

Explainable Automated Fact-Checking: A Survey

Neema Kotonya and Francesca Toni
Department of Computing
Imperial College London, United Kingdom
{nk2418, ft}@ic.ac.uk

Abstract

A number of exciting advances have been made in automated fact-checking thanks to increasingly larger datasets and more powerful systems, leading to improvements in the complexity of claims which can be accurately fact-checked. However, despite these advances, there are still desirable functionalities missing from the fact-checking pipeline. In this survey, we focus on the *explanation* functionality – that is fact-checking systems providing *reasons for their predictions*. We summarize existing methods for explaining the predictions of fact-checking systems and we explore trends in this topic. Further, we consider what makes for good explanations in this specific domain through a comparative analysis of existing fact-checking explanations against some desirable properties. Finally, we propose further research directions for generating fact-checking explanations, and describe how these may lead to improvements in the research area.

1 Introduction

The harms of false claims, which, through their verisimilitude, masquerade as statements of fact, have become widespread as a result of the self-publishing popularized on the modern Web 2.0, without verification. It is now widely accepted that if we are to efficiently respond to this problem, computational approaches are required (Fallis, 2015). Falsehoods take many forms, including audio and graphical content, however a significant quantity of misinformation exists in natural language. For this reason, computational linguistics has a significant role to play in addressing this problem.

Steady progress has been made in the area of computational journalism, and more specifically automated fact-checking and its orthogonal tasks. A number of fact-checking tasks have been formulated. Among these tasks are the following: identifying the severity or nature of false claims, i.e., distinguishing between hoaxes, rumors and satire (Rubin et al., 2016; Rashkin et al., 2017); detecting check-worthy claims (Hassan et al., 2015; Hassan et al., 2017); distinguishing between hyper-partisan and mainstream news (Potthast et al., 2018); stance detection and opinion analysis in fake news (Hanselowski et al., 2018a); fauxtography, i.e., verifying the captions or claims which accompany published images (Zlatkova et al., 2019); table-based fact verification (Chen et al., 2020); credibility assessment (Mitra and Gilbert, 2015; Derczynski et al., 2017); and an end-to-end pipeline which encompasses both the document-level evidence retrieval and sentence-level evidence selection sub-tasks of claim verification (Thorne et al., 2018). Established orthogonal tasks such as claim detection (Levy et al., 2014; Lippi and Torroni, 2015) and deception detection (Conroy et al., 2015) also play a key role in automated fact-checking. The fact-checking pipeline is shown in Figure 1.

The focus of this survey is *explainable* automated fact-checking. Indeed, over the last few years strong progress has been made both in terms of the performance of deep learning models for fact-checking and also in terms of the availability of comprehensive datasets for training these models. However, not nearly as much work has been devoted to acquiring explanations for these systems, either in the form of post-hoc explanations for the outputs of fact-checking models or by hard-wiring explanation methods into

these fact-checking models, even though the justification of claim verification judgments is arguably a very important part of journalistic processes when performing manual fact-checking.




Figure 1: Pipeline for automated fact-checking with the explanation sub-task.

This survey has the following structure. We begin (in Section 2) by laying out the current state-of-the-art in automated fact-checking. We discuss the rapid ground which has been covered in recent years with regard to both dataset construction and development of increasingly more accurate predictive models for fact-checking. We then explore (in Section 3) explanations for the fact-checking domain. We discuss human explanations offered by trained journalists and computationally derived fact-checking explanations. Further (in Section 4), we explore limitations of the current state-of-the-art in explainable fact-checking systems. Finally (in Section 5), we offer research directions which could be explored in order to extend this topic area and possibly mitigate some limitations we identify in the literature.

2 Automated Fact-Checking: State of the Art

Engineering accurate and resilient computational methods for automated fact-checking is challenging. However, steady progress has been made, and the current state-of-the-art with respect to a number of fact-checking datasets achieves promising results.

2.1 Datasets

A major bottleneck associated with automated fact-checking is the shortage of available training data. However, recently there has been an increase in the volume of data available due to the rise of fact-checking websites (Graves and Cherubini, 2016; Lowrey, 2017), corpus annotation platforms, and generative language models (Niewinski et al., 2019). One of the first fact-checking datasets is a small corpus of 106 claims from Politifact¹ (Vlachos and Riedel, 2014). A later dataset, EMERGENT, provides 300 claims retrieved from Twitter and Snopes² (Ferreira and Vlachos, 2016).

These initial corpora for fact-checking are evidence of positive steps towards a rigorous formulation of the fact-checking task, as they introduced the task as a multi-class or graded problem (Vlachos and Riedel, 2014) and furthermore related stance detection to claim verification (Ferreira and Vlachos, 2016). However, due to their small size, none was sufficient for training machine learning models. Since then, a larger number of corpora have been constructed with numbers of claims which are orders of magnitude larger, i.e., in the thousands and tens of thousands, and therefore are suitable for training classifiers.

The LIAR dataset of fact-checked claims was the first to be released of a sufficient size for training machine learning models (Wang, 2017). LIAR consists of 12.8K claims retrieved using the Politifact API, as well as meta-data for these claims (i.e., speaker of the claim, political affiliations of the speaker, medium through which the claim was first published). However, one element not present in the dataset is the evidence used by fact-checkers to verify the claims.

Since the publication of LIAR, huge strides in progress have been made, and now there are substantially larger datasets, which also include evidence and rich with meta-data to contextualize the claims. Among these large fact-checking corpora is the manually-constructed FEVER dataset (Thorne et al., 2018) which consists of 185K claims extracted from Wikipedia. There are also corpora which consist of real-world or naturally occurring claims, including claims from multiple domains, such as MultiFC (Augenstein et al., 2019), FakeNewsNet (Shu et al., 2018), and a corpus with fine-grained evi-

¹<https://www.politifact.com/>

²<https://www.snopes.com/>

dence (Hanselowski et al., 2019). Recently, datasets have been made available which examine fact-checking outside the politic context, e.g., in the scientific (Wadden et al., 2020) and health domains (Kotonya and Toni, 2020), which includes natural language explanations (see Table 1).

