OpenFraming: We brought the ML; you bring the data. Interact with your data and discover its frames

Alyssa Smith*, **David Assefa Tofu***, **Mona Jalal***, **Edward E. Halim, Yimeng Sun, Vidya Akavoor, Margrit Betke, Prakash Ishwar, Lei Guo, and Derry Wijaya**

Boston University

* asmithh@alum.mit.edu, {davidat, jalal}@bu.edu

Abstract

When journalists cover a news story, they can cover the story from multiple angles or perspectives. A news article written about COVID-19 for example, might focus on personal preventative actions such as mask-wearing, while another might focus on COVID-19's impact on the economy. These perspectives are called "frames," which when used may influence public perception and opinion of the issue. We introduce a Web-based system for analyzing and classifying frames in text documents. Our goal is to make effective tools for automatic frame discovery and labeling based on topic modeling and deep learning widely accessible to researchers from a diverse array of disciplines. To this end, we provide both state-of-the-art pre-trained frame classification models on various issues as well as a user-friendly pipeline for training novel classification models on user-provided corpora. Researchers can submit their documents and obtain frames of the documents. The degree of user involvement is flexible: they can run models that have been pre-trained on select issues; submit *labeled* documents and train a new model for frame classification; or submit *unlabeled* documents and obtain potential frames of the documents. The code making up our system is also open-sourced and well documented, making the system transparent and expandable. The system is available online at http://www.openframing.org and via our GitHub page https://github.com/davidatbu/openFraming.

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, researchers have employed computational methods based on both unsupervised and supervised machine learning (ML) techniques. This use of artificial intelligence (AI) has enabled users to detect frames automatically and robustly (Akyürek et al. 2020; Liu et al. 2019; Tsur, Calacci, and Lazer 2015). These state-of-the-art AI tools, however, are not readily accessible to social sciences scholars who typically do not have machine learning training. This hampers their ability to glean valuable insights from unprecedentedly large media datasets. Communication scholars without an AI background generally cannot benefit from automatic AI tools that support their analysis of media framing, an important research objective, since framing defines how news media coverage shapes mass opinion.

Our goal is to make AI-based framing analysis accessible to researchers from a diverse array of disciplines. We present OpenFraming (www.openframing.org), a user-friendly and interactive 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-run commercial software, but these tools often pose issues for researchers by their lack of transparency into their inner computational mechanisms. In contrast, our system is based on state-of-the-art research and our code is publicly available. While the focus of the project is on news media framing, the proposed system can also be used to implement other tasks such as sentiment detection or process other data types like social media data.

Specifically, our proposed OpenFraming system can perform two types of AI-based framing analysis:

1. Unsupervised topic modeling based on Latent Dirichlet Allocation (LDA; Blei, Ng, and Jordan (2003)), and

* These authors contributed equally.

2. Supervised learning using deep neural network Bidirectional Encoder Representations from Transformers (BERT; Devlin et al. (2018)).

Both approaches have been applied to media framing in communication research and are proven to be efficient and valid (Guo et al. 2016; Liu et al. 2019).

When encountering a large set of unknown media data, researchers can employ the LDA-based approach to make sense of the data inductively (Guo et al. 2016). Using the LDA output, researchers can find the main threads of discourse in a corpus by examining the LDA "topics" associated with keywords that are most indicative of that particular thread of discourse. Ultimately, the "topics" can prove useful for frame discovery. However, the LDA output may not produce a useful framing model on its own. Because the method is unsupervised, the "topics" it creates may overlap with each other; appear to be irrelevant to the phenomenon being studied; or seem so ill-defined to the trained researcher that the results would not contribute to the framing literature.

Given the limitations of unsupervised topic modeling, our system also provides an alternative approach that allows domain experts (i.e., the users such as communication scholars) to intervene in building the framing model. In this setting, the user can first employ the LDA-based approach to discover potential frames in the corpus. Then, using their domain-specific knowledge, they can manually label and upload a dataset to the system with frames suggested by the LDA model or uncovered from other explorations, whether machine-guided or not. We employ a BERT-based classification model to create a state-of-the-art frame classifier. Researchers can upload unlabeled documents, e.g., large media corpora, to wwww.openframing.org and use the trained classifier to extract the frames.

