

Multi-Hop Fact Checking of Political Claims

Wojciech Ostrowski*, Arnav Arora, Pepa Atanasova and Isabelle Augenstein

Department of Computer Science, University of Copenhagen, Denmark

qnj566@alumni.ku.dk, {aar, pepa, augenstein}@di.ku.dk

Abstract

Recent work has proposed multi-hop models and datasets for studying complex natural language reasoning. One notable task requiring multi-hop reasoning is fact checking, where a set of connected evidence pieces leads to the final verdict of a claim. However, existing datasets either do not provide annotations for gold evidence pages, or the only dataset which does (FEVER) mostly consists of claims which can be fact-checked with simple reasoning and is constructed artificially. Here, we study more complex claim verification of naturally occurring claims with multiple hops over interconnected evidence chunks. We: 1) construct a small annotated dataset, PolitiHop¹, of evidence sentences for claim verification; 2) compare it to existing multi-hop datasets; and 3) study how to transfer knowledge from more extensive in- and out-of-domain resources to PolitiHop. We find that the task is complex and achieve the best performance with an architecture that specifically models reasoning over evidence pieces in combination with in-domain transfer learning.

1 Introduction

Recent progress in machine learning has seen interest in automating complex reasoning, where a conclusion can be reached only after following logically connected arguments. To this end, multi-hop datasets and models have been introduced, which learn to combine information from several sentences to arrive at an answer. While most of them concentrate on question answering, fact checking is another task that often requires a combination of multiple evidence pieces to predict a claim’s veracity.

Existing fact checking models usually optimize only the veracity prediction objective and assume that the task requires a single inference step. Such models ignore that often several linked evidence chunks have to be explicitly retrieved and combined to make the correct veracity prediction. Moreover,

*Contact Author

¹We make the PolitiHop dataset and the code for the experiments publicly available on <https://github.com/copenlu/politihop>.

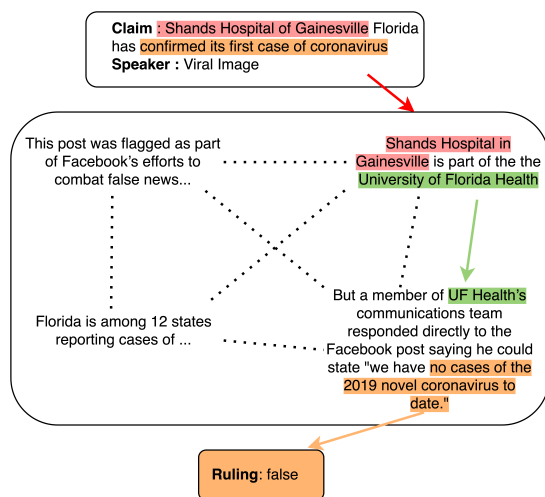


Figure 1: An illustration of multiple hops over an instance from PolitiHop. Each instance consists of a claim, a speaker, a veracity label, and a PolitiFact article the annotated evidence sentences. The highlighted sentences represent the evidence sentences a model needs to connect to arrive at the correct veracity prediction.

they do not provide explanations of their decision-making, which is an essential part of fact checking.

Atanasova *et al.* [2020] note the importance of providing explanations for fact checking verdicts, and propose an extractive summarization model, which optimizes a ROUGE score metric w.r.t. a gold explanation. Gold explanations for this are obtained from the LIAR-PLUS [Alhindi *et al.*, 2018] dataset, which is constructed from PolitiFact² articles written by professional fact checking journalists. However, the dataset does not provide guidance on the several relevant evidence pieces that have to be linked and assume that the explanation requires a single reasoning step. FEVER [Thorne *et al.*, 2018] is another fact checking dataset, which contains annotations of evidence sentences from Wikipedia pages. However, it consists of manually augmented claims, which require limited reasoning capabilities for verification as the evidence mostly consists of one or two sentences.

To provide guidance for the multi-hop reasoning process

²<https://www.politifact.com/>

of a claim’s verification and facilitate progress on explainable fact checking, we introduce **PolitiHop**, a dataset of 500 real-world claims with manual annotations of sets of interlinked evidence chunks from PolitiFact articles needed to predict the claims’ labels. We provide insights from the annotation process, indicating that fact checking real-world claims is an elaborate process requiring multiple hops over evidence chunks, where multiple evidence sets are also possible.

