opinion similarity does not indicate that the model presents diverse opinions. While both attributes are important for evaluating bias in opinion summarisation we suggest future studies look into introducing metrics that can evaluate summaries from both perspectives.

References

Md Hijbul Alam, Woo-Jong Ryu, and SangKeun Lee. 2016. Joint multi-grain topic sentiment: modeling semantic aspects for online reviews. *Information Sciences*, 339:206–223.

Stefanos Angelidis and Mirella Lapata. 2018. Summarizing opinions: Aspect extraction meets sentiment prediction and they are both weakly supervised. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3675–3686.

Roy Bar-Haim, Lilach Eden, Roni Friedman, Yoav Kantor, Dan Lahav, and Noam Slonim. 2020a. From arguments to key points: Towards automatic argument summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4029–4039.

Roy Bar-Haim, Lilach Eden, Yoav Kantor, Roni Friedman, and Noam Slonim. 2021. Every bite is an experience: Key point analysis of business reviews. 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 3376–3386.

Roy Bar-Haim, Yoav Kantor, Lilach Eden, Roni Friedman, Dan Lahav, and Noam Slonim. 2020b. Quantitative argument summarization and beyond: Cross-domain key point analysis. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 39–49.

Iman Munire Bilal, Bo Wang, Adam Tsakalidis, Dong Nguyen, Rob Procter, and Maria Liakata. 2022. [Template-based abstractive microblog opinion summarization](#). *Transactions of the Association for Computational Linguistics*, 10:1229–1248.

Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.

Su Lin Blodgett, Lisa Green, and Brendan O’Connor. 2016. [Demographic dialectal variation in social media: A case study of African-American English](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1119–1130, Austin, Texas. Association for Computational Linguistics.

Arthur Bražinskas, Mirella Lapata, and Ivan Titov. 2020. Unsupervised opinion summarization as copycat-review generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5151–5169.

Jiaao Chen and Diyi Yang. 2020. Multi-view sequence-to-sequence models with conversational structure for abstractive dialogue summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4106–4118.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.

Eric Chu and Peter Liu. 2019. Meansum: a neural model for unsupervised multi-document abstractive

- summarization. In *International Conference on Machine Learning*, pages 1223–1232. PMLR.
- Abhisek Dash, Anurag Shandilya, Arindam Biswas, Kripabandhu Ghosh, Saptarshi Ghosh, and Abhijnan Chakraborty. 2019. Summarizing user-generated textual content: Motivation and methods for fairness in algorithmic summaries. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–28.
- 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)*, pages 4171–4186.
- John W Du Bois. 2007. The stance triangle. *Stancetaking in discourse: Subjectivity, evaluation, interaction*, 164(3):139–182.
- Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research*, 22:457–479.
- Alexander Fabbri, Faiyaz Rahman, Imad Rizvi, Borui Wang, Haoran Li, Yashar Mehdad, and Dragomir Radev. 2021. **ConvoSumm: Conversation summarization benchmark and improved abstractive summarization with argument mining**. 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 6866–6880, Online. Association for Computational Linguistics.
- Kyle Glandt, Sarthak Khanal, Yingjie Li, Doina Caragea, and Cornelia Caragea. 2021. Stance detection in covid-19 tweets. 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 1596–1611.
- David Inouye and Jugal K Kalita. 2011. Comparing twitter summarization algorithms for multiple post summaries. In *2011 IEEE Third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*, pages 298–306. IEEE.
- Timo Johner, Abhik Jana, and Chris Biemann. 2021. Error analysis of using bart for multi-document summarization: A study for english and german language. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 391–397.
- Vijay Keswani and L Elisa Celis. 2021. Dialect diversity in text summarization on twitter. In *Proceedings of the Web Conference 2021*, pages 3802–3814.
- Lun-Wei Ku, Yu-Ting Liang, Hsin-Hsi Chen, et al. 2006. Opinion extraction, summarization and tracking in news and blog corpora. In *AAAI spring symposium: Computational approaches to analyzing weblogs*, volume 100107, pages 1–167.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Yang Liu. 2019. Fine-tune bert for extractive summarization. *arXiv preprint arXiv:1903.10318*.
- Xinfan Meng, Furu Wei, Xiaohua Liu, Ming Zhou, Sujian Li, and Houfeng Wang. 2012. Entity-centric topic-oriented opinion summarization in twitter. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 379–387.
- 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.
- Amita Misra, Brian Ecker, and Marilyn Walker. 2016. **Measuring the similarity of sentential arguments in dialogue**. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 276–287, Los Angeles. Association for Computational Linguistics.
- Martin Müller, Marcel Salathé, and Per E Kummervold. 2020. Covid-twitter-bert: A natural language processing model to analyse covid-19 content on twitter. *arXiv preprint arXiv:2005.07503*.
- Olubusayo Olabisi, Aaron Hudson, Antonie Jetter, and Ameeta Agrawal. 2022. **Analyzing the dialect diversity in multi-document summaries**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6208–6221, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Nils Reimers, Benjamin Schiller, Tilman Beck, Johannes Daxenberger, Christian Stab, and Iryna Gurevych. 2019. **Classification and clustering of arguments with contextualized word embeddings**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 567–578, Florence, Italy. Association for Computational Linguistics.

