

TILFA: A Unified Framework for Text, Image, and Layout Fusion in Argument Mining

Qing Zong¹, Zhaowei Wang², Baixuan Xu², Tianshi Zheng², Haochen Shi²,
Weiqi Wang², Yangqiu Song², Ginny Y. Wong³, Simon See³

¹Harbin Institute of Technology (Shenzhen), Guangdong, China

²Department of Computer Science and Engineering, HKUST, Hong Kong SAR, China

³NVIDIA AI Technology Center (NVAITC), NVIDIA, Santa Clara, USA

zongqing0068@gmail.com, {bxuan, tzhengad, hshiah}@connect.ust.hk,
{zwanggy, wwangbw, yqsong}@cse.ust.hk

Abstract

A main goal of Argument Mining (AM) is to analyze an author’s stance. Unlike previous AM datasets focusing only on text, the shared task at the 10th Workshop on Argument Mining introduces a dataset including both text and images. Importantly, these images contain both visual elements and optical characters. Our new framework, **TILFA**¹ (A Unified Framework for Text, Image, and Layout Fusion in Argument Mining), is designed to handle this mixed data. It excels at not only understanding text but also detecting optical characters and recognizing layout details in images. Our model significantly outperforms existing baselines, earning our team, **KnowComp**, the **1st** place in the leaderboard² of Argumentative Stance Classification subtask in this shared task.

1 Introduction

Argument mining (AM) aims to automatically analyze the structure and components of arguments in text. Persuasiveness analysis is a crucial aspect of it, which has gained significant attention in the NLP community (Habernal and Gurevych, 2017; Carlile et al., 2018). However, previous works focus solely on text, overlooking other modalities like images which can also impact an argument’s persuasiveness. To fill this gap, Liu et al. (2022) introduces **ImageArg**, a dataset going beyond text to include also images. It features tweets centered on contentious topics like gun control and abortion. These associated images contain not only objects but also optical characters (e.g., slogans, tables).

The 10th Workshop on Argument Mining in EMNLP 2023 introduces a shared task (Liu et al., 2023b) called ImageArg Shared Task 2023, centering around this dataset. It is divided into two subtasks: Argumentative Stance (AS) Classification and Image Persuasiveness (IP) Classification.

¹The code and data are available at <https://github.com/HKUST-KnowComp/TILFA>.

²<https://imagearg.github.io/>

Tweet Image:



Tweet Text: Background checks save lives. Full stop. We deserve #MoreThanThoughtsAndPrayers from our elected leaders to end gun violence.

Topic: gun control
Stance: support

Tweet Image:



Tweet Text: A child deserves a chance at life. A child deserves a future. A child deserves to be loved. We will always fight for the innocent unborn. #ProLife

Topic: abortion
Stance: oppose

Figure 1: Examples of positive (support) and negative (oppose) tweets of different topics. The images also contain a lot of information crucial to stance identification.

We primarily focus on the former, which aims to identify the stance of a given tweet towards a specific topic. Examples can be found in Fig. 1.

After scrutinizing the dataset, we found several challenges: (1) Imbalanced data distribution (Ratio of positive to negative examples on abortion topic is about 1 : 2.65); (2) Limited data size (Neither of the two topics has more than 1000 entries); (3) Presence of both objects and optical characters in images (They are difficult to be handled by a single model at the same time). To address these challenges, we have made the following contributions:

- To tackle data imbalance, we employ back-translation to enrich data in the fewer class, as described by Yu et al. (2018); Wieting and Gimpel (2018).
- For data augmentation, we utilize WordNet (Miller, 1994) with GlossBERT (Huang et al., 2019) to create additional data by replacing synonyms of nouns in original instances.
- We introduce **TILFA** which can understand both text and image well, especially adept at detecting optical characters and discerning layout details in images.

2 Related work

Data Augmentation: Data augmentation enhances a model’s performance and increases its generalization capabilities in Natural Language Processing (NLP). At the word level, Wang et al. (2019) used databases like WordNet (Miller, 1994), to replace certain words with their synonyms, while Rizos et al. (2019) implemented embedding replacement to find contextually fitting words; At the document level, Yu et al. (2018) used back-translation, translating data to another language and then back to the source one.