To assess the difficulty of the task, we conduct experiments with lexical baselines, as well as a single-inference step model – BERT [Devlin *et al.*, 2019], and a multi-hop model – Transformer-XH [Zhao *et al.*, 2020]. Transformer-XH allows for the sharing of information between sentences located anywhere in the document by eXtra Hop attention and achieves the best performance. We further study whether multi-hop reasoning learned with Transformer-XH can be transferred to PolitiHop. We find that the model cannot leverage any reasoning skills from training on FEVER, while training on LIAR-PLUS improves the performance on PolitiHop. We hypothesize that this is partly due to a domain discrepancy, as FEVER is constructed from Wikipedia and consists of claims requiring only one or two hops for verification. In contrast, LIAR-PLUS is based on PolitiFact, same as PolitiHop.

Finally, we perform a detailed error analysis to understand the models’ shortcomings and recognize possible areas for improvement. We find that the models perform worse when the gold evidence sets are larger and that, surprisingly, named entity (NE) overlap between evidence and non-evidence sentences does not have a negative effect on either evidence retrieval or label prediction. The best results for Transformer-XH on the dev and test sets are for a different number of hops – 2 and 6, indicating that having a fixed parameter for the number of hops is a downside of Transformer-XH; this should instead be learned for each claim. Overall, our experiments constitute a solid basis to be used for future developments.

To summarise, our **contributions** are as follows:

- We document the first study on multi-hop fact checking of political claims
- We create a dataset for the task
- We study whether reasoning skills learned with a multi-hop model on similar datasets can be transferred to PolitiHop
- We analyze to what degree existing multi-hop reasoning methods are suitable for the task

2 Multi-Hop Fact Checking

A multi-hop fact checking model $f(X)$ receives as an input $X = \{(claim_i, document_i) | i \in [1, |X|]\}$, where $document_i = [sentence_{i,j} | j \in [1, |document_i|]]$ is the corresponding PolitiFact article for $claim_i$ and consists of a list of sentences. During the training process, the model learns to (i) select which sentences from the input contain evidence needed for the veracity prediction $y_i^S = [y_{i,j}^S \in \{0, 1\} | j \in [1, |document_i|]]$ (**sentence selection task**), where 1 indicates that the sentence is selected as an evidence sentence; and (ii) predict the veracity label of the claim $y_i^L \in \{True, False, Half - True\}$, based on the extracted evidence (**veracity prediction task**). The sentences selected by

Statistic	Test	Train
#Words per article	569 (280.8)	573 (269.1)
#Sent. per article	28 (12.8)	28 (12.8)
#Evidence sent. per article	11.75 (5.56)	6.33 (2.98)
#Evidence sent. per set	2.88 (1.43)	2.59 (1.51)
#Sets per article	4.08 (1.83)	2.44 (1.28)
Label Distribution		
False	149	216
Half-true	30	47
True	21	37

Table 1: PolitiHop dataset statistics. Test set statistics are calculated for a union of two annotators; train instances are annotated by one annotator only, which makes some measures different across splits. We report the mean and standard deviation (in parentheses).

the model as evidence provide *sufficient* explanation, which allows to verify the corresponding claim *efficiently* instead of reading the whole article. Each evidence set consists of k sentences, where $k \in [1, \max_{i \in [1, |X|]}(|document_i|)]$ is a hyperparameter of the model. Figure 1 illustrates the process of multi-hop fact checking, where multiple evidence sentences provide different information, which needs to be connected in logical order to reach the final veracity verdict.

2.1 Dataset

We present PolitiHop, the first dataset for multi-hop fact checking of real-world claims. It consists of 500 manually annotated claims in written English, split into a training (300 instances) and a test set (200 instances). For each claim, the corresponding PolitiFact article was retrieved, which consists of a discussion of each claim and its veracity, written by a professional fact checker. The annotators then selected sufficient sets of evidence sentences from said articles. As sometimes more than one set can be found to describe a reason behind the veracity of a claim independently, we further take each set in the training split as a separate instance, resulting in 733 training examples. Each training example is annotated by one annotator, whereas each test example is annotated by two. We split the training data into train and dev datasets, where the former has 592 examples and the latter – 141. For veracity prediction, we arrived at Krippendorff’s α and Fleiss’ κ agreement values of 0.638 and 0.637, respectively. By comparison, Thorne *et al.* [2018] reported Fleiss’ κ of 0.684 on FEVER. For the sentence prediction, we attain Krippendorff’s α of 0.437.

