IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 16 June 2023, accepted 17 July 2023, date of publication 25 July 2023, date of current version 7 August 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3298877

RESEARCH ARTICLE

An Inferential Commonsense-Driven Framework for Predicting Political Bias in News Headlines

SWATI SWATI[1,2], **DUNJA MLADENIĆ**[1,2], **AND MARKO GROBELNIK**[1]

[1]Jožef Stefan Institute, 1000 Ljubljana, Slovenia
[2]Jožef Stefan International Postgraduate School, 1000 Ljubljana, Slovenia

Corresponding author: Swati Swati (swati@ijs.si)

This work was supported in part by the enRichMyData European Union's Horizon Europe Project under Agreement 101070284, in part by the Slovenian Research Agency under Project J2-1736 Causalify, and in part by the European Union's Horizon 2020 Research and Innovation Program under the Marie Skłodowska-Curie Grant under Agreement 812997.

ABSTRACT Identifying political bias in news headlines holds significant importance as it influences the dissemination and consumption of news stories. However, employing conventional methods to do so poses a formidable challenge, as the short headline text is often complex and lacks sufficient syntactic and semantic context. Existing approaches fail to acknowledge the potential of commonsense reasoning in facilitating text comprehension, although it has been shown to aid numerous downstream applications. To this end, to facilitate comprehension and compensate for the lack of context, we propose leveraging inferential commonsense knowledge to simplify, interpret, and explain events that are not explicitly stated in headlines. Furthermore, to fully utilise its potential and deal with the unnecessary noise it may introduce, we present a method for emphasising significant inferences. Using this knowledge, we introduce a novel framework, **IC-BAIT**, short for **I**nferential **C**ommonsense aware **B**i**A**s **I**den**T**ifier, which is a flexible neural network framework designed to enhance political bias prediction in news headlines. We also present two bias-annotated datasets: MediaBias and GoodNews. Experiments on both datasets demonstrate that IC-BAIT significantly enhances the performance of the baseline models used in the framework. Experiments on the datasets show that IC-BAIT improves the baseline models in terms of accuracy (2.0-10.0%), macro-averaged F_1 (2.2-22.2%), Jaccard-score (up to 15.1%), and micro-averaged F_1 (up to 18.6%). Our in-depth qualitative analysis reveals the scenarios in which the selected knowledge is beneficial and when it is detrimental. Datasets and scripts are available at https://github.com/Swati17293/IC-BAIT.

INDEX TERMS Deep learning, inferential commonsense knowledge, media bias, news bias, NLP, short text classification.

I. INTRODUCTION

Media outlets often publish news stories that benefit the political party they endorse [1], [2]. Their bias is particularly reflected in news headlines, as readers are more likely to be swayed if the headline is interesting and catchy [3], [4], [5]. However, identifying bias solely based on the headline can pose a challenge. It is because they are typically short and may not contain the context of bias embedded in the story [6]. Furthermore, syntactic and semantic information are difficult to capture in a short headline. For example, consider the three headlines in Fig. 1 from our dataset MediaBias reporting on the same event with conflicting political ideologies. Unless we understand the context of the event and how the bias manifests itself, an automatic bias predictor will perform poorly [7].

To deal with this lack of context, models that use Inferential Commonsense Knowledge (IC_Knwl) have the opportunity to substantially enhance their performance. This enhancement is achieved in a manner nearly similar to the way people use commonsense inferences to carry out their daily routine tasks. For instance, the inclusion of IC_Knwl has been

The associate editor coordinating the review of this manuscript and approving it for publication was Arianna Dulizia.

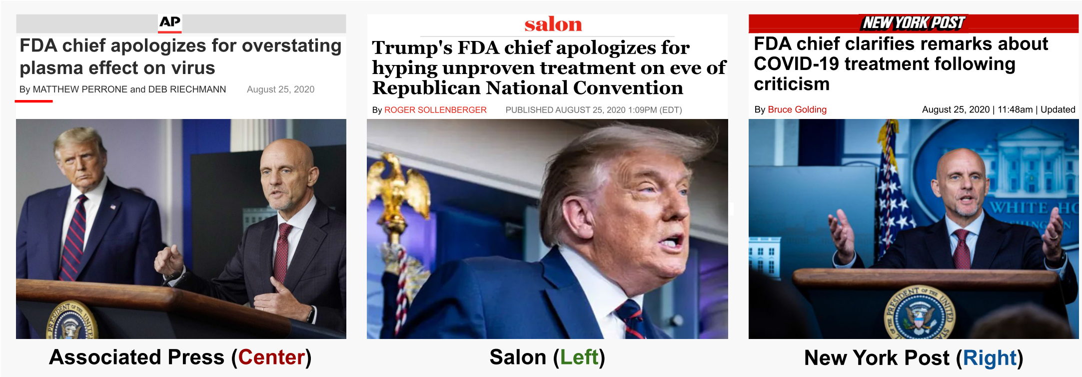

FIGURE 1. News headlines reporting on the event "*FDA Commissioner Acknowledges Misrepresenting Convalescent Plasma Data*" from opposing political ideologies. (Image source: allsides.com).

shown to improve the model's performance in a wide range of tasks such as emotion inference [8], sarcasm detection [9], sarcasm generation [10], and reading comprehension [11], among others [12].

Although IC_Knwl has been demonstrated to be promising, existing research neglects to acknowledge its potential in aiding text comprehension for the task of bias prediction. To this end, we propose leveraging it to simplify, interpret, and explain events that are not explicitly stated in headlines. We hypothesise that incorporating such knowledge would ideally help the learning model identify political bias when it is not readily evident from the headlines.

However, without proper emphasis, the additional inferential context is prone to introduce unnecessary noise. This noise can prevent models from fully exploiting the acquired knowledge. We thus present a method that facilitates emphasising significant inferences. We then use this knowledge to introduce a novel framework, **IC-BAIT** (**I**nferential **C**ommonsense aware **B**i**A**s **I**den**T**ifier). It leverages the IC_Knwl acquired using the neural knowledge model *COMET* [13] trained on *ATOMIC*$^{20}_{20}$ [13] knowledge graphs. For example, for the headline "Trump's awful advice on voting twice", the acquired commonsense knowledge provides an inferential context, as shown in Fig. 2.

Additionally, given the scarcity of large-scale datasets that adequately capture the unique challenges posed by the task of bias prediction in news headlines, we present two datasets, *MediaBias* and *GoodNews*, in Section III. We conduct experiments on both datasets to test the robustness of IC-BAIT. Furthermore, to assess its effectiveness, we compare the performance of several state-of-the-art models incorporated into it. We also conduct a comprehensive error analysis on its prediction results to determine when the selective inclusion of IC_Knwl is advantageous and when it is not. To summarise, we make the following contributions:

- Proposing to leverage Inferential Commonsense Knowledge to aid in the comprehension of news headlines by simplifying, interpreting, and explaining events that are not explicitly stated in the headlines.

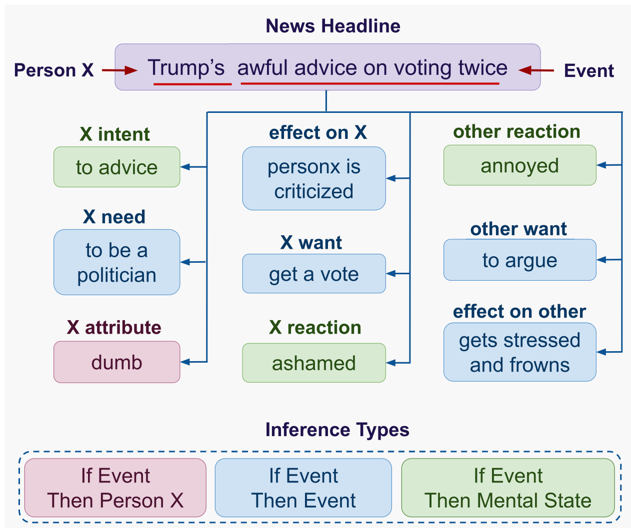

FIGURE 2. An illustration of IC_Knwl acquired using COMET for a sample news headline. "PersonX is dumb, needed to be a politician, intended to advice, wanted to get a vote, is criticized, feels ashamed. Others want to argue, gets stressed and frowns, feel annoyed" represents the set of inferred causes and effects for the headline "Trump's awful advice on voting twice."

- Introducing IC-BAIT, a neural network framework designed to enhance political bias prediction in news headlines by incorporating commonsense knowledge with an emphasis on important inferences.
- Presenting two datasets of news headlines annotated with political bias. We conduct experiments on both of them to demonstrate IC-BAIT's reliability.
- Evaluating the effectiveness of IC-BAIT with several state-of-the-art language models.
- Analysing the impact of selective knowledge augmentation on overall performance and conducting an in-depth error analysis.

We believe that our proposed framework could serve as an effective method for copy editors [14]. It could help them avoid any accidental or intentional bias in news headlines written by journalists. It could also be useful in practical applications such as e-journalism and manual news-bias prediction portals (ex: allsides.com, adfontesmedia.com), where it could be used to automatically classify headlines into different bias types. Additionally, it could help reduce the number of articles that require manual examination, which is a time-consuming process that is often susceptible to annotator bias [15]. Furthermore, it could be ideal for social scientists and those interested in the analysis of political bias, as well as for automated systems such as the "bias flipper" [6], [16], [17].

II. RELATED WORK

A news report is deemed politically biassed if it appears to favour one political ideology over another [18]. Such bias can appear in a variety of ways, such as in the selection of news stories, the tone and language of reporting, the sources of citation, etc [19]. Such bias in news reporting can result in an incomplete or inaccurate portrayal of events and issues, as well as the spread of misinformation or propaganda [20]. It also has the potential to influence public opinion on political issues and events [21]. Owing to their immense significance, bias prediction and mitigation have been the focus of research at multiple levels of granularity. In this section, we present a description of the datasets used in the related studies, followed by a discussion of the associated methodologies.

A. DATASETS

Many datasets have been designed to study political bias at different granularity levels: news sources [22], [23], articles [24], paragraphs [25], sentences [26], [27], and headlines [6] with one or more bias types [28].

Political bias annotation at the source level is typically obtained from online platforms, such as allsides.com and adfontesmedia.com. These platforms are dedicated to assessing and rating the bias of media sources to provide balanced news, perspectives, and issues across the political spectrum. They have highly qualified teams that rigorously follow predefined guidelines for rating.

At the sentence level, numerous datasets with manual bias annotations have been developed, with an emphasis on individual sentences of the articles [26], [29]. Manual annotations of individual sentences require substantial time and effort, and they do not scale well [30]. These annotations are even more time-consuming at the article level. As a result, many datasets with article-level annotations are generated by collecting articles from news sources with a known bias, where individual news articles reflect their source's political leanings [31], [32].

For example, authors in [33] used data from one such platform, Media Bias/Fact Check (MBFC), which contains manual annotations of political bias for over 2,000 news websites. Their dataset includes 1,066 websites with explicit bias labels ranging from 'extreme-left' to 'extreme-right' on a seven-point scale. Manual inspection of their dataset revealed that the 'left-center' and 'right-center' labels were ill-defined and ambiguous. Therefore, they opted to exclude news websites with these labels. They also merged the labels 'extreme-left' with 'left' and 'extreme-right' with 'right' to mitigate the impact of potentially subjective annotator decisions. As a result, their labels were reduced to a three-point scale (left, center, and right), resulting in a total of 864 websites [34]. However, the political leanings of news articles do not always correspond to the political leanings of the sources [35]. Besides that, the publisher's political leanings may shift while reporting on various topics [36].