Document AI: Document AI refers to the extraction and comprehension of information from scanned documents, web pages, ads, posters, or images with textual content. Previous works like Hao et al. (2016); Liu et al. (2019) all missed the integrated pre-training of text and layout details, which are vital for document image comprehension. To fix this, Xu et al. (2020) proposed LayoutLM. Its updated versions, LayoutLMv2 (Xu et al., 2021) and LayoutLMv3 (Huang et al., 2022), encapsulated text, layout and also image interactions within a unified multimodal framework.

3 Methods

We employ back-translation (Yu et al., 2018) to address data imbalance and apply WordNet (Miller, 1994) for data augmentation, assisted by GlossBERT (Huang et al., 2019). Our new framework, **TILFA**, uses DeBERTa (He et al., 2021) as the text encoder and LayoutLMv3 (Huang et al., 2022) as the image encoder, thus excels at not only understanding text but also detecting optical characters and recognizing layout details in images. We also experiment with several multimodal fusion mechanisms. Consequently, we achieve the **highest** F1-score in the Argumentative Stance Classification subtask of the ImageArg Shared Task 2023.

3.1 Addressing Data Imbalance

Label imbalance is serious in the training set, particularly concerning the abortion topic (The ratio of positive to negative examples is about 1 : 2.65). To address this, we preprocess the data through back-translation (Yu et al., 2018). We translate the English tweet text belonging to the underrepresented label (e.g., positive in abortion topic) to a random language (e.g., French, German) and then back to English. This maintains the tweet’s mean-

ing and thus the stance label. The translated text is finally paired with its original image.

3.2 Data Augmentation

More data usually leads to better model performance (Bayer et al., 2022; Fang et al., 2022). We employ data augmentation methods since our ImageArg training set is limited: only 918 entries for gun control and 888 for abortion. We first utilize spaCy to tokenize the tweet text and extract the nouns in it. Then we find all their synonym sets in WordNet (Miller, 1994). We determine these nouns’ meanings in context by Word Sense Disambiguation (WSD) using GlossBERT (Huang et al., 2019), thus getting their correct synonym set. Finally, we replace these nouns with their synonyms to create new data.

3.3 Model

To solve this task, we introduce a model: **TILFA** (A Unified Framework for **T**ext, **I**mage, and **L**ayout **F**usion in **A**rgument **M**ining). The structure of **TILFA** is illustrated in Fig. 2, comprising three components: Image Encoder, Text Encoder, and Multimodal Fusion. We will discuss details of them one by one.

Image Encoder: As highlighted by Liu et al. (2022), traditional image encoders like ResNet50, ResNet101 (He et al., 2016), VGG (Simonyan and Zisserman, 2015) are good at identifying objects but fall short in recognizing optical characters in images, which may hurt performance. However, as shown in Fig. 1, many of the images in our dataset contain significant amount of characters. So we reasonably believe that using models which can capture the text in the images will have better results. And we notice that those characters with more prominent position and larger size are relatively more important. So considering the importance of this layout information of characters in images, we employ LayoutLMv3 (Huang et al., 2022) to encode the images, favoring it over sole OCR tools.

Text Encoder: To encode tweet texts, we employ DeBERTa (He et al., 2021). It has shown great performance in various NLP tasks. Our experimental results confirm its effectiveness in the ImageArg task as well.

Multimodal Fusion: We explore three multimodal fusion methods, illustrated in Fig. 2. The

