Generative Entity-to-Entity Stance Detection with Knowledge Graph Augmentation

[1]Computer Science and Engineering, University of Michigan, Ann Arbor, MI
[2]Department of Political Science, Northeastern University, Boston, MA
[1]{xlfzhang,wangluxy}@umich.edu
[2]n.beauchamp@northeastern.edu

code:https://github.com/launchnlp/SEESAW

2023. 3. 9 • ChongQing

2022_EMNLP

Reported by Junhao Cao

Leibniz-Institut für Sozialwissenschaften

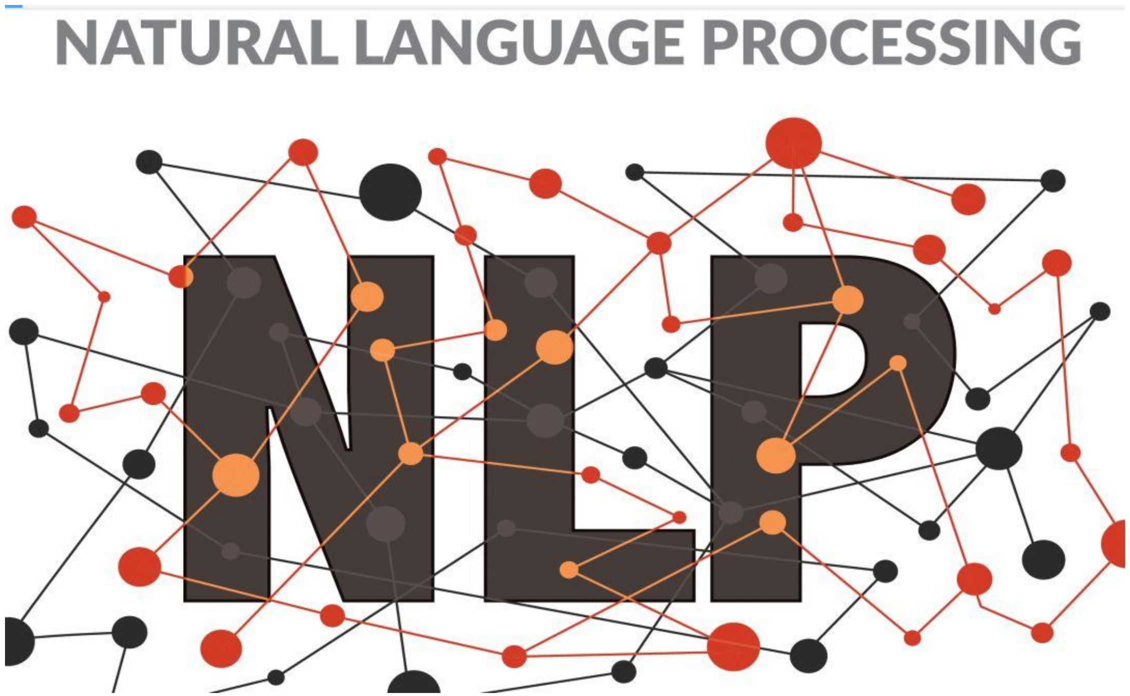

1.Introduction

2.Method

3.Experiments

Introduction

<s> [preceding con text] <s> [target text] </s> [succeeding context]

Trump's rhetoric, including calling Central Americans trying to enter the United States "an invasion," and his hard-line immigration policies have exposed him to condemnation since the El Paso shooting. "How far is it from Trump's saying this 'is an invasion' to the shooter in El Paso declaring 'his attack is a response to the Hispanic invasion of Texas?' Not far at all," Biden was due to say, according to an advance copy of his speech. "In both clear language and in code, this president has fanned the flames of white supremacy in this nation."

[0]	Joe Biden	NEG	Donald Trump
[1]	Joe Biden	NEG	white supremacy
[2]	Donald Trump	POS	white supremacy

<source, sentiment, target>

	Tweet	Target	Stance
Example 1	Pregnant people have feelings, and the ability to make decisions about their health	Legalization of abortion	Favor
	They have not the ability and shouldn't make decisions that involve their health	Legalization of abortion	Favor

Stance detection is typically framed as predicting the sentiment in a given text towards a target entity. However, this setup overlooks the importance of the source entity, i.e., who is expressing the opinion.

