

SUMMVIS: Interactive Visual Analysis of Models, Data, and Evaluation for Text Summarization

Jesse Vig[‡] Wojciech Kryściński[‡] Karan Goel[†] Nazneen Fatema Rajani[‡]

[‡] Salesforce Research [†] Stanford University

{jvig, kryscinski, nazneen.rajani}@salesforce.com

kgoel@cs.stanford.edu

Abstract

Novel neural architectures, training strategies, and the availability of large-scale corpora have been the driving force behind recent progress in abstractive text summarization. However, due to the black-box nature of neural models, uninformative evaluation metrics, and scarce tooling for model and data analysis, the true performance and failure modes of summarization models remain largely unknown. To address this limitation, we introduce SUMMVIS, an open-source tool for visualizing abstractive summaries that enables fine-grained analysis of the models, data, and evaluation metrics associated with text summarization. Through its lexical and semantic visualizations, the tool offers an easy entry point for in-depth model prediction exploration across important dimensions such as factual consistency or abstractiveness. The tool together with several pre-computed model outputs is available at <https://summviz.com>.

1 Introduction

The field of Natural Language Processing has seen substantial progress in recent years driven by the availability of large-scale corpora (Brown et al., 2020; Raffel et al., 2020), developments in neural architectures (Vaswani et al., 2017; Zaheer et al., 2020) and training strategies (Devlin et al., 2019; Zhang et al., 2020a). Despite the promising results on benchmarks and recent findings in model analysis, the true performance, generalizability, and failure modes of modern neural models are not yet fully understood, due to the black-box nature of neural models and the unmanageable scale of recent datasets for manual analysis. Software tooling for NLP research provides a plethora of mature and easy-to-use libraries for model development, such as PyTorch (Paszke et al., 2019) or Transformers (Wolf et al., 2020a), but offers disproportionately fewer tools for visual analysis and debugging,

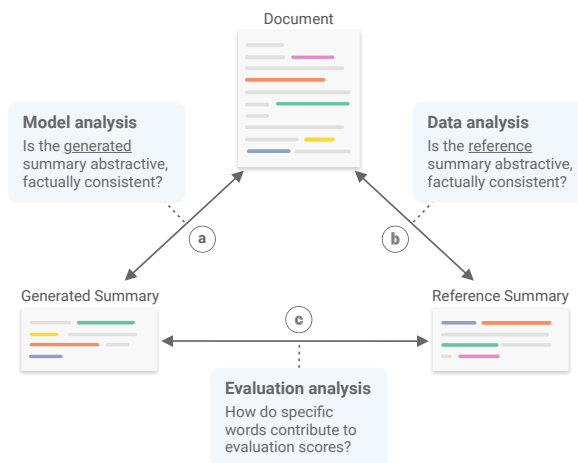


Figure 1: SUMMVIS supports fine-grained comparison between (a) source document and generated summary, (b) source document and reference summary, and (c) generated summary and reference summary, enabling analysis of models, data, and evaluation metrics.

which further hinders the understanding of model performance.

Within NLP, Automatic Text Summarization is a task that aims to convert long documents into short textual snippets that contain the most important information from the source document. To successfully summarize documents, models must first build an understanding of the source text that will allow them to evaluate the saliency of presented facts and then select only the most important details for the output summary. In case of abstractive approaches, the neural networks are also expected to paraphrase the selected content to generate novel sentences that fuse together the facts extracted from different sections of the document into coherent and factually consistent text.

Progress in the field is measured primarily using automatic metrics, such as ROUGE (Lin, 2004) or BERTScore (Zhang et al., 2020b), which quantify the lexical and semantic overlap between reference