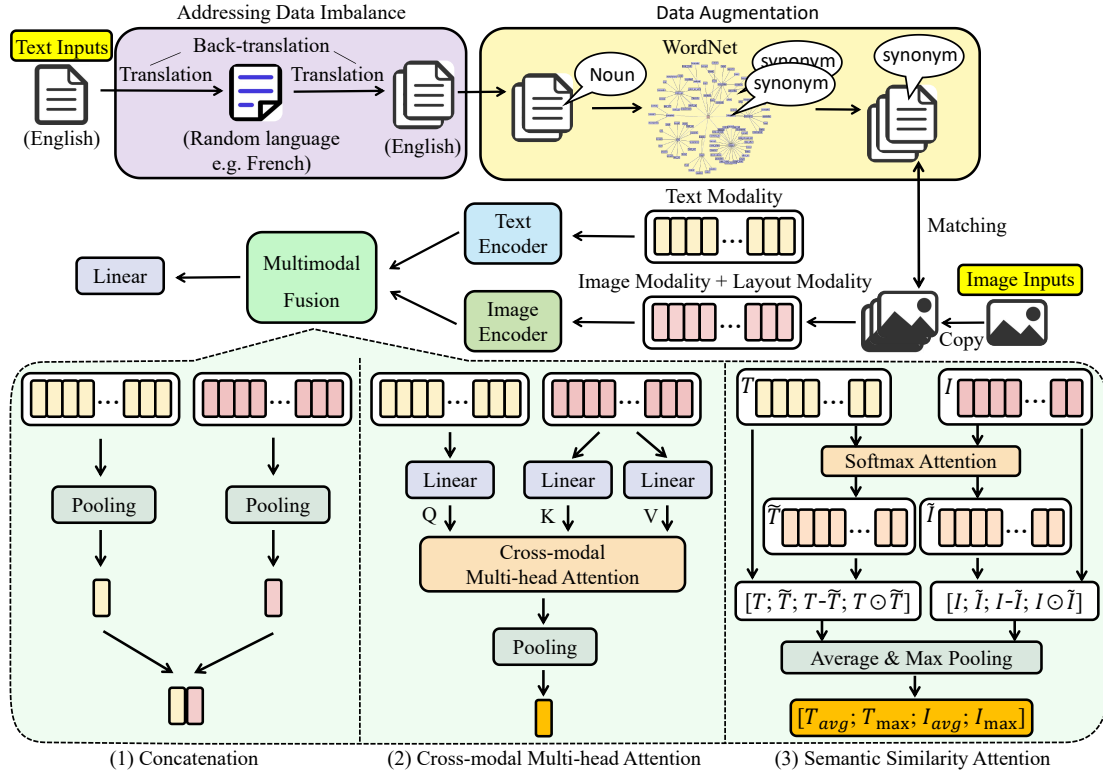


Figure 2: Our model, **TILFA**, includes three main parts: a text encoder, an image encoder, and a multimodal fusion module. In this fusion module, we experiment with three different methods: (1) Concatenation; (2) Cross-modal Multi-head Attention; (3) Semantic Similarity Attention.

first simply concatenates the hidden states from text and image inputs. The second method, named cross-modal multi-head attention, is adapted from Yu et al. (2021). And the third is a new approach adapted from ESIM (Chen et al., 2017). ESIM is a sequential natural language inference model used to predict the logic relationship between two sentences. We adapt it for text and image fusion, and name our version: Semantic Similarity Attention.

4 Experiments

4.1 Experiment Settings

Metrics: We use F1 score, Macro F1, AUC (Area Under Curve) and accuracy scores to evaluate baselines and our models. Models with the best F1-score on validation set are chosen.

Baselines: For images, we use ResNet50, ResNet101, VGG16, and LayoutLMv3. For text, we use DeBERTa. For combined image and text input, DeBERTa serves as text encoder, while ResNet50, ResNet101, and VGG16 act as image encoder. All models use the original dataset.

Implementation Details: To be more specific, we report scores within topics. Since the hyperpa-

rameters have a non-negligible effect on the scores, we conduct the experiments at a learning rate of $1e-4$, $1e-5$, $5e-6$ and a batch size of 16, 8, 4. More implementation details can be found in Appendix A.

4.2 Results and Analysis

Table 1 shows the results on both topics in ImageArg dataset. (Experimental results of more models can be seen in Appendix B.)

Our model, **TILFA**, outperforms the baselines (those above the double horizontal line) on all the four evaluation metrics by a large margin. Also, our model achieves the SOTA performance on the leaderboard of Argumentative Stance Classification subtask in ImageArg Shared Task 2023, which demonstrate the effectiveness of our methods.

For combined text and image inputs, models utilizing LayoutLMv3 for image encoding perform much better than those using traditional ones, which verifies our belief in Section 3.3 that a better understanding of the text in images is beneficial.