Table 1 presents statistics of the dataset. The average number of evidence sentences per set is above 2, which already indicates that the task is more complex than the FEVER dataset. In FEVER, 83.2% of the claims require one sentence, whereas in PolitiHop, only 24.8% require one sentence.

3 Models

We compare the performance of five different models to measure the difficulty of automating the task.

Majority. Label prediction only. The majority baseline labels all claims as false.

Random. We pick a random number $k \in [1, 10]$ and then randomly choose k sentences from the document as evidence. For label prediction, we randomly pick one of the labels.

TF-IDF. For each instance x_i we construct a vector $v_i^C = [v_{il}^C | l \in [0, |N^C|]]$ with TF-IDF scores v_{il}^C for all n-grams N^C found in all of the claims; and one vector $v_i^D = [v_{im}^D | m \in [0, |N^D|]]$ with TF-IDF scores v_{im}^D for all n-grams N^D found in all of the documents, where $n \in [2, 3]$. We then train a Naive Bayes model $g(V)$, where $V = \{v_i = (v_i^C \cdot v_i^D) | i \in [0, |X|]\}$ is the concatenation of the two feature vectors.

BERT. We first train a Transformer model [Vaswani *et al.*, 2017], which does not include a multi-hop mechanism, but applies a single inference step to both the evidence retrieval and the label prediction tasks. We employ BERT [Devlin *et al.*, 2019] with the base pre-trained weights. Each sentence from a fact checking document is encoded separately, combined with the claim and the author of the claim. We refer to the encoded triple as *node* τ . The tokens of one node $x_\tau = \{x_{\tau,j} | j \in [0, |x_\tau|]\}$ are encoded with the BERT model into contextualized distributed representations: $h_\tau = \{h_{\tau,j} | j \in [0, |x_\tau|]\}$. The encoded representations of all nodes are passed through two feed-forward layers:

$$p(y^L | \tau) = \text{softmax}(\text{Linear}(h_{\tau,0})) \quad (1)$$

$$p(y^S | \tau) = \text{softmax}(\text{Linear}(h_{\tau,0})) \quad (2)$$

$$p(y^L | X) = \sum_{\tau} p(y^L | \tau) p(y^S | \tau) \quad (3)$$

The first layer predicts the veracity of the claim given a particular node τ by using the contextual representation of the “[CLS]” token, located at the first position (Eq. 1). The second feed-forward layer learns the importance of each node in the graph (Eq. 2). The outputs of these two layers are combined for the final label prediction (Eq. 3). For evidence prediction, we choose k most important sentences, as ranked by the second linear layer. In our experiments, we set $k = 6$ since this is the average number of evidence sentences selected by a single annotator. The implementation of the feed-forward prediction layers is the same as in Transformer-XH, described below, and can be viewed as an ablation of Transformer-XH removing the eXtra Hop attention layers.

Transformer-XH. Transformer-XH is a good candidate for a multi-hop model for our setup as it has previously achieved the best multi-hop evidence retrieval results on FEVER. It is also inspired by and improves over other multi-hop architectures [Liu *et al.*, 2020; Zhou *et al.*, 2019], and we conjecture that the results should be generalisable for its predecessors as well. Not least, its architecture allows for ablation studies of the multi-hop mechanism. Following previous work on applying Transformer-XH to FEVER [Zhao *et al.*, 2020], we encode node representations as with the BERT model and construct a fully connected graph with them. Transformer-XH uses eXtra hop attention layers to enable information sharing between the nodes. An eXtra hop attention layer is a Graph Attention layer (GAT) [Veličković *et al.*, 2018], which receives as input a graph $\{X, E\}$ of all evidence nodes X and the edges between them E , where the edges encode

the attention between two nodes in the graph. Each eXtra hop layer computes the attention between a node and its neighbors, which corresponds to one hop of reasoning across nodes. Transformer-XH applies L eXtra hop layers to the BERT node encodings H_0 , which results in new representations H_L that encode the information shared between the nodes, unlike BERT, which encodes each input sentence separately. We use three eXtra hop layers as in [Zhao *et al.*, 2020], which corresponds to three-hop reasoning, and we experiment with varying the number of hops. The representations H_L are passed to the final two linear layers for label and evidence prediction as in BERT. The final prediction of the veracity label $p(y^L | \{X, E\})$ now can also leverage information exchanged in multiple hops between the nodes through the edges E between them.

