

Open-Domain, Content-based, Multi-modal Fact-checking of Out-of-Context Images via Online Resources

Sahar Abdelnabi, Rakibul Hasan, and Mario Fritz
 CISPA Helmholtz Center for Information Security
 {sahar.abdelnabi, rakibul.hasan, fritz}@cispa.de

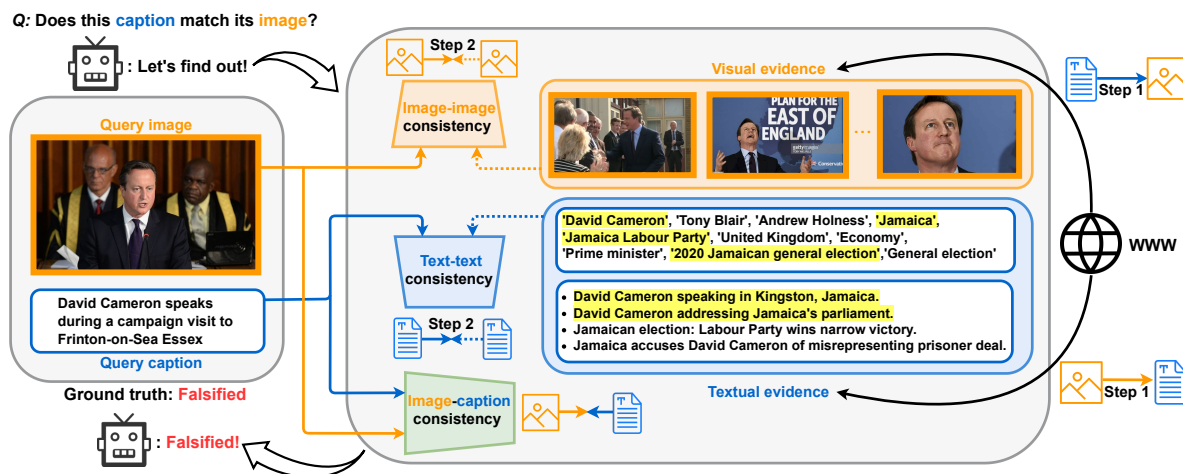



Figure 1. To evaluate the veracity of **image-caption** pairings, we leverage **visual** and **textual** evidence gathered by querying the Web. We propose a novel framework to detect the consistency of the claim-evidence (**text-text** and **image-image**), in addition to the **image-caption** pairing. Highlighted evidence represents the model’s highest attention, showing a difference in location compared to the query **caption**.

Abstract

Misinformation is now a major problem due to its potential high risks to our core democratic and societal values and orders. **Out-of-context** misinformation is one of the easiest and effective ways used by adversaries to spread viral false stories. In this threat, a real image is **re-purposed** to support other narratives by misrepresenting its context and/or elements. The internet is being used as the go-to way to verify information using different sources and modalities. Our goal is an inspectable method that automates this time-consuming and reasoning-intensive process by fact-checking the **image-caption** pairing using Web evidence. To integrate evidence and cues from both modalities, we introduce the concept of ‘**multi-modal cycle-consistency check**’ ; starting from the **image/caption**, we gather **textual/visual** evidence, which will be compared against the other paired **caption/image**, respectively. Moreover, we propose a novel architecture, **Consistency-Checking Net-**

work (CCN), that mimics the layered human reasoning across the same and different modalities: the **caption** vs. **textual evidence**, the **image** vs. **visual evidence**, and the **image** vs. **caption**. Our work offers the first step and benchmark for **open-domain, content-based, multi-modal fact-checking**, and significantly outperforms previous baselines that did not leverage external evidence¹.

1. Introduction

Recently, there has been a growing and widespread concern about ‘fake news’ and its harmful societal, personal, and political consequences [1, 18, 24], including people’s own health during the pandemic [6, 7, 30]. Misusing generative AI technologies to create deepfakes [13, 21, 36] further fuelled these concerns [5, 14]. However, **image-repurposing**— where a real image is misrepresented and used **out-of-context** with another false or unrelated narrative



¹For code, checkpoints, and dataset, check: <https://s-abdelnabi.github.io/OoC-multi-modal-fc/>

to create more credible stories and mislead the audience—is still one of the easiest and most effective ways to create realistically-looking misinformation. Image-repurposing does not require profound technical knowledge or experience [2, 27], which potentially amplifies its risks. Images usually accompany real news [41]; thus, adversaries may augment their stories with images as ‘supporting evidence’ to capture readers’ attention [15, 27, 47].

Image re-purposing datasets and threats. Gathering large-scale labelled out-of-context datasets is hard due to the scarcity and substantial manual efforts. Thus, previous work attempted to construct synthetic out-of-context datasets [20, 39]. A recent work [27] proposed to automatically, yet non-trivially, match images accompanying real news with other real news captions. The authors used trained language and vision models to retrieve a close and convincing image given a caption. While this work contributes to misinformation detection research by automatically creating datasets, it also highlights the threat that *machine-assisted* procedures may ease creating misinformation at scale. Furthermore, the authors reported that both defense models and humans struggled to detect the out-of-context images. In this paper, we use this dataset as a challenging benchmark; we leverage external evidence to push forward the automatic detection.

Fact-checking. To fight misinformation, huge fact-checking efforts are done by different organizations [33, 34]. However, they require substantial manual efforts [43]. Researchers have proposed several automated methods and benchmarks to automate fact-checking and verification [32, 42]. However, most of these works focus on textual claims. Fact-checking multi-modal claims has been under-explored.

