

Neural Deepfake Detection with Factual Structure of Text

Wanjun Zhong^{1*}, Duyu Tang², Zenan Xu^{1*}, Ruize Wang³, Nan Duan², Ming Zhou²
Jiahai Wang¹ and Jian Yin¹

¹ The School of Data and Computer Science, Sun Yat-sen University,
Guangdong Key Laboratory of Big Data Analysis and Processing, Guangzhou, P.R.China

² Microsoft Research ³ Fudan University, Shanghai, P.R.China

{zhongwj25@mail2, xuzn@mail2}.sysu.edu.cn

{wangjiah@mail, issjyin@mail}.sysu.edu.cn

{dutang, nanduan, mingzhou}@microsoft.com

rzwang18@fudan.edu.cn

Abstract

Deepfake detection, the task of automatically discriminating machine-generated text, is increasingly critical with recent advances in natural language generative models. Existing approaches to deepfake detection typically represent documents with coarse-grained representations. However, they struggle to capture factual structures of documents, which is a discriminative factor between machine-generated and human-written text according to our statistical analysis. To address this, we propose a graph-based model that utilizes the factual structure of a document for deepfake detection of text. Our approach represents the factual structure of a given document as an entity graph, which is further utilized to learn sentence representations with a graph neural network. Sentence representations are then composed to a document representation for making predictions, where consistent relations between neighboring sentences are sequentially modeled. Results of experiments on two public deepfake datasets show that our approach significantly improves strong base models built with RoBERTa. Model analysis further indicates that our model can distinguish the difference in the factual structure between machine-generated text and human-written text.

1 Introduction

Nowadays, unprecedented amounts of online misinformation (e.g., fake news and rumors) spread through the internet, which may misinform people’s opinions of essential social events (Faris et al., 2017; Thorne et al., 2018; Goodrich et al., 2019; Kryściński et al., 2019). Recent advances in neural generative models, such as GPT-2 (Radford et al.,

* Work done while this author was an intern at Microsoft Research.



Figure 1: An example of machine-generated fake news. We can observe that the factual structure of entities extracted by named entity recognition is inconsistent.

2019), make the situation even severer as their ability to generate fluent and coherent text may enable adversaries to produce fake news. In this work, we study deepfake detection of text, to automatically discriminate machine-generated text from human-written text.

Previous works on deepfake detection of text are dominated by neural document classification models (Bakhtin et al., 2019; Zellers et al., 2019; Wang et al., 2019; Vijayaraghavan et al., 2020). They typically tackle the problem with coarse-grained document-level evidence such as dense vectors learned by neural encoder and traditional features (e.g., TF-IDF, word counts). However, these coarse-grained models struggle to capture the fine-grained factual structure of the text. We define the factual structure as a graph containing entities mentioned in the text and the semantically relevant relations among them. As shown in the motivating example in Figure 1, even though machine-generated text seems coherent, its factual structure is inconsistent. Our statistical analysis further reveals the difference in the factual structure between human-written and machine-generated text (detailed in Section 3). Thus, modeling factual structures is essential for detecting machine-generated text.

Based on the aforementioned analysis, we propose **FAST**, a graph-based reasoning approach uti-

lizing **F**Actual **S**tructure of **T**ext for deepfake detection. With a given document, we represent its factual structure as a graph, where nodes are automatically extracted by named entity recognition. Node representations are calculated not only with the internal factual structure of a document via a graph convolution network, but also with external knowledge from entity representations pre-trained on Wikipedia. These node representations are fed to produce sentence representations which, together with the coherence of continuous sentences, are further used to compose a document representation for making the final prediction.

We conduct experiments on a news-style dataset and a webtext-style dataset, with negative instances generated by GROVER (Zellers et al., 2019) and GPT-2 (Radford et al., 2019) respectively. Experiments show that our method significantly outperforms strong transformer-based baselines on both datasets. Model analysis further indicates that our model can distinguish the difference in the factual structure between machine-generated text and human-written text. The contributions are summarized as follows:

- We propose a graph-based approach, which models the fine-grained factual structure of a document for deepfake detection of text.
- We statistically show that machine-generated text differs from human-written text in terms of the factual structures, and injecting factual structures boosts detection accuracy.
- Results of experiments on news-style and webtext-style datasets verify that our approach achieves improved accuracy compared to strong transformer-based pre-trained models.

