

# INTERPRETABLE MULTIMODAL OUT-OF-CONTEXT DETECTION WITH SOFT LOGIC REGULARIZATION

Huanhuan Ma<sup>1,2,\*</sup>, Jinghao Zhang<sup>1,2,\*</sup>, Qiang Liu<sup>1,2</sup>, Shu Wu<sup>1,2,†</sup>, Liang Wang<sup>1,2</sup>

<sup>1</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>2</sup> Center for Research on Intelligent Perception and Computing (CRIPAC)

State Key Laboratory of Multimodal Artificial Intelligence Systems (MAIS)

Institute of Automation, Chinese Academy of Sciences

{*huanhuan.ma, jinghao.zhang*}@cripac.ia.ac.cn, {*qiang.liu, shu.wu, wangliang*}@nlpr.ia.ac.cn

## ABSTRACT

The rapid spread of information through mobile devices and media has led to the widespread of false or deceptive news, causing significant concerns in society. Among different types of misinformation, image repurposing, also known as out-of-context misinformation, remains highly prevalent and effective. However, current approaches for detecting out-of-context misinformation often lack interpretability and offer limited explanations. In this study, we propose a logic regularization approach for out-of-context detection called LOGRAN (LOGic Regularization for out-of-context ANalysis). The primary objective of LOGRAN is to decompose the out-of-context detection at the phrase level. By employing latent variables for phrase-level predictions, the final prediction of the image-caption pair can be aggregated using logical rules. The latent variables also provide an explanation for how the final result is derived, making this fine-grained detection method inherently explanatory. We evaluate the performance of LOGRAN on the NewsCLIPPings dataset, showcasing competitive overall results. Visualized examples also reveal faithful phrase-level predictions of out-of-context images, accompanied by explanations. This highlights the effectiveness of our approach in addressing out-of-context detection and enhancing interpretability.

**Index Terms**— out-of-context detection, interpretability, soft logic regularization

## 1. INTRODUCTION

Due to the rise of mobile devices and the fast-paced development of media, information spreads more quickly than ever before. However, this rapid dissemination also enables the rapid spread of false or misleading news, causing significant concern about its negative consequences for society, individuals, and politics [1, 2, 3, 4, 5, 6, 7, 8]. Some news make a false or misleading claim by using a real image to create more credible stories [9, 10]. Specifically, the real images of people and events get reappropriated and used out-of-context to illustrate false events and misleading narratives by misrepresenting who is in the image, what is the context in which they appear, or where the event takes place. Image repurposing is unlike using deep fake models [11, 12] to generate an image unexist, which potentially amplifies its damage because of non-expertise need [13]. Such out-of-context misinformation can't be detected solely based on texts or

images separately, thus automated detection methods are in high demand.

Existing research predominantly focuses on developing “black box” models for predicting whether news articles involve the misuse of images in an out-of-context manner [9, 14, 15]. However, these approaches often lack interpretability, providing only a holistic verdict without detailed explanations. Inspired by Chen et al. [16], we can detect a caption of a news image by its composing phrases, e.g., subject, verb, and object phrases. An image can be identified as out-of-context if it contains one or more unsupported phrases, which can be referred to as the culprits. Conversely, an image is deemed valid only when all the phrases in its caption are supported by the visual content it presents.

It is also challenging due to the lack of phrase-level labels in existing out-of-context datasets, such as the COSMOS dataset [9] and the NewsCLIPPings dataset [10]. This absence of labels makes it difficult for us to determine which specific aspect of the image caption renders it out-of-context. Manually annotating such fine-grained data is unrealistic and requires tremendous human labor. The main concern is how to supervise a model to reach meaningful phrasal veracity predictions. There are symbolic logical rules that can be utilized to create weak supervision for intermediate phrasal predictions. Empirically, all phrases of the image caption should not be detected as out-of-context use if the image is not out-of-context used, and a caption is detected as out-of-context use if there exists at least one out-of-context phrase.