Dataset	Claims	Sources	Labels	Evidence	Explanation
Vlachos and Riedel (2014)	106	Various	5	✓	
Mitra and Gilbert (2015)	1,049	Twitter	5		
Pomerleau and Rao (2017)	300	Emergent ³	4	✓	
Wang (2017)	12,836	PolitiFact	6		
Pérez-Rosas et al. (2018)	980	News, GossipCop	2		
Shu et al. (2018)	23,921	PolitiFact, GossipCop	2	✓	
Alhindi et al. (2018)	12,836	PolitiFact	6	✓	
Thorne et al. (2018)	185,445	Wikipedia	3	✓	
Augenstein et al. (2019)	36,534	Various	2-40	✓	
Hanselowski et al. (2019)	6,422	Snopes	5	✓	
Wadden et al. (2020)	1,409	S2ORC ⁴	3	✓	
Kotonya and Toni (2020)	11,832	Various	4	✓	✓

Table 1: Comparison of fact-checking corpora. Vlachos and Riedel (2014) take claims from PolitiFact and Channel 4 News, Augenstein et al. (2019) extract claims from 26 fact-checking websites, and Kotonya and Toni (2020) take claims from fact-checking, news review and news websites.

Decisions in dataset annotation inform formulations of the fact-checking task, e.g., the choice between binary (Pérez-Rosas et al., 2018) and graded classification (Vlachos and Riedel, 2014; Wang, 2017) informs the definition of veracity. Various annotation schemes are employed for constructing these datasets, and thus veracity labels vary considerably across datasets. For example, LIAR uses a six-point veracity system borrowed from Politifact; FEVER uses a three-way classification, with the label NOT-ENOUGH-INFO to account for claims for which there is not insufficient data to arrive at a true or false judgment. The FactBank dataset (Saurí and Pustejovsky, 2009) defines the factuality for claims, which are all events in the dataset, as having two dimensions: polarity (i.e., accurate, inaccurate, uncertain) and degree of certainty (i.e., certainly, probably, uncertain).

One last point to mention regarding the current available resources for fact-checking is the limitations associated with synthetic or human-crafted data for automated fact-checking. These limitations include model bias and lackluster performance on new data. Schuster et al. (2019) examine the FEVER dataset, creating a symmetric dataset to show that FEVER consists of idiosyncrasies which can be exploited in order produce high accuracy fact-checking models. In order to debias the dataset, they propose a regularization method, which down weighs words and phrases which bias models.

2.2 Fact-checking Methods

There has been steady progress in engineering accurate systems for automated fact-checking. Systems have mainly been developed in two contexts: (i) shared-task, focused on a single dataset, and (ii) stand-alone, where new datasets have been built as part of the research.

Example of shared tasks are the first and second Fact Extraction and Verification (FEVER) challenges (Thorne et al., 2018). The first iteration saw twenty-four fact-checking systems submitted, with the highest performing achieving accuracy greater than 60%. The second iteration was framed as a *Build it, Break it, Fix it* problem: as well as inviting participants to build fact-checking systems, participants were also offered the opportunity to develop adversarial claims to break the builders’ systems. Fixers were then invited make systems resilient to future adversarial attacks (Thorne et al., 2019b).

³<https://www.emergent.info>

⁴<https://allenai.org/data/s2orc>

Through close examination of system architectures for automated fact-checking – both in the case of shared-tasks and standalone automated fact-checking systems – we observe that a combination of deep neural networks (DNNs), non-DNNs and heuristic approaches are employed. The technique employed is also strongly related to the sub-task. For example, the FEVER shared-task consists of document retrieval, evidence selection (i.e., extraction of the most pertinent sentences from the retrieved documents), relation prediction (recognizing textual entailment between claim and selected evidence sentences), and finally the implied task of label aggregation. These sub-tasks may benefit from different approaches.

Typically heuristic-based approaches are used for document retrieval. Hanselowski et al. (Hanselowski et al., 2018b) devise a document retrieval module based on TFIDF, Google Search API and named entity linking. This document retrieval module is re-used in a number of works (Alonso-Reina et al., 2019; Soleimani et al., 2020; Liu et al., 2020). DNNs (e.g. Enhanced Sequential Inference Model (ESIM) (Chen et al., 2017) and decomposable attention (Parikh et al., 2016)) are used for evidence selection and in some cases for joint evidence selection and relation prediction. Similar approaches are taken in the FNC1 shared-task (Hanselowski et al., 2018a), with heuristics used for evidence selection and DNN techniques (e.g., stacked LSTMs) employed for relation prediction. The only exception is Nie et al. (2019), whose method uses a DNN approach, i.e., neural semantic matching, for the entire pipeline (i.e., for all subtasks). DNN are a popular choice for the relation prediction sub-task, which is consistent with the state of the art for this NLP task. For relation aggregation, all apart from Nie et al. (2019) make use of a rule-based approach or aggregate via a multilayer perceptron (MLP). The three fact-checking systems submitted for the Builders phase of FEVER 2.0 employ similar approaches for the document retrieval sub-task, in that they all use DNNs. However, unlike first FEVER shared-task, FEVER 2.0 models put larger emphasis on jointly learning evidence selection and relation prediction subtasks.