Although the majority of these datasets are concerned with news articles or news sources, a few related datasets are focused on news headlines. Authors in [6], compiled a list of article pairs reporting on the same event from sources with opposing political views. They used it for the task of flipping the bias of news headlines. They later started with this dataset and extended it with new article pairs to analyse political bias and unfairness in news articles at different levels of granularity [7]. Event-centred pairing is considered necessary for tasks such as bias flipping and examining how different news outlets cover the same event. Any dataset devoid of such pairings, on the other hand, presents unique challenges for political bias prediction in news headlines and helps determine whether the prediction method is truly generic.

To deal with the scarcity of publicly available large-scale datasets for determining the political bias of news headlines, we present the datasets GoodNews III-B and MediaBias III-A. GoodNews includes mapping from the corpus GoodNewsEveryone [37] and MediaBias consists of headlines paired as per the events, similar to [6]. We provide data generation frameworks to aid expansion and data customization for both datasets.

B. POLITICAL BIAS PREDICTION

In general, news bias is identified by studying linguistic attributes such as keywords and syntactic features. But, with the advent of machine learning and neural networks [38], [39], [40], [41], advanced algorithmic methods for bias analysis in news texts have been developed [42], [43]. For instance, leveraging recent advancements in deep learning, the authors in [42] developed a sentence-level factuality and bias prediction model fine-tuned on BERT [44]. Authors in [45] used the Gaussian Mixture Model, whereas authors in [43] used BERT and ELMo [46] to classify political bias in news articles, respectively.

There is another set of studies that aim to identify political bias based on explicitly defined features, such as quotes from think tanks. For example, authors in [18] studied news bias by estimating the ideological scores of major outlets. For the study, they made a list of influential policy groups and active think tanks and counted how many times a specific outlet and members of Congress used quotes from those groups.

Based on the pattern of citations, they estimated the bias score. However, articles might not cite a specific policy group but still exhibit political favouritism. For instance, one party may be disproportionately criticised for building up another party's figure. Several features are therefore required to be considered to determine political leaning.

The majority of related research seeks to strengthen methods for determining bias in news content as opposed to the headline. There are only a few methods concerning bias in headlines. For instance, authors in [47] used a headline attention network consisting of a headline encoder, an article encoder, and a headline attention layer to detect political bias in news content. They used headlines, articles, and the bias of the news to generate the dataset. They concluded that headlines alone weren't enough to figure out if an article was biassed or not. To analyse the bias induced at different granularity levels, authors in [6] extracted the most representative, discriminative, and sentiment-inducing words as the features of their classifiers. Their analysis revealed that named entities are very important to discriminate the left or right orientation of the text, whereas sentiment words play a crucial role in subjectivity identification.

Although deep learning is believed to be superior at learning complex and dense text representations through the capture of semantic and syntactic information, extending it from bias classification at the article level to the headline level is not plausible [6]. This is because headlines are more complicated and ambiguous than news articles, and they lack contextual and conceptual information.

To address this bottleneck in similar tasks, prior studies have focused on utilising external knowledge. For instance, Mihaylov and Frank [48] used key-value memory to represent commonsense facts and employ word-to-knowledge attention; Chen et al. [49] used semantic relations present in WordNet [50] to enhance attention and inference capabilities; Bauer et al. [51] proposed a mutual information-based knowledge selection method and integrated knowledge using gated attention; Zhang et al. [52] proposed using triplet-based knowledge to resolve coreferences.

Although these studies incorporate knowledge into pre-trained language models, they prioritise entity-centric facts in knowledge bases over commonsense, limiting their capacity to intuitively reason about events and situations. Empirical evidence suggests that the incorporation of commonsense knowledge is a valuable resource for language inference, leading to significant improvements in a variety of tasks [53], [54]. Consequently, recent research has sought to integrate external commonsense knowledge into pre-trained models to enhance linguistic representation for knowledge-reliant NLP tasks [55].

For generating such knowledge, commonsense knowledge graphs are commonly employed as standard tools to provide models with inferential background knowledge [56], [57]. ATOMIC is an example of such a knowledge graph, with an emphasis on event representation and enhanced relation types. As evidenced by its performance in comparison to human evaluation, it demonstrates competitive performance in the if-then reasoning task. ATOMIC$^{20}_{20}$ [13] which is an extension of ATOMIC [58] is another example. It incorporates additional event-centric relations and ConceptNet [59] facts that are not easily captured by language models, resulting in a comprehensive repository of complex entities.

The commonsense knowledge bases aid in training neural language models to generate commonsense descriptions [60]. In contrast to extractive methods, the models trained on commonsense knowledge bases exhibit a substantial advantage in their ability to generate knowledge about unseen events. This attribute is particularly essential for tasks that necessitate the use of commonsense [9], [12].

Our study diverges from prior research by demonstrating the potential of leveraging implicit knowledge acquired from pre-trained language models and explicit knowledge in the form of inferential commonsense knowledge at the same time. We show that they can be used in conjunction to enhance the task of bias prediction in complex news headlines. We present our framework that utilises the commonsense knowledge acquired using the neural knowledge model COMET [13] trained on ATOMIC$^{20}_{20}$ knowledge graphs. It employs an attentive knowledge selection strategy for emphasising significant inferences in the generated inferential statements. To ensure future adaptability, we design our framework to be model-agnostic, allowing it to be packaged with cutting-edge linguistic models as they become available.

III. DATASET DESCRIPTION

To address the lack of standardised, annotated datasets for the task of bias prediction in news headlines, we present two distinct datasets. We generate the first dataset, MediaBias (Section III-A), by utilising bias labels from the bias rating portal allsides.com. We generate the second dataset, GoodNews (Section III-B), by extracting the headlines and their bias values from the dataset GoodNewsEveryone [37], which is based on the bias rating portal adfontesmedia.com.

In Table 1, we summarise dataset statistics, including the number of headlines for each political ideology, the average number of words in the headline, and so on. We document our datasets following the requirements of the FAIR Data Principles.[1]

A. MediaBias

1) RAW DATA SOURCE: ALLSIDES

To generate the dataset MediaBias, we use the bias rating portal allsides as the raw data source. The team at *allsides* consists of news experts with political leanings from all sides of the political spectrum. Every week, they identify the most important news stories and write about twenty roundups. They cover stories ranging from major national events to niche opinions. Their goal is to cover the most important news stories rather than the most clickable, viral, or interesting ones. For each roundup, they scan over 800 media outlets

[1] https://www.nature.com/articles/sdata201618/

TABLE 1. Dataset Statistics. Avg. length: average number of words in the headline.

			Size of the split				Bias label		
Dataset	**Raw data source**	**Dataset size**	**Train**	**Valid**	**Test**	**Avg. length**	**Left**	**Center**	**Right**
MediaBias	allsides	11,031	8,825	1,102	1,104	13	3,082	4,221	3,728
GoodNews	adfontesmedia	3,058	2,446	306	306	13	877	1,153	1,028

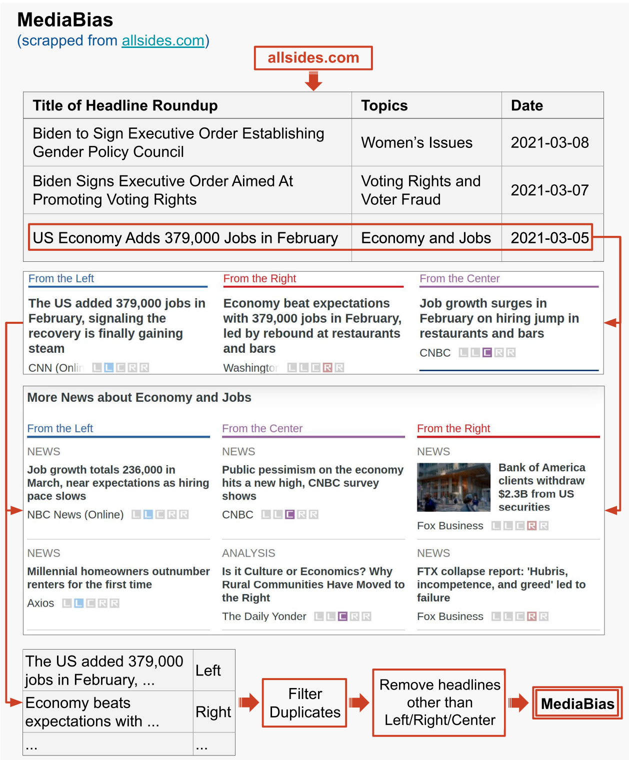

FIGURE 3. Overview of MediaBias generation process.

and release bias ratings for three sources that have published articles about it from different political perspectives. They then compile three main headlines from the *left*, *center*, and *right* sides of the same story and present them side by side to compare news coverage, identify bias, and get the entire story. Since it provides annotations for individual stories based on a rigorous process [61], we use it as a source of high-quality gold standard data for our dataset generation process.

2) MEDIABIAS GENERATION PROCESS

We generate the dataset *MediaBias* by using the existing bias labels from the featured headline roundups of the news aggregator *allsides*. We begin by crawling through all of the featured stories, regardless of their category. For each story, we then crawl through the three major stories and all the stories related to their topic from different political perspectives. Following that, we record the headlines of each story along with their bias values. We remove duplicate

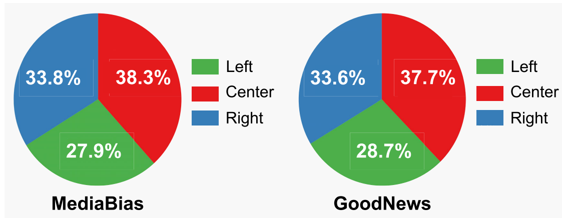

FIGURE 4. Bias label distribution. The graphs demonstrate the imbalance in both datasets.

headlines as well as those with labels *Lean Left*, *Lean Right*, *Not Rated*, and *Mixed*. The generation process is demonstrated in Fig. 3.

We end up with 11, 031 headlines with an average length of 13 words. Since the bias label distribution is not uniform, as shown in Fig. 4, we use a stratified split to replicate this imbalance across the generated train-valid-test sets.

B. GoodNews

1) RAW DATA SOURCE: ADFONTESMEDIA

We use the bias rating portal Adfontesmedia as the raw data source to generate the dataset, GoodNews. It scores news outlets by using bias and reliability as coordinates on its chart [62]. It assigns ratings to articles based on the opinions of a panel of expert analysts. Before labelling an outlet, the panels read at least three to thirty articles. They calculate the reliability score by adding individual ratings for an article's correctness, use of fact or opinion, and appropriateness of its headline and graphic content. The political bias score is influenced by topic selection or omission, the language used in the article, and the extent to which a political stance from left to right is endorsed. Individual ratings for each reviewed article are then aggregated to calculate an outlet's overall bias and reliability score, with popular articles receiving more weight. This average determines the placement of the outlet on the chart. Markers in the chart range from *Most Extreme Left* to *Most Extreme Right* along the bias axis.