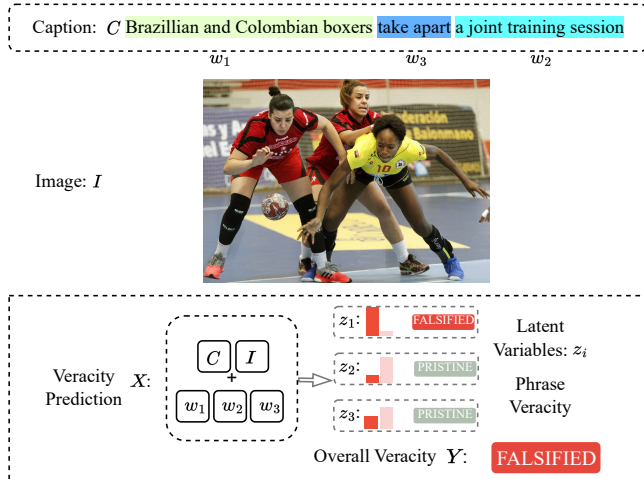
Motivated by the preceding description, we propose LOGRAN (LOGic Regularization for out-of-context ANalysis), a multimodal soft logic regularization approach designed to detect out-of-context instances for predicting the overall label, as well as to provide phrase-level predictions as explanations, thereby offering a comprehensive understanding of the detected instances.

Figure 1 illustrates that the fundamental concept of LOGRAN is the use of a latent model. This method decomposes the given caption into a sequence of detectable phrases, and then employs a phrase-level prediction module as a latent model to generate the phrase predictions, treating these phrase predictions as latent variables. This approach is anchored in the framework of variational EM [17, 18], which contributes to generating meaningful predictions for the individual caption phrases. Furthermore, LOGRAN incorporates logical rules into these latent variables by constructing a teacher module that aggregates logical reasoning across all latent variables. This process serves to distill logical knowledge into the target module.

In the subsequent sections, we will conduct an in-depth exploration of our innovative approach, LOGRAN, followed by a com-

\* The first two authors contributed equally.

† Corresponding author



**Fig. 1.** An example of the pipeline of how LOGRAN works. LOGRAN not only predicts the overall veracity of the image-caption pair, but also uses latent variables to indicate the veracity at phrase level. In this example, we can easily find the “Culprit” of this caption is  $w_1$ . It manipulates the subject and the pristine caption is “Tunisian and Angolan players fight for the ball on Sunday during a handball tournament in Spain Angola go on to win”.

prehensive evaluation of its performance using various backbone architectures on the NewsCLIPPings dataset [10].

## 2. RELATED WORK

### 2.1. Multimodal Misinformation

In response to the growing number of false claims associated with images, recent efforts have emerged to address this challenge and have introduced a few methods and datasets [19, 10, 9].

For instance, Luo et al. [10] proposed a methodology for generating falsified examples by re-matching real images with real captions [20]. They created the extensive NewsCLIPPings dataset, which encompasses both authentic and convincingly falsified examples. This matching process was achieved automatically using the trained vision-language model CLIP [21]. The falsified examples generated could distort the context, location, or individuals depicted in the image, introducing inconsistencies in terms of entities and semantic context. The authors demonstrated that detection performances are limited, underscoring the challenging nature of the task. Similarly, Abdelnabi et al. [14] proposed a method named CCN, which leverages external web evidence to verify image-caption claims.

In contrast, our approach takes a distinct direction towards interpretable out-of-context detection. We aim to identify where and how a caption has been falsified using a phrase-level approach. Our method does not rely on external content evidence for making predictions, and instead emphasizes explanations, thereby advancing the goal of being correct for the right reasons.

### 2.2. Symbolic Logic Reasoning

Previous researchers have attempted to combine symbolic logic and neural networks [15, 22, 23, 24]. One prominent approach in this area relies on the variational EM framework [17, 18]. Another

commonly employed strategy involves incorporating neural network components to loosen the constraints of logic, enabling end-to-end training [25, 26].

Our approach follows a similar trajectory to Chen et al. [16], drawing inspiration from both of these approaches, where we combine elements from both research lines. Specifically, we utilize latent variables within a latent space to represent the intermediate predictions of truthfulness, mirroring the principles of the variational EM framework. These latent variables are subsequently subjected to regularization through a softened logic mechanism.

## 3. THE METHOD

In this section, we introduce the proposed method LOGRAN for detecting whether an image-caption pair is out-of-context used. The possible prediction results are categorized as either *Pristine* or *Falsified*. Unlike most previous approaches that provide an overall prediction, our goal is to predict the final result as well as faithful phrase predictions for explanations. We define two tasks: overall caption detection and fine-grained phrase detection.

