

TWOWINGOS: A Two-Wing Optimization Strategy for Evidential Claim Verification

Wenpeng Yin, Dan Roth

University of Pennsylvania

{wenpeng, danroth}@seas.upenn.edu

Abstract

Determining whether a given claim is supported by evidence is a fundamental NLP problem that is best modeled as Textual Entailment. However, given a large collection of text, finding evidence that could support or refute a given claim is a challenge in itself, amplified by the fact that different evidence might be needed to support or refute a claim. Nevertheless, most prior work decouples evidence identification from determining the truth value of the claim given the evidence.

We propose to consider these two aspects jointly. We develop TWOWINGOS (**two-wing optimization strategy**), a system that, while identifying appropriate evidence for a claim, also determines whether or not the claim is supported by the evidence. Given the claim, TWOWINGOS attempts to identify a subset of the evidence candidates; given the predicted evidence, it then attempts to determine the truth value of the corresponding claim. We treat this challenge as coupled optimization problems, training a joint model for it. TWOWINGOS offers two advantages: (i) Unlike pipeline systems, it facilitates flexible-size evidence set, and (ii) Joint training improves both the claim verification and the evidence identification. Experiments on a benchmark dataset show state-of-the-art performance.¹

1 Introduction

A claim, e.g., “Marilyn Monroe worked with Warner Brothers”, is an assertive sentence that may be true or false. While the task of claim verification *will not* tell us the absolute truth of this claim, it is expected to determine whether the claim is supported by evidence in a given text collection. Specifically, given a claim and a text corpus, *evidential claim verification*, demonstrated in

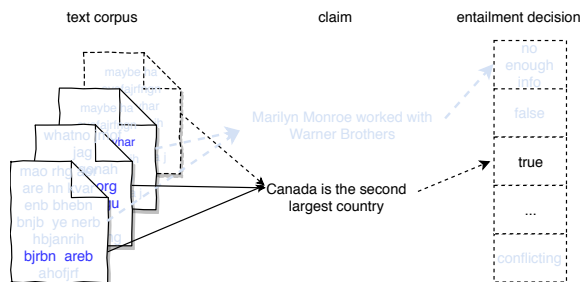


Figure 1: Illustration of evidential claim verification task. For a claim, we determine its truth value by evidence identified from a text corpus.

Figure 1, aims at identifying text snippets in the corpus that act as evidence that supports or refutes the claim.

This problem has broad applications. For example, knowledge bases (KB), such as Freebase (Bollacker et al., 2008), YAGO (Suchanek et al., 2007), can be augmented with a new relational statement such as “(Afghanistan, is_source_of, Kushan Dynasty)”. This needs to be first verified by a claim verification process and supported by evidence (Roth et al., 2009; Chaganty et al., 2017). More broadly, claim verification is a key component in any technical solution addressing recent concerns about the trustworthiness of online content (Vydiswaran et al., 2011; Pasternack and Roth, 2013; Hovy et al., 2013). In both scenarios, we care about whether or not a claim holds, and seek reliable evidence in support of this decision.

Evidential claim verification requires that we address three challenges. First, to locate text snippets in the given corpus that can potentially be used to determine the truth value of the given claim. This differs from the conventional textual entailment (TE) problem (Dagan et al., 2013) as here we first look for the premises given a hypothesis. Clearly, the evidence one seeks depends on the claim, as well as on the eventual entailment

¹cogcomp.org/page/publication_view/847

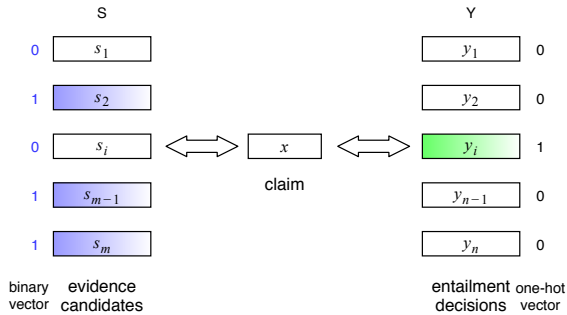


Figure 2: TWOWINGOS, a generic two-wing optimization framework. A subset of the evidence candidates $S_e = \{s_1, \dots, s_{m-1}, s_m\}$ is chosen via a binary vector (left), and an n -valued entailment decision $y_i \in Y$ is chosen (right), with respect to the claim x .

decision – the same claim would require different supporting than refuting evidence. This motivates us to develop an approach that can transfer knowledge from claim verification to evidence identification. Second, the evidence for a claim might require aggregating information from multiple sentences and even multiple documents (rf. #3 in Table 4). Therefore, a *set*, rather than a collection of independent text snippets, should be chosen to act as evidence. And, finally, in difference from TE, given a set of evidence sentences as a premise, the truth value of the claim should depend on *all of the evidence*, rather than on a single sentence there.

The discussion above suggests that claim verification and evidence identification are tightly coupled. Claim should influence the identification of appropriate evidence, and “trusted evidence boosts the claim’s veracity” (Vydiswaran et al., 2011). Consequently, we propose TWOWINGOS, a two-wing optimization strategy², to support this process. As shown in Figure 2, we consider a set of sentences S as the candidate evidence space, a claim x , and a decision space Y for the claim verification. In the optimal condition, a one-hot vector over Y indicates which decision to make towards the claim, and a binary vector over S indicates a subset of sentences S_e (in blue in Figure 2) to act as evidence.