Our approach. People frequently use the Internet to verify information. We aggregate evidence from images, articles, different sources, and we measure their consensus and consistency. Our goal is to design an inspectable framework that automates this multi-modal fact-checking process and assists users, fact-checkers, and content moderators.

More specifically, we propose to gather and reason over evidence to judge the veracity of the **image-caption** pair. First , we use the **image** to find its other occurrences on the internet, from which, we crawl **textual evidence** (e.g., captions), which we compare against the paired **caption**. Similarly , we use the **caption** to find other images as **visual evidence** to compare against the paired **image**. We call this process: ‘**multi-modal cycle-consistency check**’. Importantly, we retrieve evidence in a fully automated and flexible **open-domain** manner [8]; no ‘golden evidence’ is pre-identified or curated and given to the model.

To evaluate the claim’s veracity, we propose a novel architecture, the **Consistency-Checking Network (CCN)**, that consists of 1) memory networks components to evaluate the consistency of the claim against the evidence (de-

scribed above), 2) a CLIP [35] component to evaluate the consistency of the **image** and **caption** pair themselves. As the task requires machine comprehension and visual understanding, we perform different evaluations to design the memory components and the evidence representations. Moreover, we conduct two user studies to 1) measure the human performance on the detection task and, 2) understand if the collected evidence and the model’s attention over the evidence help people distinguish true from falsified pairs. **Figure 1** depicts our framework, showing a falsified example from the dataset along with the retrieved evidence.

Contributions. We summarize our contributions as follows: 1) we formalize a new task of multi-modal fact-checking. 2) We propose the ‘**multi-modal cycle-consistency check**’ to gather evidence about the multi-modal claim from both modalities. 3) We propose a new inspectable framework, *CCN*, to mimic the aggregation of observations from the claims and world knowledge. 4) We perform numerous evaluations, ablations, and user studies and show that our evidence-augmented method significantly improves the detection over baselines.

2. Related Work

Multi-modal Misinformation. Previous work has studied multi-modal misinformation [22, 29, 46]. For instance, Khattar et al. [22] studied multi-modal fake news on Twitter by learning representations of images and captions which were used in classification. The images in the dataset could be edited. In contrast, we focus on out-of-context real news images and verifying them using evidence.

Moreover, Zlatkova et al. [50] studied the factuality of the image-claim pairs using information from the Web. They collected features about the claim image, such as its URL. The actual content of the claim image is not considered against evidence. Our work is different in how we collect both **visual** and **textual** evidence to perform the **cycle-consistency check**. In addition, they only calculate features from the claim text such as TF-IDF, while we use memory networks with learned representations.

Related to the out-of-context threat, Aneja et al. [2] constructed a large, yet unlabelled, dataset of different contexts of the same image. They propose a self-supervised approach to detect whether two captions (given an image) are having the same context. However, unlike our work, they do not judge the veracity of a single image-caption claim. Also, the unlabelled dataset collected in this work does not allow the veracity detection training and evaluation.

In order to produce labelled out-of-context images, previous work created synthetic datasets by changing the captions, either by naive swapping or named entities manipulations [20, 39], however, the falsified examples were either too naive or contained linguistic biases that are easy to detect even by language-only models [27].


Therefore, Luo et al. [27] proposed to create falsified examples by matching real images with real captions [26]. They created the large-scale NewsCLIPPings dataset that contains both *pristine* and convincing *falsified* examples. The matching was done automatically using trained language and vision models (such as SBERT-WK [45], CLIP [35], or scene embeddings [49]). The falsified examples could misrepresent the context, the place, or people in the image, with inconsistent entities or semantic context. The authors show that both machine and human detection are limited, indicating that the task is indeed challenging. Thus, to improve the detection, we propose to use external Web evidence to verify the **image-caption** claim.

Open-domain QA and Fact-verification. Our work is similar to textual work in open-domain QA [8] and fact-verification [42] (from Wikipedia) in having a large-scale and open-domain task that involves automatic retrieval and comprehension. We do not assume that the input to the model is already labelled and identified as relevant, simulating real-life fact-checking. Moreover, we do not restrict the evidence to be from a specific curated source only, such as fact-checking websites, in contrast to [44].

Similar to our work, Popat et al. [32] built a credibility assessment end-to-end method of textual claims using external evidence. However, to the best of our knowledge, no previous work attempted to verify multi-modal claims using both modalities. Also, their model is designed to predict the per-source credibility of claims, while we learn the aggregated consistency from multiple sources.


3. Dataset and Evidence Collection

Dataset. We use the NewsCLIPPings [27] that contains both pristine and falsified (‘out-of-context’) images. It is built on the VisualNews [26] corpus that contains news pieces from 4 news outlets: The Guardian, BBC, USA Today, and The Washington Post. The NewsCLIPPings dataset contains different subsets depending on the method used to match the images with captions (e.g., text-text similarity, image-image similarity, etc.). We use the ‘balanced’ subset that has representatives of all matching methods and consists of 71,072 train, 7,024 validation, and 7,264 test examples. To kick-start our *evidence-assisted detection*, we use the **image-caption** pairs as queries to perform Web search, as depicted in Figure 1.

Textual evidence. We use the query **image** in an inverse search mode using Google Vision APIs [4] to retrieve **textual evidence** . The API returns a list of **entities** that are associated with that image, which we collect as part of the textual evidence. They might describe the content of the image and, further, the contexts of where these images appeared, such as the entities’ list in Figure 1.