**Caption Detection** Given a caption sentence  $c$  and its image  $I$ , our goal is to model the probability distribution  $p(\mathbf{y}|c, I)$ , where  $\mathbf{y} \in \{\text{Pristine}, \text{Falsified}\}$  is a two-valued variable indicating the veracity of the caption’s image. We use **bold** letters to indicate variables in this paper.

**Phrase Detection** We decompose the caption into phrases and predict the out-of-context label  $z_i$  for each caption phrase  $w_i \in \mathcal{W}_c$  using the probability  $p(z_i|c, w_i, I)$ , where  $z_i$  is treated as a binary latent variable  $z_i \in \{\text{Pristine}, \text{Falsified}\}$ . We utilize the off-the-shelf tools provided by Flair and implement a phrase chunking model [27] to extract caption phrases. The extracted caption phrases include named entities (NEs), verbs, adjectives, and noun phrases (APs and NPs).

### 3.1. Overview of LOGRAN

The LOGRAN approach offers more than just an overall prediction at the caption level; it also uncovers the veracity of individual phrases, providing insights into which part of the caption contributes to the final prediction. We begin by following the variational EM framework [17, 18] to build a latent model for phrase prediction. Each phrase prediction, denoted as  $w_i \in \mathcal{W}_c$ , is treated as a binary latent variable  $z_i$ . We implement the negative Evidence Lower Bound (ELBO) loss  $\mathcal{L}_{\text{var}}(t, l)$  to optimize the latent model. At this stage, the latent variables can provide some degree of classification between *Pristine* and *Falsified* but still lack logical knowledge and constraints. Let  $\mathbf{z} = (z_1, z_2, \dots, z_{|\mathcal{W}_c|})$ . The final prediction of a caption  $\mathbf{y}$  is then logically regularized based on the latent variables  $\mathbf{z}$ . To integrate the logical knowledge into the latent variables, we propose a distillation method involving both a teacher model and a student model, along with the introduction of the logical loss  $\mathcal{L}_{\text{logic}}(t, l)$ . The final loss function is a weighted sum of the  $\mathcal{L}_{\text{var}}(t, l)$  and the  $\mathcal{L}_{\text{logic}}(t, l)$ . In the following sections, we will provide a detailed introduction to our method.

### 3.2. Latent Variable Model

Given an input  $x = (c, E)$  comprising a textual caption  $c$  and its corresponding image  $E$ , we define the target distribution  $p_t(\mathbf{y}|x)$  as follows:

$$p_t(\mathbf{y}|x) = \sum_{\mathbf{z}} p_t(\mathbf{y}|\mathbf{z}, x) p(\mathbf{z}|x) \quad (1)$$

Here, the target distribution  $p_t(\mathbf{y}|x)$  as the sum of the probabilities of  $\mathbf{y}$  given the latent variable  $\mathbf{z}$  and  $x$ , multiplied by the prior distribution  $p(\mathbf{z}|x)$ . The independence assumption is made for  $\mathbf{z}_i$ , where  $p(\mathbf{z}|x) = \prod_i p(\mathbf{z}_i|x, w_i)$ .

The objective function is designed to minimize the negative log-likelihood of the gold label  $y^*$  given  $x$ . For optimization, we utilize variational inference and employ neural networks to amortize the variational posterior distribution. This is necessary due to the vast space of  $\mathbf{z}$ . Computing the exact posterior  $p_t(\mathbf{z}|\mathbf{y}, x)$  poses a challenge for the EM algorithm. This leads to minimizing the negative Evidence Lower BOund (ELBO), which consists of the expected log probability of  $p_t(y^*|\mathbf{z}, x)$  under the variational posterior distribution  $q_t(\mathbf{z}|\mathbf{y}, x)$  and the Kullback-Leibler divergence between  $q_t(\mathbf{z}|\mathbf{y}, x)$  and  $p(\mathbf{z}|x)$ . The variational bound is denoted as  $\mathcal{L}_{\text{var}}(t, l)$ :

$$-\mathbb{E}_{q_t}[\log p_t(y^*|\mathbf{z}, x)] + D_{\text{KL}}(q_t(\mathbf{z}|\mathbf{y}, x) \parallel p(\mathbf{z}|x)) \quad (2)$$

In our experiments, we utilize fine-tuned visual-language models as the prior distribution  $p(\mathbf{z}|x)$ . The features obtained from visual-language models are fed into a classifier, which provides some degree of classification between `Pristine` and `Falsified`.