4 Experiments

We address the following research questions:

- Can multi-hop architectures successfully reason over evidence sets on `POLITIHOP`?
- How do multi-hop vs. single inference architectures fare in an adversarial evaluation, where named entities (NE) in evidence and non-evidence sentences overlap?
- Does pre-training on related small in-domain or large out-of-domain datasets improve model performance?

We further perform ablation studies to investigate the influence of different factors on performance (see Section 6). We also include additional ablation studies, experimental details, and annotations guidelines in the supplemental material³.

4.1 Experimental Setup

Metrics. We use macro-F1 score and accuracy for the veracity prediction task and F1 and precision for the evidence retrieval task. To calculate the performance on both tasks jointly, we use the FEVER score [Thorne *et al.*, 2018], where the model has to retrieve at least one full evidence set and predict the veracity label correctly for the label prediction to count as correct. We consider a single evidence set to be sufficient for correct label prediction. As each example from train and dev sets in `POLITIHOP`, and every example from `LIAR-PLUS`, has one evidence set, all evidence sentences need to be retrieved for these. The employed measures for evidence retrieval allow for comparison to related work and for relaxing the requirements on the models. We consider the FEVER score to be the best for evaluating explainable fact checking.

Dataset settings. We consider three settings: full article, even split and adversarial. For full, the whole article for each claim is given as input. For even split, we pick all sentences from the same article, but restrict the number of non-evidence sentences to be at most equal to the number of evidence sentences. Non-evidence sentences are picked randomly. This results in a roughly even split between evidence and non-evidence sentences for the test set. Since we divide train and dev datasets into one evidence set per example, but keep all non-evidence for each, the number of non-evidence sentences for instances in these splits is usually 2-3 times

³<https://github.com/copenlu/politihop>

larger than the number of evidence sentences. To examine if the investigated multi-hop models overfit on named entity (NE) overlaps, we further construct an adversarial dataset from the even split dataset by changing each non-evidence sentence to a random sentence from any PolitiFact article, which contains at least one NE present in the original evidence sentences. While such sentences can share information about a relevant NE, they are irrelevant for the claim. We argue that this is a good testbed to understand if a fact checking model can successfully reason over evidence sets and identify non-evidence sentences, even if they contain relevant NEs, which are rather surface features not indicating whether the sentence is relevant to the claim.

Training settings. We perform transfer learning, training on in-domain data (LIAR-PLUS, PolitiHop), out-of-domain data (FEVER), or a combination thereof.

Note that the measures do not consider the order of the sentences in the evidence set, and the systems do not predict that as well. We believe other measures and models that take that into account should be explored in future work. Here we consider them to appear in the same order as in the document. This also corresponds to the way they were annotated.

5 Results

Full article setting. From the results in Table 2, we can observe that both BERT and Transformer-XH greatly outperform the Random and TF-IDF baselines. Out of BERT and Transformer-XH, neither model clearly outperforms the other on our dataset. This is surprising as Transformer-XH outperforms the BERT baselines by a significant margin on both FEVER and the multi-hop dataset HotpotQA [Zhao *et al.*, 2020]. However, we observe that the best performance is achieved with Transformer-XH trained on LIAR-PLUS, then fine-tuned on PolitiHop. It also achieves the highest FEVER scores on PolitiHop in that setting. Further, very low FEVER scores of both Transformer-XH and BERT indicate how challenging it is to retrieve the whole evidence set.

Adversarial setting. We train the Transformer-XH models on the even split setting, then evaluate it on both adversarial and even split datasets (see Table 3). The model performs similarly in both settings. When compared on test sets, it achieves a higher FEVER score on the adversarial, but dev sets FEVER score is higher on the even split setting. Overall, the results show Transformer-XH is robust towards NE overlap.

Out-of-domain pre-training on FEVER. In this experiment, we examine whether pre-training Transformer-XH on the large, but out-of-domain dataset FEVER, followed by fine-tuning on LIAR-PLUS, then on PolitiHop improves results on PolitiHop. As can be seen from Table 4, it does not have a positive effect on performance in the full setting, unlike pre-training on LIAR-PLUS. We hypothesize that the benefits of using a larger dataset are outweighed by the downsides of it being out-of-domain. We further quantify the domain differences between datasets. We use Jensen-Shannon divergence [Lin, 1991], commonly employed for this purpose [Ruder and Plank, 2017]. The divergence between FEVER

and PolitiHop is 0.278, while between LIAR-PLUS and PolitiHop is 0.063, which further corroborates our hypothesis. Another reason might be that PolitiHop has several times more input sentences compared to FEVER. Labelling difference might matter as well: FEVER uses ‘true’, ‘false’ and ‘not enough info’, while PolitiHop uses ‘true’, ‘false’ and ‘half-true’.

