BU-NEmo: An Affective Dataset of Gun Violence News

Carley Reardon, Sejin Paik, Ge Gao, Meet Parekh, Yanling Zhao, Lei Guo, Margrit Betke, Derry Wijaya

Boston University AIEM Research Group

Research Objectives

In our work, we pose the following research questions:

1. How do different multimedia content types (text vs. image) affect the way people emotionally respond to gun-violence-related news?
2. Based on demographic variables, how do people's emotional responses vary to a given multimedia content type?
3. What is the scalability and contribution of a dataset at this intersection?

Introduction

A wide-range of research has arisen in recent years in response to the increasingly large quantities of online, multimodal news content and its emotional impact on viewers. One example of such work includes analyzing the influence of partisanship in affective political news [1]. Within machine learning, the task of affective emotion recognition aims to predict the emotional response that a given piece of media will elicit. Affective emotion recognition could be hugely influential to understanding the emotional impact of hard news content, yet this relationship is largely still uninvestigated.

There is a lack of news media data for this task, with most notable datasets either collecting abstract emotional responses through art, such as in [2], or neglecting to fully annotate the data with detailed, human emotional responses, as in [3].

This study aims to facilitate future research in this intersection by presenting BU-NEmo, a **multimodal dataset of gun violence news images and headlines**, annotated with viewers' elicited emotional responses to the media content. This study focuses on gun violence and gun control news, as these topics continue to be heavily debated in the U.S. despite their devastating impacts. This data is, to our knowledge, the first of its kind to:

1. Offer *high-quality affective annotations* on news data from various news sources.
2. Include both *categorical and written* emotional responses on *multimodal news data.*
3. Explore the *differences in emotional response* when given *various modalities of news.*

Additionally, we make our dataset, along with documentation, publicly available to facilitate further progress in this area of research.

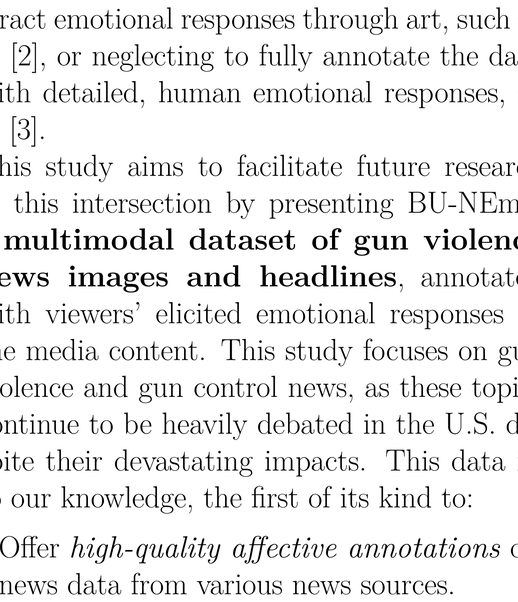

Figure 1:Overview of categorical emotional responses given by experimental condition.

Data

In this study, we make use of the Gun Violence Frame Corpus (GVFC), a gun-violence news headline corpus from [4], later extended by [5] to include the lead images corresponding to each article. GVFC ultimately fits our goals of extending emotion annotations to a multimodal news dataset with a diversity of news sources.

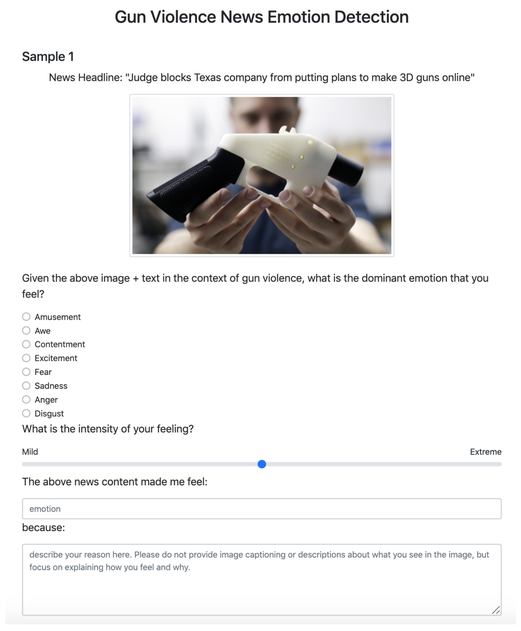

Figure 2:Our annotation procedure for a sample of corresponding news image and headline (*TI* condition).