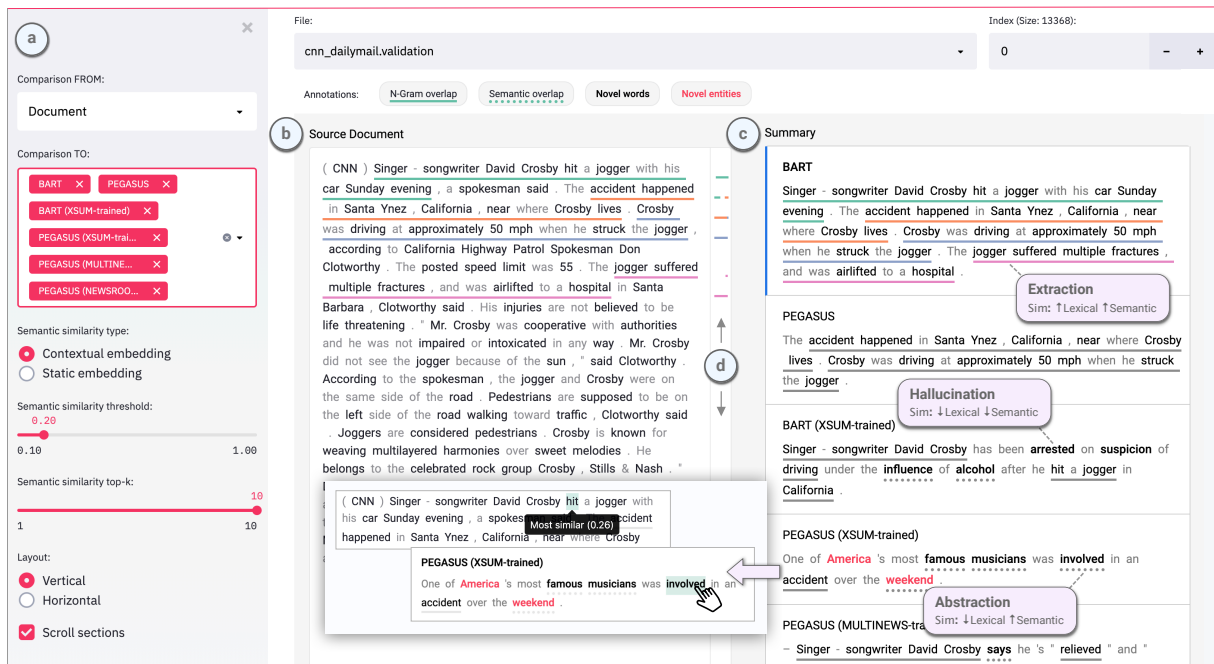


Figure 2: **Warning: Contains factually incorrect information.** SUMMVIS interface showing the first example from the CNN/DailyMail validation split. Interface components: (a) configuration panel, (b) source document (or reference summary, depending on configuration), (c) generated summaries (and/or reference summary, depending on configuration), (d) scroll bar with global view of annotations. Colored underlines align n-grams between source document and the selected summary (BART); colors are determined by the position of the containing sentence in the summary. Novel words in the summary that do not appear in the source document are **bolded**, while novel entities are bolded in **red**. Stopwords are **grayed out** and are not used in the matching algorithms. **Dotted underlines** indicate tokens that are semantically similar to a token in the source document (above the threshold specified in the configuration panel). The user may hover over a token to see the most semantically similar tokens in the source document (see inset image), or click on the token to auto-scroll the source document to the most similar token.

and generated summaries. While automatic metrics are convenient for model evaluation, they have been shown to be mismatched with human judgements (Fabbri et al., 2020) and only offer high-level insights while failing to pinpoint particular shortcomings of models. In-depth debugging across the different modes of analysis (Fig. 1) must be conducted through expensive and time-consuming human-based studies, where the substantial length of texts makes such efforts more labor-intensive.

Recent work in summarization analysis has looked at the problems of the field in isolation, focusing on: models (Kedzie et al., 2018; Kryściński et al., 2019, 2020), data (Zhong et al., 2019; Jung et al., 2019), and evaluation (Fabbri et al., 2020; Steen and Markert, 2021). However, these modes of analysis are strongly interconnected and isolating them could skew the broader view of the current state of the task and delay progress.

To address the mentioned challenges, we introduce SUMMVIS, an open-source interactive visu-

alization tool for analyzing text summarization. SUMMVIS was designed to offer fine-grained insights into the models, data, and evaluation metrics, both in isolation and jointly, thus compensating for the shortcomings of automatic evaluation metrics and shortage of dedicated debugging tooling. SUMMVIS scaffolds human analysis by offering clear visual indicators of the semantic and lexical relationships between texts and intelligent navigation within text. The tool comes pre-loaded with a set of state-of-the-art model predictions for a quick starting point for model analysis and comparison and offers out-of-the-box integration with the HuggingFace Dataset API for custom use-cases. Through a case study of state-of-the-art summarization models we show how SUMMVIS can be used to quickly conduct non-trivial analysis, debugging, and comparison of model performance across important dimensions such as factual consistency or abstractiveness. A video demonstration of the tool is available at <https://vimeo.com/540429745>.