6 Analysis and Discussion

In Section 5, we documented experimental results on multi-hop fact checking of political claims. Overall, we found that multi-hop training on Transformer-XH gives small improvements over BERT, that pre-training on in-domain data helps, and that Transformer-XH deals well with an adversarial test setting. Below, we aim to further understand the impact of modeling multi-hop reasoning explicitly with a number of ablation studies:

- How the evidence set size affects performance
- Varying the hops’ number in Transformer-XH
- The impact of evidence set size on performance
- How NE overlap affects performance
- To what extent Transformer-XH pays attention to relevant evidence sentences

Varying the number of hops in Transformer-XH. We train Transformer-XH with a varying number of hops to see if there is any pattern in how many hops result in the best performance. Zhao *et al.* [2020] perform a similar experiment and find that 3 hops are best, similar for 2-5 hops, while the decrease in performance is noticeable for 1 and 6 hops. We experiment with hops between 1 and 7 (see Table 5). Evidence retrieval performance is quite similar in each case. There are some differences for the label prediction task: 1 and 2 hops have slightly worse performance, the 4-hop model has the highest test score and the lowest dev score, while the exact opposite holds for the 5-hop model. Therefore, no clear pattern can be found. One reason for this could be the high variance of the annotated evidence sentences in PolitiHop.

Evidence set size vs. performance. Not surprisingly, larger number of evidence sentences leads to higher precision and lower recall, resulting in a lower FEVER score. This is true for both models, as Table 6 (top) indicates. We also notice that the smaller the number, the smaller the ratio of evidence to non-evidence sentences. For instance, if a claim has two sets of evidence, one of size 1 and the other of size 3, then after splitting into one example per set, there are 4 non-evidence sentences in each of the two examples, but the one with set of size 1 has only one evidence sentence – which decreases the evidence to non-evidence ratio and makes it more difficult to achieve high precision.

Named entity overlap vs. performance. To measure the effect of having the same NEs in evidence and non-evidence sentences, we computed NE overlap – a measure of the degree to which evidence and non evidence sentences share NEs. We compute the overlap as $|E \cap N|/|E \cup N|$; E and N are sets of NEs in evidence and non-evidence sentences, respectively. Table 6 (bottom) shows that a higher NE overlap results in more confusion when retrieving evidence sentences, but it

	Dev					Test				
	L-F1	L-Acc	E-F1	E-Prec	FEVER	L-F1	L-Acc	E-F1	E-Prec	FEVER
Random	34.1	38.5	22.9	30.2	4.5	24.2	27.7	14.7	12.2	0.7
Majority	27.3	69.5	-	-	-	28.6	75.0	-	-	-
Annotator	-	-	-	-	-	76.3	-	52.4	49.2	-
TF-IDF	34.4	69.5	-	-	-	34.0	76.0	-	-	-
LIAR-PLUS full articles dataset										
BERT	45.4	70.9	18.4	13.7	14.9	57.0	76.0	32.9	38.9	13.0
Transformer-XH	56.2	74.5	17.1	12.8	14.2	56.3	79.5	30.3	35.8	12.0
PolitiHop full articles dataset										
BERT	54.7	69.5	32.0	23.6	31.9	44.8	76.0	47.0	54.2	24.5
Transformer-XH	61.1	76.6	30.4	22.3	34.8	43.3	75.5	44.7	51.7	23.5
LIAR-PLUS and PolitiHop full articles										
BERT	64.4	75.9	29.6	21.7	28.4	<u>57.8</u>	79.5	45.1	52.2	23.5
Transformer-XH	64.6	78.7	32.4	23.8	38.3	57.3	80.5	47.2	54.5	24.5

Table 2: PolitiHop results for label (L), evidence (E) and joint (FEVER) performance in the full setting. Best results with a particular training dataset (LIAR-PLUS/PolitiHop/LIAR-PLUS and PolitiHop) are emboldened and the best results across all set-ups are underlined.