2 Task Definition

We study the task of deepfake detection of text in this paper. This task discriminates machine-generated text from human-written text, which can be viewed as a binary classification problem. We conduct our experiments on two datasets with different styles: a news-style dataset with fake text generated by GROVER (Zellers et al., 2019) and a large-scale webtext-style dataset with fake text generated by GPT-2 (Radford et al., 2019). The news-style dataset consists of 25,000 labeled documents, and the webtext-style dataset consists of 520,000 labeled documents. With a given document, systems are required to perform reasoning about the

content of the document and assess whether it is “human-written” or “machine-generated”.

3 Factual Consistency Verification

In this part, we conduct a statistical analysis to reveal the difference in the factual structure between human-written and machine-generated text. Specifically, we study the difference in factual structures from a consistency perspective and analyze entity-level and sentence-level consistency.

Through data observation, we find that human-written text tends to repeatedly mention the same entity in continuous sentences, while machine-written continuous sentences are more likely to mention irrelevant entities. Therefore, we define **entity consistency count (ECC)** of a document as the number of entities that are repeatedly mentioned in the next w sentences, where w is the sentence window size. **Sentence consistency count (SCC)** of a document is defined as the number of sentences that mention the same entities with the next w sentences. For instance, if entities mentioned in three continuous sentences are “ A and B ; A ; B ” and $w = 2$, then $ECC = 2$ because two entities A and B are repeatedly mentioned in the next 2 sentences. $SCC = 1$ because only the first sentence has entities mentioned in the next 2 sentences. We use all 5,000 pairs of human-written and machine-generated documents from the news-style dataset and each pair of documents share the same metadata (e.g., title) for statistical analysis. We plot the kernel density distribution of these two types of consistency count with sentence window size $w = \{1, 2\}$.

As shown in Figure 2, human-written documents are more likely to have higher entity-level and sentence-level consistency count. This analysis indicates that human-written and machine-generated text are different in the factual structure, thus modeling consistency of factual structures is essential in discriminating them.

4 Methodology

In this section, we present our graph-based reasoning approach, which considers factual structures of documents, which is used to guide the reasoning process for the final prediction.

Figure 3 gives a high-level overview of our approach. With a document given as the input, our system begins by calculating the contextual word representations by RoBERTa (§ 4.1). Then, we

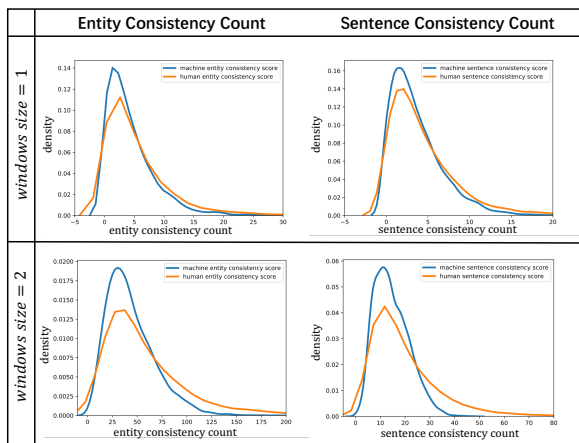


Figure 2: Statistical analysis about entity-level and sentence-level consistency. Orange curve and blue curve indicate kernel density estimation curve for human-written document and machine-generated document respectively. X-axis indicates the value of consistency count and y-axis indicates probability density.

build a graph for capturing the internal factual structure of the whole document (§ 4.2). With the constructed graph, we initialize node representations utilizing internal and external factual knowledge and propagate and aggregate information by a graph neural network to learn graph-enhanced sentence representations (§ 4.3). Then, to model the consistent relations of continuous sentences and compose a document representation for making the final prediction, we employ a sequential model with help of coherence scores from a pre-trained next sentence prediction (NSP) model (§ 4.4).

4.1 Word Representation

In this part, we present how to calculate contextual word representations by a transformer-based model. In practice, we employ RoBERTa (Liu et al., 2019).

Taking a document d as the input, we employ RoBERTa to learn contextual semantic representations for words¹. RoBERTa encoder \mathcal{B} maps document \mathbf{x} with length $|\mathbf{x}|$ into a sequence of following hidden vectors.

$$\mathbf{h}(\mathbf{x}) = [\mathbf{h}(\mathbf{x})_1, \mathbf{h}(\mathbf{x})_2, \dots, \mathbf{h}(\mathbf{x})_{|\mathbf{x}|}] \quad (1)$$

where each $\mathbf{h}(\mathbf{x})_i$ indicates the contextual representation of word i

4.2 Graph Construction

In this part, we present how to construct a graph to reveal the internal factual structure of a docu-

¹In practice, “words” may indicate tokens or token-pieces, we use “words” for a better illustration here.

ment. In practice, we observe that selecting entities, the core participants of events, as arguments to construct the graph leads to less noise to the representation of the factual structure. Therefore, we employ a named entity recognition (NER) model to parse entities mentioned in each sentence. Specifically, taking a document as the input, we construct a graph in the following steps.

