

2021

OpenFraming: open-sourced tool for computational framing analysis of multilingual data

V. Bhatia, V.P. Akavoor, S. Paik, L. Guo, M. Jalal, A. Smith, D.A. Tofu, E.E. Halim, Y. Sun, M. Betke, P. Ishwar, D.T. Wijaya. 2021. "OpenFraming: Open-sourced Tool for Computational Framing Analysis of Multilingual Data." Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. 2021-11 - 2021-11. <https://doi.org/10.18653/v1/2021.emnlp-demo.28>
<https://hdl.handle.net/2144/44172>

"Downloaded from OpenBU. Boston University's institutional repository."

OpenFraming: Open-sourced Tool for Computational Framing Analysis of Multilingual Data

Vibhu Bhatia*
Alyssa Smith†

Vidya Akavoor*
David Tofu
Prakash Ishwar

Sejin Paik*
Edward Halim

Lei Guo
Yimeng Sun
Derry Tanti Wijaya‡

Mona Jalal
Margrit Betke

Boston University
wijaya@bu.edu

Abstract

When journalists cover a news story, they can cover the story from multiple angles or perspectives. These perspectives are called “frames”, and usage of one frame or another may influence public perception and opinion of the issue at hand. We develop a web-based system for analyzing frames in multilingual text documents. We propose and guide users through a five-step end-to-end computational framing analysis framework grounded in media framing theory in communication research. Users can use the framework to analyze multilingual text data, starting from the exploration of frames in user’s corpora and through review of previous framing literature (step 1-3) to frame classification (step 4) and prediction (step 5). The framework combines unsupervised and supervised machine learning and leverages a state-of-the-art (SoTA) multilingual language model, which can significantly enhance frame prediction performance while requiring a considerably small sample of manual annotations. Through the interactive website, anyone can perform the proposed computational framing analysis, making advanced computational analysis available to researchers without a programming background and bridging the digital divide within the communication research discipline in particular and the academic community in general. The system is available online at <http://www.openframing.org>¹, via an API <http://www.openframing.org:5000/docs/>, or through our GitHub page <https://github.com/vibss2397/openFraming>.

1 Introduction

We live in a world saturated with media. Any major public issue, such as the ongoing COVID-19 pandemic and the Black Lives Matter protests, attracts

tremendous attention from hundreds of thousands of news media outlets — traditional and emerging — around the world. The reporting angles on a single issue are often varied across different media outlets. In covering COVID-19, for example, some media outlets focus on government response and actions while others emphasize the economic consequences. Social science scholars call this process *media framing*. To define, or to *frame*, is “to select some aspects of a perceived reality and make them more salient in a communicating text” (Entman, 1993). When used in news articles, frames can strongly impact public perception of the topics reported and lead to different assessments by readers (Hamborg, 2020), or even reinforce stereotypes and project explicit and implicit social and racial biases (Drakulich, 2015; Sap et al., 2019).

Frame discovery in media text has been traditionally accomplished using methods such as *quantitative content analysis* (Krippendorff, 2018), which is a manual method widely used by social scientists. However, in the emerging media environment, the sheer volume and velocity with which content is generated makes manual labeling increasingly intractable. To overcome this “Big Data” challenge, our cross-disciplinary team, which consists of computer science and communication researchers have employed methods based on SoTA machine learning (ML) techniques to detect frames automatically and robustly (Akyürek et al., 2020; Liu et al., 2019).

However, these SoTA ML models are not readily accessible to social sciences scholars who typically do not have machine learning and programming training. The current ecosystem around “Big Data” creates new digital divides between the Big Data rich and the Big Data poor (Boyd and Crawford, 2012). Among other barriers, the limited access to computational resources and skillsets prevents many communication scholars from taking advantage of a large number of unprecedented ML models of the day and hampers their ability to glean

* Contributed equally

† While at Summer research in Boston University

‡ Corresponding author

¹Best viewed with Google Chrome browser




Figure 1: Our proposed five-step computational framing analysis framework that is grounded in media framing theory in communication research. Users can use our framework to analyze frames in text, starting from the *inductive* exploration of frames in user’s corpora via topic modeling (step 1) and *deductively* through review of previous framing literature (step 2), to labeling frames for training data with content analysis (step 3), to training the frame classification model (step 4) and to prediction on unlabeled data (step 5).

valuable insights from unprecedentedly large media datasets since framing, especially at scale, defines how news media coverage shapes mass opinion.