To summarize, our system OpenFraming has the following advantages:

1. OpenFraming can process textual media data and detect frames automatically.
2. OpenFraming is accessible to researchers without computational backgrounds.
3. OpenFraming produces valid media frames based on peer-reviewed, state-of-the-art computational models.
4. OpenFraming provides many options for users to perform unsupervised ML, supervised ML, or both. In the supervised setting, the model trained on user-provided labeled data can be used to label a much larger dataset than would be feasible for human workers.

Related Work

A typical task in the field of communication research is the identification of topics, attributes, and frames in document collections to understand, for example, news media messages, elite discourse, and public opinion. Traditionally, scholars rely on content analysis approaches, both qualitative and quantitative, to manually annotate the data (Krippendorff 2018; Lindlof and Taylor 2017). In recent years, a group of communication researchers has taken advantage of advances in computational sciences and applied both unsupervised and supervised ML, to analyze large-scale communication text. In light of the growing importance of media and communication in our lives concerning agenda setting, framing, and biases, more and more computer scientists also joined this line of research and consider media framing to be a domain to apply their algorithms (Tsur, Calacci, and Lazer 2015; Field et al. 2018; Liu et al. 2019; Akyürek et al. 2020; Hamborg 2020; Sap et al. 2019).

Within the world of unsupervised ML for text analysis, LDA-based topic modeling is one of the most widely used approaches in communication research (see Maier et al. (2018) for a systematic review). The LDA algorithm generates a set number of "topics" associated with a list of terms. Researchers then review the terms and decide the label for each topic. Consider the news coverage of COVID-19 as an example. An LDA topic may include the terms *pandemic*, *job*, *million*, *economy*, and *unemployment*, which can be labeled as the topic "economic consequences". Another topic may include the terms *season*, *player*, *sport*, *game*, and *return*, and can be labeled as "the impact on the sports industry". Guo et al. (2016) made the first attempt to assess the efficacy and validity of the LDA-based approach in the context of journalism and mass communication research; furthermore, they prove that it is useful and efficient to obtain initial ideas about the data.

Since a frame, explicitly defined, is "a central organizing idea for news content that supplies a context and suggests what the issue is through the use of selection, emphasis, exclusion, and elaboration" (Reese, Gandy Jr, and Grant 2001), LDA-generated topics related to frames may elide the abstraction and nuance that the frames themselves contain. Framing scholars have identified a list of *generic* and *issue-specific* frames and argued that framing analysis should be built on the existing work to make a meaningful contribution to the literature (Guo, Holton, and Jeong 2012; Nisbet 2010; Semetko and Valkenburg 2000). This suggests that not all LDA-generated topics can be productively considered as frames.

Using the running example of the COVID-19 coverage, while the LDA topic "economic consequences" corresponds to one of the generic frames identified earlier, it is debatable whether the topic discussing the impact on the sports industry can be interpreted as a frame. The LDA-based approach has other imperfections as well: it may generate meaningless "topics" or produce "topics" that contain unrelated or even conflicting information. 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 infer potential frames for the next step of supervised frame analysis.

Unlike unsupervised ML, the supervised approach is a deductive research method and is used to identify predetermined frames based on the literature. In communication research, scholars have used supervised ML algorithms such as support vector machines and deep learning models to identify frames in a media text. Two recent studies used BERT to identify frames in the news coverage of gun violence in the US; the studies both demonstrate a high level of accuracy (Akyürek et al. 2020; Liu et al. 2019).