Back-translation and WordNet also greatly improve performance across all metrics (e.g. an improvement of 1.68 on gun control and 3.98 on abortion for **TILFA** in F1-score), confirming the value

Models	Modality	Gun Control				Abortion			
		F1	Ma-F1	AUC	Acc	F1	Ma-F1	AUC	Acc
ResNet50	I	65.00	62.50	62.90	62.67	35.71	60.42	59.55	75.84
ResNet101	I	67.43	60.91	60.86	62.00	38.89	59.71	60.44	70.47
VGG16	I	67.44	61.85	61.81	62.67	38.46	58.32	59.80	67.79
LayoutLMv3	I+L	56.98	46.67	46.92	48.67	39.47	59.38	60.66	69.13
DeBERTa	T	86.32	81.34	80.54	82.67	71.43	80.11	86.40	83.89
DeBERTa+ResNet50	T+I	88.04	84.54	83.80	85.33	70.59	79.43	85.97	83.22
DeBERTa+ResNet101	T+I	87.50	82.64	81.72	84.00	73.17	81.49	87.26	85.23
DeBERTa+VGG16	T+I	87.96	83.43	82.49	84.67	74.07	82.20	87.70	85.91
<hr/>									
DeBERTa+ResNet50	T+I ⁺	90.32	87.27	86.33	88.00	75.95	83.64	88.56	87.25
DeBERTa+ResNet101	T+I ⁺	87.43	83.89	83.21	84.67	70.73	79.81	85.32	83.89
DeBERTa+VGG16	T+I ⁺	89.73	86.60	85.75	87.33	77.50	84.62	90.07	87.92
<hr/>									
TILFA (DeBERTa+LayoutLMv3)	t1 T+I+L	89.13	85.94	85.16	86.67	76.54	83.89	89.64	87.25
TILFA (DeBERTa+LayoutLMv3)	t1 T+I+L ⁺	90.81	88.01	87.10	88.67	80.52	86.87	91.37	89.93
TILFA (DeBERTa+LayoutLMv3)	t2 T+I+L ⁺	90.32	87.27	86.33	88.00	77.50	84.62	90.07	87.92
TILFA (DeBERTa+LayoutLMv3)	t3 T+I+L ⁺	88.89	84.98	84.03	86.00	79.01	85.59	91.59	88.59

Table 1: Performance of all frameworks on the testing set of both topics. Those below the double horizontal line use our methods, and the above are baselines. For models that have both base and large sizes, we use the large one. We abbreviate F1-score, Macro F1-score, Accuracy to F1, Ma-F1, Acc, respectively. T, I and L are short for text modality, image modality and layout modality. Three multimodal fusion methods are named t1, t2 and t3 here. Those with a superscript + use both back-translation and WordNet, while others don’t use either.

of our data preprocessing and augmentation strategies. When it comes to multimodal fusion methods, the simplest Concatenation works best. We think it may be because the second method is initially applied in video field (Yu et al., 2021), and the third one in pure text field (Chen et al., 2017). So, neither of them is suitable to be migrated to this task.

We merge the answers belonging to different topics together and report the Micro F1-score. With a 90.32 F1 on gun control and a 77.50 F1 on abortion, we get a Micro F1-score of 86.47, which is the top on the leaderboard, 1.41 higher than the second best team. Our score improves even further after changing the hyperparameters, up to 87.79 (90.81 on gun control and 80.52 on abortion).

4.3 Ablation Study

To fully understand the impact of different components, we conduct an ablation study in Table 2.

Both back-translation and WordNet do help to the improvement of model performance, with WordNet having a larger impact. Models using only text inputs outperform just image inputs. This suggests that information in the text is more effective in determining the author’s stance than images. However, the best performance is achieved when both text and image inputs are used, showing that images also do contribute to stance determination.

LayoutLMv3 performs better than Resnet50 on

Text	Image	T	W	F1	
				Gun	Abortion
–	ResNet50	–	–	65.00	35.71
–	LayoutLMv3	–	–	56.98	39.47
DeBERTa	–	–	–	86.32	71.43
DeBERTa	ResNet50	–	–	88.04	70.59
DeBERTa	ResNet50	✓	–	88.42	71.11
DeBERTa	ResNet50	–	✓	88.77	75.61
DeBERTa	ResNet50	✓	✓	90.32	75.95
DeBERTa	LayoutLMv3	–	–	89.13	76.54
DeBERTa	LayoutLMv3	✓	–	89.73	76.92
DeBERTa	LayoutLMv3	–	✓	90.22	77.50
DeBERTa	LayoutLMv3	✓	✓	90.81	80.52

Table 2: Ablation studies in ImageArg. The first two columns illustrate text and image encoders. T and W represent back-translation and WordNet. All models use the Concatenation method for multimodal fusion.

abortion topic when based solely on image inputs, but on both topics when text inputs are added. This indicates that image encoders which can take the text and layout information in the images into account can really work better.