Method

$<ENT>$ *[source]* $<ENT>$ *[target]* $<STANCE>$

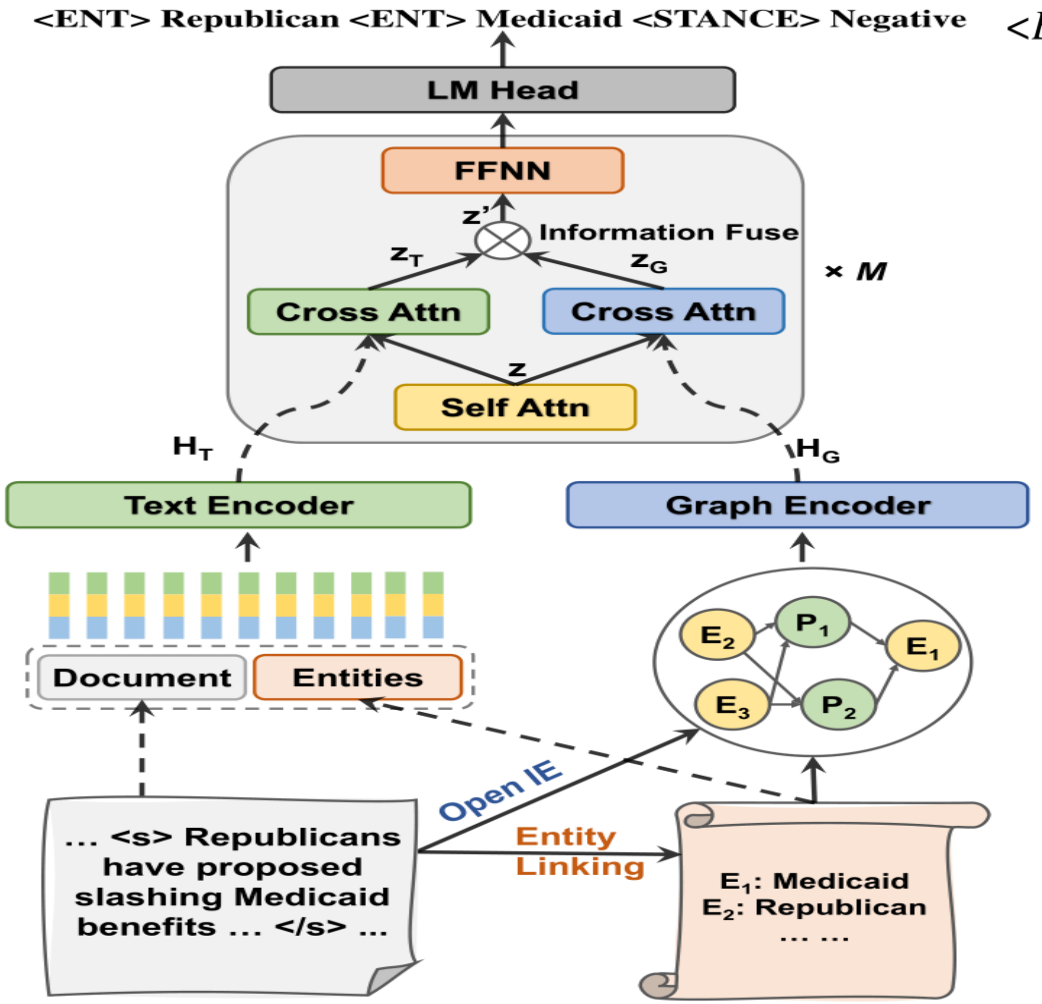

x $<s>$ $<ENT>$ E_1 $<ENT>$ E_2 …

<subject, predicate, object>.