The implementation of both unsupervised ML and supervised ML discussed above requires a computational background. Some social science scholars explore the methods themselves, and others choose to collaborate with colleagues in computer science. However, due to a lack of formal computer science training, it is often difficult for social science scholars to apply the computational models appropriately on their own. Also, not all scholars have the opportunities and resources for cross-disciplinary collaboration. Commercial software programs exist for this type of analysis, but most are costly and the algorithms they provide remain a black box. To overcome these challenges, we present OpenFraming, a free and open-sourced Web-based system specialized in AI-based framing analysis.

System Architecture

We made a Web server, www.openframing.org that runs OpenFraming publicly available. It was set up on an EC2 instance on Amazon Web Services (AWS) with minimal additional configuration. We also enable users to run a copy of the system locally. This is possible through our release of a Docker container that orchestrates the various technologies used by our system. Anyone ranging from the user who would like to have their own version of the system on their personal computers, to bigger organizations who would like to host and extend the system on more capable hardware, can get OpenFraming up and running in minutes.

The software that makes up the system includes Gensim's (Řehůřek and Sojka 2010) Python interface to Mallet (McCallum 2002) for LDA topic modeling; the transformers library for supervised classification (Wolf et al. 2019), Redis for queuing the jobs, SQLite for a database solution, Flask for the Web application backend, and jQuery and Bootstrap for the frontend.

Data Cleaning and Pre-processing for the LDA While there is some flexibility regarding the format of the dataset (the system currently supports .xls, .xlsx, and .csv), it is nonetheless necessary that it at least contain a column labeled as "Example." This column will hold the text examples, with one document or, broadly speaking, textual entity, per row. LDA employs a bag-of-words model, where each document is understood as an unordered collection of words; to make the analysis more conducive to the discovery of useful topics, the system filters out extremely common and extremely rare words. The pre-processing steps we employ include the following:

- **Removing punctuation and digits.** this is a standard step in natural language processing (NLP) applications.
- **Removing the stopwords**: stopwords are extremely common words, usually filtered out by default in NLP applications.
- **Lemmatizing the content**: this groups together different inflected forms of a word into a single entity.
- **Setting minimum word length**: the system removes words shorter than 2 characters.

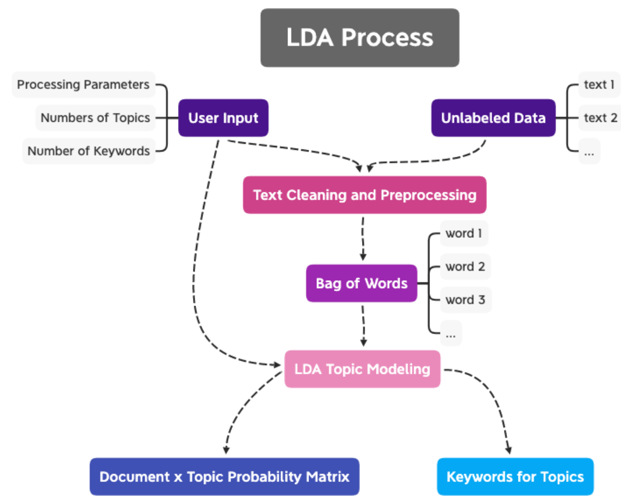

Figure 1: LDA pipeline for topic discovery

LDA for topic discovery The system runs LDA using the Mallet (McCallum 2002) implementation and its preset parameter tuning. The random seed is set deterministically so that subsequent runs of the algorithm will yield the same results. LDA models each document as a probabilistic mixture of topics. A topic is defined as a probability distribution over keywords. LDA iteratively updates the topic-keyword distributions to maximize the log-likelihood of the entire corpus. The system uses LDA to create a matrix mapping documents to weight vectors which quantify the contribution (weight) of each topic to the document; we can think of the weight vector as a probability distribution over topics for a particular document. Our system also produces a list of the most relevant keywords for each topic; the user can specify how many keywords they would like to be given before runtime. Because running the LDA over a large corpus can be time-consuming, the user's part in monitoring the modeling finishes when they hit the "submit" button. The system then sends them an e-mail with a link to download the results of the analysis when it is ready. We also provide topic quality metrics, namely coherence, and perplexity, to aid researchers in refining the number of topics they choose to use to further analysis. Figure 1 provides a more detailed explanation of the LDA pipeline.