### 3.3. Distilling Logical Knowledge

The ELBO (equation 2) alone does not ensure that the latent variables accurately represent the veracity of caption phrases without intermediate supervision. To address this issue, we need to integrate logical rules into the latent variables. Here we declare the definition of the logical rules we used. Our insight is that certain natural logical consistencies need to be met between phrase detection and caption detection. Specifically, a caption is considered: 1) `Falsified` if there is inconsistency in at least one caption phrase; 2) `Pristine` if all caption phrases are consistent.

Having established the logical constraints, we propose a distillation method involving a teacher module and a student module to integrate the above logical rules into the latent variables. The student module, denoted as  $p_t(\mathbf{y}|\mathbf{z}, x)$ , is the module we aim to optimize. Meanwhile, the teacher module is constructed by projecting the variational distribution  $q_t(\mathbf{z}|\mathbf{y}, x)$  into a subspace called  $q_t^T(\mathbf{y}_z|\mathbf{y}, x)$ , which adheres to logical rules. This subspace, represented by  $\mathbf{y}_z$ , captures the logical aggregation of  $\mathbf{z}$ . By simulating the outputs of  $q_t^T$ , we can transfer logical knowledge into  $p_t$ . The distillation loss is formulated as:

$$\mathcal{L}_{\text{logic}}(t, l) = D_{\text{KL}}\left(p_t(\mathbf{y}|\mathbf{z}, x) \parallel q_t^T(\mathbf{y}_z|\mathbf{y}, x)\right). \quad (3)$$

The behavior of  $q_t^T$  during predictions reflects the information captured by the rule-regularized subspace, indicating the uncertain and probabilistic nature of the predictions [28]. By minimizing the distillation loss  $\mathcal{L}_{\text{logic}}$  in Equation 3, the predictions of phrasal veracity are regularized by the aggregation logic, even without specific supervisions for caption phrases.

For the training process, we adopt soft logic [25] to relax the strict adherence to logic. We use product t-norms for differentiability during training and regularization of the latent variables. As discussed in Section 3.2, given the probability of the caption phrase veracity  $\mathbf{z}$ , we logically aggregate them into  $\mathbf{y}_z$  using the following

equations (ignoring the input  $x$  for simplicity):

$$\begin{aligned} q_t^T(\mathbf{y}_z = \top) &= \prod_{i=1}^{|\mathbf{z}|} q_t(\mathbf{z}_i = \top), \\ q_t^T(\mathbf{y}_z = \perp) &= 1 - \prod_{i=1}^{|\mathbf{z}|} q_t(\mathbf{z}_i = \top) \end{aligned} \quad (4)$$

where  $\sum_{\mathbf{y}_z} q_t^T(\mathbf{y}_z) = 1$  and  $\sum_{\mathbf{z}_i} q_t(\mathbf{z}_i) = 1$ . Here, the symbols  $\top$  and  $\perp$  represent the labels `Pristine` and `Falsified`, respectively.

The final loss function is a weighted sum of two objectives, where  $(1 - \lambda)$  and  $\lambda$  represent the relative importance of the variational loss and the logical loss, respectively:

$$\mathcal{L}_{\text{final}}(t, l) = (1 - \lambda)\mathcal{L}_{\text{var}}(t, l) + \lambda\mathcal{L}_{\text{logic}}(t, l), \quad (5)$$

Here,  $\lambda$  is a hyperparameter that determines the trade-off between the two objectives.

### 3.4. Classifier

We utilize neural networks to parameterize  $p_t(\mathbf{y}|\mathbf{z}, x)$  and the variational distribution  $q_t(\mathbf{z}|\mathbf{y}, x)$  for veracity prediction. These models are optimized using the variational EM algorithm and decoded iteratively. In this approach, a fine-tuned visual-language model serves as the representation decoder for both text and image inputs.

To generate joint representations for local hidden representations  $\{\mathbf{h}_{\text{local}}^{(i)} \in \mathbb{R}^d\}$  and global hidden representation  $\mathbf{h}_{\text{global}} \in \mathbb{R}^d$ , we compute the Hadamard product ( $\odot$ ) between each phrase  $w_i \in W_c$  and the image  $I$  for the local hidden representations. For the global hidden representation, we take the Hadamard product ( $\odot$ ) between the overall caption  $c$  and the image  $I$ .