2 SUMMVIS

In this section, we present SUMMVIS, an interactive visualization tool that provides rich text comparison in summarization systems, enabling fine-grained analysis of models, data, and evaluation metrics. It comes pre-loaded with model outputs for state-of-the-art models over common benchmark datasets, as well as scripts for loading data for any dataset provided by the Datasets API (Wolf et al., 2020b) and any HuggingFace-compatible model.

2.1 Analysis Modes

SUMMVIS supports three modes of analysis, depending on the type of text being compared:

1. **Model Analysis** (Fig. 1a). By comparing the source document with *generated* summaries, SUMMVIS provides insights into a model’s ability to abstract and faithfully retain information present in the document.
2. **Data Analysis** (Fig. 1b). By comparing the source document with the *reference* summary, SUMMVIS helps determine the degree to which the reference summary itself is abstractive and factually consistent with the source document.
3. **Evaluation Analysis** (Fig. 1c). By comparing the *reference* summary with the *generated* summary, SUMMVIS surfaces the word- and phrase-level relationships that form the basis of automated evaluation metrics such as ROUGE and BERTScore.

These analyses are interdependent with one another; for example, the behavior of a *model* depends on the *data* on which it was trained. By providing a unified interface for all modes of analyses, the user may also draw conclusions about the relationships between model, data, and evaluation, as we’ll demonstrate in Section 3.

2.2 Text Comparison

Understanding abstractive summaries requires comparing not only surface similarities but also building a semantic understanding of the source document and summaries. Therefore SUMMVIS incorporates similarity measures based on both lexical and semantic overlap, as described below.

Lexical Overlap. The ability to quickly compare the lexical form of source document and summary is an important first step in analyzing a generated

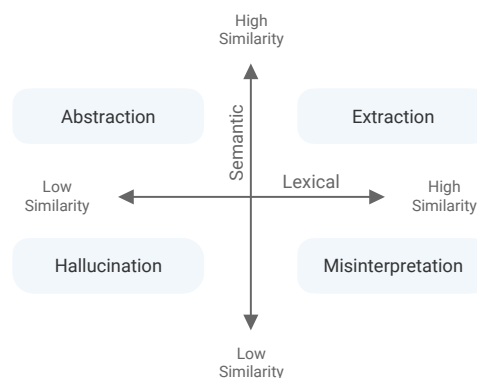


Figure 3: Taxonomy of textual relationships across lexical and semantic dimensions.

summary. For example, it is well known that many abstractive reference summaries are in fact largely extractive, copying long spans of text from the source document (Grusky et al., 2018). Other summaries might contain significant hallucinations, including words that are not found in the source document (Kryściński et al., 2020; Maynez et al., 2020). In order to identify these phenomena, SUMMVIS provides a lexical alignment based on shared n-grams between the two texts, which is also the basis for many automated metrics such as ROUGE. **Semantic Overlap.** Lexical overlap is incomplete as a measure of similarity between texts since it only considers the surface form of words. For example, a summary that is highly abstractive may share few common words with the source article, despite having a similar meaning. To address such limitations, the tool also identifies semantically-related tokens by computing the cosine similarity between word embeddings, with the option of using static word embeddings provided by spaCy (Honnibal et al., 2020), or contextual embeddings from a pretrained RoBERTa (Liu et al., 2019) model. In the later case, we apply the same default embeddings¹ used in BERTScore, a common evaluation metric for abstractive summarization systems that correlates strongly with human evaluations (Zhang et al., 2020b). As we’ll discuss in Section 3, the visualized semantic similarities can also help to interpret BERTScore values. We note that the BERTScore library² used in the tool also supports other models of semantic similarity, for example, models trained on scientific or non-English text.

Taxonomy. Considering both lexical and semantic measures of similarity provides a natural way to

¹RoBERTa-large layer 17

²https://github.com/Tiiiger/bert_score

chart out summarization datasets for further analysis. By comparing a source document to any summary along these two dimensions, four quadrants of behavior can be mapped out (Fig. 3):

1. **Extraction:** high lexical and high semantic similarity. The summary quotes text from the document verbatim.
2. **Abstraction:** low lexical and high semantic similarity. The summary consolidates and paraphrases information from the document.
3. **Hallucination:** low lexical and low semantic similarity. The summary is factually inconsistent, and includes information that is absent in the document.
4. **Misinterpretation:** high lexical and low semantic similarity. The summary misinterprets and uses information from the document, such as misunderstanding homonyms.