Prior work mostly approached this problem as a pipeline procedure – first, given a claim x , determine S_e by some similarity matching; then, conduct textual entailment over (S_e, x) pairs. Our framework, TWOWINGOS, optimizes the two

²By “two-wing optimization”, we mean that the same object, i.e., the claim, is mapped into two target spaces in a joint optimization scheme.

subtasks jointly, so that both claim verification and evidence identification can enhance each other. TWOWINGOS is a generic framework making use of a shared representation of the claim to co-train evidence identification and claim verification.

TWOWINGOS is tested on the FEVER benchmark (Thorne et al., 2018), showing $\approx 30\%$ F_1 improvement for evidence identification, and $\approx 23\%$ accuracy increase in claim verification. Our analysis shows that (i) entity mentions in claims provide a strong clue for retrieving relevant passages; (ii) composition of evidence clues across sentences helps claim verification; and that (iii) the joint training scheme provides significant benefits of a pipeline architecture.

2 Related Work

Most work focuses on the dataset construction while lacking advanced models to handle the problem. Vlachos and Riedel (2014) propose and define the “fact checking” problem, without a concrete solution. Ferreira and Vlachos (2016) release the dataset “Emergent” for rumor debunking. Each claim is accompanied by an article headline as evidence. Then a three-way logistic regression model is used over some rule-based features. No need to search for evidence. Wang (2017) release a larger dataset for fake news detection, and propose a hybrid neural network to integrate the statement and the speaker’s meta data to do classification. However, the presentation of evidences is ignored. Kobayashi et al. (2017) release a similar dataset to (Thorne et al., 2018), but they do not consider the evaluation of evidence reasoning.

Some work mainly pays attention to determining whether the claim is true or false, assuming evidence facts are provided or neglecting presenting evidence totally, e.g., (Angeli and Manning, 2014) – given a database of true facts as premises, predicting whether an unseen fact is true and should belong to the database by natural logic inference. Open-domain question answering (QA) against a text corpus (Yin et al., 2016; Chen et al., 2017; Wang et al., 2018) can also be treated as claim verification problem, if we treat (question, correct answer) as a claim. However, little work has studied how well a QA system can identify all the answer evidence.

Only a few works considered improving the evidence presentation in claim verification problems.

Roth et al. (2009) introduce the task of Entailed Relation Recognition – given a set of short paragraphs and a relational fact in the triple form of (argument₁, relation, argument₂), finding the paragraphs that can entail this fact. They first use Expanded Lexical Retrieval to rank and keep the top- k paragraphs as candidates, then build a TE classifier over each (candidate, statement) pair. The work directly related to us is by Thorne et al. (2018). Given claims and a set of Wikipages, Thorne et al. (2018) use a retrieval model based on TF-IDF to locate top-5 sentences in top-5 pages as evidence, then utilize a neural entailment model to classify (evidence, claim) pairs.

In contrast, our work tries to optimize the claim verification as well as the evidence identification in a joint training scheme, which is more than just supporting or refuting the claims.

3 The TwoWINGOS Model

Figure 2 illustrates the two-wing optimization problem addressed in this work: given a collection of evidence candidates $S = \{s_1, s_2, \dots, s_i, \dots, s_m\}$, a claim x and a decision set $Y = \{y_1, \dots, y_n\}$, the model TwoWINGOS predicts a binary vector p over S and a one-hot vector o over Y against the ground truth, a binary vector q and a one-hot vector z , respectively. A binary vector over S means a subset of sentences (S_e) act as evidence, and the one-hot vector indicates a single decision (y_i) to be made towards the claim x given the evidence S_e . Next, we use two separate subsections to elaborate the process of evidence identification (i.e., optimize p to q) and the claim verification (i.e., optimize o to z).

3.1 Evidence identification

A simple approach to identifying evidence is to detect the top- k sentences that are lexically similar to the claim, as some pipeline systems (Roth et al., 2009; Thorne et al., 2018) do. However, a claim-unaware fixed k is less optimal, adding noise or missing key supporting factors, consequently limiting the performance.

In this work, we approach the evidence by modeling sentences $S = \{s_1, \dots, s_i, \dots, s_m\}$ with the claim x as context in a supervised learning scheme. For each s_i , the problem turns out to be learning a probability: how likely s_i can entail the claim conditioned on other candidates as context, as shown by the blue items in Figure 2.

To start, a piece of text t ($t \in S \cup \{x\}$) is represented as a sequence of l hidden states, forming a feature map $\mathbf{T} \in \mathbb{R}^{d \times l}$, where d is the dimensionality of hidden states. We first stack a vanilla CNN (convolution & max-pooling) (LeCun et al., 1998) over \mathbf{T} to get a representation for t . As a result, each evidence candidate s_i has a representation \mathbf{s}_i , and the claim x has a representation \mathbf{x} . To get a probability for each s_i , we need first to build its *claim-aware representation* \mathbf{r}_i .