4.4 Case Study

We conduct a case study to better understand the behavior of our data augmentation method, with an example presented in Table 3. In the original text, we select the noun "risk". Then we find its differ-

Original Text	SCOTUS has balanced rights w/ public safety, ruling that gun safety laws essential & constitutional Rushing through a replacement to RBG could undermine that balance and put life-saving laws at risk.
Selected Noun	"risk"
Word Senses	"hazard.n.01": a source of danger; a possibility of incurring loss or misfortune "risk.n.02": a venture undertaken without regard to possible loss or injury "risk.n.03": the probability of becoming infected given that exposure to an infectious agent has occurred "risk.n.04": the probability of being exposed to an infectious agent "risk.v.01": expose to a chance of loss or damage "gamble.v.01": take a risk in the hope of a favorable outcome
Disambiguation	"hazard.n.01": a source of danger; a possibility of incurring loss or misfortune
Synonyms	"hazard.n.01" = ["peril", "jeopardy", "endangerment", "hazard"]
New Text	SCOTUS has balanced rights w/ public safety, ruling that gun safety laws essential & constitutional Rushing through a replacement to RBG could undermine that balance and put life-saving laws at peril/jeopardy/endangerment/hazard.

Table 3: An example of our data augmentation method.

ent meanings and corresponding synonym sets in WordNet. Using GlossBERT, we determine its exact meaning "hazard.n.01" and thus get the correct synonym set ["peril", "jeopardy", "endangerment", "hazard"]. Finally, we replace the noun "risk" in the original text with these synonyms to form new text.

5 Conclusion

We present **TILFA**, a new framework for multimodal argumentative stance classification. Unlike existing methods, **TILFA** considers not only the text and images in tweets but also the characters and their layout information in those images. Back-translation and WordNet also contribute to our SOTA performance. Our results reveal that better handling of images is essential to model improvement, and suggest that more effective methods for multimodal fusion are yet to be found.

Limitations

We have experimented with three multimodal fusion methods, but the simplest one, Concatenation, turned out to be the best. So the other two methods that we use are not suitable for this task actually. But we believe that there are more effective multimodal fusion methods (Liu et al., 2023a; Li et al., 2023; Yang et al., 2022) waiting to be discovered.

Also, we notice that images in the dataset vary widely, some feature only objects, but others con-

tain significant text. Further research is needed to effectively handle these differences, and we expect that better image encoders will improve performance in future works.

Moreover, in the data augmentation part, we only explore the methods related to text, but there are also many ways to augment images. Whether these methods (Shorten and Khoshgoftaar, 2019; Xu et al., 2023) are effective for images containing lots of characters is a question worth studying.

For the text modality, we found that most instances are a piece of text containing a few events, such as the example in Table 3. With the recent advances in event understanding (Lin et al., 2023), we can incorporate different relations among events, including temporal (Fang et al., 2023), causal (Zhang et al., 2022; Wang et al., 2023c; Gao et al., 2023), sub-event (Wang et al., 2022; Zhang et al., 2020), hierarchical (Wang et al., 2023b,a).

Acknowledgements

The authors of this paper were supported by the NSFC Fund (U20B2053) from the NSFC of China, the RIF (R6020-19 and R6021-20), and the GRF (16211520 and 16205322) from RGC of Hong Kong. We also thank the support from NVIDIA AI Technology Center (NVAITC) and the UGC Research Matching Grants (RMGS20EG01-D, RMGS20CR11, RMGS20CR12, RMGS20EG19, RMGS20EG21, RMGS23CR05, RMGS23EG08).

