Toward Automated Factchecking: Developing an Annotation Schema and Benchmark for Consistent Automated Claim Detection

LEV KONSTANTINOVSKIY, Full Fact
OLIVER PRICE, University of Warwick
MEVAN BABAKAR, Full Fact
ARKAITZ ZUBIAGA, Queen Mary University of London

In an effort to assist factcheckers in the process of factchecking, we tackle the claim detection task, one of the necessary stages prior to determining the veracity of a claim. It consists of identifying the set of sentences, out of a long text, deemed capable of being factchecked. This article is a collaborative work between Full Fact, an independent factchecking charity, and academic partners. Leveraging the expertise of professional factcheckers, we develop an annotation schema and a benchmark for automated claim detection that is more consistent across time, topics, and annotators than are previous approaches. Our annotation schema has been used to crowdsource the annotation of a dataset with sentences from UK political TV shows. We introduce an approach based on universal sentence representations to perform the classification, achieving an F1 score of 0.83, with over 5% relative improvement over the state-of-the-art methods ClaimBuster and ClaimRank. The system was deployed in production and received positive user feedback.

CCS Concepts: • **Computing methodologies** → **Natural language processing**; • **Information systems** → *Clustering and classification*; • **Applied computing** → *Document management and text processing*;

Additional Key Words and Phrases: Claim detection, factchecking, debates, classification

ACM Reference format:

Lev Konstantinovskiy, Oliver Price, Mevan Babakar, and Arkaitz Zubiaga. 2021. Toward Automated Factchecking: Developing an Annotation Schema and Benchmark for Consistent Automated Claim Detection. *Digit. Threat.: Res. Pract.* 2, 2, Article 14 (April 2021), 16 pages.
<https://doi.org/10.1145/3412869>

1 INTRODUCTION

Misinformation has recently become more central in public discourse [6, 37, 46]. As a consequence, interest has increased in the scientific community to further natural language processing (NLP) approaches that can help alleviate the burdensome and time-consuming human activity of factchecking [38, 40]. Factchecking is known as the task of producing an informed assessment of the veracity of a claim [13, 14]. The main mission of factcheckers

This work has been partially funded by Open Society Foundation (Grant No. OR2017-35395) and Omidyar Network.

Authors' addresses: L. Konstantinovskiy and M. Babakar, Full Fact, 2 Carlton Gardens, St. James's, London SW1Y 5AA, UK; emails: lev.konst@gmail.com, mevan@fullfact.org; O. Price, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, UK; email: op.oliverprice@gmail.com; A. Zubiaga, Queen Mary University of London, Mile End Road, London E1 4NS, UK; email: a.zubiaga@qmul.ac.uk. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2576-5337/2021/04-ART14 \$15.00

<https://doi.org/10.1145/3412869>

is to give citizens information to make political choices, improve the quality of public political discourse, and to hold politicians accountable [14]. Misinformation and misperceptions can undermine this goal [9, 11].

However, there is a very small number of factchecking organisations in the world, about 160,¹ compared to the volume of media items produced daily. The speed at which information flows online means there is less time to verify the claims made and myths spread further before being factchecked, if they are factchecked at all. Automating any parts of the factchecking process could cut down the time it takes to respond to a claim. It could also protect human factchecker's time to work on the more complicated checks that need careful human judgement.

When considered inside a factchecking organisation, the factchecking process consists of a series of tasks, which Reference [3] defined as a process consisting of four tasks: (1) monitor media, i.e., capturing content such as articles, videos, images, (2) detect claims, i.e., spotting when an item contains a checkable claim, (3) check claims, i.e., doing the research to check and verify that claim, and (4) publish, i.e., creating a piece of content that encapsulates the results of the check.

The vast majority of the scientific research has focused on the determination of veracity [20, 39, 45] or checking claims; despite often having been given different names in the scientific literature, such as fake news detection [26, 41, 44] or arguably rumour detection [2, 24, 27, 43].

The component checking the veracity of a claim cannot function in isolation, it needs a list of claims to check as an input. The majority of previous work on factchecking has started from a readily available list of claims and has omitted the previous steps of monitoring and spotting in the factchecking pipeline [20, 40]. This act of spotting claims within a corpus is known as claim detection, i.e., monitoring news sources and identifying if a sentence constitutes a claim that could be factchecked [18, 21]. Monitoring and spotting within a factchecking organisation is a time consuming semi-manual task, which is inevitably limited by the resources available. Automating the process of claim detection could mean that factcheckers can monitor a greater set of media, extract the claims made, and hopefully make smarter choices about what the most valuable items to be checking that day are. If deployed in a live factchecking situation, then it could also help separate out claims made in real-time during a ministerial speech, for example. This could help factcheckers quickly skim transcripts when time is limited. Despite the importance of this component in the factchecking pipeline, automation is still in its infancy and there is a dearth of scientific literature.