After obtaining these representations, two-layer MLPs are employed to parameterize  $p_t(\cdot)$  and  $q_t(\cdot)$

$$\begin{aligned} p_t(\mathbf{y}|\mathbf{z}, x) &= \text{MLP}(\mathbf{h}_{\text{global}} \odot \mathbf{h}_{\text{local}}^{(i)}), \\ q_t(\mathbf{z}|\mathbf{y}, x) &= \text{MLP}(\mathbf{y} \odot \mathbf{h}_{\text{local}}^{(i)}) \end{aligned} \quad (6)$$

The label embeddings of  $\mathbf{y}$  (ground truth  $y^*$  during training),  $\mathbf{h}_{\text{phrase}}^{(i)}$ , and  $\mathbf{h}_{\text{caption}}$  are concatenated as input to  $q_t(\mathbf{z}_i|\mathbf{y}, x)$ , which outputs the probability of  $\mathbf{z}_i$ . Similarly,  $(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{\text{max}})$  (padded to a maximum length) along with  $\mathbf{h}_{\text{caption}}$  and  $\mathbf{h}_{\text{phrase}}$  are concatenated as input to  $p_t(\mathbf{y}|\mathbf{z}, x)$ , which provides the distribution of  $\mathbf{y}$ . During training, both  $q_t(\cdot)$  and  $p_t(\cdot)$  are jointly optimized using the loss function described in Equation 5.

During inference, we *randomly* initialize  $\mathbf{z}$  and iteratively decode  $\mathbf{y}$  and  $\mathbf{z}$  using  $p_t(\mathbf{y}|\mathbf{z}, x)$  and  $q_t(\mathbf{z}|\mathbf{y}, x)$  until convergence. As a result, we obtain both the final prediction  $\mathbf{y}$  and the latent variables  $\mathbf{z}$ , which serve as the phrasal explanation.

## 4. EXPERIMENTS

### 4.1. Dataset

We use the NewsCLIPPings dataset [10], comprising both pristine and falsified images. It employs automation to match captions and images from the VisualNews [20] corpus, offering various subsets based on matching methods.

**Table 1.** Classification accuracy on the test set for the following models: (I) VisualBERT, (II) VisualBERT with LOGRAN, (III) Multimodal CLIP, and (IV) CLIP with LOGRAN. The underlined portions represent improvements from LOGRAN

	VisualBERT	VisualBERT-LOGRAN	CLIP	CLIP-LOGRAN
(a) Semantics/CLIP Text-Image	55.12	<u>56.88</u>	58.59	<u>59.03</u>
(b) Semantics/CLIP Text-Text	53.47	<u>55.62</u>	68.36	<u>70.81</u>
(c) Person/SBERT-WK Text-Text	63.32	<u>65.27</u>	66.57	<u>71.42</u>
(d) Scene/ResNet Place	60.72	<u>62.41</u>	69.64	<u>73.14</u>
Merged/Balanced	61.32	<u>63.18</u>	67.27	<u>70.51</u>

## 4.2. Backbone models

We use both CLIP [21] and VisualBERT[29] state-of-art models to evaluate the LOGRAN.

- **CLIP** utilizes distinct encoders for processing images and text, which are trained to produce comparable representations for associated concepts. CLIP is pretrained on a large-scale dataset consisting of 400 million image-text pairs collected from the web. During training, a contrastive loss function is utilized to maximize the cosine similarity between true image-text pairs. This approach enables the model to learn meaningful associations between images and corresponding textual descriptions.
- **VisualBERT** is another multimodal model that integrates visual and textual information. It includes a sequence of Transformer layers that use self-attention to automatically align components of a given text input with specific regions in a corresponding image input. Following the setting of NewsCLIPPING [10], we use a Faster-RCNN model[30] to extract image bounding box features.

## 4.3. Implementation Details

We first fine-tune the two backbone models on the NewsCLIPPINGs dataset, then train our model LOGRAN. We implemented our method in PyTorch. We fine-tune both backbone models as we train the classifiers on the NewsCLIPPINGs dataset and optimize two models with the Adam [31] optimizer, where the batch size is fixed at 64. When fine-tuning the backbone model, we use a learning rate of  $5e-5$  for the classifier and  $5e-7$  for other layers. When training our model LOGRAN, we freeze the two backbone models and use a learning rate of  $5e-7$ .