Labeling Procedure For the LDA Results When the LDA algorithm has completed, the user will receive its output, which contains a set of "topics", each of which is associated with a list of keywords. Communications researchers recommend that at least two researchers manually review the keywords and decide on a label for each topic. Ideally, for framing analysis, each label should correspond to one of the frames — generic or issue-specific — identified in the relevant literature. New labels may be created to signify topics or frames related to the specific issue. Given the limitations

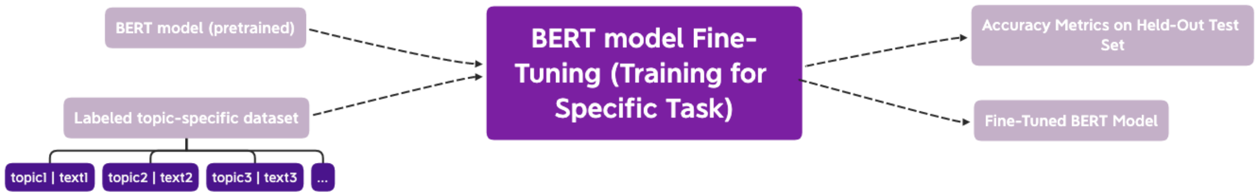

Figure 2: BERT training and fine-tuning pipeline.

of the LDA approach, it is also possible that some "topics" may not be meaningful.

Text classification using BERT BERT's masked language model (Devlin et al. 2018), which builds on a deep Transformer's encoder architecture that relies on multi-layer self-attention to compute contextual representations of its input (Vaswani et al. 2017), has shown impressive performance across a wide range of tasks, including framing analysis (Liu et al. 2019), when fine-tuned on labeled data for the task. However, there remains a significant access barrier for those with a non-computational background to truly make use of BERT's wide-ranging applicability. To our knowledge, all publicly available Web services and software packages that make use of BERT either constrain the end-user to one specific fine-tuned model (for example, fine-tuned on a specific sentiment analysis dataset), or, they require their users to be prepared to write code to fine-tune and further predict on a custom dataset. OpenFraming makes it possible for those without a computational background to take advantage of BERT's impressive fine-tuned performance on a custom dataset of their own.

When the user uploads *labeled* data for training and testing or *unlabeled* data for inference, our system either fine-tunes a new BERT model or uses our existing fine-tuned BERT for classifying the frame labels in the data. For fine-tuning, our system uses the standard configuration of BERT's internal architecture, and uses one set of training parameters recommended for BERT: a learning rate of 5e-5, 3 epochs of fine-tuning training, and a batch size of 8.

Once training or inference are completed, the user receives an e-mail with a download link to the frame prediction results on their data. In the case of fine-tuning a new BERT model on user-provided labeled data, we also provide accuracy on user-provided test data and the newly fine-tuned BERT model that the user can download. Figure 2 provides a more detailed description of the BERT training pipeline.

Labeling Procedure for Training a New BERT Model With the feature "Training BERT – Do-it-yourself method," users can train a new BERT classification model using their own labeled data. In social science research, quantitative content analysis is one of the most widely used methods for labeling visual and textual content (Kripendorff 2004; Riffe et al. 2019). The approach involves drawing a representative sample of data; training two or more human coders on a labeling protocol to identify patterns in content, and measuring inter-coder reliability between their coding results. Once the coders reach a certain degree of inter-coder reliability, they can start labeling the remaining data independently. Communications researchers have recently suggested that crowdsourcing, if appropriately implemented, can be a valid alternative to annotating media messages (Guo et al. 2019; Lind, Gruber, and Boomgaarden 2017). The labeled data can then be uploaded to our system to train a new BERT model.