Previous works on claim detection [12, 18] rely on definitions of claims that incorporate the concept of check-worthiness and/or importance. Our objective is instead to avoid these subjective concepts to come up with an objective way of determining what constitutes a claim that is checkable; rather than being important or worthy of checking, which is bound to subjective interpretation. This article describes the iterative process we followed together with factcheckers to come up with an annotation schema that would effectively capture claims and non-claims. This annotation schema avoids factors that can be affected by personal biases, such as importance, in the manual annotation to produce an objective outcome. Following this annotation schema through a crowdsourcing methodology, we generated a dataset of 5,571 sentences labelled as claims or non-claims. Further, we set out to present the development of the first stage in the automated factchecking pipeline. It constitutes the first automated claim detection system developed by an independent factchecking charity, Full Fact,² along with academic partners. The main contributions of our work are as follows:

- We introduce the first annotation schema for claim detection, iteratively developed by experts at Full Fact, comprising seven different labels.
- We describe a crowdsourcing methodology that enabled us to collect a dataset with 5,571 sentences labelled according to this schema.

¹<https://reporterslab.org/fact-checking/>.

²<http://fullfact.org/>.

- We develop a claim detection system that leverages universal sentence representations, as opposed to previous work that was limited to word-level representations. Our experiments show that our claim detection system outperforms the state-of-the-art claim detection systems, ClaimBuster and ClaimRank.
- With the annotation schema, crowdsourcing methodology and task definition, we set forth a benchmark methodology for further development of claim detection systems.

2 RELATED WORK

Most work around automatically factchecking claims has focused on the later stages of determining the veracity of claims, usually by building knowledge graphs out of knowledge bases, such as Wikidata [5, 7, 35, 36, 42]. Less work has been documented on the preceding stage of claim detection.

One of the best-known approaches to claim detection is ClaimBuster [18]. They collected a large annotated corpus of televised debates in the USA. Their model combines Term Frequency-Inverse Document Frequency (TF-IDF), part-of-speech (POS) tags, and named entity recognition (NER) features on an Support Vector Machines (SVM) classifier and produces a score of how important a claim is to factcheck. This has the caveat of then having to choose a cut-off score to determine the claims that will be considered worthy or important enough for factchecking. Our approach is instead to define an annotation schema that is binary, determining checkable claims rather than check-worthy claims, and is built on several types of claims. It is a simpler fit for the use case in this factchecking pipeline, i.e., in a live stream of subtitles, we are unable to know in advance which sentence will make it to the top ranking until the end of the entire programme.

In Reference [15] annotations were collected using ClaimBuster-inspired annotation guidance from volunteers together with their age, gender and education. With possible labels being verifiable check-worthy (VCW), verifiable not check-worthy (VNCW), and not verifiable (NV), they obtained 2,100 labelled sentences (reduced to 264 high-quality labelled sentences). They find that annotators with a natural sciences background agreed internally about what constitutes a check-worthy claim, whereas those with humanities, medicine, and ontology backgrounds saw more internal disagreement on check-worthiness. Furthermore, using 35 labelled control claims to test annotator skill, they find that the age group 40–49 obtains a higher average score, and label more claims, than that of 30–39, which correspondingly scores higher than age group 20–29.

Another recent approach to claim detection is ClaimRank [12]. They compiled a dataset by taking the outcome of factchecking a political debate, published by nine organisations simultaneously. Models were created to predict if the claim would be highlighted by at least one or by a specific organisation. The modelling is done with a large variety of features from both the individual sentence and the wider context around it. A subsequent version of the dataset [21, 29] includes a larger set of sentences in two languages, English and Arabic. These datasets are similar to ours in order of magnitude, however, use a different definition of claim, as is the case with others tackling the determination of check-worthiness of claims [4, 17, 31, 47]. We further elaborate on this in the next subsection.

A slightly different approach to claim detection is that of context-dependent claim detection (CDCD) [25]. This study proposes identifying claims *given a specific context*. Articles relevant to a topic are used to detect claims on that topic.

Another piece of work worth mentioning is the development of the FEVER (Fact Extraction and VERification) dataset [39]. Whilst this is primarily aimed at work regarding claim veracity, the mere presence of a vast quantity of claims in the dataset allow it to be extended for claim detection in the future.