	Dev					Test				
	L-F1	L-Acc	E-F1	E-Prec	FEVER	L-F1	L-Acc	E-F1	E-Prec	FEVER
even split	58.1	71.6	47.7	35.5	52.5	62.9	82.0	58.2	66.7	31.0
adversarial	56.5	70.9	49.9	38.7	46.8	56.4	77.0	63.6	76.0	33.5

Table 3: PolitiHop adversarial vs even split dataset results for label (L), evidence (E) and joint (FEVER) performance for Transformer-XH trained on LIAR-PLUS and PolitiHop on the even split setting. Best result emboldened.

does not have a significant influence on label prediction in the case of Transformer-XH. For BERT, higher NE overlap leads to a bigger, negative effect on both tasks. This suggests Transformer-XH is more robust to NE overlaps.

Attention over evidence sentences. We investigate what attention patterns Transformer-XH learns. Ideally, attention flowing from evidence sentences should be higher than from non-evidence ones since this determines how much they contribute to the final representations of each sentence. To do this, we inspect the weights in the final eXtra hop layer. We normalize results by measuring the ratio of the given attention to the average attention for the given graph: 1 means average, over/under 1 means more/less than average. Table 7 shows average ratios for evidence vs. non-evidence sentences. One notable finding is that attention weights from evidence sentences are higher than average, and attention from non-evidence sentences is lower. The Welch t-test indicates that the difference is significant with a p -value lower than 10^{-30} . So, attention weights get more importance on average, but the magnitude of this effect is quite limited. This shows the limitations of using Transformer-XH for this task.

7 Related Work

Fact checking. Several datasets have been released to assist in automating fact checking. Vlachos and Riedel [2014] present a dataset with 106 political claim-verdict pairs. The FakeNewsChallenge⁴, provides 50K headline-article pairs

⁴<http://www.fakenewschallenge.org/>

and formulates the task of fact checking as stance detection between the headline and the body of the article. The relationship between these two tasks is further explored in Hardalov *et al.* [2021]. Wang [2017] extract 12.8K claims from PolitiFact constituting the LIAR dataset. Alhindi *et al.* [2018] introduce the LIAR-PLUS dataset extending the latter with automatically extracted summaries from PolitiFact articles. These are, however, high-level explanations that omit evidence details. LIAR-PLUS also does not provide annotation of particular evidence sentences from the article leading to the final verdict and the possible different evidence sets. Augenstein *et al.* [2019] present a real-world dataset constructed from 26 fact checking portals, including PolitiFact, consisting of 35k claims paired with crawled evidence documents. Thorne *et al.* [2018] present the FEVER dataset, consisting of 185K claims produced by manually re-writing Wikipedia sentences. Furthermore, Niewinski *et al.* [2019] from the FEVER’2019 shared task [Thorne *et al.*, 2019] and Hidey *et al.* [2020] use adversarial attacks to show the vulnerability of models trained on the FEVER dataset to claims that require more than one inference step. Unlike prior work, we construct a dataset with annotations of the different reasoning sets and the multiple hops that constitute them.

Multi-hop datasets. Multi-hop reasoning has been mostly studied in the context of Question Answering (QA). Yang *et al.* [2018] introduce HotpotQA with Wikipedia-based question-answer pairs requiring reasoning over multiple documents and provide gold labels for sentences supporting the answer. Welbl *et al.* [2018] introduce MedHop and WikiHop datasets for reasoning over multiple documents. These are

	Dev					Test				
	L-F1	L-Acc	E-F1	E-Prec	FEVER	L-F1	L-Acc	E-F1	E-Prec	FEVER
FEVER+LIAR-PLUS+PolitiHop	48.6	70.2	30.5	22.2	32.6	59.9	83.0	45.1	52.7	21.5
LIAR-PLUS+PolitiHop	64.6	78.7	32.4	23.8	38.3	57.3	80.5	47.2	54.5	24.5

Table 4: PolitiHop full results for label (L), evidence (E) and joint (FEVER) performance for Transformer-XH trained on different datasets. Best model emboldened.