Available Pre-Trained BERT Models for Frame Classification For the feature "Using BERT – off-the-shelf classifier," users can use models that we have fine-tuned on benchmark frame datasets to classify their unlabeled data. We make available models that can label frames on issues that include (1) immigration, (2) tobacco-use, (3) same-sex marriage (fine-tuned on Media Frame Corpus dataset (Card et al. 2015)), (4) US Gun Violence issue (fine-tuned on Gun Violence Frame Corpus (Liu et al. 2019)), or (5) COVID-19. To validate the performance of our fine-tuned model and the quality of its predictions, users can label a sample of their documents using the aforementioned approaches — quantitative content analysis and crowdsourcing — and compare the manual and machine-generated labels.

User Interface and Site Design

Our demo Website includes framing analysis as well as the LDA topic discovery utilities. Additionally, our landing page provides an introduction to the user explaining what various building blocks of our Website are (Figure 4)

Our framing analysis page (Figure 5) is created to accommodate two use cases. Either the user inputs a file for framing classification and chooses one of the policy issues for which we already have pre-trained models (e.g. *Immigration*), or picks one of the policy issues of their choosing (e.g. *Labor Market Inequality*). If the user chooses their own policy issue for which we don't have a pre-trained model, they are required to also upload a sizable dataset labeled with frames (containing approximately 100 documents for each frame) so that the system can train a new BERT-based framing classifier for the issue in the backend.

Once the backend has completed running inference on the pre-defined and pre-trained policy issues or completed the training and inference on user-defined policy issue, the

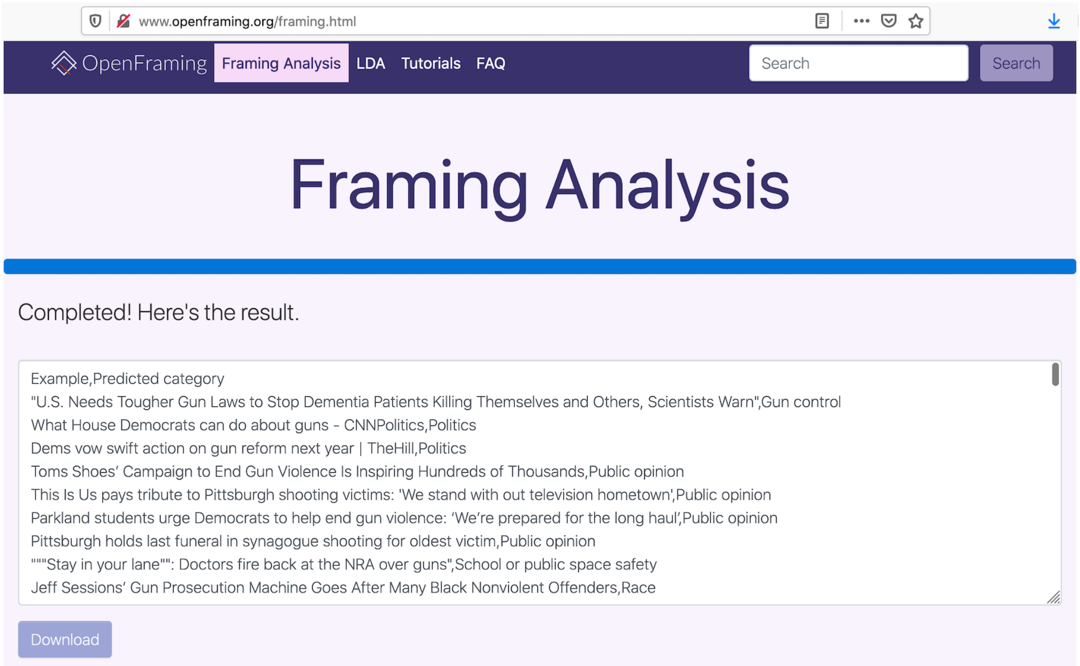

Figure 3: A snapshot of framing classification results: Headlines of news articles about gun violence and predicted frames (here, gun control, politics, public opinion, school safety, race)