2.1 Previous Attempts at Defining Claims

There is a body of work on claim detection that has not formalised the definition of a claim, e.g., [12, 21]. Instead it directly relies on what has been identified by external organisations. The lack of a formal definition prevents others from replicating or extending their work. These studies used claims identified by nine organisations in a

political speech as a proxy. The annotations were sourced from publicly available online articles. This is different from our approach, where we crowdsourced annotations following our definition of the task. The authors of Reference [12] acknowledge this limitation, which led to high number of false positives in their experiments. For example, an article consisting of a debate transcript with editorial comments will not highlight repeated instances of claims. This creates inconsistent annotation—during a TV debate, popular claims are discussed on repeated occasions. This had to be re-annotated by researchers in Reference [29]. Another caveat is that only three of the nine annotator organisations contributing to those online articles sign up to be neutral and transparent in their selection of claims as verified signatories of the International Factchecking Network’s Code of Principles.³

ClaimBuster [18] provides a definition of a claim that revolves around the question: “Will the general public be interested in knowing whether this sentence is true or false?” Claims are considered to be those sentences for which the answer to this question is yes. Their aim was for anyone to be able to feed in a source, e.g., a political speech, and for the system to produce a list of claims ranked by importance, which could directly feed into the editorial process. This definition of a claim includes the judgement of “importance,” which we avoid in our work. We believe it is an editorial judgement best left to factcheckers. ClaimBuster annotators were journalists, students and professors. Annotations that agreed with the authors of that study were selected to ensure good agreement and shared understanding of the assumptions. Researchers from the ClaimBuster team also defined an annotation schema called PolitiTax,⁴ a taxonomy of political claims that we considered. However, the categories were not useful for the downstream task of checking the veracity of the claim by routing it to the right dataset or team at Full Fact, in part, due to the level of granularity in the taxonomy and, in part, because the team at Full Fact is split across topics.

There was also a taxonomy defined by factcheckers during the HeroX factchecking challenge [10], which is less granular than PolitiTax. It has four claim types—numerical, political stance, quotes, objects. During this work, we discovered that the latter three categories are rare and intersect with others, so we did not use them in our schema.

3 DATASET

3.1 Our Claim Definition and Process

Writing the annotation guidance was a long process. Full Fact’s formal definition of a claim during the 2015 UK election, was “an assertion about the world that can be checked.” Media monitoring volunteers were encouraged to ask a factchecker if they had doubts on whether something was assessable. We worked on codifying some of this thinking in conversations with the factcheckers. However, as we captured more and more claims, this definition proved insufficient. We wanted to understand whether a claim could be better defined by breaking it down into sub-categories for better consistency across time, topics and annotators. We opted for defining a typology that would capture the different types of claims, which would be more comprehensive than the previous short definition. Asking annotators to identify the category that a sentence belongs to would encourage more critical thinking. Likewise, choosing the right category for a sentence would significantly reduce the personal bias with respect to judging whether it is a claim, which is bound to personal judgement. Annotators would choose the category pertaining to a sentence, and we would then simplify the schema by mapping those types to binary labels, claim or not a claim.

We chose to decouple the importance of the claim from the claim itself. We felt that importance was heavily subjective, reliant on context and best left to factcheckers. Importance is a subtle, and forever changing feature.

³<https://perma.cc/BM43-SJ4N>.

⁴“PolitiTax A Taxonomy of Political Claims” by IDIR Lab <https://perma.cc/4RQF-FCPV>.

Even though the most “important” issues in the view of the UK public are often about the economy, immigration and health, their relative positions change.⁵ In some cases new issues become important, e.g., in the UK, the importance of claims about the EU increased significantly after the 2016 EU referendum.⁶

We also chose to decouple the topic from the definition. By making our definition descriptive of the claim and not, by proxy, the topic, we would have a more consistent final dataset. In some cases the selection of topics is an inherently political choice, e.g., it varies across the population whether “drugs” relate to the topic of “crime” or “health.” This kind of classification was avoided.

To come up with the schema that would capture what constitutes a claim, we followed an iterative process. In the first step, factcheckers identified sentences that were definitely not a claim. They iterated on potential rules and found examples that broke them. They also identified some constraints, for example, a claim needs to be checkable with more readily available evidence, which means that a personal claim like “I woke up at 7 a.m. today” is not a claim capable of being checked. We were most concretely able to exclude claims based on an individual’s personal experience, as more often than not they were un-checkable. This is similar to ‘verifiable experiential’ statements [30].