- We parse each sentence to a set of entities with an off-the-shelf NER toolkit built by AllenNLP², which is an implementation of Peters et al. (2017). Each entity is regarded as a node in the graph.
- We establish links between inner-sentence and inter-sentence entity node pairs to capture the structural relevance. We add inner-sentence edges to entity pairs in the same sentence for they are naturally relevant to each other. Moreover, we add inter-sentence edges to literally similar inter-sentence entity pairs for they are likely to be the same entity.

After this process, the graph reveals the fine-grained factual structure of a document.

4.3 Graph Neural Network

In this part, we introduce how to initialize node representations and exploit factual structure utilizing multi-layer graph convolution network (GCN) to propagate and aggregate information and finally produce sentence representations.

4.3.1 Node Representation Initialization

We initialize node representations with contextual word representations learnt by RoBERTa and external entity representations pre-trained on Wikipedia.

Contextual Representation Since each entity node is naturally a span of words mentioned in the document, we calculate the contextual representation of each node by the contextual words representations $\mathbf{h}(\mathbf{x})$. Supposing an entity e consists of n words, then the contextual representation $\varepsilon_{\mathcal{B}}$ is calculated with the following formula:

$$\varepsilon_{\mathcal{B}} = ReLU(\mathbf{W}_{\mathcal{B}} \frac{1}{n} \sum_{i=0}^n \mathbf{h}(\mathbf{x})_{e^i}) \quad (2)$$

where $\mathbf{W}_{\mathcal{B}}$ is a weight metric, e^i is the absolute position in the document of the i^{th} word in the span of entity e , and $ReLU$ is an activation function.

²<https://demo.allennlp.org/named-entity-recognition>

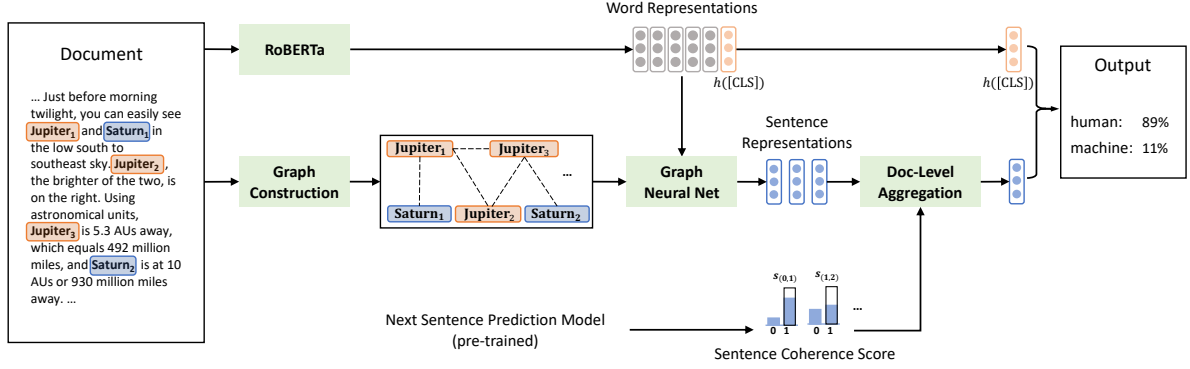


Figure 3: An overview of our approach. Taking a document as the input, we first calculate contextual word representations via RoBERTa (§ 4.1) and represent the factual structure as a graph (§ 4.2). After that, we employ graph neural network to learn sentence representations (§ 4.3). Then, sentence representations are composed to a document representation considering coherence of continuous sentences before making the final prediction (§ 4.4).

Wikipedia-based Entity Representation To model external factual knowledge about entities in the knowledge base, we further represent entity e with a projected wikipedia2vec entity representation (Yamada et al., 2018), which embeds words and entities on Wikipedia pages in a common space. The Wikipedia-based entity representation ε_w is :

$$\varepsilon_w = \text{ReLU}(\mathbf{W}_w \mathbf{v}_e) \quad (3)$$

where \mathbf{v}_e is the wikipedia2vec representation of entity e and \mathbf{W}_w is a weight metric.

The initial representation $\mathbf{H}_e^{(0)} \in \mathbf{R}^d$ of entity node e is the concatenation of contextual representation ε_B and Wikipedia-based entity representation ε_w , with dimension d .

4.3.2 Multi-layer GCN

In order to propagate and aggregate information through multihop neighbouring nodes, we employ multi-layer Graph Convolution Network (GCN) (Kipf and Welling, 2016).