Aside from the solutions developed for the FEVER shared tasks, other attempts have been made to engineer end-to-end fact-checking systems. Hassan et al. (2017) develop a system to assess the veracity of claims. They source evidence from the Web and obtain knowledge bases using keyword searches. No DNN is employed in their system, instead features, e.g., part-of-speech tags and sentiment, are employed to learn a 3-way classification for claims with random forests, naive Bayes classifiers, and support vector machines. Popat et al. (2018) and Shu et al. (2019), whose methods are also some of the first examples of explainable fact-checking, both employ DNNs for evidence selection and natural language inference. A summary of the automated fact-checking systems which are comprised of both a predictive model and explanation model is given in Table 2. In Section 3.2, we discuss the task formulations of these explainable fact-checking in greater depth.

System	Formulation	Explanation methods
Popat et al. (2017) [†]	Attention	BiLSTM + CRF
Popat et al. (2018) [†]	Attention	BiLSTM + attention
Shu et al. (2019) [†]	Attention	BiLSTM + attention
Gad-Elrab et al. (2019)	Rule discovery	Knowledge graphs + Horn rules
Ahmadi et al. (2019) [†]	Rule discovery	Knowledge graphs, Horn rules + probabilistic answer set programming
Yang et al. (2019)	Attention	CNN + self-attention
Atanasova et al. (2020a) [†]	Summarization	BERT fine-tuned for extractive summarization
Lu and Li. (2020) [†]	Attention	Graph-based co-attention networks
Wu et al. (2020) [†]	Attention	Co-attention self-attention networks
Kotonya and Toni (2020)	Summarization	BERT fine-tuned for joint extractive and abstractive summarization

Table 2: Methods employed by explainable automated fact-checking systems. We identify where joint models are employed for prediction and explanation. † indicates the model performs joint veracity prediction and explanation extraction.

Observing the trends, in particular the strong reliance on DNNs for textual entailment and the use of joint training for sub-tasks in the pipeline, there is evidently a trade off between system complexity and transparency. The increase in complexity of automated fact-checking methods is an even bigger incentive for acquiring explanations for these state-of-the-art systems. The observed trends also highlight that there are multiple fact-checking sub-tasks which should be considered for explanation.

3 Explanations for fact-checking

There is increasing research interest in explanations for DNNs. These approaches include model-agnostic methods for explaining the outputs of machine learning classifiers, e.g., LIME (Ribeiro et al., 2016) and SHAP (Lundberg and Lee, 2017); work on evaluating explanations (Narayanan et al., 2018); and the formalization of desiderata (i.e., functional, operational, and user requirements) for explanations (Sokol and Flach, 2019). For the purpose of this survey, we choose to examine explanations specifically with respect to the domain of automated fact-checking. To this end we will overview existing methods for explainable fact-checking and perform a comparative analysis thereof.

Before we begin our analysis, it is important to make a distinction between model *interpretability* and model *explainability*, which we understand as follows.

Interpretability can be understood as the ability of a machine learning model to offer a mechanism by which its decision-making can be analyzed, and possibly visualized. An interpretation is not necessarily readily understood by those who do not have specific technical expertise or else that are not well-versed in the architecture of the model, e.g., as may be the case with attention weights.

Explainability can be understood as the ability of a machine learning model to deliver a rationale for its decisions, e.g., for a model which verifies claims through cross-checking against rules in a knowledge base, the explanation could be the full set of rules used to arrive at the final judgment.

In this survey, we are primarily interested in natural language, human-readable explanations of machine learning models. However, at the same time we acknowledge that model interpretations are very much a prerequisite through which human-readable explanations are achieved, but they are neither on their own sufficient nor necessarily human comprehensible.

3.1 Fact-checking explanations by journalists

In order to better understand explanations in the fact-checking domain, it is important that we have a background on the mechanisms of fact-checking in journalism and the explanations reported by journalists. There are a number of studies which outline an epistemology of fact-checking (Uscinski and Butler, 2013; Amazeen, 2015; Graves, 2017). With regard to explanations, a few points are mentioned. One is that fact-checking is typically achieved by a team of journalists who will aim to come to a consensus regarding the veracity of a claim (Graves, 2017). Graves (2017) makes the point that fact-checkers also work with a set of guidelines (e.g., Politifact uses the *Words matter*, *Context matters* and *Timing* principles to form their judgments⁵, see Table 3). These guidelines are analogous to explainability desiderata (Narayanan et al., 2018; Sokol and Flach, 2019). Graves (2018) states that fact-checking has three parts: *identification*, *verification*, and *correction*. Here *correction* includes flagging falsehoods and providing contextual data and fact checks, which are components of explanations.

Words matter	We pay close attention to the specific wording of a claim. Is it a precise statement? Does it contain mitigating words or phrases?
Context matters	We examine the claim in the full context, the comments made before and after it, the question that prompted it, and the point the person was trying to make.
Timing	Our rulings are based on when a statement was made and on the information available at that time.

Table 3: Three of the Politifact fact-checking principles.

⁵<https://www.politifact.com/article/2013/may/31/principles-politifact/>

Another important consideration is how to gather these human explanations for fact-checking. As a certain level of expertise is required in order to produce competent fact-checking explanations, it would be difficult to use crowd-sourcing techniques, as has been done for acquiring human explanations for other NLP tasks (Camburu et al., 2018). One potential source for extracting fact-checking explanations is online fact-checking platforms, as there are a considerable number of such websites (Graves and Cherubini, 2016), and it is from similar websites that some automated fact-checking datasets have been acquired (Wang, 2017; Augenstein et al., 2019). Figure 2 shows explanations provided by journalists on two fact-checking websites. Both are justifications for the veracity label attributed to the claim by the journalist, however they take two very different formats. Indeed, Figure 2(a) provides explanatory text which identifies *what’s true* and *what’s false* in the claim, whereas the explanation provided in Figure 2(b) takes the form of overall conclusion.