Examples of such cases will be discussed in the following sections.

2.3 Interface

The main components of the SUMMVIS interface are described in detail in Figure 2. The interface supports analysis of the model, data, and evaluation (Sec. 2.1) based on which types of text are selected by the user for comparison (Figs. 2b, 2c). The annotations provided by the tool highlight both lexical and semantic relations between the text (Sec. 2.2) and are designed to be lightweight, allowing users to quickly grasp the relationship between texts while still being able to clearly read the text.

The joint lexical and semantic annotations enable the user to understand the summaries according to the taxonomy in Figure 3. Examples of extraction, abstraction, and hallucination are highlighted in Figure 2. Since measures of semantic similarity may be unreliable, the tool also enables users to hover over tokens for additional details on the semantically matched tokens in the source document, which are highlighted based on their semantic similarity scores (Fig. 2, inset image). Additionally the score of the closest match is displayed, following the BERTScore algorithm, which computes the maximum semantic similarity score for each token before averaging the results over the full text. These features enable users to manually assess whether the tokens are in fact semantically similar.

The tool supports two additional features to accommodate long source documents: a global view

and auto-scrolling functionality. The global view, embedded in the scroll bar region of the source document (Fig. 2d), displays a compressed view of the full document’s annotations that is visible even when the document exceeds the viewable region. The user may also directly navigate to matched portions of the source documents not currently visible by clicking on related annotations in the summary.

2.4 System Architecture

The interface is implemented as a Streamlit³ application with a highly customized HTML/JavaScript component that handles most interactions in the tool. The custom component enables a much richer interaction than a vanilla Streamlit app, while the Streamlit infrastructure allows for adapting or extending some components in the tool without necessarily writing additional HTML or JavaScript.

We provide pre-processing scripts to generate and cache all data required by SUMMVIS to ensure fast response times in the interface. These scripts are implemented using Robustness Gym (Goel et al., 2021) and integrate with the HuggingFace Datasets API (Wolf et al., 2020b) so that any summarization dataset available in the dataset repository or provided by the user as a `jsonl` file may be viewed in the tool. We additionally include scripts for caching outputs for any HuggingFace summarization model, and share pre-computed outputs of state-of-the-art summarization models: PEGASUS (Zhang et al., 2020a) and BART (Lewis et al., 2020). To increase the variety of outputs, we chose model checkpoints fine-tuned on multiple popular summarization datasets: CNN/DailyMail (Hermann et al., 2015), XSum (Narayan et al., 2018), Newsroom (Grusky et al., 2018), and MultiNews (Fabbri et al., 2019), and decoded on the validation splits of two benchmark datasets: CNN/DailyMail and XSum.

3 Case Study: Debugging Hallucination

As discussed earlier, SUMMVIS supports joint analysis of the model, data, and evaluation metrics. We now demonstrate how we can draw from all three modes of analysis to study the problem of hallucination in summarization systems. Through the unified view of SUMMVIS, we analyze the example shown in Figure 4 and demonstrate the existence of hallucination, suggest a possible cause, and show