In addition, the API returns the images’ URLs and the containing pages’ URLs. In contrast to previous work [50]

that only considered the containing pages’ titles, we also collect the images’ captions. We designed a Web crawler that visits the page, searches for the image’s tag using its URL or by image content matching (using perceptual hashing), then retrieves the **captions** if found. We scrape the `<figcaption>` tag, as well as the `` tag’s textual attributes such as *alt*, *image-alt*, *caption*, *data-caption*, and *title*. In addition, we observed the returned pages for a few hundreds of the API calls and implemented other strategies to scrape the captions based on them. We also save the **titles** of the pages. From each page, we collect all the non-redundant text snippets that we found. The API returns up to 20 search results. We discard a page if the detected language of the title is not English, using the fastText library [12] for language identification. We collect the **domains** of each evidence item as metadata.

Visual evidence. Second, we use the **caption** as textual queries to search for **images** . We use the Google custom search API [10] to perform the image search. We retrieve up to 10 results, while also saving their **domains**. It is important to note that, unlike the inverse image search, the search results here are not always corresponding to the exact match of the textual query. Therefore, the **visual evidence** might be more loosely related to the query **image**. However, even if it is not exactly related to the event, it works as a useful baseline of the type of images that could be associated with that topic.

Dataset decomposition. We summarize the dataset components and task as follows:

Dataset. Unless no search results were found, a single example in the dataset consists of the following:

- A query **image** I^q .
- A query **caption** C^q .
- **Visual evidence:**
 - A list of **images**: $I^e = [I_1^e, \dots, I_K^e]$.
- **Textual evidence:**
 - A list of **entities**: $ENT = [E_1, \dots, E_M]$.
 - A list of **captions/sentences**: $S = [S_1, \dots, S_N]$.

Task. Classify $\{I^q, C^q\}$ to: *Pristine* or *Falsified*.

4. The Consistency-Checking Network

We introduce the task of evidence-assisted veracity assessment of **image-caption** pairing. As shown in Figure 1, we perform the ‘**multi-modal cycle-consistency check**’ by comparing the **textual evidence** against the query **caption**, and the **visual evidence** against the query **image**.

Challenges. The task is significantly more complex than the merely one-to-one matching of the query against the evidence. First, many search results may be unrelated to the query (neither falsify nor support) and act as noise. Second, comparing the query against the evidence requires further




Figure 2. Overview of our Consistency-Checking Network, *CCN*.

comprehension and reasoning. For pristine examples, the **textual evidence** might range from being paraphrases of the query **caption** to distantly related but supporting. For falsified examples, they might range from having different named entities to having the same ones but in a different context, such as the example in **Figure 1**. Similarly, comparing the **visual evidence** against the query **image** requires visual and scene understanding or regions comparison.

We propose a novel architecture, the Consistency-Checking Network (*CCN*), to meet these challenges. We show an overview of the method in **Figure 2**. At the core of our approach is the memory networks architecture [9, 23, 28, 40], which selectively compares the claim to the relevant items of the possibly large list of evidence. In addition, the attention mechanism allows inspecting which evidence items were most relevant to the decision. The model consists of a **visual reasoning** component, a **textual reasoning** component, and a ‘CLIP’ component.

4.1. Visual Reasoning

Figure 3 outlines the visual reasoning component that inspects the consistency between the query **image** and the **visual evidence**. First, we represent the images using ResNet152 [16], pretrained on the ImageNet dataset. Each image is represented as: $I^q/I^e \in \mathbb{R}^{2048}$, where q denotes the query representation and e denotes evidence. Moreover, to reason over the overlap of regions and objects in the query image vs. evidence images, we used the label detection Google API [3] to get a list of labels for each image. Then, for each evidence image, we compute the number of *overlapping labels* between it and the query. We use this




Figure 3. Visual evidence reasoning component.

number as an additional feature, and we concatenate it with the evidence images’ representations.

The memory holds the evidence images. Each input to the memory is embedded into input and output memory representations [40], denoted by a and c , respectively. The image memory vectors $m_i \in \mathbb{R}^{1024}$ are represented by:

$$m_i^a = \text{ReLU}(W_i^a I^e + b_i^a), \quad (1)$$

$$m_i^c = \text{ReLU}(W_i^c I^e + b_i^c) \quad (2)$$

The learned parameters are W_i^a and $W_i^c \in \mathbb{R}^{2048 \times 1024}$, and b_i^a and $b_i^c \in \mathbb{R}^{1024}$. The query image I^q is also projected into a 1024-dimension vector (\hat{I}^q) by another linear layer for modelling convenience. The matching between \hat{I}^q and the memory vectors m_i^a are then computed by:

$$p_{i,j} = \text{Softmax}(\hat{I}^q T m_{i,j}^a), \quad (3)$$

where i denotes the image memory, j is a counter for the memory items, and p_i is a probability vector over the items. The output of the memory is the sum of the query and the average of the output representations m_i^c , weighted by p_i :

$$o_i = \sum_j p_{i,j} m_{i,j}^c + \hat{I}^q \quad (4)$$

In addition, for some mismatched examples, there could be context discrepancies based on the place. To make the model aware of scenes and places similarity, we also represent the images using a ResNet50 trained on the Places365 dataset [49]. We form a separate memory for the scene representations to allow more flexibility. Similar to the previous formulation, each image is represented as: $P^q/P^e \in \mathbb{R}^{2048}$, and the scenes memory vectors $m_p \in \mathbb{R}^{1024}$ are represented by:

$$m_p^{a/c} = \text{ReLU}(W_p^{a/c} P^e + b_p^{a/c}) \quad (5)$$

Similar to Eqn. 3 and Eqn. 4, we get the output of the scenes (places) memory o_p .