2) GoodNews GENERATION PROCESS

We generate the dataset GoodNews by extracting the headlines and their bias values from the corpus GoodNewsEveryone [37] which is based on adfontesmedia. To do so, we first crawl through all the headlines in the corpus. For each headline, we then extract the headline text and its associated bias values. Following the range of bias values defined by

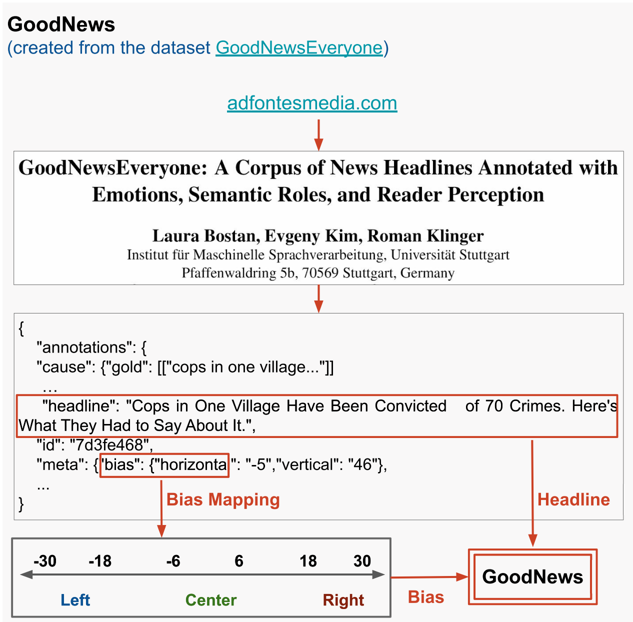

FIGURE 5. Overview of GoodNews generation process.

adfontesmedia.com, we then map the horizontal bias values from -30 to -18 to the *left*, -6 to 6 to the *center*, and 18 to 30 to the *right*. For mapping, we use the horizontal axis, which represents the political orientation of the news. Fig. 5 depicts the dataset generation process.

We end up with 3,058 headlines with an average length of 13 words. Similar to MediaBias, we use a stratified split to mimic the imbalance in the resultant train-valid-test sets, as illustrated in Fig. 4.

IV. MATERIALS AND METHODS

In this section, we first formulate the task of predicting political bias in news headlines. We then present a brief discussion of the baseline models that we use to evaluate IC-BAIT. We then introduce our proposed framework and its key components. Finally, we discuss the metrics that we use for evaluation.

A. PROBLEM FORMULATION

For a given headline text H, the task begins with the acquisition of its associated inferential commonsense knowledge IC_Knwl,

$$IC\_Knwl = c(H, \alpha) \tag{1}$$

Here, c is the commonsense knowledge modelling function, and α denotes the model parameters. Given the pair (H, IC_Knwl), our final task is to train a classifier that maps this extended feature space of short texts into the political bias label set B. Mathematically, it can be formulated as,

$$b = f(H, IC\_Knwl, \theta) \tag{2}$$

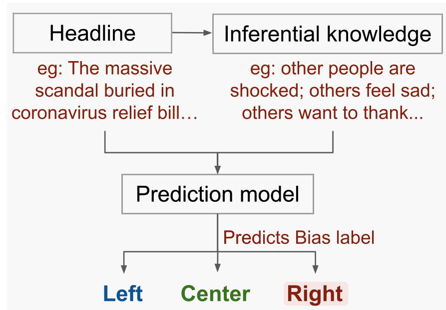

FIGURE 6. Outline of the task of predicting political bias in news headlines.

where f is the bias prediction function, θ denotes the model parameters, and b indicates the bias labels in the set B. Fig. 6 depicts an outline of the formulated problem.

B. BASELINE MODELS

We examine the following state-of-the-art language models in IC-BAIT for a comprehensive evaluation:

- **ALBERT (A Lite BERT)** [63]: a lite BERT [44] for self-supervised language representation learning. It uses a self-supervised loss that models inter-sentence coherence and the parameter reduction technique to decrease memory usage and accelerate training. Pre-trained model: https://tfhub.dev/google/albert_base/
- **DistilRoBERTa (Distilled Robustly optimized BERT approach)** [64]: a knowledge-distilled version of the robustly optimised BERT-based model, RoBERTa [65]. It significantly outperforms BERT as its model is trained longer with larger batches on longer sequences over more data, the objective of predicting the next sentence is removed, and the masking pattern on training data is modified dynamically.
- **MPNet (Masked and Permuted Net)** [66]: a masked and permuted method that inherits the benefits of BERT and XLNet [67] without suffering their drawbacks. It employs permuted language modelling to take advantage of the dependency among predicted tokens and takes auxiliary position information to present the model with a complete sentence, thereby reducing position discrepancy. Pre-trained model: https://github.com/microsoft/MPNet
- **MiniLMv2 (Mini Language Model v2)** [68]: a generalised and simplified deep self-attention distillation in MiniLM [69]. It introduces multi-head self-attention relation distillation to provide students with a flexible number of attention heads while also improving fine-grained self-attention knowledge. Distilled model: https://aka.ms/minilm
- **CMLM (Conditional Masked Language Modeling)** [70]: a conditional masked language model built on a 12-layer BERT architecture. It combines sentence representation learning with MLM training by using encoded vectors of adjacent

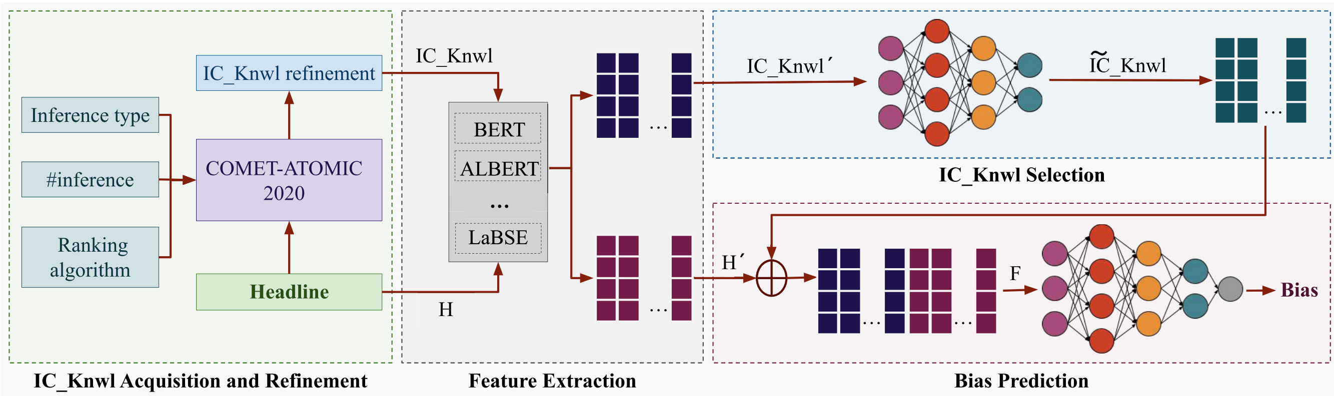

FIGURE 7. **Abstract representation of the bias prediction framework.**

sentences as a conditioning factor. Pre-trained model: https://tfhub.dev/google/universal-sentence-encoder-cmlm/en-base/

- **GPL (Generative Pseudo Labelling)** [71]: a method for training dense retrievers using unsupervised domain adaptation. It relies on pseudo-labelling in conjunction with robust cross-encoders and query generators. Pre-trained model: https://huggingface.co/GPL
- **GPL+BPR (Generative Pseudo Labelling + Binary Passage Retriever)** [72]: a method that optimises BPR [73] using GPL. It achieves memory efficiency by employing compact binary codes instead of continuous vectors, and domain adaptation efficiency without requiring domain-specific annotated training data. Model: https://github.com/NThakur20/income
- **LaBSE (Language-agnostic BERT Sentence Embedding)** [74]: a BERT-based, language-agnostic model. It employs a novel combination of masked language modelling, translation language modelling, dual encoder translation ranking, and additive margin softmax to establish a new benchmark for bi-text mining. Pre-trained model: https://tfhub.dev/google/LaBSE

C. METHODOLOGY

Our proposed framework (IC-BAIT), as illustrated in Fig. 7, is built around the following key components: IC_Knwl acquisition and refinement, feature extraction, IC_Knwl selection, and bias prediction. We sketch the pseudocode for IC-BAIT in Algorithm 1 and present a detailed description of each component in the following subsections.

1) IC_Knwl ACQUISITION AND REFINEMENT

IC_Knwl assists in simplifying, comprehending, and explaining events that are not explicitly stated in the headline. To acquire it for each headline H, we use the neural knowledge model *COMET* (COMmonsensE Transformers)[2] [13] trained on the $ATOMIC_{20}^{20}$ (ATlas Of MachIne Commonsense-2020)[3] [13] knowledge graphs:

$$IC\_Knwl = COMET(H, I_{type}, r, k) \tag{3}$$

where I_{type}, r, and k denote the inference type, ranking algorithm, and the number of returned references respectively. We set I_{type} = {*oEffect*, *oReact*, *oWant*, *xAttr*, *xEffect*, *xIntent*, *xNeed*, *xReact*, and *xWant*}, r = *beam*, $k = 3$. Fig. 8 defines the inference types.

For each I_{type} inference in *IC_Knwl*, we first determine whether or not the returned list is empty. We then assign *'none'* to *IC_Knwl* if the list is empty or if all of the entries are *'blank'* or *'garbage'*. If the list isn't empty and at least one

Algorithm 1 **bias prediction algorithm.**

Input: headline text H, true bias label b, inference type I_{type}, ranking algorithm r, number of references k
Output: predicted bias label $\hat{b}$

while not converge **do**
  for each headline H **do**
    extract *IC_Knwl* using (H, I_{type}, r, k)
      ▷ ref. Eq. 3
    refine *IC_Knwl*
      ▷ ref. Section 3
    extract feature vectors H' and IC_Knwl'
      ▷ ref. Section IV-C2
    filter and select relevant knowledge $\widetilde{IC}\_Knwl$
      ▷ ref. Eq. 4
    generate fused feature vector F using H' and $\widetilde{IC}\_Knwl$
      ▷ ref Eq. 5
    get predicted label $\hat{b}$
      ▷ ref Eq. 6
    calculate prediction *Loss*
      ▷ ref. Eq. 7
    backpropagate *Loss* to update IC-BAIT
  end for
end while

[2]https://github.com/allenai/comet-atomic-2020/

[3]https://allenai.org/data/atomic-2020

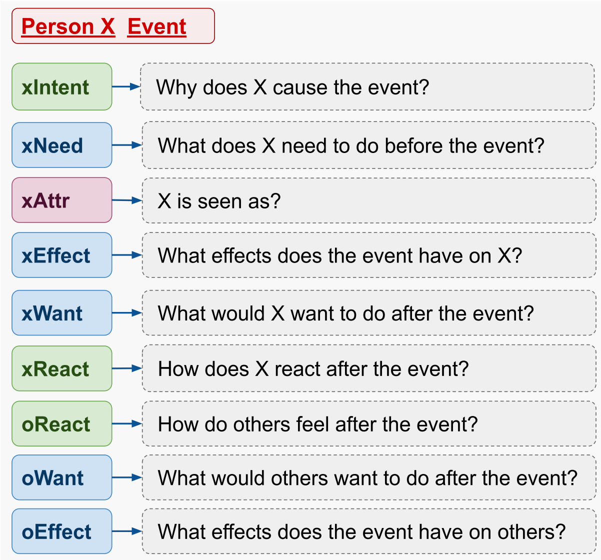

FIGURE 8. **Interpretation of ATOMIC$^{20}_{20}$ inference types.**

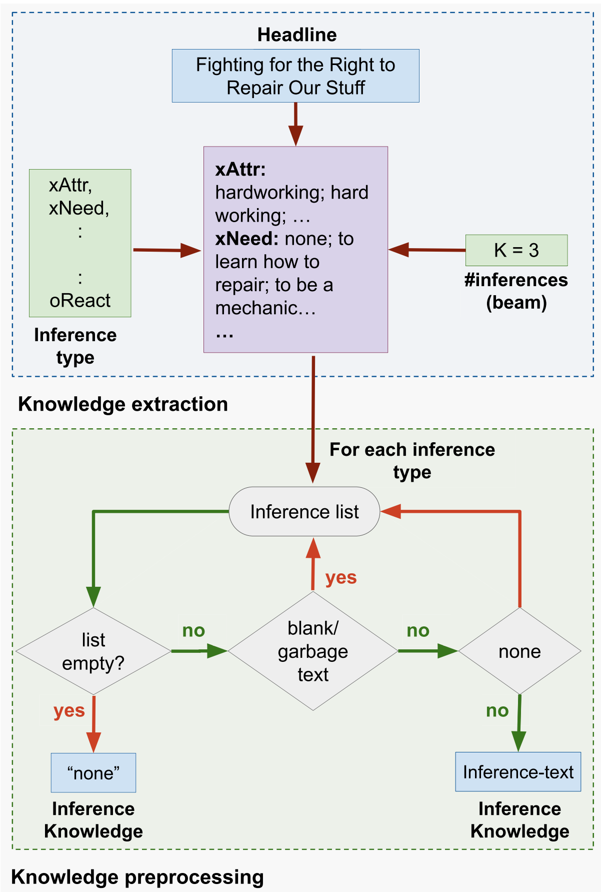

FIGURE 9. **Flowchart of inferential knowledge acquisition process.**

item isn't *blank* or *garbage*, we add the first such item to the *IC_Knwl* list. Fig. 9 illustrates the flowchart of the IC_Knwl acquisition process.