Coarse-grained representation. We directly concatenate the representation of s_i and x , generated by the vanilla CNN, as:

$$\mathbf{r}_i = [\mathbf{s}_i, \mathbf{x}, \mathbf{s}_i \cdot \mathbf{x}^T] \quad (1)$$

This coarse-grained approach makes use of merely the sentence-level representations while neglecting more fine-grained interactions between the sentences and the claim.

Fine-grained representation. Instead of directly employing the sentence-level representations, here we explore claim-aware representations for each word in sentence s_i , then compose them as the sentence representation \mathbf{r}_i , inspired by the Attentive Convolution (Yin and Schütze, 2017).

For each word s_i^j in s_i , we first calculate its matching score towards each word x^z in x , by dot product over their hidden states. Then the representation of the claim, as the context for the word s_i^j , is formed as:

$$\mathbf{c}_i^j = \sum_z \text{softmax}(\mathbf{s}_i^j \cdot (\mathbf{x}^z)^T) \cdot \mathbf{x}^z \quad (2)$$

Now, word s_i^j has left context s_i^{j-1} , right context s_i^{j+1} in s_i , and the claim-aware context \mathbf{c}_i^j from x . A convolution encoder generates its claim-aware representation \mathbf{i}_i^j :

$$\mathbf{i}_i^j = \tanh(\mathbf{W} \cdot [\mathbf{s}_i^{j-1}, \mathbf{s}_i^j, \mathbf{s}_i^{j+1}, \mathbf{c}_i^j] + \mathbf{b}) \quad (3)$$

where parameters $\mathbf{W} \in \mathbb{R}^{d \times 4d}$, $\mathbf{b} \in \mathbb{R}^d$.

To compose those claim-aware word representations as the representation for sentence s_i , we use a max-pooling over $\{\mathbf{i}_i^j\}$ along with j , generating \mathbf{i}_i .

We use term $f_{\text{int}}(s_i, x)$ to denote this whole process, so that:

$$\mathbf{i}_i = f_{\text{int}}(s_i, x) \quad (4)$$

At this point, the fine-grained representation for evidence candidate s_i is:

$$\mathbf{r}_i = [\mathbf{s}_i, \mathbf{x}, \mathbf{s}_i \cdot \mathbf{x}^T, \mathbf{i}_i] \quad (5)$$

Loss function. With a claim-aware representation \mathbf{r}_i , the sentence s_i subsequently gets a probability, acting as the evidence, $\alpha_i \in (0, 1)$ via a non-linear sigmoid function:

$$\alpha_i = \text{sigmoid}(\mathbf{v} \cdot \mathbf{r}_i^T) \quad (6)$$

where parameter vector \mathbf{v} has the same dimensionality as \mathbf{r}_i .

In the end, all evidence candidates in S have a ground-truth binary vector q and the predicted probability vector α ; then loss l_{ev} (“ev”: evidence) is implemented as a binary cross-entropy:

$$l_{ev} = \sum_{i=1}^m -(q_i \log(\alpha_i) + (1 - q_i) \log(1 - \alpha_i)) \quad (7)$$

As the output of this evidence identification module, we binarize the probability vector α by $p_i = [\alpha_i > 0.5]$ (“ $[x]$ ” is 1 if x is true or 0 otherwise). p_i indicates s_i is evidence or not. All $\{s_i\}$ with $p_i = 1$ act as evidence set S_e .

3.2 Claim verification

As shown in Figure 2, to figure out an entailment decision y_i for the claim x , the evidence S_e possibly consists of more than one sentence. Furthermore, those evidence sentences are not necessarily in textual order nor from the same passage. So, we need a mechanism that enables each evidence or even each word inside to be aware of the content from other evidence sentences. Similar to the aforementioned approach to evidence identification, we come up with three methods, with different representation granularity, to learn a representation for (S_e, x) , i.e., the input for claim verification, shown in Figure 3.

Coarse-grained representation. In this case, we treat S_e as a whole, constructing its representation \mathbf{e} by summing up the representations of all sentences in S_e in a weighted way:

$$\mathbf{e} = \sum_{i=1}^m \alpha_i \cdot p_i \cdot \mathbf{s}_i \quad (8)$$

where α_i , from Equation 6, is the probability of s_i being the evidence.

Then the (S_e, x) pair gets a coarse-grained concatenated representation: $[\mathbf{e}, \mathbf{x}]$. It does not model the interactions within the evidence nor the interactions between the evidence and the claim. Based on our experience in evidence identification