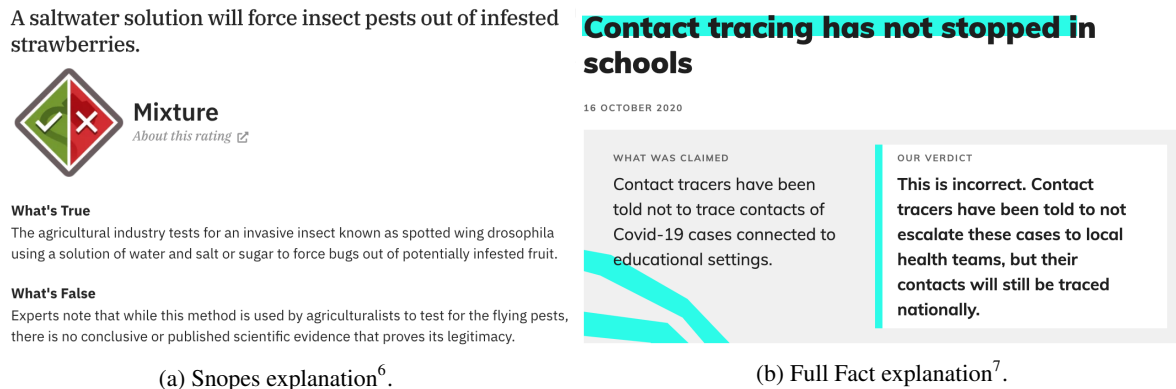


Figure 2: Explanations provided by reporters on two fact-checking websites. Figure 2(a) relates to a claim about using saltwater to remove insects from strawberries. Figure 2(b) is about a claim concerning contact tracing in UK schools.

3.2 Task Formulations

Progress has been made where accuracy of fact-checking systems is concerned, but a relatively small number of automated fact-checking systems have explainability components. Furthermore, various task formulations exist for generating explanations for automated fact-checking systems, e.g., explanations as text summaries. However, one aspect that the majority of methods we surveyed have in common is that their approaches to explainability are *extractive*, i.e., the explanations produced by these systems in some way involve extracting the components of the input which are most pertinent to the prediction (see Figure 3). All explainable systems surveyed explain individual predictions as opposed to the fact-checking model.

3.3 Attention-based Explanations

These explanations take the form of some type of visualization of neural attention weights. There is a lot of interest with regards to the relationship between attention and explanation, for example the works of Jain and Wallace (2019) and Wiegrefe and Pinter (2019) take opposing views with regards to whether attention constitutes explanation.

Popat et al. (2017), Popat et al. (2018), Shu et al. (2019), Yang et al. (2019), Lu and Li (2020) and Wu et al. (2020) all present DNN-based methods with attention mechanisms to extract explanations. An example of explanation returned by Popat et al. (2018) is given in Figure 3(a). Popat et al. (2018) use a bidirectional LSTM architecture to assess the credibility of a claim with respect to a series of articles crawled from the Web. The output of the system is a credibility score for the claim based on how well is it corroborated by the articles. The trustworthiness of the claim’s speaker and evidence articles, as assessed

⁶<https://www.snopes.com/fact-check/soaking-strawberries-in-saltwater/>

⁷<https://fullfact.org/health/schools-contact-tracing-not-escalate/>

using GloVe word embeddings, are also input features used for the final output. The explanations offered by Popat et al. (2018) are in the form of tokens highlighted in the articles, trained via attention.

Shu et al. (2019) similarly use attention weights, developing a sentence-comment co-attention sub-network to highlight the most salient excerpts from the evidence articles and link them to explanatory comments offered by readers of these articles. Unlike Popat et al. (2018), they are able to contextualize the evidence highlighted in the articles by relating them to human assessments of those articles. However, the credibility of comments by users is not verified, as it is assumed that comment authors are trustworthy. Yang et al. (2019) also use, somewhat similarly, self-attention to extract n-gram explanations and linguistic analysis to extract features, e.g., verb ratio.

Lu and Li (2020) take a slightly different approach to generating explanations. Like Shu et al. (2019) they make use of a co-attention mechanism, however the work, which looks at fact-checking tweets looks at explainability from three perspectives source tweets, retweet propagation, and retweeter characteristics i.e., suspicious users (see Figure 3(b)).

Wu et al. (2020) also perform fact-checking on Twitter data. They combine decision-trees (for evidence selection), and a co-attention mechanism composed of two hierarchical self-attention networks for explainable fake news detection. The Decision Tree-based Co-Attention model (DTCA) outputs semantically relevant words from the evidence comment tweets used for fact-checking. Furthermore, the authors provide user and comment credibility scores for evidence.

3.4 Explanation as rule discovery

Unlike the aforementioned works, the systems by Gad-Elrab et al. (2019) and Ahmadi et al. (2019) use Horn rules and knowledge graphs to mine explanations (see Figure 3(c)). A clear advantage of these systems over the attention-based explanations is that the explanations produced are more comprehensive (see Figure 3), as discussed later when we present a comparison of the explanations produced by each system in Table 4. One disadvantage that these two systems have over the attention-based DNNs presented by Shu et al. (2019) and Popat et al. (2018) is the issue of scalability. Indeed, both systems are reliant on mining rules for each of the claim triples from a knowledge base, e.g., DBpedia, which limits the statements which can be fact-checked. However, there has been recent work which looks at alleviating this problem by creating an extensible framework for rule mining (Ahmadi et al., 2020).

3.5 Explanation as summarization

Atanasova et al. (2020a) generate natural language explanations for fact-checking, formulating the explanation generation task as a text summarization problem (see Figure 3(d)). Two models are trained with this in mind: one model which generates post hoc explanations, i.e., the predictive and explanation models are trained separately; and the other is jointly trained for both tasks. The jointly trained model slightly performs more poorly than the model that trains the explainer separately. The models adapt a BERT-based architecture for which the segmentation embeddings are re-purposed for text summarization (Liu and Lapata, 2019). Similarly to other models we have discussed, this systems also takes an *extractive* view of the explanation task, i.e., it employs extractive summarization so the outputs generated as explanations of predictions are subsets of the inputs. Kotonya and Toni (2020) also model the explanation generation task as one of summarization. However, their work differs from that of Atanasova et al. (2020a) in that the explanation models of the former are fine-tuned for joint extractive and abstractive summarization in order to achieve novel explanations, rather than being purely extractive summaries.