4.2. Textual Reasoning

The second component of our model evaluates the consistency between the query **caption** and the **textual evidence**. As shown in **Figure 1**, we have two types of textual evidence: sentences (captions or pages’ titles), and entities. As they have different granularities and might differ in importance, we form a separate memory for each.

As shown in **Figure 4**, we represent the query caption and each evidence item using a sentence embedding model. We experiment with state-of-the-art inference models that were trained on large corpuses such as Wikipedia and were shown to implicitly store world knowledge [25, 31, 38], making them suitable for our task. We evaluate two methods in our experiments: 1) a pre-trained sentence transformer model [37] that is trained for sentence similarity, 2)

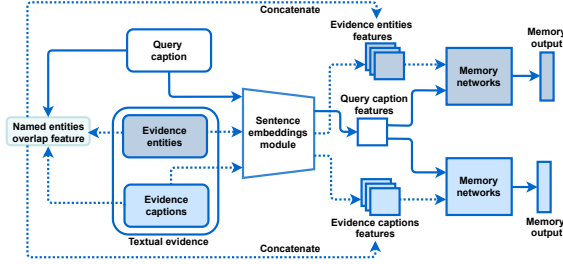


Figure 4. Textual evidence reasoning component.

using BERT [11] to get strong contextualized embeddings, in addition to an LSTM to encode the sequence. In the second method, we use the second-to-last BERT layer [48] as the tokens’ embeddings. We concatenate the last time-step LSTM output and the average of all time-steps’ outputs.

In addition, to help the model be entity-aware, we utilize a binary indicator feature to denote if there is a named entity overlap between the query caption and the evidence item. We used the spaCy NER [17] to extract the entities and concatenated the binary feature with the evidence (both captions and entities) representations.

Using either of these previously mentioned methods, we get embeddings for the query caption C^q , the evidence entities E , and the evidence captions/sentences S . The entities input and output memory representations are given by:

$$m_e^{a/c} = \text{ReLU}(W_e^{a/c} E + b_e^{a/c}), \quad (6)$$

similarly, the captions/sentences input and output memory representations are given by:

$$m_s^{a/c} = \text{ReLU}(W_s^{a/c} S + b_s^{a/c}), \quad (7)$$

where $W_e^{a/c}, W_s^{a/c} \in \mathbb{R}^{d \times d}$ and $b_e^{a/c}, b_s^{a/c} \in \mathbb{R}^d$ are trainable weights, and d is the dimension of the sentence embedding model (768 in the case of the pre-trained model, and 512 in the case of using BERT+LSTM).

As per Eqn. 3 and Eqn. 4, we compute the output of the entities and sentences memories as o_e and o_s , respectively.

Encoding the evidence’s domain. Features of websites, e.g., how frequently they appear and the types of news they feature, could help to prioritize evidence items. Thus, we learn an embedding of the evidence’s domain names. We represent the domains as one-hot vectors and project them into a 20-dimensional space. We consider the domains that appeared at least three times, resulting in 17148 unique domains, the rest are set to *UNK*. The domain embeddings are then concatenated with the evidence representations (both visual and textual, excluding entities).

4.3. CLIP

In addition to reasoning over evidence, we leverage CLIP [35], used in [27], to integrate the query **image-text** consistency into the decision. We first fine-tune CLIP

ViT/B-32 on the task of classifying image-caption pairs into pristine or falsified, without considering the evidence.

During fine-tuning, we pass the image and text through the CLIP encoders and normalize their embeddings. We produce a joint embedding that is a dot product of the image and text ones, and we add a linear classifier on top. The model is trained to classify the pair into pristine or falsified. Then, we freeze the fine-tuned CLIP and integrate the joint CLIP embeddings (J_{clip}) into the final classifier of *CCN*.

4.4. Classifier

Now that we individually evaluated the **text-text**, **image-image**, and **image-text** consistency, we aggregate these observations in order to reach a unified decision.

We found it helpful during training to apply a batch normalization layer [19] to the output of each component. We then concatenate all previous components in one feature vector o_t as follows:

$$o_t = \text{BN}(o_i) \oplus \text{BN}(o_p) \oplus \text{BN}(o_e) \oplus \text{BN}(o_s) \oplus \text{BN}(J_{\text{clip}}), \quad (8)$$

where BN denotes the batch normalization. o_t is then fed to a simple classifier that has two fully connected layers with ReLU and batch normalization after the first one (dimension: 1024), and Sigmoid after the second one that outputs a final falsified probability (p_f). The model is trained, with freezing the backbone embedding networks, to binary classify the examples using the binary cross-entropy loss:

$$L = -y_{\text{true}} \log(p_f) - (1 - y_{\text{true}}) \log(1 - p_f) \quad (9)$$

More implementation details can be found in Supp. 1.

5. Experimental Results

In this section, we show the quantitative analysis of different variants of the model and baselines. We then present our user studies, qualitative analysis, and discussion.

5.1. Quantitive Analysis

We evaluated our model and other variants of it in order to understand the effect of each component. **Table 1** shows our experiments. We summarize different aspects and highlight the most interesting observations in what follows.

Evidence types. We first show the effect of each evidence type in the first four rows. Removing the evidence **images** or the evidence **captions** dropped the performance significantly; these results indicate the importance of integrating both modalities for verification. Removing the **Entities** had less effect. This might be due to having some redundant information with the evidence captions already, or because of sometimes having generic named entities that are not helpful to verify the caption claim.