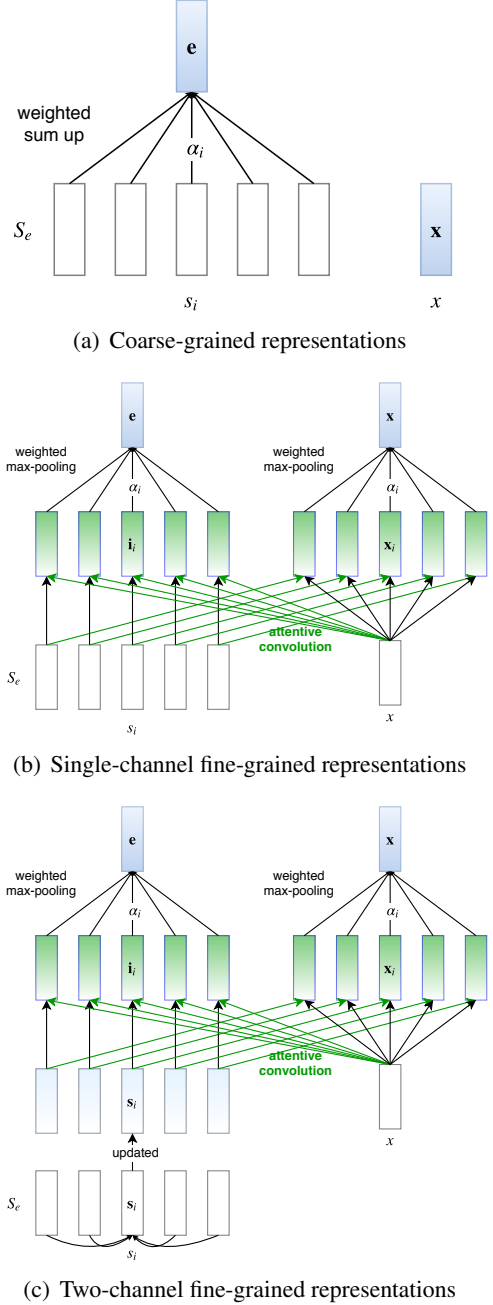


Figure 3: Three representation learning methods in claim verification. Green arrows act as context in attentive convolution.

module, the representation of a sentence is better learned by composing context-aware word-level representations. Next, we introduce how to learn fine-grained representation for the (S_e, x) pair.

Single-channel fine-grained representation. By “single-channel,” we mean each sentence s_i is aware of the claim x as its single context.

For a single pair (s_i, x) , we utilize the function $f_{\text{int}}()$ in Equation 4 to build the fine-grained representations for both s_i and x , obtaining $\mathbf{i}_i =$

$f_{\text{int}}(s_i, x)$ for s_i and $\mathbf{x}_i = f_{\text{int}}(x, s_i)$ for x .

For (S_e, x) , we compose all the $\{\mathbf{i}_i\}$ and all the $\{\mathbf{x}_i\}$ along with i , via a weighted max-pooling:

$$\mathbf{e} = \text{maxpool}_i(\alpha_i \cdot p_i \cdot \mathbf{i}_i) \quad (9)$$

$$\mathbf{x} = \text{maxpool}_i(\alpha_i \cdot p_i \cdot \mathbf{x}_i) \quad (10)$$

This weighted max-pooling ensures that the sentences with higher probabilities of being evidence have a higher chance to present their features. As a result, (S_e, x) gets a concatenated representation: $[\mathbf{e}, \mathbf{x}]$

Two-channel fine-grained representation. By “two-channel,” we mean that each evidence s_i is aware of two kinds of context, one from the claim x , the other from the remaining evidences.

Our first step is to accumulate evidence clues within S_e . To start, we concatenate all sentences in S_e as a fake long sentence \hat{S} consisting of hidden states $\{\hat{\mathbf{s}}\}$. Similar to Equation 2, for each word s_i^j in sentence s_i , we *accumulate all of its related clues* (\mathbf{c}_i^j) from \hat{S} as follows:

$$\mathbf{c}_i^j = \sum_z \text{softmax}(\mathbf{s}_i^j \cdot (\hat{\mathbf{s}}^z)^T) \cdot \hat{\mathbf{s}}^z \quad (11)$$

Then we update \mathbf{s}_i^j , the representation of word s_i^j , by element-wise addition:

$$\mathbf{s}_i^j = \mathbf{s}_i^j \oplus \mathbf{c}_i^j \quad (12)$$

This step enables the word s_i^j to “see” all related clues from S_e . The reason we add \mathbf{s}_i^j and \mathbf{c}_i^j is motivated by a simple experience: Assume the claim “Lily lives in the biggest city in Canada”, and one sentence contains a clue “... Lily lives in Toronto ...” and another sentence contains a clue “... Toronto is Canada’s largest city...”. The most simple yet effective approach to aggregating the two clues is to sum up their representation vectors (Blacoe and Lapata, 2012) (we do not concatenate them, as those clues have no consistent textual order across different s_i^j).

After updating the representation of each word in s_i , we perform the aforementioned “single-channel fine-grained representation” between the updated s_i and the claim x , generating $[\mathbf{e}, \mathbf{x}]$.

Loss function. For the claim verification input (S_e, x) , we forward its representation $[\mathbf{e}, \mathbf{x}]$ to a

	#SUPPORTED	#REFUTED	#NEI
train	80,035	29,775	35,639
dev	3,333	3,333	3,333
test	3,333	3,333	3,333

Table 1: Statistics of claims in FEVER dataset

logistic regression layer in order to infer a probability distribution o over the label space Y :

$$\mathbf{o} = \text{softmax}(\mathbf{W} \cdot [\mathbf{e}, \mathbf{x}] + \mathbf{b}) \quad (13)$$

where $\mathbf{W} \in \mathbb{R}^{n \times 2d}$, $\mathbf{b} \in \mathbb{R}^n$

The loss l_{cv} (“cv”: claim verification) is implemented as negative log-likelihood:

$$l_{\text{cv}} = -\log(\mathbf{o} \cdot \mathbf{z}^T) \quad (14)$$

where \mathbf{z} is the ground truth one-hot label vector for the claim x on the space Y .