3.6 Adversarial claims for robust fact-checking

It is worth noting the *Build it. Break it. Fix it.* theme of the FEVER 2.0 shared task (Thorne et al., 2019a), the *Break it* component of which consisted of engineering systems for generating adversarial claims to break fact-checking systems. It is clear that adversarial examples and explanations in machine learning are orthogonal (Tomsett et al., 2018). The correlation between the two has been noted in machine learning from both empirical (Thorne et al., 2019b) and theoretical (Ignatiev et al., 2019) stand points. In fact, for automated fact-checking there have even been attempts to generate task specific adversarial claims (Thorne et al., 2019a), e.g., by means of a method that uses a GPT-2 based (Radford et al., 2019) model

[False] Barbara Boxer: "Fiorina's plan
 Article Source: nytimes.com
 least of **glimmer** of **truth** while ignorin
 including this one in california democra
 and **medicare** but we found there was sk
 she has **said** **doesn't** provide much **proof**

(a) Attention mechanism (Popat et al., 2018).

FALSE : almaMater (Michael White, UT Austin)
 ← employer (Michael White, UT Austin)
 ← occupation (Michael White, UT Austin)
 ← almaMater (Michael White, Abilene
 Christian Univ.), almaMater (Michael
 White, Yale Divinity School)

(c) Rule discovery (Ahmadi et al., 2019).

kansas ku
 ks
 city
 breaking
 center
 strict

Fake news

(b) Attention mechanism & user data. (Lu and Li, 2020)

Label: Half-true **Claim:** Of the more than 1.3 million temporary mortgage m
Just: In the final full week of the U.S. Senate race, how did Rubio fare on hi
 temporary work-outs, over half have now defaulted," referring to a temporary
Explain-Extr: Over 1.3 million temporary work-outs, over half have now de
 said that more than half of those 1.3 million had defaulted." Rubio: "The temp
Explain-MT: Rubio also said that more than half of those 1.3 million had "d
 said. Of those permanent modifications, the majority survived while almost 29
 that is slightly more than half.

(d) Text summarization (Atanasova et al., 2020a)

ksdknews rt
 confirmed
 record
 irrelevant
 criminal
 ferguson

True news

Figure 3: Examples of explanations generated by various task formulations.

with input controls and original vocabulary (Niewinski et al., 2019). Furthermore, a more recent work by Atanasova et al. (2020b) extends on this theme by generating well-formed adversarial claims, which are cohesive with the veracity label, through the use of universal adversarial triggers. These are n-grams which are inserted in the input text to confuse the model into predicting a different label for the claim.

4 Comparative Analysis of Explainable Systems

In this section we offer a qualitative analysis of explainable fact-checking approaches. We evaluate the explainable fact-checking systems described in Section 3.2 against the eight desirable properties shown in Table 4. Seven of these properties are borrowed from the usability and operational requirements for machine learning explanations proposed by Sokol and Flach (2019). We also propose one new property which is informed by the journalism literature on fact-checking. This property is explanation *unbiasedness* or *impartiality*, i.e., whether the fact-checking explanation attempts to highlight or mitigate bias (e.g., hyperpartisanship (Potthast et al., 2018) in the context of automated fact-checking) in the input. For this exercise, we consider whether each property has been addressed in the design of the fact-checking systems. We do not take into consideration how effectively these properties are addressed in the systems, but rather if the systems meet the minimum required for the properties.

Actionable explanations. According to Sokol and Flach (2019), an actionable explanation a one which users can treat as a set of guidelines indicating steps towards a desired outcome. In the case of fact-checking these would most likely take the form of factual corrections, i.e., the explanation would provide a corrected version of the erroneous claim or an indication of which parts of the claim include misinformation. In addition, an actionable explanation could identify credible sources which counter the misinformation in the original claim. Explaining, correcting, and clarifying are complementary tasks, so naturally strides or progress towards one should aide in achieving the others. In his work on the limitations of automated fact-checking, Graves (2018) identifies three main stages in automated fact-checking: identification, verification and correction. Correction should be delivered "instantaneously, across different media, to audiences exposed to misinformation", and consist of the following: (1) flagging repeated falsehoods, (2) providing contextual data, (3) publishing new fact checks. None of the explainable fact-checking methods we surveyed include corrections so we do not consider the explanations actionable.

Causal explanations. For this property to be satisfied the explainable system would need to make use of a full causal model in order to infer causal relationships between inputs and outputted predictions. Also, such an explainable system would need to produce suitable explanations, e.g., counterfactual explanations. None of the explainable fact-checking systems we examined make use of a causal model and

therefore none of them can be considered to output causal explanations.

Coherent explanations. Sokol and Flach (2019) state that the user who is accessing the machine learning model’s predictions and explanations will have prior knowledge and beliefs concerning inputs. An explanation is coherent if it is consistent with this prior knowledge. They also state that there is a high degree of subjectivity concerning this property, however, for what they describe as *basic coherence*, explanations should be consistent with the natural laws. With this in mind, we have labeled the explanations which are reliant on rule-based approaches (Ahmadi et al., 2019; Gad-Elrab et al., 2019) as coherent because they offer an explanation in terms of rule satisfaction. Kotonya and Toni (2020) define three coherence properties to evaluate explanation quality. We do not consider explanations making use of attention-based and language generate methods to be coherent because they are non-deterministic.