Memory design. Adding a batch normalization layer after each component, as in Eqn. 8, improved the training and

increased the accuracy by nearly 11 percentage points. Another variant we studied had a unified memory containing images, captions, and entities. The query here was a concatenation of the image and caption pairs. As shown in row 6, this was less successful than the separate memory setup, suggesting that the explicit **text-text** and **image-image** consistency comparison aids the learning.

Evidence filtering. As the dataset is constructed from real news articles, the Google search may return the exact news as the query search (i.e., exact news with the exact webpage). While this is needed in a real fact-checking setup, it might bias the training; the model might use it/or its absence as a shortcut to predict pristine/falsified pairs, respectively, without stronger reasoning. Therefore, we filtered the evidence as follows: for pristine examples, we discard an evidence item if it *matches* the query and comes from the *same website* as the query. To detect matching, we use perceptual hashing for **images**. For **captions**, we remove punctuations and lower-case all the sentences and then check if they are an exact match. We then trained and evaluated with this filtered dataset. As shown in row 7, this did not significantly reduce the accuracy, suggesting that the model reasons about consistency beyond exact matches.

Other improvements. We show that our other enhancements, including adding CLIP and improving visual and textual representations, recovered the performance drop due to the evidence filtering. CLIP had relatively the largest effect, with around a 1.5 percentage points increase. Training the LSTM with BERT embeddings performed better than using a pre-trained sentence transformer model. This might be because it allowed the model to learn on the token level and focus on the consistency in, e.g., named entities, location, etc., which are more specific cues in our use-case than general sentence entailment tasks. Finally, the last row

#	Evidence type	Separate mem.	BN	Dataset filter	CLIP	ResNet (ImageNet)	ResNet (Scenes)	Labels	Sent. transformer	BERT+LSTM	NER	Accuracy
1	all	✓	✗	✗	✗	✓	✗	✗	✓	✗	✗	73.5%
2	all w/o Images	✓	✗	✗	✗	✓	-	-	✓	✗	✗	62.5%
3	all w/o Captions	✓	✗	✗	✗	✓	✗	✗	✓	✗	✗	57.4%
4	all w/o Entities	✓	✗	✗	✗	✓	✗	✗	✓	✗	✗	71.8%
5	all	✓	✓	✗	✗	✓	✗	✗	✓	✗	✗	84.2%
6	all	✗	✓	✗	✗	✓	✗	✗	✓	✗	✗	81.7%
7	all	✓	✓	✓	✗	✓	✗	✗	✓	✗	✗	80.3%
8	all	✓	✓	✓	✗	✓	✗	✗	✓	✗	✓	81.2%
9	all	✓	✓	✓	✓	✓	✗	✗	✓	✗	✓	82.6%
10	all	✓	✓	✓	✓	✓	✓	✗	✓	✗	✓	83.4%
11	all	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	83.9%
12	all	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	84.7%
13	all w/o domains	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	83.9%

Table 1. Classification performance on the test set for different variants of the model. Highlighted cells represent the changed factor in that experiment. The green box represents the best model.

Method	Evidence	Pair	All	Falsified	Pristine
CLIP	✗	✓	66.1%	68.1%	64.2%
Averaged	✓	✗	70.6%	72.4%	68.9%
CCN	✓	✓	84.7%	84.8%	84.5%

Table 2. Classification performance on the test set for our model in comparison with baselines.

shows that including the evidence’s domain helps to some extent, as it might help the model to attend to and prioritize evidence items. Additional experiments are in Supp. 2.

Baselines. We compare our evidence-assisted detection against the CLIP-only baseline used in [27] in Table 2. We fine-tuned CLIP [35], reaching a higher accuracy than originally reported in [27] on this dataset subset. As the dataset pairing is not trivial, this baseline achieved a relatively low performance. In contrast, we achieve a significant improvement of a nearly 19 percentage points increase, indicating that leveraging evidence is important to solve the task.

As there are no previous baselines for evidence-assisted out-of-context detection, we design a baseline that uses evidence. We use the pretrained image and text representations of ResNet-152 and sentence transformer in the same setup of **text-text** and **image-image** similarity. We compute the matching between the query and the evidence via dot product. Then, we use an average pooling layer across all evidence items, which will be used for classification. As shown in Table 2, this baseline outperforms the CLIP-only. However, our proposed model with the other improvements achieves a ~14 percentage points increase.

5.2. User Studies

We conducted user studies to estimate the human performance on the dataset and evaluate the usefulness of the evidence in detection, as well as the relevancy of the evidence items that the model highly attends to.

5.2.1 Study 1: Human Performance Baseline

We aim to establish a human baseline as an upper bound estimate of the out-of-context images detection accuracy. Due to the automatic open-world evidence retrieval, we do not have a labelled dataset to indicate if an evidence item is relevant to the claim. Furthermore, some examples might not have any relevant evidence retrieved. Also, the falsified examples could be very close to the original context, making them hard to verify even with the presence of evidence.