	Dev					Test				
	L-F1	L-Acc	E-F1	E-Prec	FEVER	L-F1	L-Acc	E-F1	E-Prec	FEVER
1	54.0	75.2	47.1	35.2	52.5	58.9	79.5	58.7	67.6	33.0
2	56.1	73.0	47.5	35.4	53.9	59.8	78.5	58.1	66.7	32.5
3	58.1	71.6	47.7	35.5	52.5	62.9	82.0	58.2	66.7	31.0
4	53.3	70.9	47.0	34.9	50.4	65.0	82.0	58.9	67.6	33.0
5	59.6	73.0	47.7	35.5	51.8	55.3	76.5	58.7	67.3	32.0
6	56.5	73.0	45.9	34.2	50.4	64.9	81.5	57.5	66.0	35.0
7	56.3	71.6	46.4	34.6	50.4	62.8	81.5	57.9	66.4	33.0

Table 5: PolitiHop Transformer-XH results for label (L), evidence (E) and joint (FEVER) performance for training on the LIAR-PLUS + PolitiHop even split datasets with a varying number of hop layers. Best sentence number emboldened.

	Transformer-XH					BERT				
	L-F1	L-Acc	E-F1	E-Prec	FEVER	L-F1	L-Acc	E-F1	E-Prec	FEVER
1 or 2 evidence sentences	63.9	76.8	43.7	29.2	74.4	53.5	72.0	41.3	27.5	62.2
3+ evidence sentences	60.9	66.1	67.5	56.2	42.4	57.8	66.1	65.8	54.8	40.7
< 40% NE overlap	62.5	77.0	59.1	46.2	62.3	62.0	75.4	57.7	45.1	59.0
≥ 40% NE overlap	63.6	71.0	48.5	35.5	60.9	47.5	66.7	46.2	33.8	49.3

Table 6: PolitiHop adversarial dev set performance vs. (top) evidence set size and (bottom) NE overlap between evidence and non-evidence sentences for label (L), evidence (E) and joint (FEVER) performance. Better model emboldened.

ev → non-ev	ev → ev	non-ev → non-ev	non-ev → ev
1.085	1.076	0.966	0.964

Table 7: Attention weights in the last eXtra hop layer of Transformer-XH. The numbers are the average ratios of the actual attention weights to average attention weight of the given graph.

constructed using Wikipedia and DrugBank as Knowledge Bases (KB), and are limited to entities and relations existing in the KB. This, in turn, limits the type of questions that can be generated. TriviaQA [Joshi *et al.*, 2017] contains multiple documents for question-answer pairs but has few examples where reasoning over multiple paragraphs from different documents is necessary.

Multi-hop models. Chen and Durrett [2019] observe that models without multi-hop reasoning are still able to perform well on a large portion of the test dataset. Hidey *et al.* [2020] employ a pointer-based architecture, which re-ranks documents related to a claim and jointly predicts the sequence of evidence sentences and their stance to the claim. Asai *et al.* [2020] sequentially extract paragraphs from the reasoning path conditioning on the documents extracted on the previous step. CogQA [Ding *et al.*, 2019] detect spans and entities of interest and then run a BERT-based Graph Convolutional Network for ranking. Nie *et al.* [2019] perform semantic retrieval

of relevant paragraphs followed by span prediction in the case of QA and 3-way classification for fact checking. Zhou *et al.* [2019], Liu *et al.* [2020], and Zhao *et al.* [2020] model documents as a graph and apply attention networks across the nodes of the graph. We use Zhao *et al.*'s [2020] model due to its strong performance in multi-hop QA on the HotpotQA dataset, in evidence-based fact checking on FEVER, and to evaluate its performance on real-world claim evidence reasoning.

8 Conclusions

In this paper, we studied the novel task of multi-hop reasoning for fact checking of real-world political claims, which encompasses both evidence retrieval and claim veracity prediction. We presented PolitiHop, the first political fact checking dataset with annotated evidence sentences. We compared several models on PolitiHop and found that the multi-hop architecture Transformer-XH slightly outperforms BERT in most of the settings, especially in terms of evidence retrieval, where BERT is easily fooled by named entity overlaps between the claim and evidence sentences. The performance of Transformer-XH is further improved when retrieving more than two evidence sentences and the number of hops larger than one, which corroborates the assumption of the multi-hop nature of the task.