Context-full explanations. According to Sokol and Flach (2019), the *contextfulness* property is satisfied by an explanation if this is presented in a way that its context can be fully understood. We decided that all the explanations generated fulfill this criterion because they are all presented in the context of larger inputs i.e., the claim and sometimes user context (Shu et al., 2019; Wu et al., 2020).

Interactive explanations. Explanations which are interactive allow users to provide feedback to the system. Only two of the explainable fact-checking systems in the literature satisfy this property: CredEye (Popat et al., 2018) and XFake (Yang et al., 2019). The CredEye systems allows users to rate (a binary *yes* or *no*) whether they agree with the fact-checking assessment of the system. XFake offers a visualization of various assessments of the fact-checking, e.g., linguistic analysis so that the user can understand the misinformation at different levels of granularity.

Unbiased or impartial explanations. In the context of automated fact-checking, an unbiased explanation is one for which biases in the evidence data and other inputs have not influenced the explanation. Typically, for fact-checking, bias is in the form of hyper-partisan language or opinions presented as evidence. None of the systems which we examine explicitly consider this property. However we consider the rule-based approaches to be a step towards this because they output triples and not text sequences.

Parsimonious explanations. This relates to brevity, i.e., an explanation should communicate all information needed with minimal redundant text. We consider the explanations which output rules, tokens, and word clouds to be parsimonious. Atanasova et al. (2020a) evaluate explanations for non-redundancy.

Chronological explanations. This property is related to the time-sensitive nature of fact-checking as described by the *timing* principle in Table 3. None of the systems considers this property.

System	ACT	CSL	CHR	CNT	INT	UNB	PRS	CHN
Popat et al. (2017)				✓	✓		✓	
Popat et al. (2018)				✓			✓	
Shu et al. (2019)				✓				
Gad-Elrab et al. (2019)			✓	✓		✓	✓	
Ahmadi et al. (2019)			✓	✓		✓	✓	
Yang et al. (2019)				✓	✓			
Atanasova et al. (2020a)				✓			✓	
Lu and Li (2020)				✓			✓	
Wu et al. (2020)				✓			✓	
Kotonya and Toni (2020)			✓	✓				

Table 4: Properties of explanations, which are actionability (ACT), causality (CSL), coherence (CHR), contextfulness (CNT), interactiveness (INT), unbiasedness or impartiality (UNB), parsimony (PRS) and chronology or timing (CHN). ✓ indicates a property is considered. There is no explainable fact-checking system which considers the actionability, causality or chronology properties.

5 Future Directions

Here we present a number of limitations in the state-of-the-art approaches to explainable fact-checking (in addition to the limitations emerging from Table 4). With each shortcoming we also present possible approaches which might be explored in order to overcome the limitation.

Unverifiable Claims. In fact-checking, there is often an issue which arises with unverifiable claims. Indeed, there is sometimes insufficient evidence to either verify or refute claims. This problem arises in both manual and computational fact-checking (Uscinski and Butler, 2013; Thorne et al., 2018). For this reason, most formulations of the fact-checking task take this scenario into account with the inclusion of an additional class, e.g., the NOT-ENOUGH-INFO label in the FEVER dataset (Thorne et al., 2018). In such cases it would be necessary to offer an explanation as to why the claim could neither be verified nor refuted or what supporting evidence was missing such as to prevent a verdict being reached (i.e., a counterfactual explanation). For many explainable systems this is particularly tricky because explanations are typically generated with respect to the inputs. The knowledge graph-based systems (Gad-Elrab et al., 2019; Ahmadi et al., 2019) take steps towards this, as their output amounts to rules that serve as both evidence and counter-evidence, but more human understandable explanations are needed which state explicitly what evidence, if any, is available to verify the claim and why it is insufficient.

Predictions vs. Models. One limitation of current explainable fact-checking systems is the emphasis on outcome-driven explanations. These typically amount to explaining the predictions with reference to relation prediction. There is no system which offers process-driven or model explanations for multiple or all sub-tasks in the fact-checking pipeline, including justifications which might shed light on why and how decisions made in earlier sub-tasks affect outputs of later sub-tasks, e.g., evidence selection and natural language inference (NLI). There is existing work on explaining NLI (Thorne et al., 2019b), with, in particular, the e-SNLI dataset containing explanation annotations for NLI (Camburu et al., 2018). Building upon these works in explainable NLI, there are a number of directions which we envision as future avenues of research in explainable fact-checking which could result in valuable gains.

Evaluating explanations. As well as extracting explanations for predictions it is important that metrics are established for evaluating these explanations. The explainable fact-checking systems surveyed make use of both computational and human evaluations for explanation quality. Shu et al. (2019), Gad-Elrab et al. (2019), Atanasova et al. (2020a), Kotonya and Toni (2020), and Yang et al. (2019) employ human annotators to assess the quality of explanations generated by their models. Yang et al. (2019) conduct their human evaluation on the largest number of annotators, this is the only work which considers age, gender and education level as factors when choosing participants for the evaluation study. Kotonya and Toni (2020) formalize three coherence properties for evaluating the quality of explanations: local coherence, and strong and weak global coherence. ExpClaim (Ahmadi et al., 2019) is evaluated against and outperforms CredEye (Popat et al., 2018), and GCAN (Lu and Li, 2020) is shown to outperform dE-FEND (Shu et al., 2019). However, these evaluations are with respect to the prediction, not explanations. No evaluation is performed on explanations generated for some systems (Popat et al., 2017; Popat et al., 2018). In order to compare performance across explainable systems, it would be helpful to establish task-specific metrics similar to those which exist for other NLP tasks, e.g., machine translation (Papineni et al., 2002) and text summarization (Lin, 2004).