Setup. We randomly selected 100 examples (48 pristine, 52 falsified) from the test dataset. Along with the **image-caption** pairs, we presented the gathered evidence (**images**, **captions**, and **entities**). For each pair, first, we asked users if the **caption** matches the **image**, considering any of: inconsistency cues between them, the evidence presented, or their prior knowledge about the subject. Then, they answered which source(s) of information helped them label

	Study	All	Falsified	Pristine
Average	1 st	81.0%±4.71	79.5%±8.31	82.3%±9.31
	2 nd , Highest	86.2%±4.9	84.5%±9.3	88.0%±7.2
	2 nd , Lowest	77.7%±6.0	76.0%±9.0	79.5%±7.5
Best worker	1 st	89.0%	92.0%	93.7%
	2 nd , Highest	94.0%	98.0%	98.0%
	2 nd , Lowest	88.0%	90.0%	86.0%

Table 3. Our two user studies. The first is to label random 100 examples. The second is to label another 100 examples using 1) the highest-attention, and 2) the lowest-attention evidence.

the pair, or indicated ‘None’ if it was hard to verify. We instructed them *not* to search for other evidence, so that both our model and humans have access to the same evidence, and to evaluate the usefulness of the evidence gathered by our framework. We recruited 8 experienced native English-speaking crowd workers through Amazon Mechanical Turk.

Results. Table 3 shows the average performance across all workers and the results of the best worker. Compared to the findings reported in [27], human performance significantly increased when presented with evidence (average detection was 65.6%, with only 35% falsified detection rate). Additionally, *CCN* achieved 80% accuracy on these 100 examples, which is lower than the best worker but on a par with the average worker.

Figure 5 shows which information helped workers to label the **image-caption** pairs during the study. We highlight the following observations: 1) In 77.2% of the examples, on average, the evidence contributed to the workers’ decision, in comparison with 59.3% only for the **image-caption** pair. In 28.3%, the evidence was the only helpful cue. 2) Among the evidence types, the **images** were the most helpful (64%), possibly because it is easier to grasp different images at a glance. 3) 12.3% of the examples were hard to verify. When checking some of them (Supp. 3), we observed that they do not have obvious cues (e.g., generic scenes with event-specific captions, an image for the same person with a similar context). Also, they sometimes had poor retrieval (the inverse search did not find the **image**, so there are no evidence **captions**, and the evidence **images** are unrelated or not conclusive). Our model struggled in detecting these examples as well. Augmenting with looser retrieval (e.g.,




Figure 5. Workers indicated the factors that helped their decision. ‘Any evid.’ means that any evidence type was helpful. ‘Evid. only’ means that only the evidence was helpful.

searching with keywords of the caption, finding captions of other similar images) might help in these cases.

5.2.2 Study 2: Evaluating the Attention

One of our main goals is to have an automated fact-checking tool while also allowing humans to be in the loop, if needed. We hypothesize that the attention weights given by the model can be used to retrieve the most relevant and useful evidence, which enables a quick inspection.

We design a second study to evaluate this hypothesis. We randomly selected 100 examples (50 each) that at least have 8 evidence items in each type². We designed two variants using the same 100 pairs; in the first, we display the highest-attention 4 items from each evidence type, in the second, we display the lowest-attention 4 ones. The two variants are labelled by non-overlapping groups (8 workers each). We follow the rest of the first study’s setup and instructions.

Results. Table 3 and Figure 5 show that the highest-attention evidence had higher performance and generally better ratings as ‘helpful’ compared to the lowest-attention evidence. These findings suggest that the model learned to prioritize the most relevant items, as intended, and can potentially be beneficial for 1) inspectability and, 2) assistive fact-checking; as workers had a higher performance with only a subset of evidence.

5.3. Qualitative Analysis

We show some successful predictions of our model in Figure 6. When inspecting the attention in the case of *pristine* examples, we found that the highest attention is on items that are most relevant to the query (e.g., a similar **image** in the first example, named **entities** that are present in or similar to the query **caption** such as cities’ names, and semantically similar **captions**). The model also predicted the second example correctly, despite not having an **image** of the same scene. For *falsified* examples, we observe that the third one is predicted correctly despite having a similar falsified topic (‘Affordable health care’ and ‘Lawsuits’). Moreover, the fourth one shows the highest attention on contradicting locations in **entities**, and on the most syntactically similar **caption**. This was predicted correctly, despite having similar-style evidence to the query. Similarly, the falsified example in Figure 1 was similar in the persons’ names and images (‘David Cameron’), but different in context and scene details. Finally, the last example shows a *pristine* example that was misclassified as *falsified*. When inspecting the **textual evidence**, we observed that although it is revolving around the same topic, there is little connection to the context of the query **caption**, in addition to having a diverse set of **visual evidence** that is not similar to the query **image**. Other examples are in Supp. 4.

²In this first study, some examples might not have enough evidence. However, we keep them to have a representative set of the dataset.