Our goal is to make computational framing analysis accessible to researchers from a diverse array of disciplines. We present OpenFraming (www.openframing.org), a free and user-friendly Web-based system that allows researchers to conduct computational framing analysis without having to write and debug complex code. There does, of course, exist click-and-play commercial software, but these tools are often costly and pose issues for researchers due to a lack of transparency into their inner computational mechanisms. This black box problem is present in various applications of data science techniques in communication research (Guo, 2018), thus hindering the open science movement in the field (van Atteveldt and Peng, 2018). In contrast, our system is based on SoTA framing research and our code is publicly available under MIT license.

Specifically, we propose a five-step analytical framework (Figure 1) allowing users to identify frames in large-scale media text by leveraging SoTA computational frame analysis research techniques. Our work is advantageous in at least five aspects. First, the framework is grounded in media framing theory, one of the most established theories in communication research (D’Angelo, 2018; Reese et al., 2001b). Second, we provide a web-based, user-friendly graphic interface where researchers with little or no computational background can perform advanced data analysis through

a click-and-play approach. Third, all algorithms used are open to users and the benefits and limitations of the algorithms are explained at each step. Fourth, given the increasing importance of understanding information flow at a global scale, our tool can be used to analyze media content in 23 languages based on a SoTA multilingual language model (Devlin et al., 2018). Lastly, with the support of a research grant, we make our tool entirely free to the academic community.

We will start with a review of the theoretical and methodological backgrounds based on which our system is developed. We will then detail the five-step framing analysis facilitated by the system. The importance of bridging the digital divide in the field of computational communication research will also be discussed.

2 Related Work

To frame is “to select some aspects of a perceived reality and make them more salient in a communicating text” (Entman, 1993). Like any type of communication, news involves framing. The ideology of a society, ownership of a news organization, media routine, as well as individual media worker’s preferences all play some role, consciously or not, in shaping the news content (Shoemaker and Reese, 1995). Reese et al. (2001a) defines media frames as “organizing principles that are socially shared and persistent over time, that work symbolically to meaningfully structure the social world”. In other words, one can frame an issue in multiple ways,

but a frame must be shared by the target audience on some level for it to be communicable and effective. Accordingly, news frame analysis should focus on frames that are “persistent over time”. These include generic frames that appear across issues, time, and space, such as human interest, conflict, attribution of responsibility, and economic consequences (Neuman et al., 1992; Nisbet, 2010; Semetko and Valkenburg, 2000). For any particular issue, journalists also apply issue-specific frames e.g., reporters often use peace- and war-oriented frames to help their audience understand the complexity of wars (Neumann and Fahmy, 2012).

Empirically, communication researchers have developed a variety of approaches to analyze frames. There are in general three computational approaches: 1) lexical-based, 2) unsupervised ML, and 3) supervised ML.

The lexical-based approach relies on predefined lists of words, known as lexicons or dictionaries, with each word associated with a certain frame (Field et al., 2018). For example, (Lind et al., 2019) develops keywords to search for frames in the news coverage of immigration. We, however, contend that the lexical-based approach is not ideal for news frame analysis. Unlike the topic-like frames in Lind et al. (2019), many enduring media frames (e.g., conflict, human interest) are abstract and involve complex meanings, which cannot be easily captured by a list of words and terms.

The other two approaches are based on ML models that learn from data. While unsupervised ML models discover patterns of frames from unlabeled data, supervised ML is done by training a model on a sample of documents that are labeled with the “correct” frame. In communication research, the “correct frame” often refers to labels provided by human coders through quantitative content analysis (Krippendorff, 2004).

Several existing news frame studies used an unsupervised ML approach. The Latent Dirichlet Allocation (LDA) topic modeling (Blei et al., 2003) is a popular example (see Maier et al. (2018) for a systematic review). In analyzing news content, the text is observed as a set of latent “topics” and these topics are distributed over words in a probabilistic order. The output of the LDA topic modeling is a “topic matrix” with a list of keywords representing each topic. Researchers will have to review the top keywords and decide on a label to represent the meaning of each topic. Some studies approach

media framing by analyzing mainly themes or topics. This approach is problematic because, again, news frame analysis should identify patterns that endure over time, which is different from thematic or topical analyses that describe themes or topics as instances reported in certain stories (Reese, 2007). Given this, the LDA approach is most useful for exploratory analysis. Although the LDA-generated topics are not necessarily equivalent to frames, the information can be used to obtain initial ideas about the data and infer potential frames for the next step of supervised frame analysis (Guo et al., 2016).