6 Conclusion

In conclusion, the emerging work on explainable machine learning in the fact-checking domain shows a great deal of promise despite the particularly challenging nature of the problem. However, there are some limitations of current methods and task formulations which we highlight in this survey: all existing methods only look to explain one component of the fact-checking pipeline (relation prediction or entailment) and the systems only explain the predictions. In order to garner the most relevant, useful and insightful fact-checking explanations, a holistic and journalistic-informed approach should be taken. We identify possible directions for future research, which will hopefully extend the work achieved so far.

References

- Naser Ahmadi, Joohyung Lee, Paolo Papotti, and Mohammed Saeed. 2019. Explainable fact checking with probabilistic answer set programming. *CoRR*, abs/1906.09198.
- Naser Ahmadi, Thi-Thuy-Duyen Truong, Le-Hong-Mai Dao, Stefano Ortona, and Paolo Papotti. 2020. Rulehub: A public corpus of rules for knowledge graphs. *J. Data and Information Quality*, 12(4), October.
- Tariq Alhindi, Savvas Petridis, and Smaranda Muresan. 2018. Where is your evidence: Improving fact-checking by justification modeling. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 85–90, Brussels, Belgium, November. Association for Computational Linguistics.
- Aimée Alonso-Reina, Robiert Sepúlveda-Torres, Estela Saquete, and Manuel Palomar. 2019. Team GPLSI. approach for automated fact checking. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 110–114, Hong Kong, China, November. Association for Computational Linguistics.
- Michelle A Amazeen. 2015. Revisiting the epistemology of fact-checking. *Critical Review*, 27(1):1–22.
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020a. Generating fact checking explanations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7352–7364, Online, July. Association for Computational Linguistics.
- Pepa Atanasova, Dustin Wright, and Isabelle Augenstein. 2020b. Generating label cohesive and well-formed adversarial claims.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4685–4697, Hong Kong, China, November. Association for Computational Linguistics.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-snli: natural language inference with natural language explanations. In *Advances in Neural Information Processing Systems*, pages 9539–9549.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. Enhanced LSTM for natural language inference. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada, July. Association for Computational Linguistics.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou, and William Yang Wang. 2020. Tabfact : A large-scale dataset for table-based fact verification. In *8th International Conference on Learning Representations, (ICLR 2020)*, Online, April.
- Niall J Conroy, Victoria L Rubin, and Yimin Chen. 2015. Automatic deception detection: Methods for finding fake news. *Proceedings of the Association for Information Science and Technology*, 52(1):1–4.
- Leon Derczynski, Kalina Bontcheva, Maria Liakata, Rob Procter, Geraldine Wong Sak Hoi, and Arkaitz Zubiaga. 2017. SemEval-2017 task 8: RumourEval: Determining rumour veracity and support for rumours. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 69–76, Vancouver, Canada, August. Association for Computational Linguistics.
- Don Fallis. 2015. What is disinformation? *Library Trends*, 63(3):401–426.
- William Ferreira and Andreas Vlachos. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1163–1168, San Diego, California, June. Association for Computational Linguistics.
- Mohamed H Gad-Elrab, Daria Stepanova, Jacopo Urbani, and Gerhard Weikum. 2019. Exfakt: a framework for explaining facts over knowledge graphs and text. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 87–95. ACM.
- L Graves and F Cherubini. 2016. The rise of fact-checking sites in europe. Technical report, Reuters Institute for the Study of Journalism, University of Oxford.
- Lucas Graves. 2017. Anatomy of a Fact Check: Objective Practice and the Contested Epistemology of Fact Checking. *Communication, Culture and Critique*, 10(3):518–537, 10.

- D Graves. 2018. Understanding the promise and limits of automated fact-checking. Technical report, Reuters Institute for the Study of Journalism, University of Oxford.
- Andreas Hanselowski, Avinesh PVS, Benjamin Schiller, Felix Caspelherr, Debanjan Chaudhuri, Christian M. Meyer, and Iryna Gurevych. 2018a. A retrospective analysis of the fake news challenge stance-detection task. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1859–1874, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Andreas Hanselowski, Hao Zhang, Zile Li, Daniil Sorokin, Benjamin Schiller, Claudia Schulz, and Iryna Gurevych. 2018b. UKP-athene: Multi-sentence textual entailment for claim verification. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 103–108, Brussels, Belgium, November. Association for Computational Linguistics.
- Andreas Hanselowski, Christian Stab, Claudia Schulz, Zile Li, and Iryna Gurevych. 2019. A richly annotated corpus for different tasks in automated fact-checking. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 493–503, Hong Kong, China, November. Association for Computational Linguistics.
- Naeemul Hassan, Chengkai Li, and Mark Tremayne. 2015. Detecting check-worthy factual claims in presidential debates. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, CIKM '15*, page 1835–1838, New York, NY, USA. Association for Computing Machinery.
- Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward automated fact-checking: Detecting check-worthy factual claims by claimbuster. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1803–1812.
- Alexey Ignatiev, Nina Narodytska, and Joao Marques-Silva. 2019. On relating explanations and adversarial examples. In *Advances in Neural Information Processing Systems*, pages 15857–15867.
- Sarthak Jain and Byron C. Wallace. 2019. Attention is not Explanation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Neema Kotonya and Francesca Toni. 2020. Explainable automated fact-checking for public health claims. In *2020 Conference on Empirical Methods in Natural Language Processing*, Online, November. Association for Computational Linguistics.
- Ran Levy, Yonatan Bilu, Daniel Hershcovich, Ehud Aharoni, and Noam Slonim. 2014. Context dependent claim detection. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1489–1500.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Marco Lippi and Paolo Torroni. 2015. Context-independent claim detection for argument mining. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China, November. Association for Computational Linguistics.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2020. Fine-grained fact verification with kernel graph attention network. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7342–7351, Online, July. Association for Computational Linguistics.
- Wilson Lowrey. 2017. The emergence and development of news fact-checking sites: Institutional logics and population ecology. *Journalism Studies*, 18(3):376–394.
- Yi-Ju Lu and Cheng-Te Li. 2020. GCAN: Graph-aware co-attention networks for explainable fake news detection on social media. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 505–514, Online, July. Association for Computational Linguistics.
- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc.