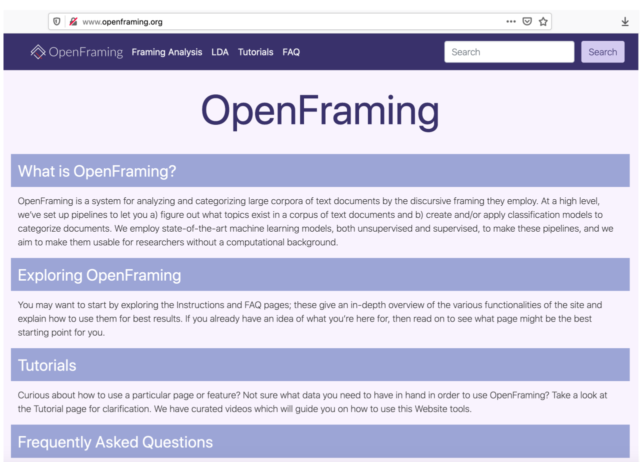

Figure 4: The landing page of openframing.org

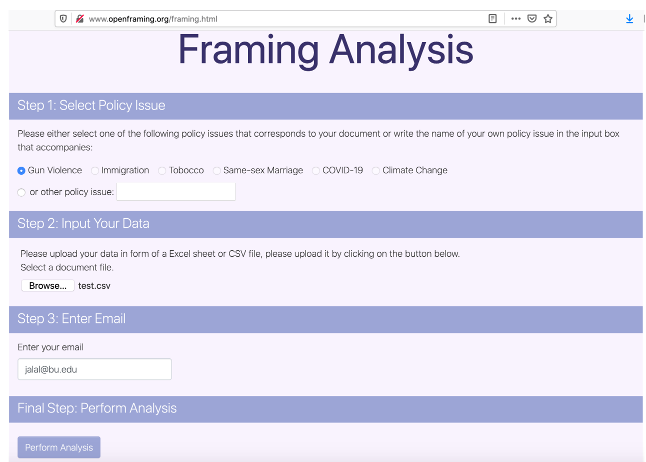

Figure 5: Framing analysis Web page

results will be shown dynamically on the same page (Figure 3). The user can then scroll through the predicted results and download the results to their local machines.

Here, we illustrate the topic discovery functionality of OpenFraming (Figure 7) using a sample from the Kaggle 'A Million News Headline' dataset[1]. Once topics are discovered, we send the topics and their keywords as well as the document topic probabilities to user's provided e-mail (Figure 8 and Figure 6).