References

- Markus Bayer, Marc-André Kaufhold, and Christian Reuter. 2022. [A survey on data augmentation for text classification](#). *ACM Comput. Surv.*, 55(7).
- Winston Carlike, Nishant Gurrupadi, Zixuan Ke, and Vincent Ng. 2018. [Give me more feedback: Annotating argument persuasiveness and related attributes in student essays](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 621–631, Melbourne, Australia. Association for Computational Linguistics.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1657–1668. Association for Computational Linguistics.
- Tianqing Fang, Zhaowei Wang, Wenxuan Zhou, Hongming Zhang, Yangqiu Song, and Muhao Chen. 2023. [Getting sick after seeing a doctor? diagnosing and mitigating knowledge conflicts in event temporal reasoning](#). *CoRR*, abs/2305.14970.
- Tianqing Fang, Wenxuan Zhou, Fangyu Liu, Hongming Zhang, Yangqiu Song, and Muhao Chen. 2022. [On-the-fly denoising for data augmentation in natural language understanding](#). *CoRR*, abs/2212.10558.
- Jinglong Gao, Xiao Ding, Bing Qin, and Ting Liu. 2023. [Is chatgpt a good causal reasoner? A comprehensive evaluation](#). *CoRR*, abs/2305.07375.
- Ivan Habernal and Iryna Gurevych. 2017. [Argumentation mining in user-generated web discourse](#). *Computational Linguistics*, 43(1):125–179.
- Leipeng Hao, Liangcai Gao, Xiaohan Yi, and Zhi Tang. 2016. [A table detection method for pdf documents based on convolutional neural networks](#). In *2016 12th IAPR Workshop on Document Analysis Systems (DAS)*, pages 287–292.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: decoding-enhanced bert with disentangled attention](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Luyao Huang, Chi Sun, Xipeng Qiu, and Xuanjing Huang. 2019. [GlossBERT: BERT for word sense disambiguation with gloss knowledge](#). 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 3507–3512, Hong Kong, China. Association for Computational Linguistics.
- Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. 2022. [Layoutlmv3: Pre-training for document AI with unified text and image masking](#). In *MM '22: The 30th ACM International Conference on Multimedia, Lisboa, Portugal, October 10 - 14, 2022*, pages 4083–4091. ACM.
- Yunxin Li, Baotian Hu, Xinyu Chen, Yuxin Ding, Lin Ma, and Min Zhang. 2023. [A multi-modal context reasoning approach for conditional inference on joint textual and visual clues](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 10757–10770. Association for Computational Linguistics.
- Zizheng Lin, Hongming Zhang, and Yangqiu Song. 2023. [Global constraints with prompting for zero-shot event argument classification](#). In *Findings of the Association for Computational Linguistics: EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*, pages 2482–2493. Association for Computational Linguistics.
- Chengliang Liu, Jie Wen, Xiaoling Luo, and Yong Xu. 2023a. [Incomplete multi-view multi-label learning via label-guided masked view- and category-aware transformers](#). In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pages 8816–8824. AAAI Press.
- Xiaoqing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. 2019. [Graph convolution for multimodal information extraction from visually rich documents](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Industry Papers)*, pages 32–39, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zhexiong Liu, Mohamed Elaraby, Yang Zhong, and Diane Litman. 2023b. [Overview of ImageArg-2023: The first shared task in multimodal argument mining](#). In *Proceedings of the 10th Workshop on Argument Mining*, Online and in Singapore. Association for Computational Linguistics.
- Zhexiong Liu, Meiqi Guo, Yue Dai, and Diane Litman. 2022. [ImageArg: A multi-modal tweet dataset for image persuasiveness mining](#). In *Proceedings of the 9th Workshop on Argument Mining*, pages 1–18, Online and in Gyeongju, Republic of Korea. International Conference on Computational Linguistics.