## 4.4. Results and Discussion

Table 1 shows our experiments on the NewsCLIPPINGs dataset. The results indicate a significant improvement achieved by LOGRAN across all splits of NewsCLIPPINGs when compared to baseline methods without LOGRAN. Our proposed LOGRAN method demonstrates a significant improvement, achieving a three-percentage-point increase when utilizing the two backbone models. This finding demonstrates that our method can significantly improve overall classification accuracy while providing fine-grained explainability and robustness across various backbone models.

Additionally, we select specific cases to demonstrate the explainability of LOGRAN for phrase prediction, shown in figure 2. The results also demonstrate that the phrase veracity predictions generated by LOGRAN exhibit interpretability, enabling us to identify the ‘Culprit’ responsible for the falsified instance.



**Fig. 2.** Some examples showcasing the outputs of LOGRAN. Each example includes a pair of images with matching captions. The left image is labeled as ‘Pristine’, while the right image is labeled as ‘Falsified’. We use the color red to highlight the phrases identified as ‘Culprit’.

## 5. CONCLUSION

This paper introduces LOGRAN, an interpretable Multimodal Out-of-context Detection method that not only makes overall predictions but also offers phrase predictions as explanations. We achieve this by utilizing a variational EM framework to get latent variables and decomposing the out-of-context detection task at the phrase level. To incorporate logical rules into the latent variables, we propose a distillation method that involves a teacher model and a student model. This ensures that the latent variables accurately capture the veracity of caption phrases without relying on intermediate supervision. To assess the efficacy of our approach, we have instantiated it in conjunction with two state-of-the-art visual-language models as the underlying backbone. We conducted extensive experiments on the NewsCLIPPINGs dataset, yielding compelling results and the robustness of our methodology. Furthermore, we visualize some prediction results, which also demonstrate the interpretability of the LOGRAN. Thereby emphasizing the contribution of our LOGRAN advancing the interpretability of out-of-context detection.

## 6. ACKNOWLEDGEMENTS

This work is supported by National Natural Science Foundation of China (62372454, 62206291, 62236010).