We have also created a screencast video demonstrating

[1]https://www.kaggle.com/therohk/million-headlines

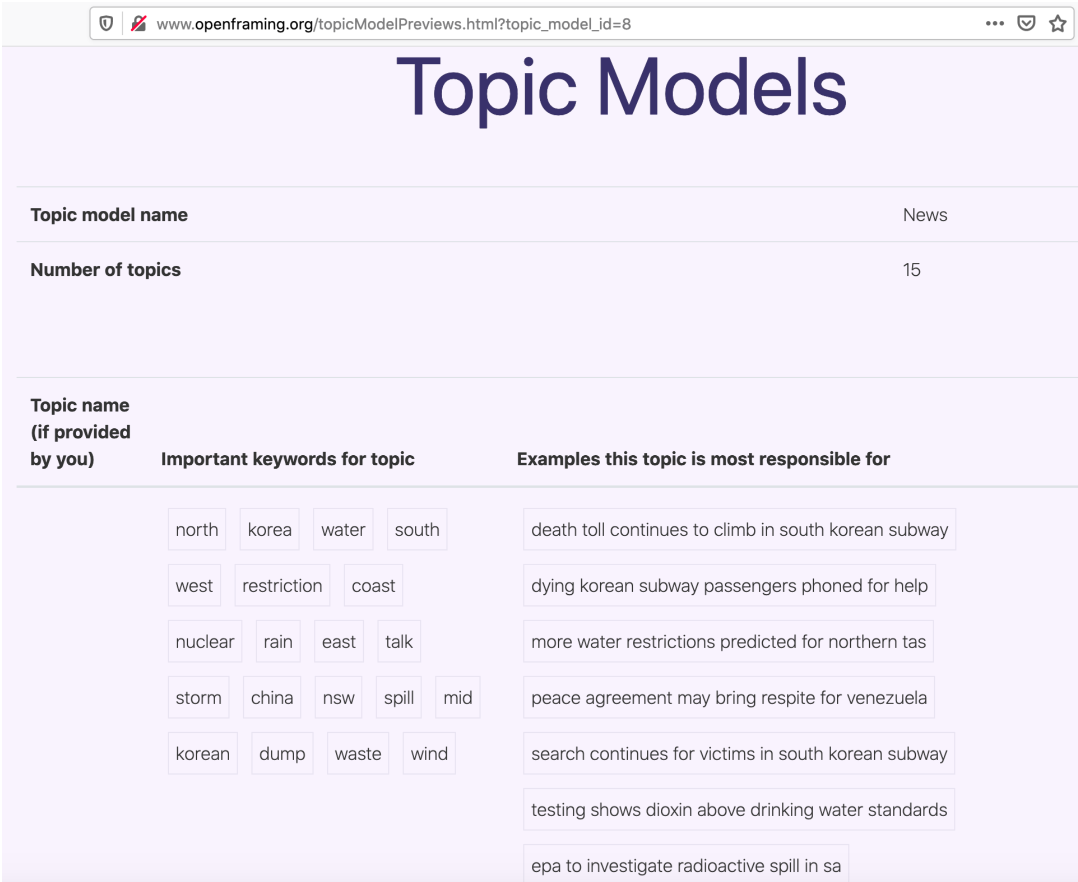

Figure 6: A snapshot of one of the topics discovered by LDA on 'A Million News Headline' dataset, the keywords for the topic, and the headlines labeled with the topic.

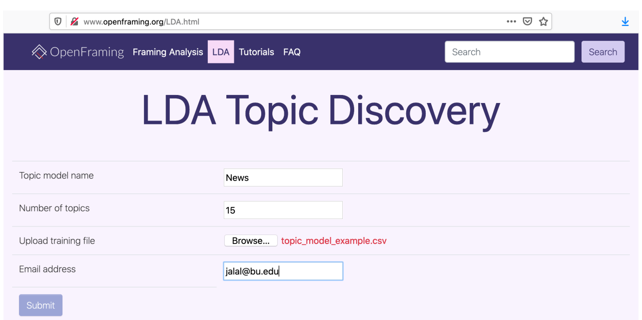

Figure 7: LDA topic discovery page

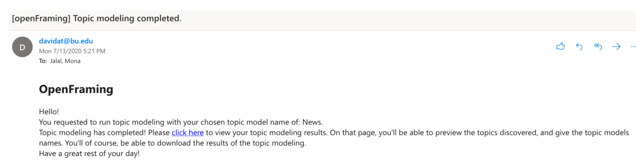

Figure 8: LDA results are ready and e-mailed to the user.

the use of the system, which can be accessed at https://www.youtube.com/watch?v=u8SJAZ-EbgU.

Conclusion and Future Work

We have introduced OpenFraming, a Web-based system for analyzing and classifying frames in the text documents. OpenFraming is designed to lower the barriers to applying machine learning for frame analysis, including giving researchers the capability to build models using their own labeled data. Its architecture is designed to be user-friendly and easily navigable, empowering researchers to comfortably make sense of their text corpora without specific machine learning knowledge.

In future work, we hope to incorporate semi-supervised machine learning methods to allow researchers to iterate quickly on models; if a researcher submits a dataset with a relatively small number of labels, for example, the system will eventually be able to generate labels for the much larger unlabeled dataset, creating a synthetic training set for the BERT supervised model to train on.