- George A. Miller. 1994. [WordNet: A lexical database for English](#). In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Georgios Rizos, Konstantin Hemker, and Björn W. Schuller. 2019. [Augment to prevent: Short-text data augmentation in deep learning for hate-speech classification](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, pages 991–1000. ACM.
- Connor Shorten and Taghi M. Khoshgoftaar. 2019. [A survey on image data augmentation for deep learning](#). *Journal of Big Data*, 6(1):1–48.
- Karen Simonyan and Andrew Zisserman. 2015. [Very deep convolutional networks for large-scale image recognition](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- WeiQi Wang, Tianqing Fang, Wenxuan Ding, Baixuan Xu, Xin Liu, Yangqiu Song, and Antoine Bosselut. 2023a. [CAR: conceptualization-augmented reasoner for zero-shot commonsense question answering](#). *CoRR*, abs/2305.14869.
- WeiQi Wang, Tianqing Fang, Baixuan Xu, Chun Yi Louis Bo, Yangqiu Song, and Lei Chen. 2023b. [CAT: A contextualized conceptualization and instantiation framework for commonsense reasoning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13111–13140, Toronto, Canada. Association for Computational Linguistics.
- X. Wang, Y. Sheng, H. Deng, and Z. Zhao. 2019. [Charcnn-svm for chinese text datasets sentiment classification with data augmentation](#). *International Journal of Innovative Computing, Information and Control*, 15:227–246.
- Zhaowei Wang, Quyet V. Do, Hongming Zhang, Jiayao Zhang, WeiQi Wang, Tianqing Fang, Yangqiu Song, Ginny Y. Wong, and Simon See. 2023c. [COLA: contextualized commonsense causal reasoning from the causal inference perspective](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 5253–5271. Association for Computational Linguistics.
- Zhaowei Wang, Hongming Zhang, Tianqing Fang, Yangqiu Song, Ginny Y. Wong, and Simon See. 2022. [Subeventwriter: Iterative sub-event sequence generation with coherence controller](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 1590–1604. Association for Computational Linguistics.
- John Wieting and Kevin Gimpel. 2018. [Paranmt-50m: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 451–462. Association for Computational Linguistics.
- Mingle Xu, Sook Yoon, Alvaro Fuentes, and Dong Sun Park. 2023. [A comprehensive survey of image augmentation techniques for deep learning](#). *Pattern Recognition*, 137:109347.
- Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei A. F. Florêncio, Cha Zhang, Wanxiang Che, Min Zhang, and Lidong Zhou. 2021. [Layoutlmv2: Multi-modal pre-training for visually-rich document understanding](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 2579–2591. Association for Computational Linguistics.
- Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2020. [Layoutlm: Pre-training of text and layout for document image understanding](#). In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pages 1192–1200. ACM.
- Qian Yang, Yunxin Li, Baotian Hu, Lin Ma, Yuxin Ding, and Min Zhang. 2022. [Chunk-aware alignment and lexical constraint for visual entailment with natural language explanations](#). In *MM '22: The 30th ACM International Conference on Multimedia, Lisboa, Portugal, October 10 - 14, 2022*, pages 3587–3597. ACM.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V. Le. 2018. [Qanet: Combining local convolution with global self-attention for reading comprehension](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Tiezheng Yu, Wenliang Dai, Zihan Liu, and Pascale Fung. 2021. [Vision guided generative pre-trained language models for multimodal abstractive summarization](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3995–4007, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hongming Zhang, Muhao Chen, Haoyu Wang, Yangqiu Song, and Dan Roth. 2020. [Analogous process structure induction for sub-event sequence prediction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP*

2020, Online, November 16-20, 2020, pages 1541–1550. Association for Computational Linguistics.

Jiayao Zhang, Hongming Zhang, Weijie J. Su, and Dan Roth. 2022. [ROCK: causal inference principles for reasoning about commonsense causality](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 26750–26771. PMLR.

A Implementation Details

In this appendix, we introduce the implementation details of every component in our framework **TILFA**.

We use the dataset divisions provided in the shared task, and the dataset sizes are detailed in Table 4. Each tweet text is cut by a maximum length of 512, and each tweet image is resized to 224×224 dimension. For back-translation, we use Youdao translation API³. For layout information, we follow Xu et al. (2020) and use Tesseract⁴, an open-source OCR engine, to get the recognized words and their 2-D positions in the images. Our models are implemented with Pytorch, and trained on a NVIDIA A6000 GPU. AdamW optimizer is used for those networks with LayoutLMv3 and Adam optimizer for others.

Topic	Train	Validation	Test
Gun control	918	96	150
Abortion	888	100	149

Table 4: Dataset scale of both topics. Following the shared task, one unavailable tweet in the abortion testing set is removed. And due to the downloading issues, our downloaded train and dev sets have little difference from the original one.

B Experimental Results

Our full experiment results are shown in Table 5.