Formally, we denote the constructed graph as G and representation of all nodes as $\mathbf{H} \in \mathbf{R}^{N \times d}$, where N denote the number of nodes. Each row $\mathbf{H}_e \in \mathbf{R}^d$ in \mathbf{H} indicates a representation of node e . We denote the adjacency matrix of graph G as \mathbf{A} and degree matrix as \mathbf{D} . We further calculate $\tilde{\mathbf{A}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$. Then, the formula of multi-layer GCN is described as follows:

$$\mathbf{H}_e^{(i+1)} = \sigma(\tilde{\mathbf{A}} \mathbf{H}_e^{(i)} \mathbf{W}_i) \quad (4)$$

where $\mathbf{H}_e^{(i)}$ denotes the representation of node e calculated by i^{th} layer of GCNs, \mathbf{W}_i is the weight matrix of layer i . σ is an activation function. Specially, $\mathbf{H}_e^{(0)}$ is the initialized node representations.

Finally, through m layers of GCN, we obtain the graph-enhanced node representations based on the structure of the factual graph.

4.3.3 Sentence Representation

According to compositionality, we believe that global representation should come from partial representations. Therefore, we calculate sentence-level representations based on graph-enhanced node representations. Supposing sentence i has N_i corresponding entities, we calculate the representation \mathbf{y}_i of sentence i as follows:

$$\mathbf{y}_i = \frac{1}{N_i} \sum_{j=0}^{N_i} \sigma(\mathbf{W}_s \mathbf{H}_{(i,j)} + \mathbf{b}_s) \quad (5)$$

where σ is an activation function, \mathbf{W}_s is a weight matrix, \mathbf{b}_s is a bias vector and $\mathbf{H}_{(i,j)}$ indicates the representation of j^{th} node in sentence i . The compositionality can also be implemented in other ways, which we leave to future work.

4.4 Aggregation to Document Representation

In this part, we present how to compose a document representation for the final prediction utilizing graph-enhanced sentence representations and coherence score calculated by a pre-trained next sentence prediction (NSP) model.

Coherence Tracking LSTM With graph-enhanced sentence representations given as the input, the factual consistency of continuous sentences is automatically modeled by a sequential model. Specifically, We employ LSTM to track the consistent relations and produce representations $\tilde{\mathbf{y}}_i$ for sentence i

$$\tilde{\mathbf{y}}_i = \text{LSTM}([\mathbf{y}_i]) \quad (6)$$

Next Sentence Prediction Model In order to further model contextual coherence of neighbouring sentence pairs as an additional information, we pre-train an NSP model to calculate the contextual coherence score for each neighbouring sentence pair. We employ RoBERTa (Liu et al., 2019) as the backbone, which receives pairs of sentences as the input and assesses whether the second sentence is a subsequent sentence of the first. Further training details are explained in Appendix A. The outputs S are described as follows.

$$S = [S_{(0,1)}, \dots, S_{(s-1,s)}] \quad (7)$$

where $s+1$ is the number of sentences in document x and each $S_{(i-1,i)}$ is the positive probability score for sentence pair $(i-1, i)$, which indicates how likely it is that sentence i is a subsequent sentence of sentence $i-1$.

Prediction with NSP Score We generate a document-level representation by composing sentence representations before making the final prediction. To achieve this, we take NSP scores as weights and calculate the weighted sum of representations of sentence pairs with the assumption that sentence pairs with higher contextual coherence score should also carry more importance in making the final prediction. The final document representation D is calculated as follows.

$$D = \sum_{j=1}^s S_{(j-1,j)} * [\tilde{y}_{j-1}, \tilde{y}_j] \quad (8)$$

Finally, we make the final prediction by feeding the combination of D and the last hidden vector $h([CLS])$ from RoBERTa through an classification layer. The goal of this operation is to maintain the complete contextual semantic meaning of the whole document because some linguistic clues are left out during graph construction.

5 Experiment

5.1 Experiment Settings

In this paper, we evaluate our system on the following two datasets:

- News-style GROVER-generated dataset provided by Zellers et al. (2019). The human-written instances are collected from Real-News, and machine-generated instances are generated by GROVER-Mega, a large state-of-the-art transformer-based generative model

developed for neural fake news. We largely follow the experimental settings as described by Zellers et al. (2019) and adopt two evaluation metrics: **paired accuracy** and **unpaired accuracy**. In the paired setting, the system is given human-written news and machine-generated news with the same meta-data. The system needs to assign higher machine probability to the machine-generated news than the human-written one. In the unpaired setting, the system is provided with single news document and states whether the document is human-written or machine-generated.