Annotation Method

To understand the emotional response triggered by news samples, we designed a crowdsourcing survey to collect emotional annotations on our data and deployed it on the BU SONA System in October 2021. Annotators were presented with gun-violence news samples in one of three experimental conditions: only the news headlines (T), only the news images (I), or both the news headline and image (TI) are presented. In each condition, the annotators were required to answer the following questions:

1. Choose the emotional category they feel among Amusement, Awe, Contentment, Excitement, Fear, Sadness, Anger, and Disgust
2. Choose the intensity of their emotional response on a 5-point scale from mild to extreme
3. Describe why they feel the emotion in written text

The eight emotional categories were chosen due to their use in the prior emotional language corpus, ArtEmis [2]. As a result of this process, we collected 15K high-quality, crowdsourced annotations on 840 news articles. An example of the interface we used for the three steps of our annotation procedure can be found in figure 2.

Results

The distribution of emotional responses in each condition can be seen in figure 1 on the left. We observed higher frequency of negative emotions, notably "fear" and "sadness", especially in the I condition. Negative emotions also tended to garner higher intensity scores.

We assessed the consensus of annotators by measuring entropy of the response distributions. Examples of news with low and high response entropy can be found in figure 3.

	Top Emotion	Entropy	Headline
Text Only			
Minimum	Sadness (10/10)	0.0	"The Latest: Victim's family marks anniversary of shooting"
Maximum	Disgust (3/12)	1.9792	"Pittsburgh synagogue-shooting suspect wheeled into courtroom; Trump to visit city Tuesday"
Image Only			
Minimum	Sadness (12/12)	0.0	"California bar massacre leaves another American community reeling after 12 gunned down by Marine veteran"
Maximum	Anger (3/12)	1.9792	"Department of Education considers letting schools buy guns for teachers with federal grant funds"
Text & Image			
Minimum	Sadness (10/10)	0.0	"The Latest: Victim's family marks anniversary of shooting"
Maximum	Awe (4/14)	1.9459	"'Stay in your lane': Doctors fire back at the NRA over guns"

Figure 3:Examples of news samples per experimental condition that yielded the minimum and maximum response entropy.

All three minimum-entropy examples have headlines which tend to center victims/human interest, while the maximum-entropy headlines are more politically-charged. These results offer support for the hypothesis that news framing and content can have an effect on emotional consensus.

We tested the significance of our findings within each observed relationship and present the results in table 1. In each case, we received significantly small p-values. These significance tests are strong support for our hypothesis that the presentation of news media can have a **significant impact** on how viewers experience this content emotionally, even when considering any political predisposition.

Ind. Variable	Dep. Variable	Significance Test	p-value
Condition	Emotion	Chi Square	$1.9873e-14$
Political Leaning	Emotion	Chi Square	0.03534
Condition	Emotional Intensity	Two-way ANOVA	$8.3474e-12$
Emotion	Emotional Intensity	Two-way ANOVA	$4.2729e-46$
Condition	Response Entropy	One-way ANOVA	$5.6743e-05$

Table 1:Results of our significance tests on the relationships in our dataset.

In figure 4, we present some sample written responses, along with their chosen emotion, to a specific sample of gun-violence news. This news set is a concrete example of when the T condition allowed for emotional responses that were more informed by the context of the news story. The news content itself was meant to evoke an action being done, and while the image sends a certain message, knowledge of the policy in the headline and its political context have the potential to inform more readers emotionally.