Method

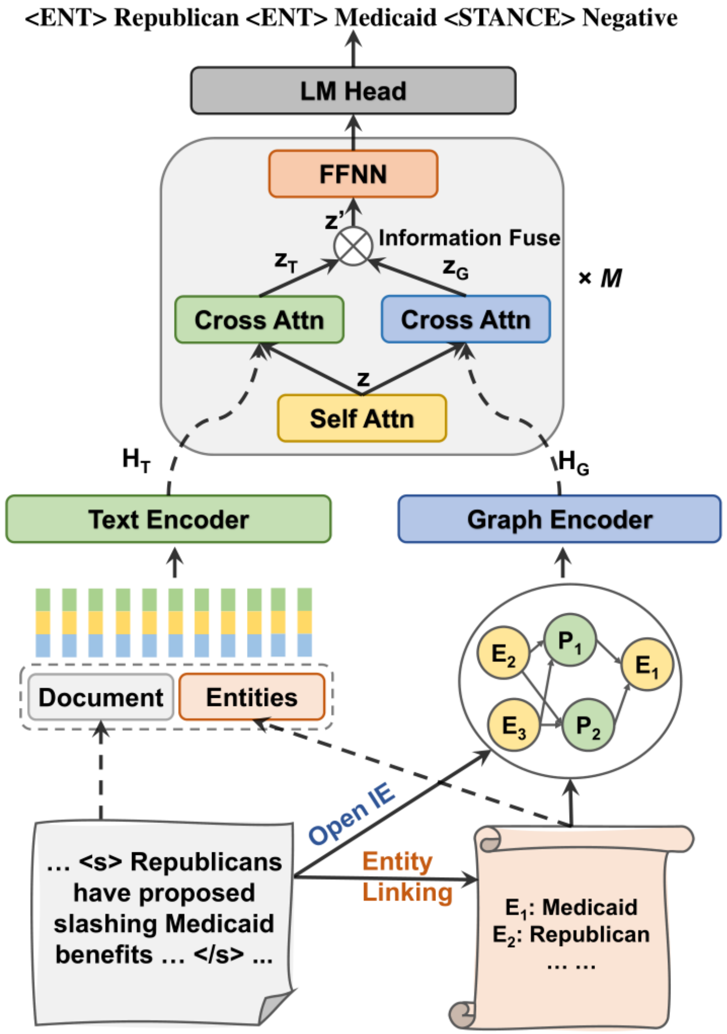

$$\mathbf{z}_T = \text{LayerNorm}(\mathbf{z} + \text{Attn}(\mathbf{z}, \mathbf{H}_T))$$
$$\mathbf{z}_G = \text{LayerNorm}(\mathbf{z} + \text{Attn}(\mathbf{z}, \mathbf{H}_G)) \quad (1)$$

$$\mathbf{z_f} = \text{GELU}(\mathbf{W^f}[\mathbf{z_T}; \mathbf{z_G}] + \mathbf{b_f}) \quad (4)$$

$$\lambda = \text{sigmoid}(\mathbf{W}^\lambda[\mathbf{z_T}; \mathbf{z_G}] + \mathbf{b}_\lambda) \quad (5)$$

$$\mathbf{z} = \boldsymbol{\lambda} \odot \mathbf{z_f} + (1 - \boldsymbol{\lambda}) \odot \mathbf{z_T} \quad (6)$$

$$\mathcal{L}_{stance} = -\sum_{(\mathbf{x},\mathbf{y})\in D} \log p(\mathbf{y}|\mathbf{x}) \quad (7)$$