To render *IC_Knwl* more meaningful, we combine the retrieved inferences into a statement describing the properties of the person involved, for example,

1) **Headline:** Fighting for the Right to Repair Our Stuff
2) **Raw IC_Knwl:** xAttr: *hardworking*, xNeed: *to learn how to repair*, xIntent: *to be able to repair*, xEffect: *fighting for the right to repair*, xWant: *to make sure it's fixed*, xReact: *proud*, oWant: *to help others*, oEffect: *fight for rights*, oReact: *happy*
3) **Processed IC_Knwl:** PersonX is *hardworking*, needed *to learn how to repair*, intended *to be able to repair*, *fighting for the right to repair*, *to make sure it's fixed*, feels *proud*. Others want *to help others*, *fight for rights*, feel *happy*.

2) FEATURE EXTRACTION

To acquire the feature vectors H' and IC_Knwl' from the headline (H) and Inferential Commonsense Knowledge (IC_Knwl), we use the state-of-the-art pre-trained language models that we describe in Section IV-B. IC-BAIT is model-agnostic, and thus models of any depth can be easily integrated.

3) KNOWLEDGE SELECTION

Using only one of the inference types to predict the bias may not be ideal. For instance, if we rely exclusively on one of the inference types, such as Person X's reaction, the resulting inferential knowledge may aid in sentiment prediction but not in comprehending the context of the event for bias prediction. As a result, we use all of the inference types for our task. However, not all of them are equally important. Consequently, we utilise a method that resembles an established and efficacious approach to selecting relevant knowledge [12], [75].

To filter the relevant knowledge, we first apply the Sigmoid function [76] to IC_Knwl', to measure the importance of each inference. Specifically, we employ it as a cumulative density function of the logistic distribution to assign probabilities to the inferences [77]. We then multiply IC_Knwl' by the resulting importance scores and feed the resultant vector to a Multi-Layer Perceptron (MLP) network that learns to mix the inferences of different I_{type} to generate $\widetilde{IC}\_Knwl$:

$$\widetilde{IC}\_Knwl = MLP(sigmoid(IC\_Knwl') \odot IC\_Knwl') \quad (4)$$

where $\odot$ denotes element-wise multiplication.

4) BIAS PREDICTION

For the task of bias prediction, we first generate the fused vector F by concatenating H' and $\widetilde{IC}\_Knwl$:

$$F = H' \oplus \widetilde{IC}\_Knwl \quad (5)$$

where $\oplus$ represents the concatenation operation. We then feed F to the MLP network. To finally predict the bias label $\hat{b}$, we pass the resultant vector to the Fully Connected layer (FC) having Softmax (σ) activation with three output neurons that represent bias labels in B.

$$\hat{b} = FC(\sigma(MLP(F)) \tag{6}$$

For the training process, we use categorical cross-entropy as the loss function:

$$Loss = -\sum_{i=1}^{|B|}(b_i * \log(\hat{b}_i)) \tag{7}$$

where b_i and $\hat{b}_i$ denotes the actual and predicted probability of selecting the i^{th} bias label in B. We release the detailed experimental settings at https://github.com/Swati17293/IC-BAIT

D. EVALUATION METRICS

We use the following well-known metrics [78] for performance evaluation:

- **Accuracy (Acc.)** measures the proportion of correct predictions and overall predictions.
- **$\mathbf{F_1}$-score** is the weighted average of Precision (P) and Recall (R) calculated as:

$$\text{F}_1\text{-score} = 2 * ((P * R)/(P + R)) \tag{8}$$

  where P denotes the proportion of correctly predicted instances of a bias category, say (c), to the total number of predicted instances of that category, and R denotes the proportion of correctly predicted instances out of the total number of actual instances that fall under the category c.
- **Jaccard-score (J)** [79] computes the fraction of correctly predicted instances over all instances, excluding true negatives. It disregards true negatives in favour of true positives. For an imbalanced dataset like the ones used in this study, where true negatives outnumber true positives, it helps to gain a deeper understanding of the results.

We report a macro-averaged F_1-score ($\mathbf{F_1}$) to ensure that all bias classes are treated equally. In addition, we report a micro-averaged F_1-score ($\mathbf{F_{1\mu}}$) and a micro-averaged Jaccard-score ($\mathbf{J_\mu}$) to account for the problem of class imbalance.

V. RESULTS AND ANALYSIS

We begin the section by presenting the experimental results. We then thoroughly examine the experimental results to study the impact of IC-BAIT on the task of bias prediction. We conclude with a comprehensive analysis to determine the scenarios under which IC-BAIT is beneficial and when it is counterproductive.

A. EXPERIMENTAL RESULTS

Table 2 shows the overall performance distribution of the baseline models and our framework. For each metric, it reports the average score across the models. For comparison, we use Fig. 10a and Fig. 10b to illustrate the

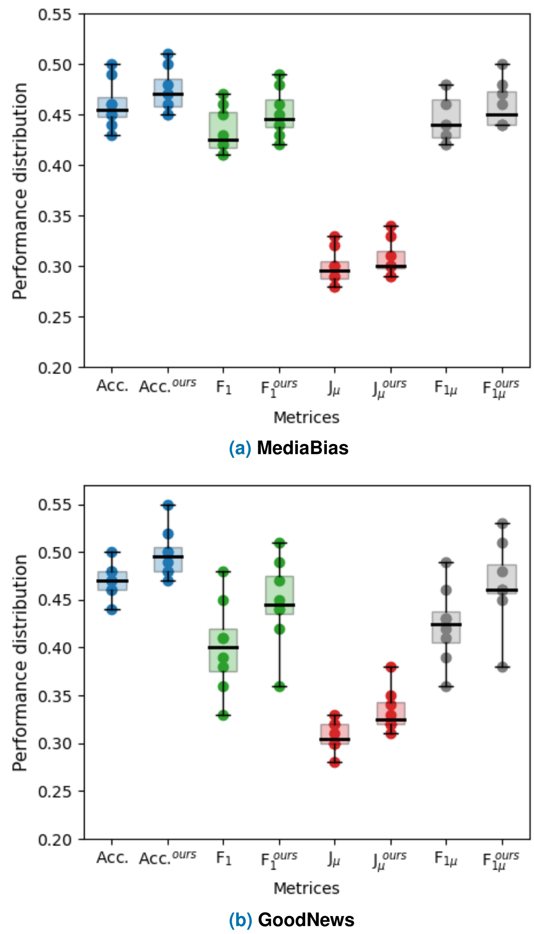

FIGURE 10. Boxplots of overall performance distributions for the datasets (a) MediaBias and (b) GoodNews, where the presence of a median near the third quartile (Q3) and a relatively high Q3 indicates superior performance.

distributional characteristics of overall performance scores for all reported metrics across the MediaBias and GoodNews datasets. We use boxplots for visualisation as they are reliable and can reveal a lot of statistical information, such as ranges, outliers, and medians. For instance, the line positioned in the middle of the box represents the median, with scores above it indicating the top 50% of results. Similarly, the line at the top part of the box signifies the third quartile (Q3) with scores above it indicating the top 25% of results. From the plots, it can be clearly seen that our framework has superior performance for each metric on both datasets, with a comparatively higher median and Q3.

In general, a model performs better when there are more training samples. However, interestingly, the results reveal that the models perform better on the smaller dataset (GoodNews) than on the larger one (MediaBias). The reason for this can be attributed to the dataset's characteristics. Whereas a team of experts selected the headlines in the MediaBias dataset to present the news stories from different political perspectives, the GoodNews dataset was assembled to classify emotions, which may

TABLE 2. Overall performance distribution of the baseline models and our proposed framework. Percentage improvement, denoted by %↑, indicates that our framework significantly improves prediction performance on both datasets.

Dataset	Acc.	Acc.^{ours}	%↑	F_1	F_1^{ours}	%↑	J_μ	J_μ^{ours}	%↑	$F1_\mu$	$F1_\mu^{ours}$	%↑
MediaBias	0.46	0.47	**2.2**	0.43	0.45	**4.6**	0.30	0.31	**3.3**	0.45	0.46	**2.2**
GoodNews	0.47	0.50	**6.4**	0.40	0.45	**12.5**	0.31	0.33	**6.4**	0.42	0.47	**11.9**

TABLE 3. Detailed performance distribution of the baseline models and our proposed framework for the datasets (a) MediaBias and (b) GoodNews, with %↑ representing the percentage improvement.

(a) MediaBias

	Acc.	Acc.^{ours}	%↑	F_1	F_1^{ours}	%↑	J_μ	J_μ^{ours}	%↑	$F1_\mu$	$F1_\mu^{ours}$	%↑
ALBERT [63]	0.43	0.45	**4.6**	0.41	0.44	**7.3**	0.28	0.29	**3.6**	0.42	0.44	**4.8**
DistilRoBERTa [64]	0.45	0.46	**2.2**	0.42	0.43	**2.4**	0.29	0.30	**3.4**	0.44	0.44	0.0
MPNet [66]	0.46	0.48	**4.3**	0.45	0.46	**2.2**	0.30	0.31	**3.3**	0.46	0.47	**2.2**
MiniLMv2 [68]	0.44	0.45	**2.3**	0.41	0.42	**2.4**	0.28	0.29	**3.6**	0.42	0.44	**4.8**
CMLM [70]	0.50	0.51	**2.0**	0.47	0.49	**4.2**	0.33	0.34	**3.0**	0.48	0.50	**4.2**
GPL [71]	0.45	0.47	**4.4**	0.43	0.44	**2.3**	0.29	0.30	**3.4**	0.44	0.44	0.0
GPL+BPR [72]	0.46	0.47	**2.2**	0.42	0.45	**7.1**	0.30	0.30	0.0	0.43	0.46	**7.0**
LaBSE [74]	0.49	0.50	**2.0**	0.46	0.48	**4.3**	0.32	0.33	**3.1**	0.48	0.48	0.0

(b) GoodNews

	Acc.	Acc.^{ours}	%↑	F_1	F_1^{ours}	%↑	J_μ	J_μ^{ours}	%↑	$F1_\mu$	$F1_\mu^{ours}$	%↑
ALBERT [63]	0.46	0.48	**4.3**	0.45	0.47	**4.4**	0.30	0.32	**6.7**	0.46	0.48	**4.3**
DistilRoBERTa [64]	0.47	0.48	**2.1**	0.39	0.45	**15.4**	0.30	0.32	**6.7**	0.42	0.46	**9.5**
MPNet [66]	0.48	0.52	**8.3**	0.41	0.49	**19.5**	0.32	0.35	**9.4**	0.43	0.51	**18.6**
MiniLMv2 [68]	0.46	0.49	**6.5**	0.41	0.42	**2.4**	0.30	0.32	**6.7**	0.43	0.45	**4.6**
CMLM [70]	0.48	0.50	**4.2**	0.38	0.44	**15.8**	0.32	0.34	**6.2**	0.41	0.46	**12.2**
GPL [71]	0.50	0.55	**10.0**	0.48	0.51	**6.2**	0.33	0.38	**15.1**	0.49	0.53	**8.2**
GPL+BPR [72]	0.44	0.47	**6.8**	0.33	0.36	**9.1**	0.28	0.31	**10.7**	0.36	0.38	**5.5**
LaBSE [74]	0.47	0.50	**6.0**	0.36	0.44	**22.2**	0.31	0.33	**6.4**	0.39	0.46	**17.9**

have contributed to improved performance. The fact that MediaBias is event-centric whereas GoodNews is not could also be a contributing factor.