³<https://github.com/streamlit/streamlit>

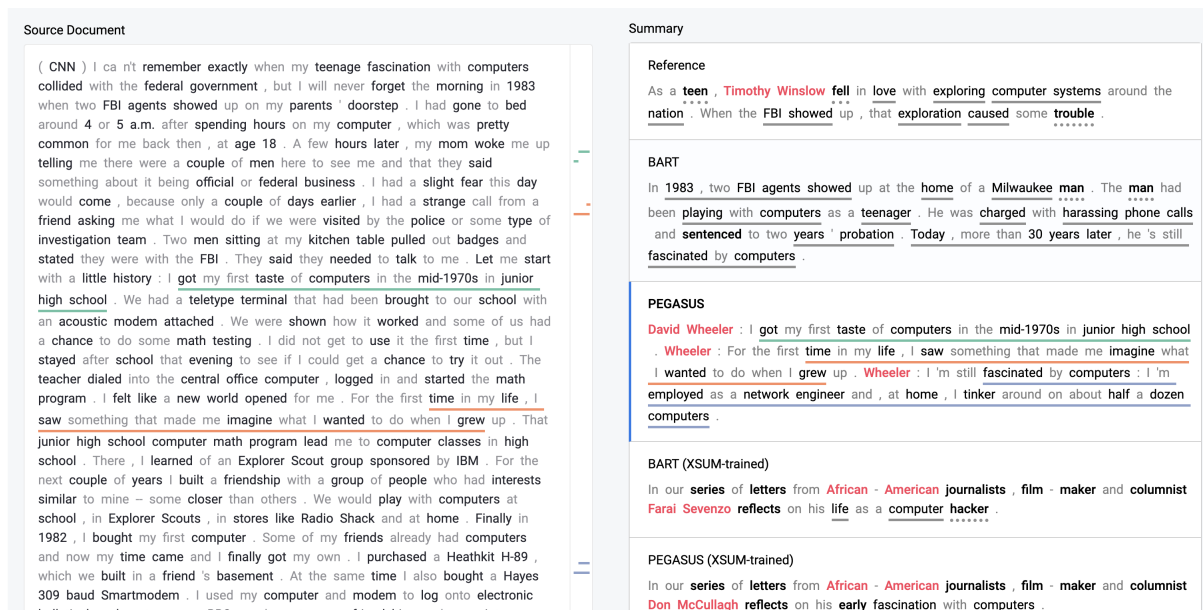


Figure 4: SUMMVIS snapshot of an example from CNN/DailyMail, showing the source document (left) and the reference summary along with generated summaries from four different models (right). The first two models are trained on CNN/DailyMail, while the last two are trained on XSum. Red text highlights entities in the summaries that are not present in the source document.

how a common evaluation metric prefers hallucinated entities over faithful descriptors in this case.

Model Analysis. SUMMVIS supports analysis of the model by visualizing the relationship between each generated summary and the source document. For the example in question (Fig. 4), this visualization reveals that three of the four models generate names of people that are absent from the source document. The XSum-trained models generate the names in the context of the phrase “*In our series of letters from African-American journalists, film-maker and columnist <person_name> reflects on ...*”. suggesting that the hallucinations for these two models may be related to artifacts in the shared XSum training set that both models have memorized. On the other hand, the summary generated by the version of PEGASUS that was trained on CNN/DailyMail is largely extractive, copying several sentences, but then also inserting the name “*David Wheeler*”, which is absent from the source document. We now show how artifacts in the *reference* summaries may explain this hallucination.

Data Analysis. We now turn to the visualization comparing source document and *reference* summary (Fig. 4, top right). We see that the reference summary also contains an entity that is missing from the source document (“*Timothy Winslow*”). This may be due to the name appearing in metadata

such as author name that was available to the person writing the summary, but was not included in the dataset. If this pattern occurs in similar types of examples in the training set (e.g., first-person written articles), then it may effectively teach the model to hallucinate, providing a possible explanation for the model behavior described earlier.

Evaluation Analysis. One remaining question is how state-of-the-art models can hallucinate but still perform well on benchmark datasets according to standard evaluation metrics. Of course, one reason is that the models only hallucinate on some fraction of examples in the dataset. However, there is also the question of how the evaluation metrics score hallucinated content. While lexical overlap metrics such as ROUGE are well-defined, semantic similarity metrics like BERTScore are less well understood as they depend on embeddings from black-box neural network models.

SUMMVIS supports fine-grained analysis of evaluation metrics through its comparison of *generated* and *reference* summaries. In particular, the token-level semantic similarity scores visualized in the tool use the same similarity measure as BERTScore (Sec. 2.2). By inspecting these token-level relationships, we can better understand how hallucinated tokens contribute to the overall BERTScore, which is computed by aggregating token-level scores.