Method

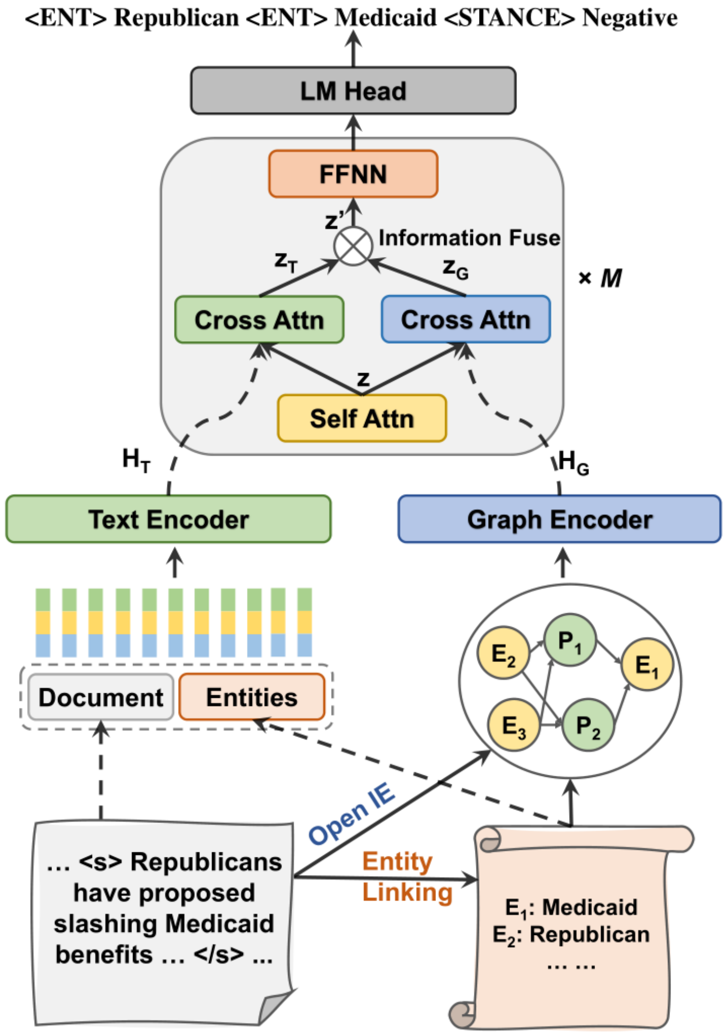

$$\hat{\mathbf{s}} = \mathrm{sigmoid}(\mathbf{u}\mathbf{H}_{\mathbf{G}}^{\mathbf{E}}) \tag{8}$$

$$\mathcal{L}_{node} = -\sum_{s_i} w * s_i \log(\hat{s}_i) + (1 - s_i)\log(1 - \hat{s}_i) \tag{9}$$

$$\mathcal{L}_{multi} = \mathcal{L}_{stance} + \mathcal{L}_{node}$$

Experiment

	SEESAW (Task A)	Park et al. (Task B)
Target Sentence length	30.3	31.0
Label ratio (pos/neg)	37.6%/62.4%	35.3%/64.7%
Splits (train/valid/test)	4505/1313/1378	4252/506/562

Table 1: Statistics of the two datasets for experiments. Data by Park et al. (2021) only contains single sentences without context. We split the SEESAW chronologically, and use the same splits as in Park et al. (2021).

	Full			Aspect		
	Acc.	F1	Acc.Any	src-s	s-tgt	src-tgt
Baselines (no graph)						
Sentence (Sen)	7.26	10.35	12.39	36.66	23.81	14.88
Sen + Context (Ctx)	9.66	14.08	16.87	45.24	27.79	20.01
Sen + Ctx + Entities	11.32	16.00	19.15	47.43	30.03	23.18
Pipeline Models (Ours)						
Graph	12.03	15.77	19.86	46.60	31.07	23.19
+ Oracle Entities	31.84	35.58	44.87	66.33	55.50	53.82
End-to-end Models (Ours)						
Graph (seq. attn.)	12.97	17.22	20.62	50.58	32.45	24.76
Graph	13.62	18.12	21.78	51.01	32.65	**26.08**
+ Multitask	13.34	18.16	21.77	**52.10**	32.06	26.07
+ Wiki	**13.74**	**18.24**	**21.87**	51.41	**32.69**	25.94

Table 2: Results on SEESAW for E2E stance detection task, and breakdown of accuracy scores by aspects (average of 5 runs). Best results without oracle entities are in **bold**. Our graph-augmented model with Wikipedia knowledge performs the best on 4 out of 6 metrics, indicating the effectiveness of encoding knowledge. Results with standard deviation are in Table A3.