- Webtext-style GPT2-generated dataset provided by OpenAI³. The human-written instances are collected from WebText. Machine-generated instances are generated by GPT-2 XL-1542M (Radford et al., 2019), a powerful transformer-based generative model trained on a corpus collected from popular webpages. For this dataset, we adopt binary classification accuracy as the evaluation metric.

We set nucleus sampling with $p = 0.96$ as the sampling strategy of generator for both datasets, which leads to better generated text quality (Zellers et al., 2019; Ippolito et al., 2019). The statistics of the two datasets are shown in the Table 1.

Dataset	Train	Valid	Test Set	
			Unpaired	Paired
News-style	10,000	3,000	8,000	8,000
Webtext-style	500,000	10,000	10,000	-

Table 1: Statistics of news-style and webtext-style datasets.

Furthermore, we adopt RoBERTa-Base (Liu et al., 2019) as the direct baseline for our experiments because RoBERTa achieves state-of-the-art performance on several benchmark NLP tasks. The hyper-parameters and training details of our model are described in Appendix B.

5.2 Model Comparison

Baseline Settings We compare our system with transformer-based baselines for DeepFake detection, including three powerful transformer-based pre-trained models: **BERT** (Devlin et al., 2018), **XLNet** (Yang et al., 2019) and **RoBERTa** (Liu

³<https://github.com/openai/gpt-2-output-dataset>

Size	Model	Unpaired Acc	Paired Acc
355M	Chance	50.0%	50.0%
	GROVER-Large	80.8%	89.0%
	BERT-Large	73.1%	84.1%
	GPT2	70.1%	78.8%
124M	GROVER-Base	70.1%	77.5%
	BERT-Base	67.2%	80.0%
	GPT2	66.2%	72.5%
	XLNet	77.1%	88.6%
	RoBERTa	80.7%	89.2%
	FAST	84.9%	93.5%

Table 2: Performance on the test set of news-style dataset in terms of unpaired and paired accuracy. Our model is abbreviated as FAST. Size indicates approximate model size. The performance of GROVER, BERT, and GPT2 are reported by [Zellers et al. \(2019\)](#)

[et al., 2019](#)), which are large bidirectional transformers achieving strong performance on multiple benchmark NLP tasks. These baselines add a simple classification layer on top of them and are fine-tuned with standard cross-entropy loss on the binary classification.

For the news-style dataset, we further compare our model with **GPT-2** ([Radford et al., 2019](#)) and **GROVER** ([Zellers et al., 2019](#)). The GROVER-based discriminator is a fine-tuned version of generator GROVER, which has three model sizes: GROVER-Base (124 million parameters), GROVER-Large (335 million parameters), and GROVER-Mega (1.5 billion parameters). Our model is not comparable with GROVER-Mega for the following reasons. Firstly, GROVER-Mega is the fake news generator, and it has a strong inductive bias (e.g., data distribution and sampling strategy) as the discriminator ([Zellers et al., 2019](#)). Secondly, GROVER-Mega has a much larger model size (1.5 billion parameters) than our model.

For the webtext-style dataset, we compare with the baselines we trained with the same hyperparameters. We don’t compare with GPT-2 because it’s the generator for machine-generated text.

Results and Analysis In Table 2, we compare our model with baselines on the test set of news-style dataset with negative instances generated by GROVER-Mega. As shown in the table, our model significantly outperforms our direct baseline RoBERTa with 4.2% improvements on unpaired accuracy and 4.3% improvements on paired accuracy. Our model also significantly outperforms GROVER-Large and other strong transformer-based baselines (i.e., GPT2, BERT, XLNet).

Model	Development Acc	Test Acc
Random	50.00%	50.00%
BERT	85.32%	85.10%
XLNet	88.79%	88.56%
RoBERTa	90.46%	90.10%
FAST	93.10%	93.17%

Table 3: Performance on the development and test set of webtext-style dataset in terms of binary classification accuracy. Our model is abbreviated as FAST.

In Table 3, we compare our model with baselines on the development set and the test set of webtext-style dataset. Our model significantly outperforms strongest transformer-based baseline RoBERTa by 2.64% on the development set and 3.07% on the test set of webtext-style GPT2-generated dataset.

This observation indicates that modeling fine-grained factual structures empower our system to discriminate the difference between human-written text and machine-generated text.

5.3 Ablation Study

Moreover, we also conduct ablation studies to evaluate the impact of each component by conducting experiments about direct baseline RoBERTa-Base and four variants of our full model.