To aid clarity and comprehension of the findings, Table 3 presents a detailed breakdown of the scores for each model. In a nutshell, IC-BAIT improves the baseline models in terms of Acc. (2.0-10.0%), F_1 (2.2-22.2%), J_μ (up to 15.1%), and $F1_\mu$ (up to 18.6%). The results reveal that CMLM for MediaBias and GPL for GoodNews stand out as the best-performing models, with impressive percentage improvements across all metrics. In a nutshell, our findings indicate that the integration of baseline models within our framework has a notable impact on prediction performance.

B. CASE STUDY

We investigate the prediction results of IC-BAIT (with and without IC_Knwl) for Left, Right, and Center oriented headlines and present a representative example from each set in Table 4.

In the study, it can be observed that IC-BAIT without IC_Knwl appears to fail in the majority of cases where the entities and/or terms mentioned in the headline are strongly correlated to an ideology other than that of the headline. In these cases, the model is more likely to learn the unjust correlation and predict it incorrectly. However, by incorporating IC_Knwl, IC-BAIT improves its understanding of the inferred meaning and makes predictions guided by that. As a result, it gains the ability to focus not only on the important entities and events mentioned in the headline but also on the explanation of unstated events for better comprehension.

1) IMPACT of IC_Knwl

To better understand the impact of IC_Knwl, it is necessary to investigate whether or not it guides the model to take a specific decision. For that, we use the LIME (Local Interpretable Model-Agnostic Explanations) [80] xAI framework, a popular qualitative interpretation tool. It provides a set of feature (word) weights as well as colour-coded text to aid in the explanation of a model prediction.

We first use the headline as the framework's sole input, and then we separate the misclassified instances, for which we

TABLE 4. Case study of bias label predictions by IC-BAIT (with and without IC_Knwl) for Left, Right, and Center oriented headlines.

Headline	**Donald Trump gets ripped to pieces over LGBT tweet**
IC_Knwl	PersonX is intolerant, needed to write a homophobic tweet, intended to be liked, gets yelled at, wants to make amends, feels upset. Others want to defend themselves, gets hurt, feel angry.
Bias Label	True: **Left**, Predicted(without IC_Knwl): **Right**, Predicted(with IC_Knwl): **Left**
Comment	Due to data bias, news related to the named entity "Trump" is heavily skewed to Right-wing media. Furthermore, idioms such as "ripped to pieces" are typically associated with the Right ideology. The model without IC_Knwl tends to learn this unjust correlation and thus ends up predicting it as "Right". However, with the additional commonsense inference, important information such as "personX is seen as intolerant" and "Others get hurt" was passed to the model, allowing it to learn the prediction correctly.
Headline	**Time to Kick the Islamizing Turkey Out of NATO**
IC_Knwl	PersonX is aggressive, needed to be a member of NATO, intended to get rid of terrorism, gets yelled at, wants to to get rid of the Islamists, feels angry. Others want to fight back, gets hurt, feel angry.
Bias Label	True: **Right**, Predicted(without IC_Knwl): **Left**, Predicted(with IC_Knwl): **Right**
Comment	An incorrect correlation between "Islam" and Left-wing media in the collected data causes the model without IC_Knwl to incorrectly predict the label as "Left". However, the acquired commonsense inferences such as, "personX is seen as aggressive and gets yelled at" provide a critical understanding of the statement, allowing the model with IC_Knwl to correctly learn the prediction.
Headline	**The refugee families caught up in a war zone in Libya**
IC_Knwl	PersonX is traumatized, needed to be deployed, intended to escape from war, gets lost in war, wants to find a new home, feels scared. Others want to get out of there, have to find a new home, feel sad.
Bias Label	True: **Center**, Predicted(without IC_Knwl): **Left**, Predicted(with IC_Knwl): **Center**
Comment	The terms "refugee" and "caught" in the headline indicate that the headline is right-oriented. However, with the given commonsense inference, such as "personX wants to find a new home" the model selects the Neutral or Central narration of the headline.

generate explanations to understand the probable reasons for misclassification. We then include IC_Knwl in the framework with the headline and regenerate the explanations for the same instances to see its impact on the prediction. Fig. 11 and Fig. 12 illustrate results for the top five significant features.

We follow the same procedure for the rest of the misclassified instances and conduct an in-depth analysis to determine the cases in which IC_Knwl is useful, and we reach the following conclusions:

- **Short headlines**: they lack the necessary contextual information to make an accurate prediction.
  Example: the short headlines, such as "*Brickbat: It's a Gas Gas Gas*" and "*Grit Won*", are insufficiently descriptive.
- **Entity**: whereas entities mentioned in headlines carry considerable weight, in the absence of contextual information, they quite often confuse the model.
  Example: although the entity "*FDA chief*" is the main subject of the headlines from different political ideologies, it has no bearing on definite ideology (ref. Table 5).
- **Metaphor**: while the words used in metaphors are significantly important, their interpretation may or may not be symbolic depending on their context. Typically, their literal and symbolic meanings are quite dissimilar, which complicates prediction.
  Example: consider the following two headlines: "*Trump: Al-Baghdadi Died Like a Dog*" and "*Ari Fleischer Lied and People Died*". The word "*died*" in the metaphor "*died like a dog*" refers to "*died in a painful and humiliating manner*" and thus carries a negative emotional pull. The second headline, on the other hand, uses the same word "*died*", implying an emotion such as grief.
- **Domain-specific/slang words**: these words are typically unimportant and provide little or no context to other significant entities (if present), and may lead to misclassification.
  Example: in the headline "*Journos Hit Hick for Pot Flip-Flop*", "*journos*" is slang for "*Journalists*" which is difficult for a model to comprehend. "*flip-flop*" in turn is a domain-specific word (political news) that refers to "*a sudden reversal of proposals*" in contrast to its general definition "*a type of footwear*".
- **Satire**: it appears to be a straightforward statement, but it contains humorous criticism.
  Example: the headline "*America's Racist Legacy From Slavery to the War on Immigrants*" is a criticism of history being whitewashed.

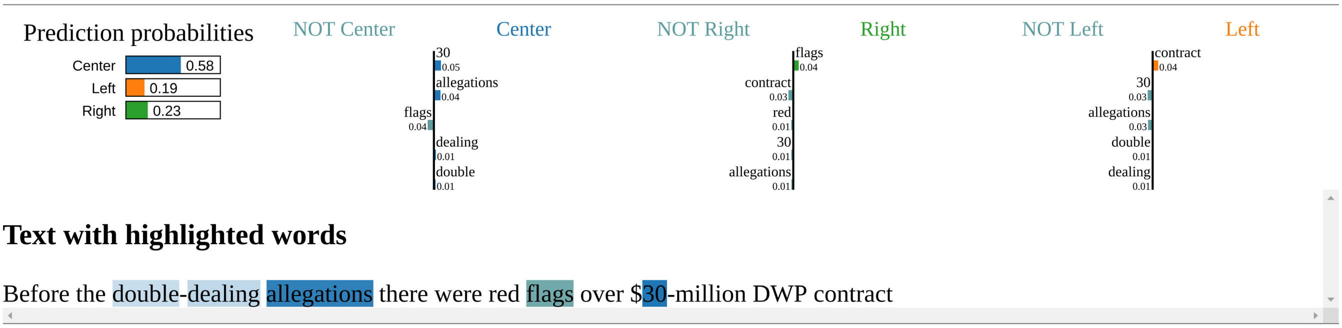

(a) Prediction framework *without* IC_Knwl, True: *Left*, Predicted: *Center*

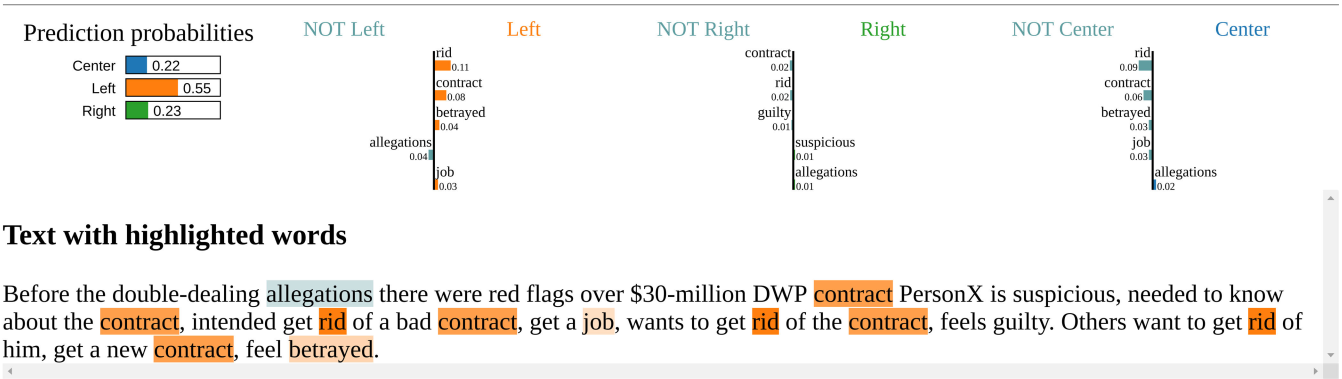

(b) Prediction framework *with* IC_Knwl, True: *Left*, Predicted: *Left*

FIGURE 11. **LIME visualization result: words like "*allegations*" in the headline 11a strongly indicates that the headline is center-oriented. However, the addition of words such as "*rid*" and "*betrayed*" from IC_Knwl results in a final prediction score of *0.55*, making the model 11b quite confident in its prediction.**

- **Sarcasm**: sarcastic headlines imply the inverse of what is stated plainly. When used sarcastically, words with a larger weight mislead the model.
  Example: the headline "*Trump: Illegal immigrants in sanctuary cities would make the 'Radical Left' happy*" refers to Trump's sarcastic remark about Left-wing supporters. Words like "*Trump*", "*Radical Left*", and "*happy*" convey a positive emotion rather than negative taunting.