Anurag Shandilya, Abhisek Dash, Abhijnan Chakraborty, Kripabandhu Ghosh, and Saptarshi Ghosh. 2020. *Fairness for whom? understanding the reader’s perception of fairness in text summarization*. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 3692–3701.

Tejpal Singh Siledar, Jigar Makwana, and Pushpak Bhat-tacharyya. 2023. Aspect-sentiment-based opinion summarization using multiple information sources. In *Proceedings of the 6th Joint International Conference on Data Science & Management of Data (10th ACM IKDD CODS and 28th COMAD)*, pages 55–61.

Jiaan Wang, Yunlong Liang, Fandong Meng, Zhixu Li, Jianfeng Qu, and Jie Zhou. 2023. Cross-lingual summarization via chatgpt. *arXiv preprint arXiv:2302.14229*.

Xianjun Yang, Yan Li, Xinlu Zhang, Haifeng Chen, and Wei Cheng. 2023. Exploring the limits of chatgpt for query or aspect-based text summarization. *arXiv preprint arXiv:2302.08081*.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.

A Appendix

A.1 Dataset

The dataset provided by Glandt et al. (2021) consists of four different stance targets related to COVID-19. The detail of the data distribution can be found in Table 3.

Table 3: COVID-related stance detection data distribution in training, validation and test subsets per target provided in Glandt et al. (2021)

Target	Train	Val	Test
Anthony S. Fauci, M.D.	1464	200	200
Keeping Schools Closed	790	200	200
Stay at Home Orders	972	200	200
Wearing a Face Mask	1307	200	200

A.2 COVID-Stance-BERT Performance

A.3 Illustrations of Opinion Diversity

Different illustrations of opinion diversity calculation can be found in Table 5. The opinion precision measures the proportion of important opinions in the generated summary. The opinion recall measures the degree of salient opinion in the source documents that the generated summary contains.

Table 4: Performance of the COVID-Stance-BERT models for stance detection on the targets in the dataset provided by Glandt et al. (2021). The performance is reported in terms of accuracy (Acc) and macro F1 score (F1).

Target	Acc	F1
Anthony S. Fauci, M.D.	0.7714	0.7557
Stay at Home Orders	0.8652	0.8340
Wearing a Face Mask	0.8257	0.8180

A.4 Stance Examples and Opinion Diversity Calculations

Table 6 provided an example of the stance expressed in the source tweets and the generated summaries using a cluster of tweets under the topic of "Stay at Home Orders". In this example, the source tweets presented three different stances towards the topic. For the generated summaries, Pegasus (Zhang et al., 2020) presented only a neutral stance towards the topic, while ChatGPT successfully captured both support and against stances. We provide how opinion diversity can be measured using the provided example and the calculation detail can be found in Table 7. Applying F1 to these sets of stances yields the results shown in Table 6; from this, we can see that ChatGPT in this example has a higher opinion diversity score.