Image-caption pair	Textual evidence	Visual evidence
 <p>The Futenma marine corps airbase on the southern Japanese island of Okinawa</p>	<p>'United States', 'Ginowan', 'Governor', 'Military base', 'Politics', 'Japan', 'Takeshi Onaga', 'Governor of Okinawa Prefecture', 'Hirokazu Nakaima', 'Shinzo Abe', 'Okinawa', 'airport'</p> <p>1- Hercules aircraft parked on the tarmac at Marine Corps Air Station Futenma in Ginowan on Okinawa. 2- Japan Decides to Stop Works on US Airbase Relocation in Okinawa. 3- Japan Decides to Restart Relocation of US Base in Okinawa Despite Protests.</p>	 <p>Prediction: Pristine</p>
 <p>The soaring number of Syrian refugees has sparked increasing resentment in Lebanon</p>	<p>'Syria', 'Lebanon', 'United Kingdom', 'Tent', 'Syrians', 'Language', 'Refugee', 'Recreation', 'Tourism', 'Camping', 'Language barrier', 'rural area'</p> <p>1- Syrian refugees at a camp in eastern Lebanon, December 2014. 2- Syrians entering Lebanon face new restrictions 3- Among those displaced, 1.6 million children have fled Syria. 4- Syrian refugees in the UK: 'We will be good people. We will build this country'</p>	 <p>Prediction: Pristine</p>
 <p>Healthcare activists say the ruling against Novartis ensures poor people will be able to access cheap versions of cancer medicines</p>	<p>'United States Capitol', 'Affordable Care Act', 'Supreme Court of the United States', 'Presidency of Donald Trump', 'President of the United States', 'United States', 'us capitol grounds'</p> <p>1- Demonstrators from Doctors for America in support of Obamacare march in front of the Supreme Court on March 4, 2015. 2-The Affordable Care Act Is Back In Court, 5 Facts You Need To Know. 3- As Court Hears Arguments in Lawsuit To Eliminate Obamacare, Conn. Senators Plead Their Case.</p>	 <p>Prediction: Falsified</p>
 <p>Smoke rises following an Israeli air strike in Gaza City</p>	<p>'Kobane', 'Kurdistan Region', 'United States', 'Peshmerga', 'Turkey', 'Kurds', 'Syria', 'Iraq', 'kobani war'</p> <p>1- Smoke rises after a U.S.-led airstrike in the Syrian town of Kobani 2- The border town of Kobani is under threat after the Islamists drove 180,000 Kurds into Turkey. 3-Former Kurdish Sniper Claims To Have Killed Around 250 ISIS Fighters.</p>	 <p>Prediction: Falsified</p>
 <p>How can our young readers persuade their parents to get them a Playstation 3</p>	<p>'Grand Theft Auto V', 'Gamer', 'Grand Theft Auto IV', 'Wii', 'Grand Theft Auto VI', 'PlayStation 3', 'Rockstar Leeds', 'terry seeborne marshall', 'Gordon Hall', 'Rockstar Games'</p> <p>1- A court order banning Sony from importing PS3s into the Netherlands has been lifted. 2- Rockstar Games, creators of the Grand Theft Auto franchise, said it was "very saddened" to hear of Mr Hall's death 3- Oakland Athletics to Begin Accepting Bitcoin for Private Suites</p>	 <p>Prediction: Falsified</p>

Figure 6. Qualitative examples of news pairs along with the collected evidence. Examples with green background are pristine, red background are falsified. Highlighted items are the ones with the highest attention. Only a subset of the evidence is shown for display purposes.

5.4. Discussion and Limitations

We propose a multi-modal fact-checking framework that significantly outperforms baselines and is comparable to human performance. However, the task has yet many challenges, and fully relying on automated tools might have dangerous consequences. Therefore, humans should still be in the loop. Thus, we offer an inspectable and assistive tool that helps to reduce the load of the otherwise fully manual process [43]. We further discuss potential risks in Supp. 5.

Moreover, our approach relies on the retrieval results of the search engine. However, as we show in our analysis (Tables 1 and 2), naively considering the evidence is not adequate, and a careful design of the model is needed to meet the challenges of the task, including the noisy open-domain setup with no relevancy supervision, and the high resemblance of evidence across pristine and falsified examples.

Finally, in some situations, some evidence items might contradict others, e.g., due to the websites' opposing politi-

cal orientations, or misinformation on the Web. We did not observe such scenarios with the used dataset; identifying and studying them might require poisoning the search results, or carefully curating claims that lead to contradicting results, which is beyond the scope of this work.

6. Conclusion

We mimic the complex fact-checking process in an automated framework, *CCN*, that aggregates consistency signals and consensus from multi-modal evidence found on the Web, and the given *image-caption* pairing. Our work significantly outperforms previous baselines and offers a new task and benchmark of multi-modal fact-checking, and an automated, inspectable tool to assist manual fact-checking.

Acknowledgment

This work was supported by the Google Cloud Research Credits program. We also thank Rebecca Weil for helpful advice and feedback.