Additionally, we identified instances where including IC_Knwl adds little or no value. For example,

- **When IC_Knwl is comprised of generic information.** There are numerous instances where IC_Knwl returns nearly identical information for different headlines. Consider the following headlines and their IC_Knwl as an example:
  - **Headline**: Bernie Sanders Wins Indiana Democratic Primary. (Bias Label: Left)
    ICKnwl: *PersonX is hopeful, needed to run for office, intended to win the election, wins the election*, wants to get elected, *feels happy*. Others want to congratulate him, *wins the election, feel happy*.
  - **Headline**: Biden Sweeps Dem Primaries in Arizona Florida Illinois. (Bias Label: Right)
    ICKnwl: *PersonX is hopeful, needed to run for office, intended to win the election, wins the election*, wants to win the election, *feels happy*. Others want to win the election, *wins election, feel happy*.
- **When IC_Knwl is devoid of significance.** There are headlines where the inferences drawn by IC_Knwl are insignificant (ref. Fig. 13).
- **When the world knowledge overpowers IC_Knwl.** There are instances where world knowledge is more important than IC_Knwl for the ideology prediction (ref. Fig. 14).

VI. LIMITATIONS AND FUTURE DIRECTIONS

The limitations of our framework originate primarily from its key components. To begin with, including IC_Knwl renders little or no significant improvement when it lacks any useful information (ref. Section V-B1). For instance, when it is comprised of generic information, is insignificant, or when world knowledge overpowers it. Similarly, there are instances where IC_Knwl alone is sufficient for accurate prediction. Identifying when to include and exclude IC_Knwl precisely and quantitatively continues to pose a challenge. Another limitation applies to its reliability. Augmenting explicable justifications for its outcomes would greatly enhance its reliability.

(a) Prediction model *without* IC_Knwl, True: *Center*, Predicted: *Left*

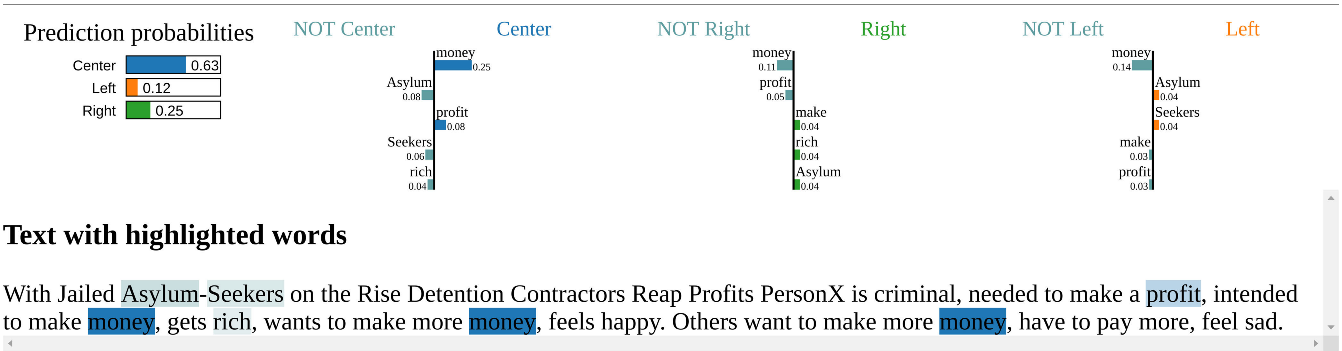

(b) Prediction model *with* IC_Knwl, True: *Center*, Predicted: *Center*

FIGURE 12. **LIME visualization result: the headline contains too many important words representing different ideologies, which confuses the model 12a. Using IC_Knwl, the model 12b, on the other hand, correctly predicts the ideology.**

TABLE 5. **IC_Knwl acquired from headlines of varying political ideologies reporting on the same event. The inclusion of IC_knwl facilitates prediction.**

	Center	**Left**	**Right**
Headline	FDA chief apologizes for overstating plasma effect on virus	Trump's FDA chief apologizes for hyping unproven treatment on eve of Republican National Convention	FDA chief clarifies remarks about COVID-19 treatment following criticism
xAttr	careless	irresponsible	helpful
xNeed	to be in charge	to be honest	to make a statement
xIntent	to be honest	to be a better representative	to clear up confusion
xEffect	gets reprimanded	has to apologize	gets asked for clarification
xWant	to make amends	to apologize to the public	to make sure everyone knows the facts
xReact	sorry	ashamed	relieved
oWant	to ask for an explanation	to ask for an apology	to thank the FDA
oEffect	they get sick	gets fired from job	they get better
oReact	angry	scared	informed

Furthermore, given that our framework is model-independent, it can be used to predict bias in headlines that are not in English. However, this scenario is only possible if the language model can process multiple languages and the extracted IC_Knwl is in the target language. For instance, the multilingual language model CMLM [70] has the potential to serve as a multilingual language model. It is a fully unsupervised learning model that can readily be extended to a wide range of languages. For knowledge extraction, MultiCOMET [81] is a viable option. However, when used in this setting, the quality of knowledge it produces will directly influence the framework's performance.

In addition, advanced methods for knowledge selection and its subsequent integration into the framework remain unexplored, despite their potential to significantly improve the framework's performance. Also, the large-scale language models used in the study require considerable resources. Consequently, there remains a need to investigate strategies

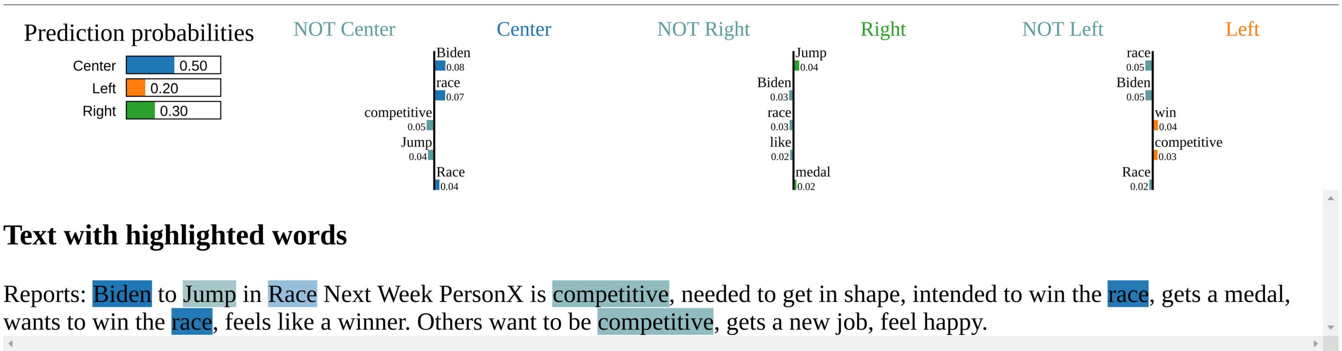

FIGURE 13. **True: *Right*, Predicted: *Center*. A flashy headline is typically associated with right-wing ideology. However, including IC_Knwl, which is insignificant for this headline, confuses the model.**

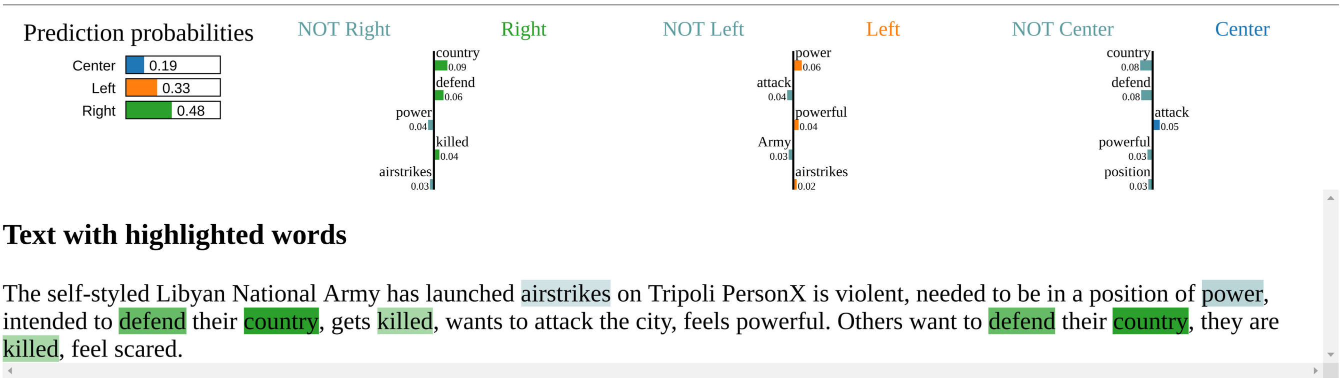

FIGURE 14. **True: *Left*, Predicted: *Right*. In order to correctly predict the ideology of this headline, the world's knowledge that the "Libyan National Army" is associated with the "Left-wing" is more important than the IC_Knwl.**

to accelerate training time [82], [83]. Finally, the limitations of the knowledge base and models employed within the framework are also inherent limitations of our proposed framework.

Besides handling the above limitations in the future, additional inferential knowledge bases and their impact on prediction performance can be investigated. Moreover, additional input data, such as information derived from external knowledge bases, can also be pursued.

VII. CONCLUSION

Detecting bias in a news headline can be difficult, as they are typically short and may lack the context of bias embedded in the underlying article. Furthermore, it is challenging to capture syntactic and semantic information in its short text. We propose our framework, *IC-BAIT* to address these challenges. It uses IC_Knwl acquired from *ATOMIC*$^{20}_{20}$ graphs using *COMET* to enhance political bias prediction in news headlines. It employs IC_Knwl in a manner that is very similar to how people use commonsense inferences to carry out their day-to-day tasks.

Due to the scarcity of large-scale datasets for the given task, we also present two datasets: *MediaBias* and *GoodNews*. We conduct experiments on both of them to determine IC-BAIT's generalisability. The results demonstrate that the baseline models, when used within IC-BAIT, had superior performance on both datasets.

We conduct an in-depth, case-by-case investigation using LIME (Local Interpretable Model-Agnostic Explanations), an xAI framework. We discover that, while IC_Knwl can be extremely beneficial in some cases, it can also be counterproductive in others. In a nutshell, when the headline is combined with a carefully chosen IC_Knwl, the model can focus not only on the important entities and events in the headline but also on the explanations for unstated events, resulting in better predictions.

Although the task of classifying political bias in news headlines is still in its early stages, we believe that our method and findings will prove beneficial for copy editors, practical application domains such as e-journalism and manual news bias prediction portals, and automated headline bias-flipping systems, among others.

REFERENCES

[1] N. Hagar, J. Wachs, and E.-Á. Horvát, "Writer movements between news outlets reflect political polarization in media," *New Media Soc.*, vol. 2021, Jan. 2021, Art. no. 14614448211027173, doi: 10.1177/14614448211027173.

[2] P. Maurer and M. Beiler, "Networking and political alignment as strategies to control the news: Interaction between journalists and politicians," *Journalism Stud.*, vol. 19, no. 14, pp. 2024–2041, Oct. 2018, doi: 10.1080/1461670X.2017.1310627.

[3] D. Dor, "On newspaper headlines as relevance optimizers," *J. Pragmatics*, vol. 35, no. 5, pp. 695–721, 2003, doi: 10.1016/S0378-2166(02)00134-0.

[4] N. Hagar, N. Diakopoulos, and B. DeWilde, "Anticipating attention: On the predictability of news headline tests," *Digital Journalism*, vol. 10, no. 4, pp. 1–22, 2021, doi: 10.1080/21670811.2021.1984266.

[5] P. H. Tannenbaum, "The effect of headlines on the interpretation of news stories," *Journalism Quart.*, vol. 30, no. 2, pp. 189–197, Mar. 1953, doi: 10.1177/107769905303000206.