Supervised ML also has become increasingly common in communication research (Colleoni et al., 2014; De Grove et al., 2020). Our recent studies use BERT language model (Devlin et al., 2018) and fine-tune it for the task of identifying frames in the news coverage about US gun violence which demonstrate a high level of accuracy for multilingual frame detection (Liu et al., 2019; Akyürek et al., 2020) with a relatively small amount of data: 1.3k English frame-labeled news headlines.

Based on the review of the literature, we propose a five-step end-to-end multilingual framing analysis framework that combines unsupervised and supervised ML. Unsupervised ML can help develop a holistic picture of large-scale text corpora, but is not sufficient for a news frame analysis. Adding a supervised approach to the research framework is essential because the goal is to identify enduring frames—generic and issue-specific frames to contribute to the media framing literature. Furthermore, since developing multilingual ground truth labels through manual content analysis is labor-intensive and time-consuming, to address this challenge in part, our proposed framework incorporates a SoTA multilingual language model such as BERT that allows transfer learning from pretrained models to the task at hand. As a result, it would require relatively fewer labeled documents to fine-tune, while still achieving robust prediction performance.

3 System Description: Five-step Multilingual Framing Analysis

Our analytical framework, available publicly via www.openframing.org, involves five steps that researchers can take to conduct an end-to-end computational framing analysis of multilingual text corpora (Figure 1). To help demonstrate the research procedure, we will accompany the description of our system by a case study that exam-

ines frames in the U.S. news coverage of the U.S. gun violence issue—also provided as a demo at <http://www.openframing.org/demos.html>.

Before using the system, researchers should first collect data related to the issue under consideration. As for the case study, we collect news headlines using the keyword combination (gun OR firearm OR nra OR “2nd amendment” OR “second amendment” OR AR15 OR “assault weapon” OR rifle OR “brady act” OR “brady bill” OR shooting) and collected a total of 42,917 U.S. English news articles from 2018 using Brandwatch Consumer Research².

Step 1: Explore Topics with Topic Modeling

To analyze how the news media frame an issue, we first suggest users come up with a list of specific “frames” that guide the discussion of the issue. This process of searching for frames should be both inductively—based on an observation of the data—and deductively—based on the review of the previous framing literature. Both steps are essential because the analysis of frames should not just aim for a full capture of the data (inductive), but also to build and further advance the prior knowledge of the media framing theory (deductive).

Step 1 (Figure 2) focuses on the inductive part of the research. The goal of this first step is to preliminarily examine the LDA “topic” information of the data, which helps researchers decide the final frames to be analyzed in Step 2.

While there is some flexibility regarding the format of the input dataset (the system currently supports .xls, .xlsx, .csv, or .tsv), the file must contain a column labeled “example”, which contains one document per row. Throughout this paper, we call the unit of analysis a “document”, which can be a news headline, a news article, or a tweet. Our system also provides users with a list of data cleaning options (some of them are shown in Figure 2). In particular, we provide the option of running the analysis in multiple languages. The Natural Language Toolkit (NLTK) provides a library of stopwords available in 23 languages, which we use to remove stopwords in text in the language specified by the user. The LDA topic modeling, which is language invariant, is then applied to generate topics from the text when user hits the “submit” button. We use the Mallet LDA implementation (McCallum, 2002). The system then sends the user an e-mail with a link to download the results of the

²<https://www.brandwatch.com/>

Figure 2: Step 1 of the 5-step framing analysis that allows users to explore topics in their corpora with LDA topic modeling.

analysis—the topics discovered and the keywords per topic.

Although the LDA topic modeling is a computational method, its implementation involves a series of human reasoning. For instance, as mentioned above, researchers should decide the data cleaning procedures, the number of LDA topics, and the number of keywords associated with each topic. Our tool not only allows users to specify their preferred settings but also provides guidance and recommendations for each decision through the information **i** buttons on Figure 2. Following previous research, we recommend users try different numbers of topics before making the final decision.