3.3 Joint optimization

Given the loss l_{ev} in evidence identification and the loss l_{cv} in claim verification, the overall training loss is represented by:

$$l = l_{\text{ev}} + l_{\text{cv}} \quad (15)$$

To ensure that we jointly train the two coupled subtasks with intensive knowledge communication instead of simply putting two pipeline neural networks together, our TOWINGOS has following configurations:

- Both subsystems share the same set of word embeddings as parameters; the vanilla CNNs for learning sentence and claim representations share parameters as well.
- The output binary vector p by the evidence identification module is forwarded to the module of claim verification, as shown in Equations 8-10.
- Though the representation of a claim’s decision y_i is not put explicitly into the module of evidence identification, the claim’s representation \mathbf{x} will be fine-tuned by the y_i , so that the evidence candidates can get adjustment from the decision y_i , since the claims are shared by two modules.

4 Experiments

4.1 Setup

Dataset. In this work, we use FEVER (Thorne et al., 2018). The claims in FEVER were generated from the introductory parts of about 50K

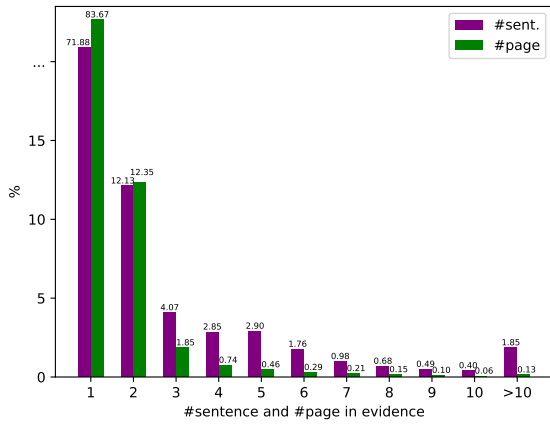


Figure 4: Distribution of #sentence and #pages in FEVER evidence

Wikipedia pages of a June 2017 dump. Annotators construct claims about a single fact of the title entity with arbitrarily complex expressions and entity forms. To increase the claim complexity so that claims would not be trivially verified, annotators adopt two routes: (i) Providing additional knowledge: Annotators can explore a dictionary of terms that were (hyper-)linked, along with their pages; (ii) Mutate claims in six ways: negation, paraphrasing, substitution of a relation/entity with a similar/dissimilar one, and making the claims more general/specific. All resulting claims have 9.4 tokens in average. Apart from claims, FEVER also provides a Wikipedia corpus in size of about 5.4 million.

Each claim is labeled as SUPPORTED, REFUTED or NOTENOUGHINFO (NEI). In addition, evidence sentences, from any wiki page, are required to be provided for SUPPORTED and REFUTED. Table 1 lists the data statistics. Figure 4 shows the distributions of sentence sizes and page sizes in FEVER’s evidence set. We can see that roughly 28% of the evidence covers more than one sentence, and approximately 16.3% of the evidence covers more than one wiki page.

This task has three evaluations: (i) NOSCOREEV – accuracy of claim verification, neglecting the validity of evidence; (ii) SCOREEV – accuracy of claim verification with a requirement that the predicted evidence fully covers the gold evidence for SUPPORTED and REFUTED; (iii) F_1 – between the predicted evidence sentences and the ones chosen by annotators. We use the officially released evaluation scorer³.

³<https://github.com/sheffieldnlp/fever-scorer>

Wiki page retrieval⁴. For each claim, we search in the given dictionary of wiki pages in the form of {title: sentence_list}, and keep the top-5 ranked pages for fair comparison with Thorne et al. (2018). Algorithm 1 briefly shows the steps of wiki page retrieval. To speed up, we first build an inverted index from words to titles, then for each claim, we only search in the titles that cover at least one claim word.

Input: A claim, $wiki=\{title: page_vocab\}$

Output: A ranked top- k wiki titles

Generate entity_mentions from the claim;

while each title do

if $claim.vocab \cap title.vocab$ is empty **then**

 | discard this title

else

 title_score = the max recall value of title.vocab
 in claim and in entity_mentions of the claim;

if title_score = 1.0 **then**

 | title.score = title_score

else

 page_score = recall of claim in

 page.vocab;

 title.score = title_score + page_score

end

end

end

Sort titles by title.score in descending order

Algorithm 1: Algorithm description of wiki page retrieval for FEVER claims.

All sentences of the top-5 retrieved wiki pages are kept as evidence candidates for claims in train, dev and test. It is worth mentioning that this page retrieval step is a reasonable preprocessing which controls the complexity of evidence searching in real-world, such as the big space – 5.4 million – in this work.

Training setup. All words are initialized by 300D Word2Vec (Mikolov et al., 2013) embeddings, and are fine-tuned during training. The whole system is trained by AdaGrad (Duchi et al., 2011). Other hyperparameter values include: learning rate 0.02, hidden size 300, mini-batch size 50, filter width 3.