- **RoBERTa-Base** is our direct baseline without considering any structural information.
- **FAST (GCN)** calculate a global document representation by averaging node representations after representation learning by GCN.
- **FAST (GCN w/o wiki)** The node representations eliminate entity representations from wikipedia2vec and the rest are the same as FAST (GCN).
- **FAST (GCN + LSTM)** takes the final hidden state from coherence tracking LSTM (§ 4.4) as the final document-level representation.
- **FAST (GCN + LSTM + NSP)** is the full model introduced in this paper.

As shown in Figure 4, adding GCN improve performance on the development the set of news-style dataset and webtext-style dataset. This verifies that incorporating fine-grained structural information is beneficial for detecting generated text. Eliminating wikipedia-based entity representation from FAST (GCN) drops performance, which indicates

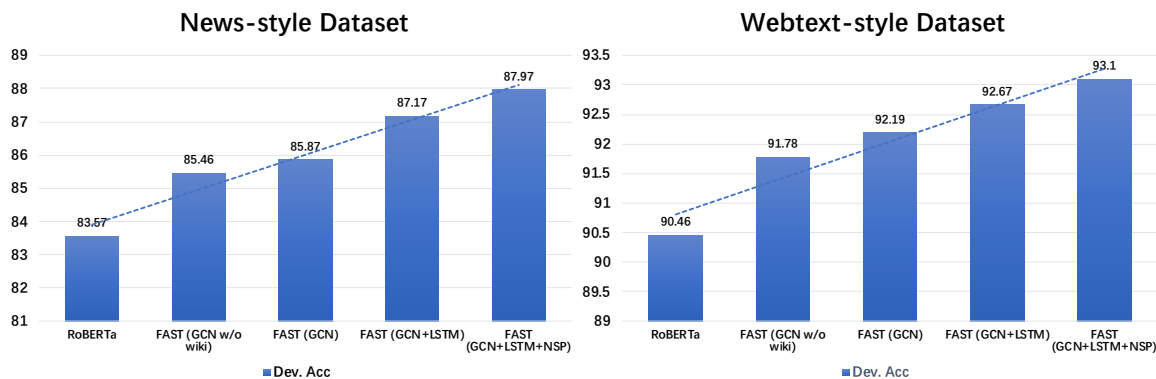


Figure 4: Ablation studies on the development set of the two datasets in terms of binary classification accuracy.

that incorporating external knowledge is also beneficial. Moreover, incorporating coherence tracking LSTM brings further improvement on two datasets, which indicates that modeling consistency of factual structure of continuous sentences is better than simply using global structural information of the document, like the setting in FAST (GCN). Lastly, results also show that incorporating semantic coherence score of pre-trained NSP model is beneficial for discriminating generated text.

5.4 Case Study

As shown in Figure 5, we conduct a case study by giving an example. This example shows human-written news and machine-generated news with the same metadata (i.e., title). The veracity of both documents are correctly predicted by our model. With the given document, our system constructs a factual graph and makes the correct predictions by reasoning over the constructed graph. We can observe that although the continuous sentences in the machine-generated news look coherent, their factual structure is not consistent as they describe events about irrelevant entities. Instead, the human-written news has a more consistent factual structure. However, without utilizing factual structure information, RoBERTa fails to discriminate between these two articles. This observation reflects that our model can distinguish the difference in the factual consistency of machine-generated text and human-written text.

5.5 Error Analysis

To explore further directions for future studies, we randomly select 200 instances and manually summarize representative error types.

The **primary** type of errors is those caused by failing to extract core entities of sentences. The

quality of a constructed graph is somehow limited by the performance of the NER model. This limitation leaves further exploratory space for extraction of internal factual structure. The **second** type of errors is caused by the weakness in the mathematical calculation of the model. For instance, a document describes that “a smaller \$5 million one-off was seized in 2016 and the National Bank of Antigua and Barbuda reclaimed \$30 million stolen in the 2015 heist last year. \$100 million, it was a massive amount. But now we are talking of \$50 million, this is extremely conservative...”. Humans can easily observe that the mentioned numbers are highly inconsistent in the generated text. A machine struggles to discern that. This error type calls for the development of a machine’s mathematical calculation abilities. The **third** error type is caused by failing to model commonsense knowledge. For example, a famous generated document mentioned “In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. ... These four-horned, silver-white unicorns were previously unknown to science.”. Although the text looks coherent, it is still problematic in terms of commonsense knowledge that “unicorn has only one horn”. This leaves space for further research on exploring commonsense knowledge in deepfake detection.