[6] W.-F. Chen, H. Wachsmuth, K. Al-Khatib, and B. Stein, "Learning to flip the bias of news headlines," in *Proc. 11th Int. Conf. Natural Lang. Gener.*, 2018, pp. 79–88, doi: 10.18653/v1/W18-6509.

[7] W.-F. Chen, K. Al Khatib, H. Wachsmuth, and B. Stein, "Analyzing political bias and unfairness in news articles at different levels of granularity," in *Proc. 4th Workshop Natural Lang. Process. Comput. Social Sci.*, 2020, pp. 149–154, doi: 10.18653/v1/2020.nlpcss-1.16.

[8] D. Li, X. Zhu, Y. Li, S. Wang, D. Li, J. Liao, and J. Zheng, "Enhancing emotion inference in conversations with commonsense knowledge," *Knowl.-Based Syst.*, vol. 232, Nov. 2021, Art. no. 107449, doi: 10.1016/j.knosys.2021.107449.

[9] S. B. R. Chowdhury and S. Chaturvedi, "Does commonsense help in detecting sarcasm?" in *Proc. 2nd Workshop Insights Negative Results NLP*, 2021, pp. 9–15, doi: 10.18653/v1/2021.insights-1.2.

[10] T. Chakrabarty, D. Ghosh, S. Muresan, and N. Peng, "R^3: Reverse, retrieve, and rank for sarcasm generation with commonsense knowledge," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics*, 2020, pp. 7976–7986, doi: 10.18653/v1/2020.acl-main.711.

[11] L. Huang, R. Le Bras, C. Bhagavatula, and Y. Choi, "Cosmos QA: Machine reading comprehension with contextual commonsense reasoning," in *Proc. Conf. Empirical Methods Natural Lang. Process. 9th Int. Joint Conf. Natural Lang. Process. (EMNLP-IJCNLP)*, 2019, pp. 2391–2401, doi: 10.18653/v1/D19-1243.

[12] S. Sabour, C. Zheng, and M. Huang, "CEM: Commonsense-aware empathetic response generation," in *Proc. Conf. Artif. Intell. (AAAI)*, Feb. 2022, vol. 36, no. 10, pp. 11229–11237, doi: 10.1609/aaai.v36i10.21373.

[13] J. D. Hwang, C. Bhagavatula, R. Le Bras, J. Da, K. Sakaguchi, A. Bosselut, and Y. Choi, "(comet-) atomic 2020: On symbolic and neural commonsense knowledge graphs," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, 2021, pp. 6384–6392. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/16792

[14] A. Einsohn and M. Schwartz, *The Copyeditor's Handbook*. Berkeley, CA, USA: Univ. California Press, 2019, doi: 10.1525/9780520972278.

[15] T. Spinde, D. Krieger, M. Plank, and B. Gipp, "Towards a reliable ground-truth for biased language detection," in *Proc. ACM/IEEE Joint Conf. Digit. Libraries (JCDL)*, Sep. 2021, pp. 324–325, doi: 10.1109/JCDL52503.2021.00053.

[16] R. Liu, L. Wang, C. Jia, and S. Vosoughi, "Political depolarization of news articles using attribute-aware word embeddings," in *Proc. Int. AAAI Conf. Web Social Media*, vol. 15, 2021, pp. 385–396. [Online]. Available: https://ojs.aaai.org/index.php/ICWSM/article/view/18069

[17] R. Liu, C. Jia, and S. Vosoughi, "A transformer-based framework for neutralizing and reversing the political polarity of news articles," *Proc. ACM Hum.-Comput. Interact.*, vol. 5, no. CSCW1, pp. 1–26, Apr. 2021, doi: 10.1145/3449139.

[18] T. Groseclose and J. Milyo, "A measure of media bias," *Quart. J. Econ.*, vol. 120, no. 4, pp. 1191–1237, Nov. 2005, doi: 10.1162/003355305775097542.

[19] F. Hamborg, K. Donnay, and B. Gipp, "Automated identification of media bias in news articles: An interdisciplinary literature review," *Int. J. Digit. Libraries*, vol. 20, no. 4, pp. 391–415, Dec. 2019, doi: 10.1007/s00799-018-0261-y.

[20] P. Navia and R. Osorio, "*El MercurioLies*, and *La TerceraLies* more. Political bias in newspaper headlines in Chile, 1994–2010," *Bull. Latin Amer. Res.*, vol. 34, no. 4, pp. 467–485, Oct. 2015, doi: 10.1111/blar.12364.

[21] M. McCluskey, "A content analysis of 2004 presidential election headlines of The Los Angeles Times and The Washington Times," Ph.D. dissertation, Dept. Commun., College Arts Sci., Univ. Central Florida Orlando, Orlando, FL, USA, 2005. [Online]. Available: https://stars.library.ucf.edu/etd/358

[22] V. P. Aires, F. G. Nakamura, and E. F. Nakamura, "A link-based approach to detect media bias in news websites," in *Proc. Companion World Wide Web Conf.*, May 2019, pp. 1–9. [Online]. Available: https://sentic.net/wisdom2020aires.pdf

[23] Z. Xu, "Examining us newspapers' partisan bias in COVID-19 news using computational methods," *Commun. Stud.*, vol. 2023, pp. 1–19, Jan. 2023, doi: 10.1080/10510974.2023.2169729.

[24] S. D'Alonzo and M. Tegmark, "Machine-learning media bias," *PLoS ONE*, vol. 17, no. 8, Aug. 2022, Art. no. e0271947, doi: 10.1371/journal.pone.0271947.

[25] B. Sinno, B. Oviedo, K. Atwell, M. Alikhani, and J. J. Li, "Political ideology and polarization of policy positions: A multi-dimensional approach," 2021, *arXiv:2106.14387*.

[26] S. Lim, A. Jatowt, M. Färber, and M. Yoshikawa, "Annotating and analyzing biased sentences in news articles using crowdsourcing," in *Proc. 12th Lang. Resour. Eval. Conf.*, 2020, pp. 1478–1484. [Online]. Available: https://www.aclweb.org/anthology/2020.lrec-1.184

[27] T. Spinde, J.-D. Krieger, T. Ruas, J. Mitrović, F. Götz-Hahn, A. Aizawa, and B. Gipp, "Exploiting transformer-based multitask learning for the detection of media bias in news articles," in *Information for a Better World: Shaping the Global Future*. Cham, Switzerland: Springer, 2022, pp. 225–235, doi: 10.1007/978-3-030-96957-8_20.

[28] F. Sakketou, J. Plepi, R. Cervero, H. J. Geiss, P. Rosso, and L. Flek, "FACTOID: A new dataset for identifying misinformation spreaders and political bias," in *Proc. 13th Lang. Resour. Eval. Conf.*, 2022, pp. 3231–3241. [Online]. Available: https://aclanthology.org/2022.lrec-1.345

[29] F. Morstatter, L. Wu, U. Yavanoglu, S. R. Corman, and H. Liu, "Identifying framing bias in online news," *ACM Trans. Social Comput.*, vol. 1, no. 2, pp. 1–18, Jun. 2018, doi: 10.1145/3204948.

[30] L. Fan, M. White, E. Sharma, R. Su, P. K. Choubey, R. Huang, and L. Wang, "In plain sight: Media bias through the lens of factual reporting," in *Proc. Conf. Empirical Methods Natural Lang. Process. 9th Int. Joint Conf. Natural Lang. Process. (EMNLP-IJCNLP)*, 2019, pp. 6343–6349, doi: 10.18653/v1/D19-1664.

[31] J. Kulshrestha, M. Eslami, J. Messias, M. B. Zafar, S. Ghosh, K. P. Gummadi, and K. Karahalios, "Search bias quantification: Investigating political bias in social media and web search," *Inf. Retr. J.*, vol. 22, nos. 1–2, pp. 188–227, Apr. 2019, doi: 10.1007/s10791-018-9341-2.

[32] B. Sinno, B. Oviedo, K. Atwell, M. Alikhani, and J. J. Li, "Political ideology and polarization: A multi-dimensional approach," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2022, pp. 231–243, doi: 10.18653/v1/2022.naacl-main.17.

[33] R. Baly, G. Karadzhov, D. Alexandrov, J. Glass, and P. Nakov, "Predicting factuality of reporting and bias of news media sources," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2018, pp. 3528–3539, doi: 10.18653/v1/D18-1389.

[34] R. Baly, G. Karadzhov, J. An, H. Kwak, Y. Dinkov, A. Ali, J. Glass, and P. Nakov, "What was written vs. who read it: News media profiling using text analysis and social media context," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics*, 2020, pp. 3364–3374, doi: 10.18653/v1/2020.acl-main.308.

[35] S. Ganguly, J. Kulshrestha, J. An, and H. Kwak, "Empirical evaluation of three common assumptions in building political media bias datasets," in *Proc. Int. AAAI Conf. Web Social Media*, vol. 14, 2020, pp. 939–943. [Online]. Available: https://ojs.aaai.org/index.php/ICWSM/article/view/7362

[36] R. Baly, G. Da San Martino, J. Glass, and P. Nakov, "We can detect your bias: Predicting the political ideology of news articles," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2020, pp. 4982–4991, doi: 10.18653/v1/2020.emnlp-main.404.

[37] L. A. M. Oberländer, E. Kim, and R. Klinger, "GoodNewsEveryone: A corpus of news headlines annotated with emotions, semantic roles, and reader perception," in *Proc. 12th Lang. Resour. Eval. Conf.*, 2020, pp. 1554–1566. https://www.aclweb.org/anthology/2020.lrec-1.194/

[38] M. A. I. Mahmud, A. A. T. Talukder, A. Sultana, K. I. A. Bhuiyan, M. S. Rahman, T. H. Pranto, and R. M. Rahman, "Toward news authenticity: Synthesizing natural language processing and human expert opinion to evaluate news," *IEEE Access*, vol. 11, pp. 11405–11421, 2023, doi: 10.1109/ACCESS.2023.3241483.

[39] Q. Meng, Y. Song, J. Mu, Y. Lv, J. Yang, L. Xu, J. Zhao, J. Ma, W. Yao, R. Wang, M. Xiao, and Q. Meng, "Electric power audit text classification with multi-grained pre-trained language model," *IEEE Access*, vol. 11, pp. 13510–13518, 2023, doi: 10.1109/ACCESS.2023.3240162.

[40] J. Shin, S.-B. Park, and H.-J. Song, "Token-level fact correction in abstractive summarization," *IEEE Access*, vol. 11, pp. 1934–1943, 2023, doi: 10.1109/ACCESS.2022.3233854.

[41] Y. Sujana and H.-Y. Kao, "LiDA: Language-independent data augmentation for text classification," *IEEE Access*, vol. 11, pp. 10894–10901, 2023, doi: 10.1109/ACCESS.2023.3234019.

[42] Y. Lei, R. Huang, L. Wang, and N. Beauchamp, "Sentence-level media bias analysis informed by discourse structures," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2022, pp. 10040–10050. [Online]. Available: https://aclanthology.org/2022.emnlp-main.682

[43] N. R. Naredla and F. F. Adedoyin, "Detection of hyperpartisan news articles using natural language processing technique," *Int. J. Inf. Manage. Data Insights*, vol. 2, no. 1, Apr. 2022, Art. no. 100064, doi: 10.1016/j.jjimei.2022.100064.

[44] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, vol. 1, 2019, pp. 4171–4186, doi: 10.18653/v1/N19-1423.