Experiment

	Accuracy	Macro F1
BART (Lewis et al., 2020)	86.32	77.53
POLITICS (Liu et al., 2022)	86.33	77.48
DSE2QA (Park et al., 2021)	87.78	**79.90**
Our Model	**87.79**	79.01

Table 3: Results on stance-only prediction for specified pairwise entities. Our model performs on par with state-of-the-art models in stance detection tasks (POLITICS and DSE2QA). Results with std. deviation in Table A4.

Ms. Harris, who is making her first trip to a battleground state since joining the Democratic ticket, is visiting with union workers and leaders as well as African-American businesspeople and pastors in Milwaukee, the Black hub of the state. Each is expected to focus on the economy, with Mr. Pence hailing **the state**'s job growth before the coronavirus pandemic and Ms. Harris critiquing the administration's handling of the virus and the resultant impact on the economy. Yet their political missions are different. The vice president is hoping to appeal to voters in a historically Democratic part of **Wisconsin**, where Mr. Trump outperformed his Republican predecessors, in hopes they abandon their political roots again.

[0] Mike Pence POS Wisconsin
[1] Kamala Harris NEG Donald Trump
[2] Kamala Harris NEG <Someone>

Sent.: [0] Mike Pence POS <Someone>
Sent. + Cxt.: [0] Mike Pence POS <Someone>
Sent. + Cxt. + Ent.: [0] Mike Pence POS <Someone>
Graph model (ours): [0] Mike Pence POS job growth; [1] Kamala Harris NEG Donald Trump
Graph model + Wiki (ours): [0] Mike Pence POS Wisconsin; [1] Kamala Harris NEG Donald Trump

Figure 3: Sample system predictions (below the dotted line) with human labeled triples (above the dotted line).

Experiment

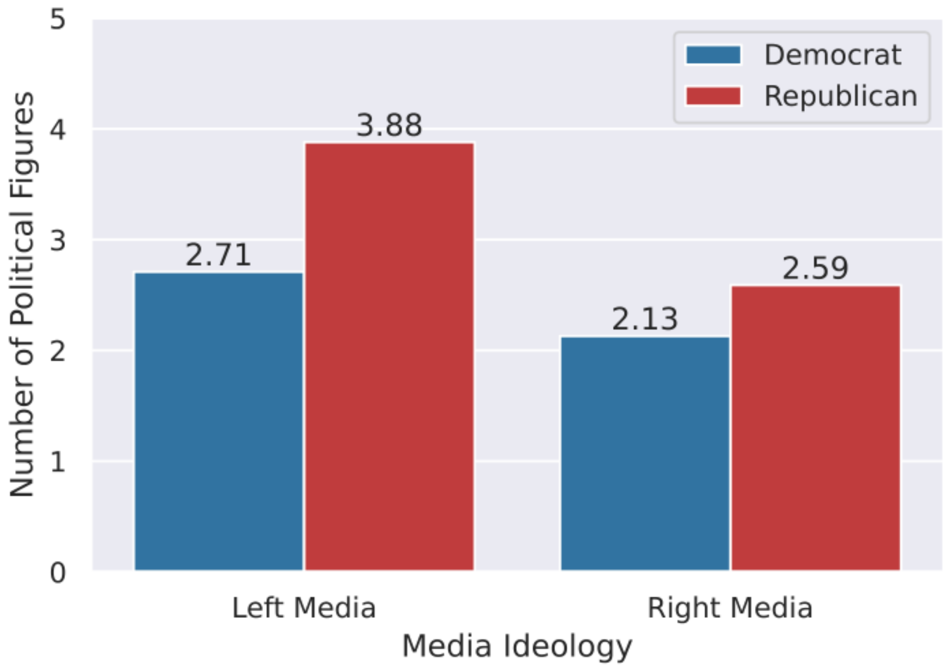

Figure 4: Media quoting Democrats vs. Republicans by counting source entities per article. Both left- and right-leaning media outlets quote Republicans more.