Baselines. In this work, we first consider the two systems reported by Thorne et al. (2018): (i) MLP: A multi-layer perceptron with one hidden layer, based on TF-IDF cosine similarity between the claim and the evidence (all evidence sentences are concatenated as a longer text piece) (Riedel et al., 2017); (ii) Decomp-Att (Parikh et al., 2016): A decomposable attention model that develops atten-

⁴Our retrieval results are released as well.

k	(Thorne et al., 2018)		ours	
	rate	acc_ceiling	rate	acc_ceiling
1	25.31	50.21	76.58	84.38
5	55.30	70.20	89.63	93.08
10	65.86	77.24	91.19	94.12
25	75.92	83.95	92.81	95.20
50	82.49	90.13	93.36	95.57
100	86.59	91.06	94.19	96.12

Table 2: Wikipedia retrieval evaluation on dev. “rate”: claim proportion, e.g., $x\%$, if its gold passages are fully retrieved (for “SUPPORT” and “REFUTE” only); “acc_ceiling”: $\frac{x\% \cdot (\#S + \#R) + \#N}{\#S + \#R + \#N}$, the upper bound of accuracy for three classes if the coverage $x\%$ satisfies.

tion mechanisms to decompose the problem into subproblems to solve in parallel. Note that both systems first employed an IR system to keep top-5 relevant sentences from the retrieved top-5 wiki pages as *static evidence* for claims.

We further consider the following variants of our own system TWOWINGOS:

- *Coarse-coarse*: Both evidence identification and claim verification adopt coarse-grained representations.

To further study our system, we test this “coarse-coarse” in three setups: (i) “pipeline” – train the two modules independently. Forward the predicted evidence to do entailment for claims; (ii) “diff-CNN” – joint training with separate CNN parameters to learn sentence/claim representations; (iii) “share-CNN” – joint training with shared CNN parameters.

The following variants are in joint training.

- *Fine&sentence-wise*: Given the evidence with multiple sentences, a natural baseline is to do entailment reasoning for each (sentence, claim), then compose. We do entailment reasoning between each predicted evidence sentence and the claim, generating a probability distribution over the label space Y . Then we sum up all the distribution vectors element-wise, as an ensemble system, to predict the label;

- Four combinations of different grained representation learning: “*coarse&fine(single)*”, “*coarse&fine(two)*”, “*fine&coarse*” and “*fine&fine(two)*”. “Single” and “two” refer to the single/two-channel cases respectively.

4.2 Results

Performance of passage retrieval. Table 2 compares our wikipedia retriever with the one in

(Thorne et al., 2018), which used a document retriever⁵ from DrQA (Chen et al., 2017).

Our document retrieval module surpasses the competitor by a big margin in terms of the coverage of gold passages: 89.63% vs. 55.30% ($k = 5$ in all experiments). Its powerfulness should be attributed to: (i) Entity mention detection in the claims. (ii) As wiki titles are entities, we have a bi-channel way to match the claim with the wiki page: one with the title, the other with the page body, as shown in Algorithm 1.

Performance on FEVER Table 3 lists the performances of baselines and the TWOWINGOS variants on FEVER (dev&test). From the *dev* block, we observe that:

- TWOWINGOS (from “share-CNN”) surpasses prior systems in big margins. Overall, fine-grained schemes in each subtask contribute more than the coarse-grained counterparts;

- In the three setups – “pipeline”, “diff-CNN” and “share-CNN” – of *coarse-coarse*, “pipeline” gets better scores than (Thorne et al., 2018) in terms of evidence identification. “Share-CNN” has comparable F_1 as “diff-CNN” while gaining a lot on NOSCOREEV (72.32 vs. 39.22) and SCOREEV (50.12 vs. 21.04). This clearly shows that the claim verification gains much knowledge transferred from the evidence identification module. Both “diff-CNN” and “share-CNN” perform better than “pipeline” (except for the slight inferiority at SCOREEV: 21.04 vs. 22.26).

- Two-channel fine-grained representations show more effective than the single-channel counterpart in claim verification (NOSCOREEV: 78.77 vs. 75.65, SCOREEV: 53.64 vs. 52.65). As we expected, evidence sentences should *collaborate* in inferring the truth value of the claims. Two-channel setup enables an evidence candidate aware of other candidates as well as the claim.

- In the last three rows of *dev*, there is no clear difference among their evidence identification scores. Recall that “sent-wise” is essentially an ensemble system over each (sentence, claim) entailment result. “Coarse-grained”, instead, first sums up all sentence representation, then performs $(\sum(sentence), claim)$ reasoning. We can also treat this “sum up” as an ensemble. Their comparison shows that *these two kinds of tricks do not*

⁵It compares passages and claims as TF-IDF weighted bag-of-bigrams.

system		claim verification		evidence identification			
		NO SCOREEV	SCOREEV	recall	precision	F_1	
dev	MLP	41.86	19.04	44.22	10.44	16.89	
	Decomp-Att	52.09	32.57	44.22	10.44	16.89	
	TWO WINGOS	coarse&coarse					
		pipeline	35.72	22.26	53.75	29.42	33.80
		diff-CNN	39.22	21.04	46.88	43.01	44.86
		share-CNN	72.32	50.12	45.55	40.77	43.03
		coarse&fine(single)	75.65	52.65	45.81	42.53	44.11
		coarse&fine(two)	78.77	53.64	45.78	39.23	42.25
		fine&sent-wise	71.02	53.43	52.70	48.31	50.40
		fine&coarse	71.48	53.17	52.75	47.30	49.87
fine&fine(two)	78.90	56.16	53.81	47.73	50.59		
test	(Thorne et al., 2018)	50.91	31.87	45.89	10.79	17.47	
	TWO WINGOS	75.99	54.33	49.91	44.68	47.15	

Table 3: Performance on *dev* and *test* of FEVER. TWO WINGOS outperforms prior systems if vanilla CNN parameters are shared by evidence identification and claim verification subsystems. It gains more if fine-grained representations are adopted in both subtasks.