A.5 Additional Results

Table 8 contains results of opinion precision and recall for the three discussed topics. Based on the result, we found that BART (Lewis et al., 2020) has the best precision score across the three discussed topics. ChatGPT has relatively strong recall scores. BART is better at capturing important opinions, while ChatGPT is relatively better at capturing salient opinions.

Table 5: Illustration of different opinion precision and recall scenarios.

Source documents: Opinion A; Opinion B			
Model generated summary:	Opinion Precision	Opinion Recall	F1
Good precision, weak recall: Opinion A	1.00	0.50	0.67
Good precision, good recall: Opinion A; Opinion B	1.00	1.00	1.00
Weak precision, weak recall: Opinion A; Opinion C	0.50	0.50	0.50
Bad precision, bad recall: Opinion C	0.00	0.00	0.00

Table 6: An example of source tweets and generated summaries with different opinions towards the "Stay at Home Orders" topic. In the source tweets, users expressed support, against and neutral stances towards the topic, forming three different opinions. In the generated summaries, ChatGPT covered two opinions, whereas Pegasus (Zhang et al., 2020) covered a single opinion. In this scenario, ChatGPT has better opinion diversity.

Source tweets	Stance in source tweets
'Hey, #billmaher, having people stay home is not "fear;" it's *public health.* @USER' '@USER It goes against every fibre of my being to stay home but I have to because my dads high risk. If you die from covid you can't support the movement long-term. Do what you can from home if you're worried'	Support
'@USER Zero protesters but we all are ordered to stay home, and this is a "free" country. What a disgrace to the veterans that died for "free".'	Against
'The UFCW told their members to "please stay home" and expressed their sympathies to members "who lost their workplaces." And this is in a union that has tens-of-thousands of Black members working in essential low-wage jobs in retail and meat packing. Cowards.'	Neutral
Generated summaries	Stance in generated summaries
ChatGPT	
'People are upset about being told to stay home during the pandemic, but it's important for public health.' 'Some are frustrated that they can't attend church or receive sacraments, while others are protesting and attending funerals.' 'There are concerns about the spread of COVID-19 and the impact on vulnerable individuals.' 'Some are critical of those who are not taking'	Support Against
Pegasus	
'I stay home to protect family, friends, strangers.' '...and this is in a union that has tens-of-thousands of black members working in essential low-wage jobs in retail and meat packing.' '...and this is in a state that has much harsher lasting restrictions depending on the needs of each state...look it up during a'	Neutral

Table 7: Illustration of opinion diversity calculation using examples provided in Table 6

Source tweets: Stance-A, Stance-B, Stance-C
System summary: ChatGPT: Stance-A, Stance-B Pegasus: Stance-C
Opinion diversity: ChatGPT: Precision = 2/2; Recall = 2/3; F1 = 0.8 Pegasus: Precision = 1/1; Recall = 1/3; F1 = 0.5

Table 8: Results of opinion precision and opinion recall for various models under different discussed topics.

Events	CDC		Stay at Home Orders		Wearing a Face Mask	
Models	Precision	Recall	Precision	Recall	Precision	Recall
BART	0.8462	0.7286	1.0000	0.6354	0.9808	0.7308
Pegasus	0.6111	0.5256	0.9931	0.6424	0.4071	0.3558
T5	0.7137	0.6560	0.9826	0.6285	0.4872	0.4615
ChatGPT	0.7863	0.7521	0.9931	0.6979	0.8077	0.5000
Copycat	0.6303	0.5192	0.9792	0.5868	0.8782	0.5929
TextRank	0.6303	0.5235	0.9931	0.6319	0.2981	0.2468
LexRank	0.6517	0.5577	1.0000	0.6389	0.3942	0.3397
Hybrid TFIDF	0.6496	0.5791	0.9792	0.5833	0.4038	0.3494