References

Akyürek, A. F.; Guo, L.; Elanwar, R.; Ishwar, P.; Betke, M.; and Wijaya, D. T. 2020. Multi-Label and Multilingual News Framing Analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8614–8624. Online: Association for Computational Linguistics. URL https://www.aclweb.org/anthology/2020.acl-main.763.

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3(Jan): 993–1022.

Card, D.; Boydstun, A.; Gross, J. H.; Resnik, P.; and Smith, N. A. 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)*, 438–444.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* .

Drakulich, K. M. 2015. Explicit and hidden racial bias in the framing of social problems. *Social Problems* 62(3): 391–418.

Entman, R. M. 1993. Framing: Toward Clarification of a Fractured Paradigm. *Journal of Communication* .

Field, A.; Kliger, D.; Wintner, S.; Pan, J.; Jurafsky, D.; and Tsvetkov, Y. 2018. Framing and agenda-setting in russian news: a computational analysis of intricate political strategies. *arXiv preprint arXiv:1808.09386* .

Guo, L.; Holton, A.; and Jeong, S. H. 2012. Transnational comparative framing: A model for an emerging framing approach. *International Journal of Communication* 6: 24.

Guo, L.; Mays, K.; Lai, S.; Jalal, M.; Ishwar, P.; and Betke, M. 2019. Accurate, Fast, But Not Always Cheap: Evaluating "Crowdcoding" as an Alternative Approach to Analyze Social Media Data. *Journalism & Mass Communication Quarterly* 1077699019891437.

Guo, L.; Vargo, C. J.; Pan, Z.; Ding, W.; and Ishwar, P. 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.

Hamborg, F. 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*, 79–87.

Kripendorff, K. 2004. Content analysis: An introduction to its methodology.

Krippendorff, K. 2018. *Content analysis: An introduction to its methodology*. Sage publications.

Lind, F.; Gruber, M.; and Boomgaarden, H. G. 2017. Content analysis by the crowd: Assessing the usability of crowdsourcing for coding latent constructs. *Communication methods and measures* 11(3): 191–209.

Lindlof, T. R.; and Taylor, B. C. 2017. *Qualitative communication research methods*. Sage publications.

Liu, S.; Guo, L.; Mays, K.; Betke, M.; and Wijaya, D. T. 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)*, 504–514.

Maier, D.; Waldherr, A.; Miltner, P.; Wiedemann, G.; Niekler, A.; Keinert, A.; Pfetsch, B.; Heyer, G.; Reber, U.; Häussler, T.; 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.

McCallum, A. K. 2002. MALLET: A Machine Learning for Language Toolkit. Http://mallet.cs.umass.edu.

Nisbet, M. C. 2010. Knowledge into action. *Doing news framing analysis: Empirical and theoretical perspectives* 43.

Reese, S. D.; Gandy Jr, O. H.; and Grant, A. E. 2001. *Framing public life: Perspectives on media and our understanding of the social world*. Routledge.

Řehůřek, R.; and Sojka, P. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 45–50. Valletta, Malta: ELRA. http://is.muni.cz/publication/884893/en.

Riffe, D.; Lacy, S.; Fico, F.; and Watson, B. 2019. *Analyzing media messages: Using quantitative content analysis in research*. Routledge.

Sap, M.; Gabriel, S.; Qin, L.; Jurafsky, D.; Smith, N. A.; and Choi, Y. 2019. Social bias frames: Reasoning about social and power implications of language. *arXiv preprint arXiv:1911.03891* .

Semetko, H. A.; and Valkenburg, P. M. 2000. Framing European politics: A content analysis of press and television news. *Journal of communication* 50(2): 93–109.

Tsur, O.; Calacci, D.; and Lazer, D. 2015. A frame of mind: Using statistical models for detection of framing and agenda setting campaigns. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 1629–1638.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.

Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; and Brew, J. 2019. HuggingFace's Transformers: State-of-the-art Natural Language Processing. *ArXiv* abs/1910.03771.