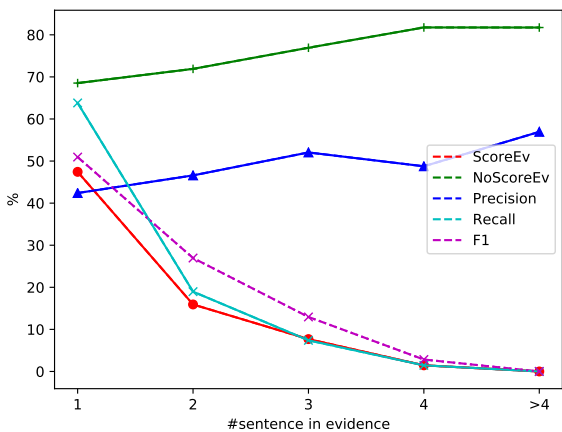


Figure 5: Performance vs. #sentence in evidence. Our system has robust precisions. The overall performance NO SCOREEV is not influenced by the decreasing recall; this verifies the fact that the truth value of most claims can be determined by a single identified evidence sentence.

make much difference.

If we adopt “two-channel fine-grained representation” in claim verification, big improvements are observed in both NO SCOREEV (+7.42%) and SCOREEV (+3%).

In the *test* block, our system (fine&fine(two)) beats the prior top system across all measurements by big margins – F_1 : 47.15 vs. 17.47; SCOREEV: 54.33 vs. 31.87; NO SCOREEV: 75.99 vs. 50.91.

In both *dev* and *test* blocks, we can observe that our evidence identification module consistently

obtains balanced recall and precision. In contrast, the pipeline system by Thorne et al. (2018) has much higher recall than precision (45.89 vs. 10.79). It is worth mentioning that the SCOREEV metric is highly influenced by the recall value, since SCOREEV is computed on the claim instances whose evidences are fully retrieved, regardless of the precision. So, ideally, a system can set all sentences as evidence, so that SCOREEV can be promoted to be equal to NO SCOREEV. Our system is more reliable in this perspective.

Performance vs. #sent. in evidence. Figure 5 shows the results of the five evaluation measures against different sizes of gold evidence sentences in test set. We observe that: (i) Our system has robust precisions across #sentence; however, the recall decreases. This is not that surprising, since the more ground-truth sentences in evidence, the harder it is to retrieve all of them; (ii) Due to the decrease in recall, the SCOREEV also gets influenced for bigger #sentence. Interestingly, high precision and worse recall in evidence with more sentences still make consistently strong overall performance, i.e., NO SCOREEV. This should be due to the fact that the majority (83.18% (Thorne et al., 2018)) of claims can be correctly entailed by a single ground truth sentence, even if any remaining ground truth sentences are unavailable.

Error analysis. The case #1 in Table 4 shows that our system identifies two pieces of evidence

#	G/P	claim	gold evidence	predicted evidence
1	0/1	Telemundo is an English-language television network.	(Telemundo, 0) (Telemundo, 1) (Telemundo, 4) (Telemundo, 5) (Hispanic_and_Latino_Americans, 0)	(Telemundo, 0) (Telemundo, 4) (Fourth_television_network, 0) (Fourth_television_network, 4)
2	1/2	Home for the Holidays stars a famous American actor.	(Anne_Bancroft, 0) (Charles_Durning, 0) (Holly_Hunter, 0) (Home_for_the_Holidays_(1995_film), 5)	\emptyset
3	0/2	Both hosts of Weekly Idol were born in 1983.	(Weekly_Idol, 0) (Weekly_Idol, 1) (Defconn, 0)	(Weekly_Idol, 1)

Table 4: Error cases of TWOWINGOS in FEVER. “G/P”: gold/predicted label (“0”: refute; “1”: support; “2”: not enough information). To save space, we use “(title, i)” to denote the i^{th} sentence in the corresponding wiki page.

(i.e., (Telemundo, 0) and (Telemundo, 4)) correctly; however, it falsely predicts the claim label. (Telemundo, 0): *Telemundo is an American Spanish-language terrestrial television . . .*. We can easily find that the keyword “Spanish-language” should refute the claim. However, both “Spanish-language” in this evidence and the “English-language” in the claim are unknown tokens with randomly initialized embeddings. This hints that a more careful data preprocessing may be helpful. In addition, to refute the claim, another clue comes from the combination of (Telemundo, 4) and (Hispanic_and_Latino_Americans, 0). (Telemundo, 4): *“The channel . . . aimed at Hispanic and Latino American audiences”*; (Hispanic_and_Latino_Americans, 0): *“Hispanic Americans and Latino Americans . . . are descendants of people from countries of Latin America and Spain.”*. Our system only retrieved (Telemundo, 4). And this clue is hard to grasp as it requires some background knowledge – people from Latin America and Spain usually are not treated as English-speaking.