- Tanushree Mitra and Eric Gilbert. 2015. Credbank: A large-scale social media corpus with associated credibility annotations. In *Ninth International AAAI Conference on Web and Social Media*.
- Menaka Narayanan, Emily Chen, Jeffrey He, Been Kim, Sam Gershman, and Finale Doshi-Velez. 2018. How do humans understand explanations from machine learning systems? an evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1802.00682*.
- Yixin Nie, Haonan Chen, and Mohit Bansal. 2019. Combining fact extraction and verification with neural semantic matching networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6859–6866.
- Piotr Niewinski, Maria Pszona, and Maria Janicka. 2019. GEM: Generative enhanced model for adversarial attacks. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 20–26, Hong Kong, China, November. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255, Austin, Texas, November. Association for Computational Linguistics.
- Verónica Pérez-Rosas, Bennett Kleinberg, Alexandra Lefevre, and Rada Mihalcea. 2018. Automatic detection of fake news. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3391–3401, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Dean Pomerleau and Delip Rao. 2017. The fake news challenge: Exploring how artificial intelligence technologies could be leveraged to combat fake news. *Fake News Challenge*.
- Kashyap Papat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2017. Where the truth lies: Explaining the credibility of emerging claims on the web and social media. In *Proceedings of the 26th International Conference on World Wide Web Companion, WWW '17 Companion*, page 1003–1012, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Kashyap Papat, Subhabrata Mukherjee, Andrew Yates, and Gerhard Weikum. 2018. DeClarE: Debunking fake news and false claims using evidence-aware deep learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 22–32, Brussels, Belgium, October-November. Association for Computational Linguistics.
- Martin Potthast, Johannes Kiesel, Kevin Reinartz, Janek Bevendorff, and Benno Stein. 2018. A stylometric inquiry into hyperpartisan and fake news. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 231–240, Melbourne, Australia, July. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2931–2937, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- Victoria L Rubin, Niall Conroy, Yimin Chen, and Sarah Cornwell. 2016. Fake news or truth? using satirical cues to detect potentially misleading news. In *Proceedings of the second workshop on computational approaches to deception detection*, pages 7–17.
- Roser Saurí and James Pustejovsky. 2009. Factbank: a corpus annotated with event factuality. *Language resources and evaluation*, 43(3):227.
- Tal Schuster, Darsh J Shah, Yun Jie Serene Yeo, Daniel Filizzola, Enrico Santus, and Regina Barzilay. 2019. Towards debiasing fact verification models. *arXiv preprint arXiv:1908.05267*.

- Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2018. Fakenewsnet: A data repository with news content, social context and dynamic information for studying fake news on social media. *arXiv preprint arXiv:1809.01286*.
- Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. defend: Explainable fake news detection. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pages 395–405, New York, NY, USA. ACM.
- Kacper Sokol and Peter Flach. 2019. Desiderata for interpretability: Explaining decision tree predictions with counterfactuals. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 10035–10036.
- Amir Soleimani, Christof Monz, and Marcel Worring. 2020. Bert for evidence retrieval and claim verification. In Joemon M. Jose, Emine Yilmaz, João Magalhães, Pablo Castells, Nicola Ferro, Mário J. Silva, and Flávio Martins, editors, *Advances in Information Retrieval*, pages 359–366, Cham. Springer International Publishing.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana, June. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2019a. Evaluating adversarial attacks against multiple fact verification systems. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2944–2953, Hong Kong, China, November. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2019b. The FEVER2.0 shared task. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 1–6, Hong Kong, China, November. Association for Computational Linguistics.
- Richard Tomsett, Amy Widdicombe, Tianwei Xing, Supriyo Chakraborty, Simon Julier, Prudhvi Gurram, Raghuvveer Rao, and Mani Srivastava. 2018. Why the failure? how adversarial examples can provide insights for interpretable machine learning. In *2018 21st International Conference on Information Fusion (FUSION)*, pages 838–845. IEEE.
- Joseph E Uscinski and Ryden W Butler. 2013. The epistemology of fact checking. *Critical Review*, 25(2):162–180.
- Andreas Vlachos and Sebastian Riedel. 2014. Fact checking: Task definition and dataset construction. In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 18–22. Association for Computational Linguistics.
- David Wadden, Kyle Lo, Lucy Lu Wang, Shanchuan Lin, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. In *EMNLP*.
- William Yang Wang. 2017. “liar, liar pants on fire”: A new benchmark dataset for fake news detection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 422–426, Vancouver, Canada, July. Association for Computational Linguistics.
- Sarah Wiegrefe and Yuval Pinter. 2019. Attention is not not explanation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China, November. Association for Computational Linguistics.
- Lianwei Wu, Yuan Rao, Yongqiang Zhao, Hao Liang, and Ambreen Nazir. 2020. DTCA: Decision tree-based co-attention networks for explainable claim verification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1024–1035, Online, July. Association for Computational Linguistics.
- Fan Yang, Shiva K. Pentylala, Sina Mohseni, Mengnan Du, Hao Yuan, Rhema Linder, Eric D. Ragan, Shuiwang Ji, and Xia (Ben) Hu. 2019. Xfake: Explainable fake news detector with visualizations. In *The World Wide Web Conference*, WWW '19, page 3600–3604, New York, NY, USA. Association for Computing Machinery.
- Dimitrina Zlatkova, Preslav Nakov, and Ivan Koychev. 2019. Fact-checking meets fauxtography: Verifying claims about images. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2099–2108, Hong Kong, China, November. Association for Computational Linguistics.