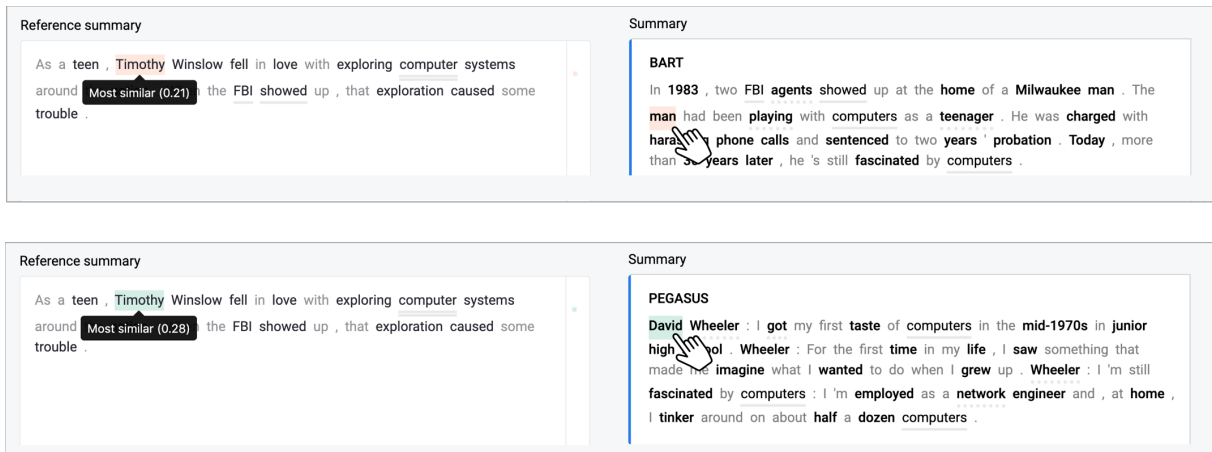


Figure 5: Snapshot from SUMMVIS showing the reference summary on the left and two of the generated summaries on the right. In the first example, the user has hovered over “man” in the generated summary, which causes the tool to highlight the most semantically similar word in the reference summary, “Timothy”, with a similarity score of 0.21. A second occurrence of “man” has an even lower semantic similarity score of just 0.02 (not shown). In the second example, the user hovers over “David”, revealing that this word is also most semantically similar to “Timothy”, but with a higher similarity score (0.28).

Figure 5 shows the comparison between the reference summary and two of the generated summaries, revealing that the factually correct “man” has a lower maximum semantic similarity score compared to the hallucinated “David”. The same is true for the corresponding hallucinated last name “Wheeler” (similarity: 0.28), and this disparity with “man” is even more pronounced for the hallucinated name “Don McCullagh” (Similarity: 0.34, 0.31) generated by the last model shown in Figure 4. Thus BERTScore does not discriminate factual consistency of proper names in this example, consistent with anecdotal evidence for other types of entities (Zhang et al., 2020b). Note that the hallucinated name “Farai Sevenzo” (Fig. 4, 4th row) has maximum similarity scores that are negative (-0.43, -0.12). This disparity may relate to name biases in word embeddings (Caliskan et al., 2017).

4 Related Work

Text Summarization requires models to be adept at both natural language understanding (NLU) and natural language generation (NLG). A gap in either of these areas has consequences on the progress of summarization as a whole. An example of this is the lack of meaningful metrics in NLG for high entropy tasks (Steen and Markert, 2021). Several recent works have realized the need for evolving benchmarks and evaluations (Goel et al., 2021; Gehrmann et al., 2021; Khashabi et al., 2021).

Existing tools support some forms of text com-

parison for summarization models. The Newsroom dataset visualization tool (Grusky et al., 2018) highlights n-grams in the summary that overlap with the source article. The LIT tool (Tenney et al., 2020) highlights words or characters that differ between reference and generated texts. However neither tool aligns (Yousef and Janicke, 2021) the matched text. The CSI framework (Gehrmann et al., 2019) and the Seq2SeqVis (Strobel et al., 2019) tool align the source document and summary, but use model-specific attention mechanisms. SUMMVIS on the other hand supports a model-agnostic comparison between source document, reference summary, and generated summary, and aligns text along lexical and semantic dimensions.

5 Conclusion

In this work we introduced SUMMVIS, an interactive visualization tool for analyzing text summarization models, datasets, and evaluation metrics. Through a case study we showed that our tool can be used to efficiently identify the shortcomings and failure modes of state-of-the-art summarization models and datasets. Together with the tool we released a set of pre-computed model outputs to enable easy, out-of-the-box use. We hope this work will positively contribute to the ongoing efforts in building tools for model evaluation and analysis and enable a deeper understanding of the performance of summarization models and the intricacies of datasets and metrics.

6 Ethics Statement

To the best of our knowledge, there is no work on ethical bias in automated text summarization. The news summarization datasets currently used by the NLP community are mainly crawled from Western news outlets and therefore are not representative of a majority of geographies. There are also biases in news reporting that can distill into parameters of models trained on such biased datasets and may even be further amplified in the generated model outputs. All datasets are in English, and all models are trained on English datasets.