References

- [1] Shruti Agarwal, Hany Farid, Yuming Gu, Mingming He, Koki Nagano, and Hao Li. Protecting world leaders against deep fakes. In *CVPR workshops*, 2019. 1
- [2] Shivangi Aneja, Christoph Bregler, and Matthias Nießner. Catching out-of-context misinformation with self-supervised learning. *arXiv*, 2021. 2
- [3] Google Vision API. Detect labels. Available at: <https://cloud.google.com/vision/docs/labels>. 4
- [4] Google Vision API. Detect web entities and pages. Available at: <https://cloud.google.com/vision/docs/detecting-web>. 3
- [5] BBC. Deepfakes: A threat to democracy or just a bit of fun?, <https://www.bbc.com/news/business-51204954>. 1
- [6] BBC. Coronavirus: The human cost of virus misinformation, <https://www.bbc.com/news/stories-52731624>. 1
- [7] BBC. Youtube to remove all anti-vaccine misinformation, <https://www.bbc.com/news/technology-58743252>. 1
- [8] Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. Reading wikipedia to answer open-domain questions. In *ACL*, 2017. 2, 3
- [9] Cesc Chunseong Park, Byeongchang Kim, and Gunhee Kim. Attend to you: Personalized image captioning with context sequence memory networks. In *CVPR*, 2017. 4
- [10] Google Developers. Programmable search engine. Available at: <https://developers.google.com/custom-search/v1/overview>. 3
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. 5
- [12] Facebook. fasttext. Available at: <https://fasttext.cc/>. 3
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *NeurIPS*, 2014. 1
- [14] The Guardian. The rise of the deepfake and the threat to democracy, <https://www.theguardian.com/technology/ng-interactive/2019/jun/22/the-rise-of-the-deepfake-and-the-threat-to-democracy>. 1
- [15] Michael Hameleers, Thomas E Powell, Toni GLA Van Der Meer, and Lieke Bos. A picture paints a thousand lies? the effects and mechanisms of multimodal disinformation and rebuttals disseminated via social media. *Political Communication*, 37(2):281–301, 2020. 2
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [17] Matthew Honnibal and Ines Montani. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. 2017. 5
- [18] Minyoung Huh, Andrew Liu, Andrew Owens, and Alexei A Efros. Fighting fake news: Image splice detection via learned self-consistency. In *ECCV*, 2018. 1
- [19] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 5
- [20] Ayush Jaiswal, Ekraam Sabir, Wael AbdAlmageed, and Premkumar Natarajan. Multimedia semantic integrity assessment using joint embedding of images and text. In *ACM MM*, 2017. 2
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 1
- [22] Dhruv Khattar, Jaipal Singh Goud, Manish Gupta, and Vasudeva Varma. Mvae: Multimodal variational autoencoder for fake news detection. In *WWW*, 2019. 2
- [23] Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. Ask me anything: Dynamic memory networks for natural language processing. In *ICML*, 2016. 4
- [24] David MJ Lazer, Matthew A Baum, Yoichai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, et al. The science of fake news. *Science*, 359(6380):1094–1096, 2018. 1
- [25] Nayeon Lee, Belinda Z Li, Sinong Wang, Wen-tau Yih, Hao Ma, and Madian Khabsa. Language models as fact checkers? In *the Third Workshop on Fact Extraction and VERification (FEVER)*, 2020. 4
- [26] Fuxiao Liu, Yinghan Wang, Tianlu Wang, and Vicente Ordonez. Visual news: Benchmark and challenges in news image captioning. In *EMNLP*, 2021. 3
- [27] Grace Luo, Trevor Darrell, and Anna Rohrbach. NewsCLIP-pings: Automatic Generation of Out-of-Context Multimodal Media. In *EMNLP*, 2021. 2, 3, 5, 6, 7
- [28] Mitra Mohtarami, Ramy Baly, James Glass, Preslav Nakov, Lluís Màrquez, and Alessandro Moschitti. Automatic stance detection using end-to-end memory networks. In *NAACL-HLT*, 2018. 4
- [29] Kai Nakamura, Sharon Levy, and William Yang Wang. Fakeddit: A new multimodal benchmark dataset for fine-grained fake news detection. In *the Language Resources and Evaluation Conference*, 2020. 2
- [30] Sophie Nightingale and Hany Farid. Examining the global spread of covid-19 misinformation. *arXiv*, 2020. 1
- [31] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. Language models as knowledge bases? In *EMNLP-IJCNLP*, 2019. 4
- [32] Kashyap Papat, Subhabrata Mukherjee, Andrew Yates, and Gerhard Weikum. Declare: Debunking fake news and false claims using evidence-aware deep learning. In *EMNLP*, 2018. 2, 3
- [33] Poynter. Politifact, <https://www.politifact.com/>. 2
- [34] Poynter. The international fact-checking network, <https://www.poynter.org/ifcn/>. 2
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 3, 5, 6

- [36] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. 1
- [37] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP*, 2019. 4
- [38] Adam Roberts, Colin Raffel, and Noam Shazeer. How much knowledge can you pack into the parameters of a language model? In *EMNLP*, 2020. 4
- [39] Ekraam Sabir, Wael AbdAlmageed, Yue Wu, and Prem Natarajan. Deep multimodal image-repurposing detection. In *ACM MM*, 2018. 2
- [40] Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. End-to-end memory networks. *NeurIPS*, 2015. 4
- [41] Reuben Tan, Bryan Plummer, and Kate Saenko. Detecting cross-modal inconsistency to defend against neural fake news. In *EMNLP*, 2020. 2
- [42] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. Fever: a large-scale dataset for fact extraction and verification. In *NAACL-HLT*, 2018. 2, 3
- [43] Vice. Facebook is literally hiring people to just google stuff, <https://www.vice.com/en/article/g5xp5j/facebook-is-literally-hiring-people-to-just-google-stuff>. 2, 8
- [44] Nguyen Vo and Kyumin Lee. Where are the facts? searching for fact-checked information to alleviate the spread of fake news. In *EMNLP*, 2020. 3
- [45] Bin Wang and C-C Jay Kuo. Sbert-wk: A sentence embedding method by dissecting bert-based word models. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2020. 3
- [46] Yaqing Wang, Fenglong Ma, Zhiwei Jin, Ye Yuan, Guangxu Xun, Kishlay Jha, Lu Su, and Jing Gao. Eann: Event adversarial neural networks for multi-modal fake news detection. In *KDD*, 2018. 2
- [47] Yuping Wang, Fatemeh Tahmasbi, Jeremy Blackburn, Barry Bradlyn, Emiliano De Cristofaro, David Magerman, Savvas Zannettou, and Gianluca Stringhini. Understanding the use of fauxtography on social media. In *International AAAI Conference on Web and Social Media*, 2021. 2
- [48] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *CVPR*, 2019. 5
- [49] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE TPAMI*. 3, 4
- [50] Dimitrina Zlatkova, Preslav Nakov, and Ivan Koychev. Fact-checking meets fauxtography: Verifying claims about images. In *EMNLP-IJCNLP*, 2019. 2, 3