References

- [Alhindi *et al.*, 2018] Tariq Alhindi, Savvas Petridis, and Smaranda Muresan. Where is Your Evidence: Improving Fact-checking by Justification Modeling. In *Proceedings of FEVER*, November 2018.
- [Asai *et al.*, 2020] Akari Asai, Kazuma Hashimoto, Hananeh Hajishirzi, Richard Socher, and Caiming Xiong. Learning to Retrieve Reasoning Paths over Wikipedia Graph for Question Answering. In *Proceedings of ICLR*, 2020.
- [Atanasova *et al.*, 2020] Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. Generating Fact Checking Explanations. In *Proceedings of ACL*, pages 7352–7364, July 2020.
- [Augenstein *et al.*, 2019] Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims. In *Proceedings of EMNLP-IJCNLP*, pages 4685–4697, November 2019.
- [Chen and Durrett, 2019] Jifan Chen and Greg Durrett. Understanding Dataset Design Choices for Multi-hop Reasoning. In *Proceedings of NAACL*, pages 4026–4032, June 2019.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL*, pages 4171–4186, June 2019.
- [Ding *et al.*, 2019] Ming Ding, Chang Zhou, Qibin Chen, Hongxia Yang, and Jie Tang. Cognitive Graph for Multi-Hop Reading Comprehension at Scale. In *Proceedings of ACL*, pages 2694–2703, July 2019.
- [Hardalov *et al.*, 2021] Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. A survey on stance detection for mis- and disinformation identification. *arXiv preprint arXiv:2103.00242*, 2021.
- [Hidey *et al.*, 2020] Christopher Hidey, Tuhin Chakrabarty, Tariq Alhindi, Siddharth Varia, Kriste Krstovski, Mona Diab, and Smaranda Muresan. DeSePtion: Dual sequence prediction and adversarial examples for improved fact-checking. In *Proceedings of ACL*, pages 8593–8606, July 2020.
- [Joshi *et al.*, 2017] Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of ACL*, pages 1601–1611, July 2017.
- [Lin, 1991] Jianhua Lin. Divergence Measures Based on the Shannon Entropy. *IEEE Transactions on Information theory*, 37(1):145–151, pages 145–151, 1991.
- [Liu *et al.*, 2020] Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. Fine-grained Fact Verification with Kernel Graph Attention Network. In *Proceedings of ACL*, pages 7342–7351, July 2020.
- [Nie *et al.*, 2019] Yixin Nie, Songhe Wang, and Mohit Bansal. Revealing the Importance of Semantic Retrieval for Machine Reading at Scale. In *Proceedings of EMNLP-IJCNLP*, pages 2553–2566, November 2019.
- [Niewinski *et al.*, 2019] Piotr Niewinski, Maria Pszona, and Maria Janicka. GEM: Generative enhanced model for adversarial attacks. In *Proceedings of FEVER*, pages 20–26, November 2019.
- [Ruder and Plank, 2017] Sebastian Ruder and Barbara Plank. Learning to select data for transfer learning with Bayesian optimization. In *Proceedings of EMNLP*, pages 372–382, September 2017.
- [Thorne *et al.*, 2018] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. FEVER: a large-scale dataset for fact extraction and VERification. In *Proceedings of NAACL*, pages 809–819, June 2018.
- [Thorne *et al.*, 2019] James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. The FEVER 2.0 Shared Task. In *Proceedings of the Second Workshop on FEVER*, pages 1–6, November 2019.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Proceedings of NeurIPS*, page 6000–6010, 2017.
- [Veličković *et al.*, 2018] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. In *Proceedings of ICLR*, 2018.
- [Vlachos and Riedel, 2014] Andreas Vlachos and Sebastian Riedel. Fact Checking: Task definition and dataset construction. In *Proceedings of LTCSS*, pages 18–22, June 2014.
- [Wang, 2017] William Yang Wang. “Liar, Liar Pants on Fire”: A New Benchmark Dataset for Fake News Detection. In *Proceedings of ACL*, pages 422–426, July 2017.
- [Welbl *et al.*, 2018] Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. Constructing datasets for multi-hop reading comprehension across documents. *TACL*, 6:287–302, 2018.
- [Yang *et al.*, 2018] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Proceedings of EMNLP*, pages 2369–2380, November 2018.
- [Zhao *et al.*, 2020] Chen Zhao, Chenyan Xiong, Corby Rosset, Xia Song, Paul N. Bennett, and Saurabh Tiwary. Transformer-XH: Multi-Evidence Reasoning with eXtra Hop Attention. In *ICRL*, 2020.
- [Zhou *et al.*, 2019] Jie Zhou, Xu Han, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. GEAR: Graph-based Evidence Aggregating and Reasoning for Fact Verification. In *Proceedings of ACL*, pages 892–901, July 2019.