6 Related Work

Recently, fake news detection has attracted growing interest due to the unprecedented amount of fake contents propagating through the internet (Vosoughi et al., 2018). Spreading of fake news arises public concerns (Cooke, 2018) as it may influence essential public events like politic elections (Allcott and Gentzkow, 2017). Online reviews can also be generated by machines, and

Title: Sky Watch: No need to blow your mind when wrapping your brain around celestial distances		
	human-written	machine-generated
Document	... Just before morning twilight, you can easily see Jupiter₁ and Saturn₁ in the low south to southeast sky. Jupiter₂ , the brighter of the two, is on the right. Using astronomical units, Jupiter₃ is 5.3 AUs away, which equals 492 million miles, and Saturn₂ is at 10 AUs or 930 million miles away.As Earth₁ tilts away from the sun and toward the Earth₂ equator, we can see how the plane and plane of the equator move away from the Sun . A 15th-century astronomer, Joseph Bagnold , wrote the 20th-century English word " Gravitational Lensing ". ...
Graph		
	Our Model: [0.02%, 99.98%] RoBERTa: [3.64%, 96.36%]	Our Model: [99.96%, 0.04%] RoBERTa: [25.58%, 74.42%]

Figure 5: A case study of our approach. Continuous words in orange indicate an entity node extracted by our system. Each green solid box indicates a sub-graph corresponding to a sentence, and a blue dashed line indicates an edge between semantically relevant entity pairs. Numbers in orange and blue indicate probability for the human-written document and the machine-generated document respectively.

can even be as fluent as human-written text (Ade-lani et al., 2020). This situation becomes even more serious when recent development of large pre-trained language models (Radford et al., 2019; Zellers et al., 2019) are capable of generating coherent, fluent and human-like text. Two influential works are GPT-2 (Radford et al., 2019) and GROVER (Zellers et al., 2019). The former is an open-sourced, large-scale unsupervised language model learned on web texts, while the latter is particularly learned for news. In this work, we study the problem of discriminating machine-generated and human-written text, and evaluate on datasets produced by both GPT-2 and GROVER.

Advances in generative models have promoted the development of detection methods. Previous studies in the field of DeepFake detection of generated text are dominated by deep-learning based document classification models and studies about discriminating features of generated text. GROVER (Zellers et al., 2019) detects generated text by a fine-tuned model of the generative model itself. Ippolito et al. (2019) fine-tune the BERT model for discrimination and explore how sampling strategies and text excerpt length affect the detection. GLTR (Gehrmann et al., 2019) develops a statistical method of computing per-token likelihoods and visualizes histograms over them to help deepfake detection. Badaskar et al. (2008) and Pérez-Rosas et al. (2017) study language distributional features including n-gram frequencies, text coherence and syntax features. Vijayaraghavan et al. (2020) study the effectiveness of different numeric representations (e.g., TFIDF and Word2Vec) and different

neural networks (e.g., ANNs, LSTMs) for detection. Bakhtin et al. (2019) tackle the problem as a ranking task and study about the cross-architecture and cross-corpus generalization of their scoring functions. Schuster et al. (2019) indicate that simple provenance-based detection methods are insufficient for solving the problem and call for development of fact checking systems. However, existing approaches struggle to capture fine-grained factual structures among continuous sentences, which in our observation is essential in discriminating human-written text and machine-generated text. Our approach takes a step towards modeling fine-grained factual structures for deepfake detection of text.

7 Conclusion

In this paper, we present FAST, a graph-based reasoning approach utilizing fine-grained factual knowledge for DeepFake detection of text. We represent the factual structure of a document as a graph, which is utilized to learn graph-enhanced sentence representations. Sentence representations are further composed through document-level aggregation for the final prediction, where the consistency and coherence of continuous sentences are sequentially modeled. We evaluate our system on a news-style dataset and a webtext-style dataset, whose fake instances are generated by GROVER and GPT-2 respectively. Experiments show that components of our approach bring improvements and our full model significantly outperforms transformer-based baselines on both datasets. Model analysis further suggests that our model can

distinguish the difference in the factual structure of machine-generated and human-written text.

Acknowledgement

Wanjun Zhong, Zenan Xu, Jiahai Wang and Jian Yin are supported by the National Natural Science Foundation of China (U1711262, U1611264, U1711261, U1811261, U1811264, U1911203, 62072483), National Key R&D Program of China (2018YFB1004404), Guangdong Basic and Applied Basic Research Foundation (2019B1515130001), Key R&D Program of Guangdong Province (2018B010107005). The corresponding author is Jian Yin.