In conducting the gun violence study, we first used the LDA to explore prominent topics in our dataset about the U.S. gun violence issue. For demonstration purposes, we tried 5, 10, and 15 topics. Based on the five-topic LDA output (Figure 3), we can manually assign labels to these topics: 1) mass shootings, 2) police officers, 3) school shootings and demonstrations, 4) gun rights and gun control, and 5) the second amendment. It is recommended that at least two researchers independently review the topics and then decide the labels collec-

	Topic 1	Topic 2	Topic 3	Topic 4	Topic 5
KW 1	time	police	student	trump	law
KW 2	life	officer	high	president	firearm
KW 3	mass	shot	florida	nra	state
KW 4	victim	man	parkland	house	weapon
KW 5	year	told	douglas	republican	rifle
KW 6	family	county	teacher	state	violence
KW 7	killed	yearold	stoneman	white	company
KW 8	friend	suspect	cruz	democrat	year
KW 9	video	report	march	bill	amendment
KW10	day	department	marjory	control	ban

Figure 3: The LDA topic modeling output for our gun violence news headlines based on five topics.

tively. When we increase the number of topics from five to 10, more information emerges such as mental health. However, when we further increase the number of topics to 15, redundancy occurs—for example, many topics are related to gun control—and certain topics contain words that are not semantically meaningful. This indicates that we may have reached the saturation point, thus further increasing the number of topics would be less likely to generate any new topical information.

After users try a series of numbers and explore the corresponding topic-keyword matrices in datasets of different languages, they will be able to develop a preliminary idea of the multilingual data. This concludes Step 1 of the framing analysis.

Step 2-3: Decide and Label Frames In Step 2, researchers are recommended to consult the LDA topic modeling results from Step 1 (inductive) and previous literature about media framing of the issue under investigation (deductive), and then decide a list of frames to be analyzed.

For our case study, based on the LDA modeling results of the news coverage of U.S. gun violence and the literature review of the media framing of this topic, we decided in Step 2 the following list of frames: 1) Gun/2nd amendment rights; 2) Gun control/regulation; 3) Politics; 4) Mental health; 5) School or public space safety; 6) Race/ethnicity; 7) Public opinion; 8) Society/culture; and 9) Economic consequences. On this list, some frames are issue-specific frames that are unique to the media coverage of gun violence such as “gun/2nd Amendment rights” and “mental health,” other frames are generic frames such as “economic consequences” that apply to all kinds of issues. It is also important to note that although the LDA topic modeling results do not have explicit reference to “society/culture,” we still include it because it is a media frame discussed in the previous literature about gun violence media coverage (Birkland and Lawrence, 2009; Schnell, 2001).

In Step 3, once the list of frames is decided, the user should draw a sample of the data and apply

quantitative content analysis (QCA) (Krippendorff, 2004) to manually label frames. This annotated sample will be used as the ground truth to train an ML model in Step 4. In our case study, we selected a random sample of 1.3k English headlines, and recruited two human coders to annotate the frames of the headlines. Following the QCA procedure, we created a codebook to explain each frame and held multiple training sessions for the coders to understand how to identify the dominant frames of the headlines. To test intercoder reliability, the two coders for each language were instructed to code a sample of news headlines independently and their results were compared. They ultimately reached a robust level of intercoder reliability (0.90 α).

Step 4: Build a Frame Classification Model with Deep Learning The goal of Step 4 is to use the documents the users have labeled from Step 3 to build supervised ML models that can then predict frames in unlabeled documents. Our analytical framework incorporates SoTA language model BERT, which stands for Bidirectional Encoder Representations from Transformers (Devlin et al., 2018). In order to analyze text in multiple languages, we use a recent multilingual extension of BERT: XLM-Roberta (Conneau et al., 2019).

BERT is one of the most successful deep learning language models in natural language processing (NLP). Trained on a large text corpus (i.e., Wikipedia pages and books), the model produces embeddings (i.e., vectors of numbers) to represent the meaning of sentences, taking into consideration the relationships between words and their context. On top of that, XLM-Roberta is further trained on a large corpus of multilingual data, that is, 2.5TB of filtered web data in 100 languages. The vector representations of text in any of the 100 languages the model is pretrained on can then be used to generate insight into any text in the given language. Further discussion of our system’s multilingual capability and other recommendations for best practices can be found in the comprehensive FAQ section of our website.³

Building a deep learning model from scratch is hard because it requires extensive training data. A common approach is to “transfer” insight from a pretrained deep learning model and use it to perform similar tasks on another dataset. This is called transfer learning. With the capability of transferring knowledge from a pretrained model to the

³<http://www.openframing.org/faq.html>

current task, one can build a model with a high level of accuracy even using a small sample of ground truth labels. In short, our system first obtains some knowledge from XLM-Roberta about how to create meaningful vector representations of text in multiple languages, and fine-tunes these representations using the provided human annotations for frame classification⁴. Our system implements 5-fold cross-validation and provides three evaluation scores—precision, recall, and F-score—to assess the performance of the trained model. These are sent via an email to the user once training is done, together with the ID of the model trained on the user’s entire labeled data that the user can use for frame prediction in Step 5. As for our gun violence case study, we use the above-discussed approach. Based on 5-fold cross-validation, the model to predict frames in the English news headlines reached 0.83 accuracy.