SUMMVIS uses spaCy for entity detection and because we did not stress test the detector, there might be biases in the system that have percolated into our tool. Similarly, the text similarity metrics used in our tool including the BERTScore and the word-embeddings carry biases of the data they were trained on. For example, they have been known to have bias associating professions with a particular gender. We request our users to be aware of these ethical issues that might affect their analyses.

Acknowledgments

We would like to thank Michael Correll for his insightful feedback.

References

- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. [Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir R. Radev. 2020. [Summeval: Re-evaluating summarization evaluation](#). *CoRR*, abs/2007.12626.
- Sebastian Gehrmann, Tosin Adewumi, Karmanya Agarwal, Pawan Sasanka Ammanamanchi, Aremu Anuoluwapo, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna Clinciu, Dipanjan Das, Kaustubh D Dhole, et al. 2021. [The gem benchmark: Natural language generation, its evaluation and metrics](#). *arXiv preprint arXiv:2102.01672*.
- Sebastian Gehrmann, Hendrik Strobelt, Robert Krüger, Hanspeter Pfister, and Alexander M Rush. 2019. [Visual interaction with deep learning models through collaborative semantic inference](#). *IEEE transactions on visualization and computer graphics*, 26(1):884–894.
- Karan Goel, Nazneen Rajani, Jesse Vig, Zachary Taschdjian, Mohit Bansal, and Christopher Ré. 2021. [Robustness Gym: Unifying the NLP evaluation landscape](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: System Demonstrations*.
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 708–719, New Orleans, Louisiana. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Advances in neural information processing systems*, pages 1693–1701.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Taehee Jung, Dongyeop Kang, Lucas Mentch, and Edward Hovy. 2019. [Earlier isn’t always better: Subaspect analysis on corpus and system biases in summarization](#). In *Proceedings of the 2019 Conference*

- on *Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3324–3335, Hong Kong, China. Association for Computational Linguistics.
- Chris Kedzie, Kathleen R. McKeown, and Hal Daumé III. 2018. [Content selection in deep learning models of summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 1818–1828. Association for Computational Linguistics.
- Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, Jungo Kasai, Yejin Choi, Noah A Smith, and Daniel S Weld. 2021. Genie: A leaderboard for human-in-the-loop evaluation of text generation. *arXiv preprint arXiv:2101.06561*.
- Wojciech Kryściński, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 540–551. Association for Computational Linguistics.
- Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 9332–9346. Association for Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Don’t give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8024–8035.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Julius Steen and Katja Markert. 2021. [How to evaluate a summarizer: Study design and statistical analysis for manual linguistic quality evaluation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1861–1875, Online. Association for Computational Linguistics.
- Hendrik Strobelt, Sebastian Gehrmann, Michael Behrisch, Adam Perer, Hanspeter Pfister, and Alexander M. Rush. 2019. [Seq2seq-vis: A visual debugging tool for sequence-to-sequence models](#). *IEEE Transactions on Visualization and Computer Graphics*, 25(1):353–363.
- Ian Tenney, James Wexler, Jasmijn Bastings, Tolga Bolukbasi, Andy Coenen, Sebastian Gehrmann, Ellen Jiang, Mahima Pushkarna, Carey Radebaugh, Emily Reif, and Ann Yuan. 2020. [The language interpretability tool: Extensible, interactive visualizations and analysis for NLP models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 107–118, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. Rush. 2020a. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pages 38–45. Association for Computational Linguistics.
- Thomas Wolf, Quentin Lhoest, Patrick von Platen, Yacine Jernite, Mariama Drame, Julien Plu, Julien Chaumond, Clement Delangue, Clara Ma, Abhishek Thakur, Suraj Patil, Joe Davison, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angie McMillan-Major, Simon Brandeis, Sylvain Gugger, François Lagunas, Lysandre Debut, Morgan Funtowicz, Anthony Moi, Sasha Rush, Philipp Schmid, Pierric Cistac, Victor Muštar, Jeff Boudier, and Anna Tordjmann. 2020b. Datasets. *GitHub. Note: <https://github.com/huggingface/datasets>*, 1.
- T. Yousef and S. Janicke. 2021. [A survey of text alignment visualization](#). *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1149–1159.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. [Big bird: Transformers for longer sequences](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020a. [PEGASUS: pre-training with extracted gap-sentences for abstractive summarization](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339. PMLR.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Ming Zhong, Danqing Wang, Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2019. [A closer look at data bias in neural extractive summarization models](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 80–89, Hong Kong, China. Association for Computational Linguistics.