References

- David Ifeoluwa Adelani, Haotian Mai, Fuming Fang, Huy H Nguyen, Junichi Yamagishi, and Isao Echizen. 2020. Generating sentiment-preserving fake online reviews using neural language models and their human-and machine-based detection. In *International Conference on Advanced Information Networking and Applications*, pages 1341–1354. Springer.
- Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–36.
- Sameer Badaskar, Sachin Agarwal, and Shilpa Arora. 2008. Identifying real or fake articles: Towards better language modeling. In *Proceedings of the Third International Joint Conference on Natural Language Processing: Volume-II*.
- Anton Bakhtin, Sam Gross, Myle Ott, Yuntian Deng, Marc’Aurelio Ranzato, and Arthur Szlam. 2019. Real or fake? learning to discriminate machine from human generated text. *arXiv preprint arXiv:1906.03351*.
- Nicole A Cooke. 2018. *Fake news and alternative facts: Information literacy in a post-truth era*. American Library Association.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Robert Faris, Hal Roberts, Bruce Etling, Nikki Bourassa, Ethan Zuckerman, and Yochai Benkler. 2017. Partisanship, propaganda, and disinformation: Online media and the 2016 us presidential election. *Berkman Klein Center Research Publication*, 6.
- Sebastian Gehrmann, Hendrik Strobelt, and Alexander M Rush. 2019. Gltr: Statistical detection and visualization of generated text. *arXiv preprint arXiv:1906.04043*.
- Ben Goodrich, Vinay Rao, Peter J Liu, and Mohammad Saleh. 2019. Assessing the factual accuracy of generated text. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 166–175. ACM.
- Daphne Ippolito, Daniel Duckworth, Chris Callison-Burch, and Douglas Eck. 2019. Human and automatic detection of generated text. *arXiv preprint arXiv:1911.00650*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Evaluating the factual consistency of abstractive text summarization. *arXiv preprint arXiv:1910.12840*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Verónica Pérez-Rosas, Bennett Kleinberg, Alexandra Lefevre, and Rada Mihalcea. 2017. Automatic detection of fake news. *arXiv preprint arXiv:1708.07104*.
- Matthew E Peters, Waleed Ammar, Chandra Bhagavathula, and Russell Power. 2017. Semi-supervised sequence tagging with bidirectional language models. *arXiv preprint arXiv:1705.00108*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Tal Schuster, Roei Schuster, Darsh J Shah, and Regina Barzilay. 2019. Are we safe yet? the limitations of distributional features for fake news detection. *arXiv preprint arXiv:1908.09805*.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. *arXiv preprint arXiv:1803.05355*.
- Sairamvinay Vijayaraghavan, Ye Wang, Zhiyuan Guo, John Voong, Wenda Xu, Armand Nasser, Jiaru Cai, Linda Li, Kevin Vuong, and Eshan Wadhwa. 2020. Fake news detection with different models. *arXiv preprint arXiv:2003.04978*.
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Yaqing Wang, Weifeng Yang, Fenglong Ma, Jin Xu, Bin Zhong, Qiang Deng, and Jing Gao. 2019. Weak supervision for fake news detection via reinforcement learning. *arXiv preprint arXiv:1912.12520*.

Ikuya Yamada, Akari Asai, Jin Sakuma, Hiroyuki Shindo, Hideaki Takeda, Yoshiyasu Takefuji, and Yuji Matsumoto. 2018. Wikipedia2vec: An efficient toolkit for learning and visualizing the embeddings of words and entities from wikipedia.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. In *Advances in Neural Information Processing Systems*, pages 9051–9062.

A Training Details of NSP Model

In this part, we describe the training details of our next sentence prediction model. The training data of the NSP model comes from the human-written component of the webtext-style dataset or the news-style dataset depending on which dataset we are running experiments on. We construct the dataset with balanced numbers of positive instances and negative instances. Supposing a positive instance is a continuous sentence pair “ $A; B$ ” from the human-written text, we construct a negative instance “ $A; C$ ”, where C is the most similar sentence in the document of B .

We tackle this problem as a binary classification task. We fine-tune the RoBERTa-Large model with standard cross-entropy loss on the binary classification task. We apply AdamW as the optimizer for model training. We set the learning rate as $1e-5$, batch size as 8, and set max sequence length as 128.

B Training Details of FAST Model

In this part, we describe the training details for our experiments. We employ cross-entropy loss as the loss function. We apply AdamW as the optimizer for model training. We employ RoBERTa-Base as the backbone of our approach. The RoBERTa network and graph-based reasoning model are trained jointly. We set the learning rate as $1e-5$, warmup step as 0, batch size as 4 per gpu, and set max sequence length as 512. The training time for one epoch takes 2 hours on 4 P40 GPUs for the webtext-style dataset, and 20 minutes for the news-style dataset. We set the dimension of the contextual node representation as 100. The dimension of the wikipedia2vec entity representation is 100.