Example 1		Top 10 Keywords	Response Samples
Shopify bans sales of certain firearms, accessories 	Text	'good', 'access', 'mean', 'decision', 'reduce', 'agree', 'help', 'start', 'actual', 'news'	1. "There is actual action being implemented for gun control." - *Contentment/Happy* 2. "This is saying that change is being made and that companies are preventing people from accessing very dangerous weapons." - *Excitement/Happy* and Excited 3. "It will help to reduce the ownership of guns I believe." - *Contentment/Agree* 4. "Reducing access to buying guns, as part of gun restrictions, makes me feel safer." - *Contentment/Good*
	Image	'trigger', 'type', 'people', 'afraid', 'access', 'sale', 'flag', 'scary', 'shelf', 'country'	1. "Guns and the country flag are always associated together, but they shouldn't be" - *Sadness/Sad* 2. "Guns are scary and those are big guns" - *Fear/Fearful* 3. "These types of guns should not just be sold on shelves" - *Disgust/Disgust* 4. "These types of guns are especially dangerous and shouldn't be accessible to citizens under any circumstances." - *Fear/Scared*
	Text + Image	'good', 'step', 'right', 'thing', 'control', 'buy', 'safety', 'help', 'small', 'purchase'	1. "This is the right thing to do." -*Amusement/joyful* 2. "It's a good move" - *Contentment/Satisfied* 3. "A small step to help gun control" - *Awe/Awe* 4. "We should look for safety rather than sales of guns." - *Contentment/Contentment*

Figure 4:News sample where the T and TI conditions aligned most closely on emotional response.

Conclusion

- The BU-NEmo dataset offers **high-quality** human emotional response annotations on **highly impactful** gun-violence news.
- Our analysis offers support for hypotheses that the inclusion of certain news media can influence the readers' emotional response.
- Our work evidences the necessity of future research to study the relationship between news content, its modalities, and the effect it has on news consumers.

Future Research

In future work, we hope to extend the BU-NEmo dataset with more annotated news data in other topics, such as climate change. Additionally, we plan to extend our current analysis to examine the effect of news framing on emotion. We hope to inform future works in this space by building a vision & language model to more intricately explore the relationship between modality, news framing, and emotion.

References

[1] Ariel Hasell.
Shared emotion: The social amplification of partisan news on twitter.
Digital Journalism, 9(8):1085–1102, 2020.

[2] Panos Achlioptas, Maks Ovsjanikov, Kilichbek Haydarov, Mohamed Elhoseiny, and Leonidas Guibas.
Artemis: Affective language for visual art.
CoRR, abs/2101.07396, 2021.

[3] Wenya Guo, Ying Zhang, Xiangrui Cai, Lei Meng, Jufeng Yang, and Xiaojie Yuan.
Ld-man: Layout-driven multimodal attention network for online news sentiment recognition.
IEEE Transactions on Multimedia, 23:1785–1798, 2021.

[4] Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya.
Detecting frames in news headlines and its application to analyzing news framing trends surrounding U.S. gun violence.
In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 504–514, Hong Kong, China, November 2019. Association for Computational Linguistics.

[5] Isidora Tourni, Lei Guo, Taufiq Husada Daryanto, Fabian Zhafransyah, Edward Edberg Halim, Mona Jalal, Boqi Chen, Sha Lai, Hengchang Hu, Margrit Betke, Prakash Ishwar, and Derry Tanti Wijaya.
Detecting frames in news headlines and lead images in U.S. gun violence coverage.
In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4037–4050, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.

[6] Xueyao Zhang, Juan Cao, Xirong Li, Qiang Sheng, Lei Zhong, and Kai Shu.
Mining dual emotion for fake news detection.
Proceedings of the Web Conference 2021, Apr 2021.

Acknowledgements & Ethical Considerations

This work is supported in part by the U.S. National Science Foundation (NSF), grant 1838193. We ensure we are not knowingly introducing bias to the data nor inflicting any emotional harm on participants or breaching their confidentiality, for which we have obtained IRB exemption approval (protocol #6280X) from Boston University's Charles River Campus on October 29th, 2021.

Contact Information

- AIEM Website: https://sites.bu.edu/aiem/
- BU-NEmo Dataset: https://github.com/Tdrinker/NEmo-dataset
- Email: {reardonc, sejin, wijaya}@bu.edu

BOSTON UNIVERSITY