³<http://fanyi.youdao.com/openapi/>

⁴<https://github.com/tesseract-ocr/tesseract>

Models	Modality	Gun Control				Abortion			
		F1	Ma-F1	AUC	Acc	F1	Ma-F1	AUC	Acc
ResNet50	I	65.00	62.50	62.90	62.67	35.71	60.42	59.55	75.84
ResNet101	I	67.43	60.91	60.86	62.00	38.89	59.71	60.44	70.47
VGG16	I	67.44	61.85	61.81	62.67	38.46	58.32	59.80	67.79
LayoutLMv3	I+L	56.98	46.67	46.92	48.67	39.47	59.38	60.66	69.13
BERT	T	81.48	74.97	74.52	76.67	64.10	75.69	79.26	81.21
RoBERTa	T	83.52	79.05	78.55	80.00	71.60	80.50	85.75	84.56
DeBERTa	T	86.32	81.34	80.54	82.67	71.43	80.11	86.40	83.89
BERT+ResNet50	T+I	81.97	76.88	76.43	78.00	65.06	75.79	81.00	80.54
RoBERTa+ResNet50	T+I	84.82	79.11	78.42	80.67	73.42	81.91	86.61	85.91
DeBERTa+ResNet50	T+I	88.04	84.54	83.80	85.33	70.59	79.43	85.97	83.22
BERT+ResNet101	T+I	80.63	73.34	72.99	75.33	64.52	74.21	82.52	77.85
RoBERTa+ResNet101	T+I	83.77	77.66	77.06	79.33	65.22	74.84	82.95	78.52
DeBERTa+ResNet101	T+I	87.50	82.64	81.72	84.00	73.17	81.49	87.26	85.23
BERT+VGG16	T+I	79.37	72.11	71.81	74.00	68.24	77.78	84.03	81.88
RoBERTa+VGG16	T+I	83.52	79.05	78.55	80.00	72.50	81.20	86.18	85.23
DeBERTa+VGG16	T+I	87.96	83.43	82.49	84.67	74.07	82.20	87.70	85.91
BERT+ResNet50	T+I ⁺	81.05	74.16	73.76	76.00	68.24	77.78	84.03	81.88
RoBERTa+ResNet50	T+I ⁺	84.66	79.26	78.60	80.67	71.60	80.50	85.75	84.56
DeBERTa+ResNet50	T+I ⁺	90.32	87.27	86.33	88.00	75.95	83.64	88.56	87.25
BERT+ResNet101	T+I ⁺	82.87	78.41	77.96	79.33	62.50	72.34	81.23	75.84
RoBERTa+ResNet101	T+I ⁺	84.21	78.47	77.83	80.00	68.97	78.08	85.11	81.88
DeBERTa+ResNet101	T+I ⁺	87.43	83.89	83.21	84.67	70.73	79.81	85.32	83.89
BERT+VGG16	T+I ⁺	81.03	72.89	72.62	75.33	68.97	78.08	85.11	81.88
RoBERTa+VGG16	T+I ⁺	83.24	78.14	77.60	79.33	72.50	81.20	86.18	85.23
DeBERTa+VGG16	T+I ⁺	89.73	86.60	85.75	87.33	77.50	84.62	90.07	87.92
BERT+LayoutLMv3	t1 T+I+L	81.32	76.25	75.84	77.33	65.75	77.32	79.47	83.22
RoBERTa+LayoutLMv3	t1 T+I+L	85.26	79.90	79.19	81.33	75.95	83.64	88.56	87.25
DeBERTa+LayoutLMv3	t1 T+I+L	89.13	85.94	85.16	86.67	76.54	83.89	89.64	87.25
BERT+LayoutLMv3	t1 T+I+L ⁺	81.72	75.95	75.48	77.33	69.14	78.81	83.80	83.22
RoBERTa+LayoutLMv3	t1 T+I+L ⁺	86.19	82.59	82.04	83.33	76.32	84.10	87.90	87.92
DeBERTa+LayoutLMv3	t1 T+I+L ⁺	90.81	88.01	87.10	88.67	80.52	86.87	91.37	89.93
BERT+LayoutLMv3	t2 T+I+L ⁺	83.52	79.05	78.55	80.00	69.88	79.13	84.89	83.22
RoBERTa+LayoutLMv3	t2 T+I+L ⁺	84.95	80.19	79.55	81.33	72.29	80.80	86.83	84.56
DeBERTa+LayoutLMv3	t2 T+I+L ⁺	90.32	87.27	86.33	88.00	77.50	84.62	90.07	87.92
BERT+LayoutLMv3	t3 T+I+L ⁺	82.16	76.73	76.24	78.00	69.05	78.45	84.46	82.55
RoBERTa+LayoutLMv3	t3 T+I+L ⁺	84.78	80.32	79.73	81.33	71.05	80.57	84.01	85.23
DeBERTa+LayoutLMv3	t3 T+I+L ⁺	88.89	84.98	84.03	86.00	79.01	85.59	91.59	88.59

Table 5: Experimental results of more models on both topics in ImageArg dataset. Compared to Table 1, the scores of the models which use BERT or RoBERTa as the text encoder are also listed here.