Step 5: Predict Frames with Deep Learning

Once the user is satisfied with the average model performance from Step 4, they can upload an unlabeled dataset and the user’s trained model can be used to predict the frames in the dataset. Our system also provides four English pretrained models on topics of gun violence trained on the gun violence frame corpus (Liu et al., 2019; Akyürek et al., 2020; Guo et al., 2021), and immigration, tobacco, and same-sex marriage trained on the media frame corpus (Card et al., 2015; Field et al., 2018).

We use the gun violence model trained from Step 4—which obtains a 5-fold cross-validation accuracy of 0.83—to predict the frames of the remaining English news headlines about the U.S. gun violence issue from different years. Figure 4 visualize the results, which can also be accessed at the demo part of our website <http://www.openframing.org/demos.html>. It is clear that the volume of coverage increased after each mass shooting case (Figure 4) and overall, the issue was largely politicized in the media discourse. The demo also allows comparison of conservative and liberal-leaning news coverage in the U.S., where the former emphasized the mental health frame more than so the latter. With results of news frames like the ones demonstrated here, users can run additional statistics to compare news framing strategies across different societies, or different types of news media within a

⁴We use one set of training parameters recommended for BERT: a learning rate of 5e-5, 3 epochs of fine-tuning, and a batch size of 8




Figure 4: Screenshot of part of the web page for the gun violence study demo, showing 2018 headline frame predictions.

certain society.

4 Discussion and Conclusion

As we have argued, applying SoTA computational methods to framing analysis can make a significant methodological contribution to the field. However, the implementation involves at least two challenges: the lack of computational resources and skills. Deep learning models such as XLM-Roberta, are extremely large ML models with millions of parameters to train. Even fine-tuning them requires computers with GPU compute capability, which is expensive and not widely available. Some cloud services such as Google Colab provide free but limited access to GPUs. The intensive computational requirements for running large ML models and the unequal access to these computing resources among researchers contribute to the digital and compute divide (Strubell et al., 2019).

For communication researchers, the divide is exacerbated due to the shortage of computational research skills. Implementing and fine-tuning topic modeling or deep learning models such as the ones introduced here requires the understanding and a considerable degree of comfortability with using a deep learning programming framework such as Pytorch, machine learning libraries such as scikit-learn for the training and evaluation setup and analysis, NLP libraries such as NLTK to clean and pre-process the text, as well as Python programming that is required for using these frameworks and libraries. Users with little experience in computer science or programming would find it challenging to run the computational analysis on their own.

With the support of a cross-disciplinary team, our system aims to make the computational fram-

ing analysis accessible to researchers with no or little prior experience in computer science and programming. Through a click-and-play web-based system, the users can follow the guidance on the website and run the advanced computational analysis step by step. Users with different levels of expertise would know where to start and how to interact with the system. Also unlike many of the similar applications in the market, our system prioritizes transparency in its data processing and algorithms. The tool is entirely open-sourced and users will have access to the raw code on our [Github](#) page.

5 Future Directions

In the future, we are extending our system to support multilingual framing analysis on noisy or informal text, such as those present in social media posts, using recent methods such as [Wibowo et al. \(2021\)](#) to convert informal to formal text prior to doing framing analysis. Other interesting future directions include extending our system to support computational multimodal (i.e., text and image) framing analysis. As journalists have been using both text and images to frame news stories ([Messaris and Abraham, 2001](#); [Coleman and Wu, 2015](#); [Dan, 2017](#); [Powell et al., 2015](#)), text and images have worked together to create a holistic perception of news and hence must be considered together when analyzing news frames ([Wessler et al., 2016](#)). Although such use of multimodal inputs has been explored in many NLP tasks such as multimodal machine translation ([Specia et al., 2016](#); [Hewitt et al., 2018](#); [Khani et al., 2021](#)) and vision-language tasks such as multilingual image retrieval or captioning ([Kim et al., 2020](#); [Burns et al., 2020](#); [Rasooli et al., 2021](#)), there is not yet a computational tool that can support multimodal framing. Furthermore, in addition to communication scholars benefiting from such tool that can analyze, on large scale, images and headlines in tandem for frames, newsroom editors would also benefit from tools that can identify images that help depict the main thrust of the story’s focus ([Caple, 2010](#)). Such tools do not yet exist, and a system that can support multimodal framing will be able to address this need.