In the case #2, our system fails to identify any evidence. This is due to the failure of our passage retrieval module: it detects entity mentions “Home”, “Holidays” and “American”, and the top-5 retrieved passages are “Home”, “Home_for_the_Holidays”, “American_Home”, “American” and “Home_for_the_Holidays_(song)”, which unfortunately cover none of the four ground truth passages. Interestingly, (i) given the falsely retrieved passages, our system predicts “no sentence is valid evidence” (denoted as \emptyset in Table 4); (ii) given the empty evidence, our system predicts “NoEnoughInfo” for this claim. Both make sense.

In the case #3, a successful classification of the

claim requires information aggregation over the three gold evidence sentences: (Weekly_Idol, 0): *“Weekly Idol is a South Korean variety show . . .”*; (Weekly_Idol, 1): *“The show is hosted by comedian Jeong Hyeong-don and rapper Defconn.”*; (Defconn, 0): *“Defconn (born Yoo Dae-joon; January 6, 1977) is a . . .”*. To successfully retrieve the three sentences as a whole set of evidence is challenging in evidence identification. Additionally, this example relies on the recognition and matching of digital numbers (1983 vs. 1977), which is beyond the expressivity of word embeddings, and is expected to be handled by rules more easily.

5 Summary

In this work, we build TWOWINGOS, a two-wing optimization framework to address the claim verification problem by presenting precise evidence. Differing from a pipeline system, TWOWINGOS ensures the evidence identification module and the claim verification module are trained jointly, in an end-to-end scheme. Experiments show the superiority of TWOWINGOS in the FEVER benchmark.

Acknowledgments

We thank group colleagues (Nitish Gupta and Jennifer Sheffield) and Dr. Mo Yu from IBM AI Foundations Lab for providing insightful comments and critiques. This work was supported by Contract HR0011-15-2-0025 with the US Defense Advanced Research Projects Agency (DARPA). Approved for Public Release, Distribution Unlimited. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

- Gabor Angeli and Christopher D. Manning. 2014. Natralli: Natural logic inference for common sense reasoning. In *Proceedings of EMNLP*, pages 534–545.
- William Blacoe and Mirella Lapata. 2012. A comparison of vector-based representations for semantic composition. In *Proceedings of EMNLP-CoNLL*, pages 546–556.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of SIGMOD*, pages 1247–1250.
- Arun Tejasvi Chaganty, Ashwin Paranjape, Percy Liang, and Christopher D. Manning. 2017. Importance sampling for unbiased on-demand evaluation of knowledge base population. In *Proceedings of EMNLP*, pages 1038–1048.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of ACL*, pages 1870–1879.
- Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzoto. 2013. Recognizing textual entailment: Models and applications.
- John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*, 12:2121–2159.
- William Ferreira and Andreas Vlachos. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of NAACL*, pages 1163–1168.
- Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard H. Hovy. 2013. Learning whom to trust with MACE. In *Proceedings of NAACL*, pages 1120–1130.
- Mio Kobayashi, Ai Ishii, Chikara Hoshino, Hiroshi Miyashita, and Takuya Matsuzaki. 2017. Automated historical fact-checking by passage retrieval, word statistics, and virtual question-answering. In *Proceedings of IJCNLP*, pages 967–975.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of NIPS*, pages 3111–3119.
- Ankur P. Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proceedings of EMNLP*, pages 2249–2255.
- Jeff Pasternack and Dan Roth. 2013. Latent credibility analysis. In *Proceedings of WWW*, pages 1009–1020.
- Benjamin Riedel, Isabelle Augenstein, Georgios P. Spithourakis, and Sebastian Riedel. 2017. A simple but tough-to-beat baseline for the fake news challenge stance detection task. *CoRR*, abs/1707.03264.
- Dan Roth, Mark Sammons, and V. G. Vinod Vydiswaran. 2009. A framework for entailed relation recognition. In *Proceedings of ACL*, pages 57–60.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of WWW*, pages 697–706.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and verification. In *Proceedings of NAACL*.
- Andreas Vlachos and Sebastian Riedel. 2014. Fact checking: Task definition and dataset construction. In *Proceedings of the Workshop on Language Technologies and Computational Social Science@ACL*, pages 18–22.
- V. G. Vinod Vydiswaran, ChengXiang Zhai, and Dan Roth. 2011. Content-driven trust propagation framework. In *Proceedings of SIGKDD*, pages 974–982.
- Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerry Tesauro, Bowen Zhou, and Jing Jiang. 2018. R³: Reinforced ranker-reader for open-domain question answering. In *Proceedings of AAAI*.
- William Yang Wang. 2017. “Liar, Liar Pants on Fire”: A new benchmark dataset for fake news detection. In *Proceedings of ACL*, pages 422–426.
- Wenpeng Yin, Sebastian Ebert, and Hinrich Schütze. 2016. Attention-based convolutional neural network for machine comprehension. In *Proceedings of the NAACL Workshop on Human-Computer Question Answering*, pages 15–21.
- Wenpeng Yin and Hinrich Schütze. 2017. Attentive convolution. *CoRR*, abs/1710.00519.