[45] W.-F. Chen, K. Al Khatib, B. Stein, and H. Wachsmuth, "Detecting media bias in news articles using Gaussian bias distributions," in *Proc. Findings Assoc. Comput. Linguistics, EMNLP*, 2020, pp. 4290–4300, doi: 10.18653/v1/2020.findings-emnlp.383.

[46] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2018, pp. 1–13, doi: 10.18653/v1/N18-1202.

[47] R. R. R. Gangula, S. R. Duggenpudi, and R. Mamidi, "Detecting political bias in news articles using headline attention," in *Proc. ACL Workshop BlackboxNLP, Analyzing Interpreting Neural Netw. NLP*, 2019, pp. 77–84, doi: 10.18653/v1/W19-4809.

[48] T. Mihaylov and A. Frank, "Knowledgeable reader: Enhancing cloze-style reading comprehension with external commonsense knowledge," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics*, 2018, pp. 821–832, doi: 10.18653/v1/P18-1076.

[49] Q. Chen, X. Zhu, Z.-H. Ling, D. Inkpen, and S. Wei, "Neural natural language inference models enhanced with external knowledge," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics*, 2018, pp. 2406–2417, doi: 10.18653/v1/P18-1224.

[50] G. A. Miller, "WordNet: A lexical database for English," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, Nov. 1995, doi: 10.1145/219717.219748.

[51] L. Bauer, Y. Wang, and M. Bansal, "Commonsense for generative multi-hop question answering tasks," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2018, pp. 4220–4230, doi: 10.18653/v1/D18-1454.

[52] H. Zhang, Y. Song, Y. Song, and D. Yu, "Knowledge-aware pronoun coreference resolution," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, 2019, p. 867, doi: 10.18653/v1/p19-1083.

[53] R. Li, Z. Jiang, L. Wang, X. Lu, M. Zhao, and D. Chen, "Enhancing transformer-based language models with commonsense representations for knowledge-driven machine comprehension," *Knowl.-Based Syst.*, vol. 220, May 2021, Art. no. 106936, doi: 10.1016/j.knosys.2021.106936.

[54] A. Lieto, G. L. Pozzato, S. Zoia, V. Patti, and R. Damiano, "A commonsense reasoning framework for explanatory emotion attribution, generation and re-classification," *Knowl.-Based Syst.*, vol. 227, Sep. 2021, Art. no. 107166, doi: 10.1016/j.knosys.2021.107166.

[55] L. Du, X. Ding, K. Xiong, T. Liu, and B. Qin, "Enhancing pretrained language models with structured commonsense knowledge for textual inference," *Knowl.-Based Syst.*, vol. 254, Oct. 2022, Art. no. 109488, doi: 10.1016/j.knosys.2022.109488.

[56] A. Bosselut, H. Rashkin, M. Sap, C. Malaviya, A. Celikyilmaz, and Y. Choi, "COMET: Commonsense transformers for automatic knowledge graph construction," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, 2019, pp. 4762–4779, doi: 10.18653/v1/P19-1470.

[57] F. Ilievski, P. Szekely, and B. Zhang, "CSKG: The commonsense knowledge graph," in *Proc. 18th Int. Conf., ESWC*. Cham, Switzerland: Springer, 2021, pp. 680–696, doi: 10.1007/978-3-030-77385-4_41.

[58] M. Sap, R. Le Bras, E. Allaway, C. Bhagavatula, N. Lourie, H. Rashkin, B. Roof, N. A. Smith, and Y. Choi, "Atomic: An atlas of machine commonsense for if-then reasoning," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 3027–3035, doi: 10.1609/aaai.v33i01.33013027.

[59] R. Speer, J. Chin, and C. Havasi, "ConceptNet 5.5: An open multilingual graph of general knowledge," in *Proc. AAAI Conf. Artif. Intell.*, vol. 31, 2017, pp. 1–10, doi: 10.1609/aaai.v31i1.11164.

[60] B. Y. Lin, W. Zhou, M. Shen, P. Zhou, C. Bhagavatula, Y. Choi, and X. Ren, "CommonGen: A constrained text generation challenge for generative commonsense reasoning," in *Proc. Findings Assoc. Comput. Linguistics, EMNLP*, 2020, pp. 1823–1840, doi: 10.18653/v1/2020.findings-emnlp.165.

[61] H. A. Brechter. (Jun. 2021). *A Step-by-Step Guide to How Allsides Creates balanced News*. AllSides. [Online]. Available: https://www.allsides.com/blog/how-does-allsides-create-balanced-news

[62] V. Otero. (Sep. 2021). *White Paper: Multi-Analyst Content Analysis Methodology September 2021*. Ad Fontes Media. [Online]. Available: https://adfontesmedia.com/white-paper-2021/

[63] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "ALBERT: A lite BERT for self-supervised learning of language representations," 2019, *arXiv:1909.11942*.

[64] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "Distilbert, a distilled version of bert: Smaller, faster, cheaper and lighter," in *Proc. NeurIPS*, 2019, pp. 1–16. [Online]. Available: https://www.emc2-ai.org/assets/docs/neurips-19/emc2-neurips19-paper-33.pdf

[65] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A robustly optimized BERT pretraining approach," 2019, *arXiv:1907.11692*.

[66] K. Song, X. Tan, T. Qin, J. Lu, and T.-Y. Liu, "MPNet: Masked and permuted pre-training for language understanding," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 16857–16867. https://proceedings.neurips.cc/paper/2020/file/c3a690be93aa602ee2dc0ccab5b7b67e-Paper.pdf

[67] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "XLNet: Generalized autoregressive pretraining for language understanding," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 1–4. [Online]. Available: https://proceedings.neurips.cc/paper/2019/file/dc6a7e655d7e5840e66733e9ee67cc69-Paper.pdf

[68] W. Wang, H. Bao, S. Huang, L. Dong, and F. Wei, "MiniLMv2: Multi-head self-attention relation distillation for compressing pretrained transformers," in *Proc. Findings Assoc. Comput. Linguistics*, 2021, pp. 2140–2151, doi: 10.18653/v1/2021.findings-acl.188.

[69] W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, and M. Zhou, "MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers," in *Proc. 34th Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 33, 2020, pp. 5776–5788. [Online]. Available: https://proceedings.neurips.cc/paper/2020/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

[70] Z. Yang, Y. Yang, D. Cer, J. Law, and E. Darve, "Universal sentence representation learning with conditional masked language model," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2021, pp. 6216–6228, doi: 10.18653/v1/2021.emnlp-main.502.

[71] K. Wang, N. Thakur, N. Reimers, and I. Gurevych, "GPL: Generative pseudo labeling for unsupervised domain adaptation of dense retrieval," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2022, pp. 2345–2360, doi: 10.18653/v1/2022.naacl-main.168.

[72] N. Thakur, N. Reimers, and J. Lin, "Injecting domain adaptation with learning-to-hash for effective and efficient zero-shot dense retrieval," 2022, *arXiv:2205.11498*.

[73] I. Yamada, A. Asai, and H. Hajishirzi, "Efficient passage retrieval with hashing for open-domain question answering," in *Proc. 59th Annu. Meeting Assoc. Comput. Linguistics 11th Int. Joint Conf. Natural Lang. Process.*, 2021, pp. 979–986, doi: 10.18653/v1/2021.acl-short.123.

[74] F. Feng, Y. Yang, D. Cer, N. Arivazhagan, and W. Wang, "Language-agnostic BERT sentence embedding," in *Proc. 60th Annu. Meeting Assoc. Comput. Linguistics*, 2022, pp. 878–891, doi: 10.18653/v1/2022.acl-long.62.

[75] N. Majumder, P. Hong, S. Peng, J. Lu, D. Ghosal, A. Gelbukh, R. Mihalcea, and S. Poria, "MIME: MIMicking emotions for empathetic response generation," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2020, pp. 8968–8979, doi: 10.18653/v1/2020.emnlp-main.721.

[76] K. Nantomah. (2019). *On Some Properties of the Sigmoid Function*. [Online]. Available: Asia Mathematika. [Online]. Available: https://hal.archives-ouvertes.fr/hal-02635089/.

[77] N. Balakrishnan, *Handbook Logistic Distribution*. Boca Raton, FL, USA: CRC Press, 2013.

[78] K. Kowsari, K. J. Meimandi, M. Heidarysafa, S. Mendu, L. Barnes, and D. Brown, "Text classification algorithms: A survey," *Information*, vol. 10, no. 4, p. 150, Apr. 2019, doi: 10.3390/info10040150.

[79] R. Real, J. M. Vargas, and R. Olmstead, ''The probabilistic basis of Jaccard's index of similarity,'' *Syst. Biol.*, vol. 45, no. 3, pp. 380–385, Sep. 1996.
[80] M. T. Ribeiro, S. Singh, and C. Guestrin, '''Why should I trust you?': Explaining the predictions of any classifier,'' in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.
[81] A. M. Grobelnik, M. Grobelnik, and D. Mladenic, ''Multicomet–multilingual commonsense description,'' in *Proc. 23rd Int. Multiconf. Inf. Soc.*, 2020, pp. 37–40.
[82] J. Zhang, Y. Li, W. Xiao, and Z. Zhang, ''Non-iterative and fast deep learning: Multilayer extreme learning machines,'' *J. Franklin Inst.*, vol. 357, no. 13, pp. 8925–8955, Sep. 2020, doi: 10.1016/j.jfranklin.2020.04.033.
[83] J. Zhang, Y. Zhao, F. Shone, Z. Li, A. F. Frangi, S. Q. Xie, and Z.-Q. Zhang, ''Physics-informed deep learning for musculoskeletal modeling: Predicting muscle forces and joint kinematics from surface EMG,'' *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 31, pp. 484–493, 2023, doi: 10.1109/TNSRE.2022.3226860.

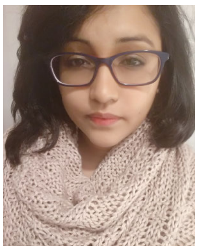

SWATI SWATI received the M.Tech. degree in mathematics and computing from the Indian Institute of Technology Patna, India. She is currently pursuing the Ph.D. degree with the Jožef Stefan International Postgraduate School, Ljubljana. She is also an Early Stage Researcher (ESR) of the EU-Funded Horizon 2020 Project CLEOPATRA with Jožef Stefan Institute, Ljubljana. Her current research work focuses on the analysis and development of deep learning-based solutions for cross-lingual news reporting bias. It includes domains such as news-value determination and computational journalism. She works on the study and development of machine learning, deep learning, data/text mining, and semantic technology techniques and their applications.

DUNJA MLADENIĆ is leading the Artificial Intelligence Laboratory, Jožef Stefan Institute, Ljubljana. From 2013 to 2017, she was with the Institute's Scientific Council, where she was the Vice President, from 2015 to 2017, and working on several research projects mainly related to machine learning, data and text mining, big data analytic, semantic technologies, and their application to real-world problems. She was with the Jožef Stefan International Postgraduate School, the University of Ljubljana, the University of Primorska, and the University of Zagreb (FERI and FOI), teaching classes related to data analytics and artificial intelligence.

MARKO GROBELNIK is currently an expert researcher in the field of artificial intelligence (AI). He co-leads the Department for Artificial Intelligence, Jožef Stefan Institute Ljubljana; and has co-founded the UNESCO International Research Center on AI (IRCAI). He is also the CEO of Quintelligence.com. He collaborates with major European academic institutions and major industries. He is the coauthor of several books and the co-founder of several start-ups. He is/was involved in over 50 EU-funded research projects in various fields of artificial intelligence. His research interests include machine learning, data/text/web mining, network analysis, semantic technologies, deep text understanding, and data visualization.