Acknowledgments

This work is supported in part by the U.S. NSF grant 1838193, DARPA HR001118S0044 (the LwLL program), and the Department of the Air Force FA8750-19-2-3334 (Semi-supervised Learn-

ing of Multimodal Representations). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes. The views and conclusions contained in this publication are those of the authors and should not be interpreted as representing official policies or endorsements of DARPA, the Air Force, and the U.S. Government.

References

- Afra Feyza Akyürek, Lei Guo, Randa Elanwar, Prakash Ishwar, Margrit Betke, and Derry Tanti Wijaya. 2020. [Multi-label and multilingual news framing analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8614–8624, Online. Association for Computational Linguistics.
- Thomas A Birkland and Regina G Lawrence. 2009. Media framing and policy change after columbine. *American Behavioral Scientist*, 52(10):1405–1425.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Danah Boyd and Kate Crawford. 2012. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society*, 15(5):662–679.
- Andrea Burns, Donghyun Kim, Derry Wijaya, Kate Saenko, and Bryan A Plummer. 2020. Learning to scale multilingual representations for vision-language tasks. In *European Conference on Computer Vision*, pages 197–213. Springer.
- Helen Caple. 2010. [What you see and what you get: The evolving role of news photographs in an Australian broadsheet](#). *Journalism and Meaning-Making: Reading the Newspaper*, pages 199–220.
- Dallas Card, Amber Boydston, Justin H Gross, Philip Resnik, and Noah A Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444.
- Renita Coleman and Denis Wu. 2015. *Image and emotion in voter decisions: The affect agenda*. Lexington Books.
- Elanor Colleoni, Alessandro Rozza, and Adam Arvidsson. 2014. Echo chamber or public sphere? predicting political orientation and measuring political homophily in twitter using big data. *Journal of communication*, 64(2):317–332.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco

- Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Viorela Dan. 2017. *Integrative framing analysis: Framing health through words and visuals*. Routledge.
- Paul D’Angelo. 2018. *Doing news framing analysis II: Empirical and theoretical perspectives*. Routledge.
- Frederik De Grove, Kristof Boghe, and Lieven De Marez. 2020. (what) can journalism studies learn from supervised machine learning? *Journalism Studies*, 21(7):912–927.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Kevin M Drakulich. 2015. Explicit and hidden racial bias in the framing of social problems. *Social Problems*, 62(3):391–418.
- Robert M Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of Communication*.
- Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. Framing and agenda-setting in russian news: a computational analysis of intricate political strategies. *arXiv preprint arXiv:1808.09386*.
- Lei Guo. 2018. Ignorance or uncertainty: How the “black box” dilemma in big data research may “misinform” political communication. In *Digital Discussions*, pages 50–67. Routledge.
- Lei Guo, Kate Mays, Yiyan Zhang, Derry Wijaya, and Margrit Betke. 2021. What makes gun violence a (less) prominent issue? a computational analysis of compelling arguments and selective agenda setting. *Mass communication and society*, pages 1–25.
- Lei Guo, Chris J Vargo, Zixuan Pan, Weicong Ding, and Prakash Ishwar. 2016. Big social data analytics in journalism and mass communication: Comparing dictionary-based text analysis and unsupervised topic modeling. *Journalism & Mass Communication Quarterly*, 93(2):332–359.
- Felix Hamburg. 2020. Media bias, the social sciences, and nlp: Automating frame analyses to identify bias by word choice and labeling. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 79–87.
- John Hewitt, Daphne Ippolito, Brendan Callahan, Reno Kriz, Derry Tanti Wijaya, and Chris Callison-Burch. 2018. Learning translations via images with a massively multilingual image dataset. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2566–2576.
- Nikzad Khani, Isidora Tourni, Mohammad Sadegh Rasooli, Chris Callison-Burch, and Derry Tanti Wijaya. 2021. Cultural and geographical influences on image translatability of words across languages. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 198–209.
- Donghyun Kim, Kuniaki Saito, Kate Saenko, Stan Sclaroff, and Bryan Plummer. 2020. Mule: Multimodal universal language embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11254–11261.
- Klaus Krippendorff. 2004. Content analysis: An introduction to its methodology.
- Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- Fabienne Lind, Jakob-Moritz Eberl, Tobias Heidenreich, and Hajo G Boomgaarden. 2019. Computational communication science when the journey is as important as the goal: A roadmap to multilingual dictionary construction. *International Journal of Communication*, 13:21.
- Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019. Detecting frames in news headlines and its application to analyzing news framing trends surrounding us gun violence. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 504–514.
- Daniel Maier, Annie Waldherr, Peter Miltner, Gregor Wiedemann, Andreas Niekler, Alexa Keinert, Barbara Pfetsch, Gerhard Heyer, Ueli Reber, Thomas Häussler, et al. 2018. Applying lda topic modeling in communication research: Toward a valid and reliable methodology. *Communication Methods and Measures*, 12(2-3):93–118.
- Andrew Kachites McCallum. 2002. Mallet: A machine learning for language toolkit. [Http://mallet.cs.umass.edu](http://mallet.cs.umass.edu).
- Paul Messaris and Linus Abraham. 2001. *The role of images in framing news stories*. In *Framing public life*, pages 231–242. Routledge.
- W Russell Neuman, Russell W Neuman, Marion R Just, and Ann N Crigler. 1992. *Common knowledge: News and the construction of political meaning*. University of Chicago Press.
- Rico Neumann and Shahira Fahmy. 2012. Analyzing the spell of war: A war/peace framing analysis of the 2009 visual coverage of the sri lankan civil war in western newswires. *Mass Communication and Society*, 15(2):169–200.
- Matthew C Nisbet. 2010. Knowledge into action: Framing the debates over climate change and poverty. In *Doing news framing analysis*, pages 59–99. Routledge.