## 7. REFERENCES

- [1] Gordon Pennycook, Adam Bear, Evan T. Collins, and David G. Rand, “The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings,” *Manag. Sci.*, vol. 66, no. 11, pp. 4944–4957, 2020.
- [2] Feng Qian, Chengyue Gong, Karishma Sharma, and Yan Liu, “Neural user response generator: Fake news detection with collective user intelligence,” in *IJCAI*. 2018, pp. 3834–3840, ijcai.org.
- [3] Weizhi Xu, Qiang Liu, Shu Wu, and Liang Wang, “Counterfactual debiasing for fact verification,” in *ACL (1)*. 2023, pp. 6777–6789, Association for Computational Linguistics.
- [4] Weizhi Xu, Junfei Wu, Qiang Liu, Shu Wu, and Liang Wang, “Evidence-aware fake news detection with graph neural networks,” in *WWW*. 2022, pp. 2501–2510, ACM.
- [5] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan, “A convolutional approach for misinformation identification,” in *IJCAI*. 2017, pp. 3901–3907, ijcai.org.
- [6] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan, “Attention-based convolutional approach for misinformation identification from massive and noisy microblog posts,” *Comput. Secur.*, vol. 83, pp. 106–121, 2019.
- [7] Junfei Wu, Qiang Liu, Weizhi Xu, and Shu Wu, “Bias mitigation for evidence-aware fake news detection by causal intervention,” in *SIGIR*. 2022, pp. 2308–2313, ACM.
- [8] Shu Wu, Qiang Liu, Yong Liu, Liang Wang, and Tieniu Tan, “Information credibility evaluation on social media,” in *AAAI*. 2016, pp. 4403–4404, AAAI Press.
- [9] Shivangi Aneja, Christoph Bregler, and Matthias Nießner, “Cosmos: Catching out-of-context misinformation with self-supervised learning,” *CoRR*, vol. abs/2101.06278, 2021.
- [10] Grace Luo, Trevor Darrell, and Anna Rohrbach, “Newsclippings: Automatic generation of out-of-context multimodal media,” in *EMNLP (1)*. 2021, pp. 6801–6817, Association for Computational Linguistics.
- [11] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer, “High-resolution image synthesis with latent diffusion models,” in *CVPR*. 2022, pp. 10674–10685, IEEE.
- [12] “Midjourney,” <https://www.midjourney.com/>.
- [13] Lisa Fazio, “Out-of-context photos are a powerful low-tech form of misinformation,” <https://theconversation.com/out-of-contextphotos-are-a-powerful-low-techform-of-misinformation-129959/>, 2020.
- [14] Sahar Abdelnabi, Rakibul Hasan, and Mario Fritz, “Open-domain, content-based, multi-modal fact-checking of out-of-context images via online resources,” in *CVPR*. 2022, pp. 14920–14929, IEEE.
- [15] Gustav Sourek, Vojtech Aschenbrenner, Filip Zelezný, and Ondrej Kuzelka, “Lifted relational neural networks,” in *CoCo@NIPS*. 2015, vol. 1583 of *CEUR Workshop Proceedings*, CEUR-WS.org.
- [16] Jiangjie Chen, Qiaoben Bao, Changzhi Sun, Xinbo Zhang, Ji-aze Chen, Hao Zhou, Yanghua Xiao, and Lei Li, “LOREN: logic-regularized reasoning for interpretable fact verification,” in *AAAI*. 2022, pp. 10482–10491, AAAI Press.
- [17] Meng Qu and Jian Tang, “Probabilistic logic neural networks for reasoning,” in *Advances in Neural Information Processing Systems*, 2019, pp. 7712–7722.
- [18] Wangchunshu Zhou, Jinyi Hu, Hanlin Zhang, Xiaodan Liang, Maosong Sun, Chenyan Xiong, and Jian Tang, “Towards interpretable natural language understanding with explanations as latent variables,” in *NeurIPS*, 2020.
- [19] Scott McCrae, Kehan Wang, and Avidah Zakhori, “Multi-modal semantic inconsistency detection in social media news posts,” in *MMM (2)*. 2022, vol. 13142 of *Lecture Notes in Computer Science*, pp. 331–343, Springer.
- [20] Fuxiao Liu, Yinghan Wang, Tianlu Wang, and Vicente Ordonez, “Visual news: Benchmark and challenges in news image captioning,” in *EMNLP (1)*. 2021, pp. 6761–6771, Association for Computational Linguistics.
- [21] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever, “Learning transferable visual models from natural language supervision,” in *ICML*. 2021, vol. 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763, PMLR.
- [22] Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt, “Deepproblog: Neural probabilistic logic programming,” in *Advances in Neural Information Processing Systems*, 2018, pp. 3749–3759.
- [23] Luís C. Lamb, Artur S. d’Avila Garcez, Marco Gori, Marcelo O. R. Prates, Pedro H. C. Avelar, and Moshe Y. Vardi, “Graph neural networks meet neural-symbolic computing: A survey and perspective,” in *IJCAI*. 2020, pp. 4877–4884, ijcai.org.
- [24] Yifan Wei, Fangyu Lei, Yuanzhe Zhang, Jun Zhao, and Kang Liu, “Multi-view graph representation learning for answering hybrid numerical reasoning question,” *CoRR*, vol. abs/2305.03458, 2023.
- [25] Tao Li, Vivek Gupta, Maitrey Mehta, and Vivek Srikumar, “A logic-driven framework for consistency of neural models,” in *Proceedings of the 2019 Conference on EMNLP-IJCNLP*, Hong Kong, China, Nov. 2019, pp. 3924–3935, Association for Computational Linguistics.
- [26] Wenya Wang and Sinno Jialin Pan, “Integrating deep learning with logic fusion for information extraction,” in *AAAI*, 2020, pp. 9225–9232.
- [27] Alan Akbik, Duncan Blythe, and Roland Vollgraf, “Contextual string embeddings for sequence labeling,” in *COLING 2018, 27th International Conference on Computational Linguistics*, 2018, pp. 1638–1649.
- [28] Tongfei Chen, Zhengping Jiang, Adam Poliak, Keisuke Sakaguchi, and Benjamin Van Durme, “Uncertain natural language inference,” in *ACL*. 2020, pp. 8772–8779, Association for Computational Linguistics.
- [29] Liunan Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang, “Visualbert: A simple and performant baseline for vision and language,” *CoRR*, vol. abs/1908.03557, 2019.
- [30] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun, “Faster R-CNN: towards real-time object detection with region proposal networks,” in *NIPS*, 2015, pp. 91–99.
- [31] Diederik P. Kingma and Jimmy Ba, “Adam: A Method for Stochastic Optimization,” in *ICLR*, 2015.