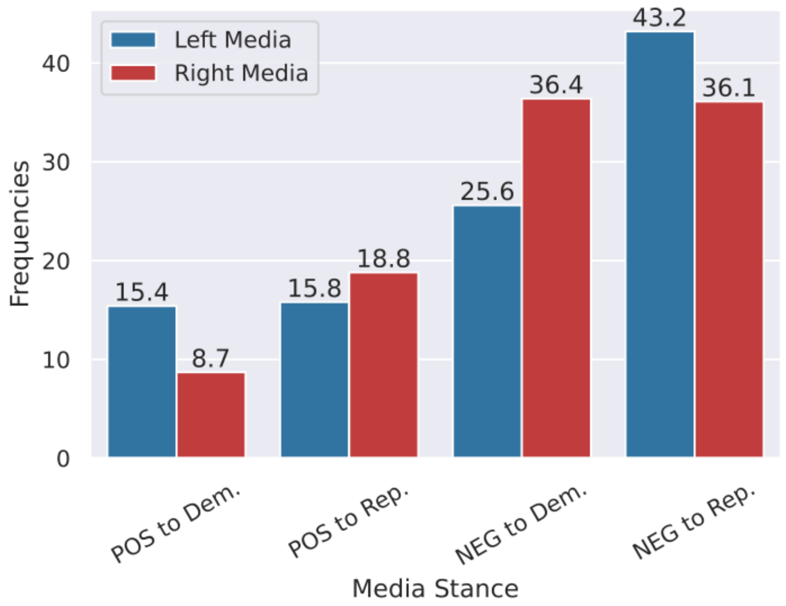

Figure 5: Percentage of stance triplets that media favoring or criticizing entities from the same or the opposite side. Media of both sides attack politicians from the opposite parties more than their own parties. Note there is a *symmetrical asymmetry* phenomenon: Left is balanced while the right is unbalanced in terms of indicated positivity, and the other way around for negativity.

Experiment

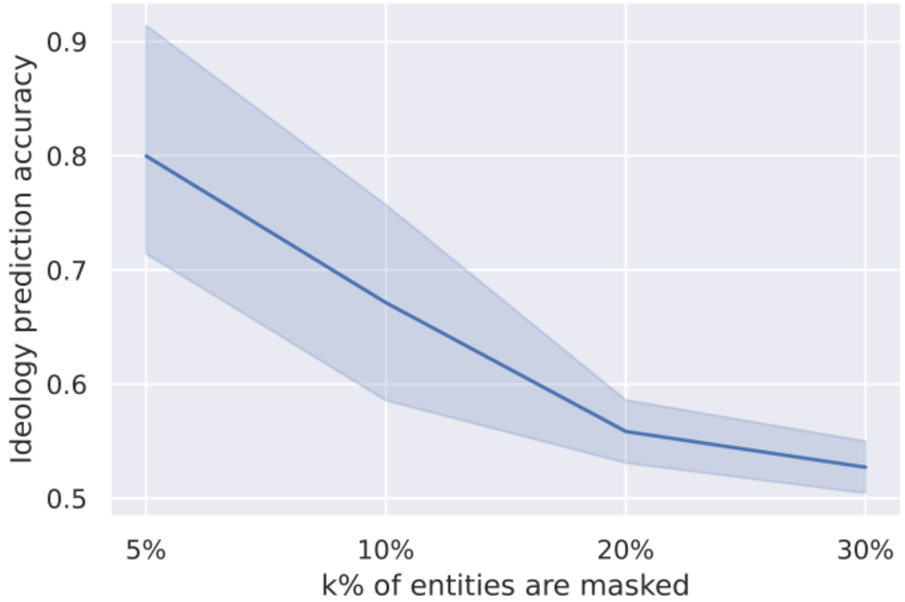

Figure 6: Entity-level ideology prediction using stances from/to their neighboring entities with known ideology. We increase the ratio of entities being masked, which decreases the ideology prediction accuracy. This implies knowing entity's support/oppose interactions with other entities is helpful for predicting their own ideology.

Thank you!