- Thomas E Powell, Hajo G Boomgaarden, Knut De Swert, and Claes H de Vreese. 2015. [A clearer picture: The contribution of visuals and text to framing effects](#). *Journal of Communication*, 65(6):997–1017.
- Mohammad Sadegh Rasooli, Chris Callison-Burch, and Derry Tanti Wijaya. 2021. "wikily" neural machine translation tailored to cross-lingual tasks. *arXiv preprint arXiv:2104.08384*.
- Stephen D Reese. 2007. The framing project: A bridging model for media research revisited. *Journal of communication*, 57(1):148–154.
- Stephen D Reese, Jr Gandy, and August E Grant. 2001a. Prologue—framing public life: A bridging model for media research. In *Framing public life*, pages 23–48. Routledge.
- Stephen D Reese, Oscar H Gandy Jr, and August E Grant. 2001b. *Framing public life: Perspectives on media and our understanding of the social world*. Routledge.
- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. 2019. Social bias frames: Reasoning about social and power implications of language. *arXiv preprint arXiv:1911.03891*.
- Frauke Schnell, Karen Callaghan. 2001. Assessing the democratic debate: How the news media frame elite policy discourse. *Political communication*, 18(2):183–213.
- Holli A Semetko and Patti M Valkenburg. 2000. Framing european politics: A content analysis of press and television news. *Journal of communication*, 50(2):93–109.
- P. Shoemaker and Stephen D. Reese. 1995. Mediating the message: Theories of influences on mass media content.
- Lucia Specia, Stella Frank, Khalil Sima'an, and Desmond Elliott. 2016. A shared task on multimodal machine translation and crosslingual image description. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 543–553.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in nlp. *arXiv preprint arXiv:1906.02243*.
- Wouter van Atteveldt and Tai-Quan Peng. 2018. When communication meets computation: Opportunities, challenges, and pitfalls in computational communication science. *Communication Methods and Measures*, 12(2-3):81–92.
- Hartmut Wessler, Antal Wozniak, Lutz Hofer, and Julia Lück. 2016. [Global multimodal news frames on climate change: A comparison of five democracies around the world](#). *The International Journal of Press/Politics*, 21(4):423–445.
- Haryo Akbarianto Wibowo, Made Nindyatama Nityasya, Afra Feyza Akyürek, Suci Fitriany, Alham Fikri Aji, Radityo Eko Prasajo, and Derry Tanti Wijaya. 2021. Indocollex: A testbed for morphological transformation of indonesian word colloquialism. *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*.