We went through several versions of the guidance with different taxonomies. They were trialled within Full Fact, and then two versions with external volunteers. The first version applied the 2015 thinking and was a binary accept/reject classification task, accompanied by a guidance. It listed several types of qualities of claims and non-claims. Claims, for example, may be explicit, implicit, or trivial. Non-claims in this version were formed of personal experience and opinion. We decided against these categories in the end as they sometimes involve explicit judgements from our annotators – these choices can sometimes be highly political. For example, in the case of “*The EU is made up of 27 [instead of 28] countries*” or “*The NHS is there for everyone*” some annotators could classify them as trivial while others might consider them explicit legal claims. The implicit/explicit categories were also removed, because whether the claim is implicit or explicit is not important for the next downstream task in the factchecking process after claim detection.

For the second version, we looked at Full Fact’s factchecks. They mostly covered statistical claims. We also identified claims around current laws or rules of operation and correlation/causation claims, e.g., “there’s no clear correlation between prisons’ performance ratings and whether they’re publicly-run or contracted out to the private sector.” This became the basis of our claim categories. Merging these categories and removing personal experience was deemed to be a good proxy for claims. There were many other types of claims that we identified, such as definitions, voting records, and expressions of support. We limited our categories to seven to make the task realistic for annotators. We also wanted to minimise the overlap between categories to make the task single-choice.

3.2 Annotation Guidance

Our annotation schema is the first to be created with a factchecking organisation. It comprises seven categories, only one of which can be assigned to each sentence. Annotators were given definitions and examples of the seven categories, including the more detailed breakdown shown in Table 1:

- **Personal experience.** Claims that are not capable of being checked using publicly available information, e.g., “I can’t save for a deposit.”
- **Quantity in the past or present.** Current value of something, e.g., “One in four wait longer than 6 weeks to be seen by a doctor.” Changing quantity, e.g., “The Coalition Government has created 1,000 jobs for

⁵“Ipsos MORI Issues Index: 2017 in review” <https://perma.cc/9SMV-CQR8>.

⁶“July 2016 Economist/Ipsos MORI Issues Index” <https://perma.cc/DPA5-4XV5>.

Table 1. Breakdown of the 4,080 Sentences with Majority Agreement

Category	Subcategory	Counts	Example
Not a claim		54.8%	"Give it all to them, I really don't mind."
Other	Other other	10.4%*	"Molly gives so much of who she is away throughout the film."
	Support/policy	5.5%*	"He has advocated for a junk food tax."
	Quote	4.7%*	"The Brexit secretary said he would guarantee free movement of bankers."
	Trivial claim	1.6%*	"It was a close call."
	Voting record	0.7%*	"She just lost a parliamentary vote."
	Public opinion	0.4%*	"A poll showed that most people who voted Brexit were concerned with immigration."
	Definition	0.0%*	"The unemployed are only those actively looking for work."
Quantity	Current value	9.9%	"1 in 4 people wait longer than 6 weeks to see a doctor."
	Changing quantity		
	Comparison		
	Ranking		
Prediction	Hypothetical statements	4.4%	"The IFS says that school funding will have fallen by 5% by 2019."
	Claims about the future		
Personal experience	Uncheckable	3.0%	"I can't save for a deposit."
Correlation/causation	Correlation	2.6%	"Tetanus vaccine causes infertility."
	Causation		
	Absence of a link		
Laws/rules of operation	Public institutional procedures	1.9%	"The UK allows a single adult to care for fewer children than other European countries."
	Rules/rule changes		

*The proportions for "Other" sub-categories are taken from a random sample of 160 claims labelled as "Other."

every day it's been in office." Comparison, e.g., "Free schools are outperforming state schools." Ranking, e.g., "The UK's the largest importer from the Eurozone."

- **Correlation or causation**, Correlation, e.g., "GCSEs are a better predictor than AS if a student will get a good degree." Causation, e.g., "Tetanus vaccine causes infertility." Absence of a link, e.g., "Grammar schools don't aid social mobility."
- **Current laws or rules of operation**, Declarative sentences, which generally have the word "must" or legal terms, e.g., "The UK allows a single adult to care for fewer children than other European countries." Procedures of public institutions, e.g., "Local decisions about commissioning services are now taken by organisations that are led by clinicians." Rules and changes, e.g., "EU residents cannot claim Jobseeker's Allowance if they have been in the country for 6 months and have not been able to find work."
- **Prediction**, Hypothetical claims about the future, e.g., "Indeed, the IFS says that school funding will have fallen by 5% in real terms by 2019 as a result of government policies."
- **Other type of claim**, Voting records, e.g., "You voted to leave, didn't you?" Public Opinion e.g. "Public satisfaction with the NHS in Wales is lower than it is in England." Support, e.g., "The party promised free childcare" Definitions, e.g., "Illegal killing of people is what's known as murder." Any other sentence that you think is a claim.
- **Not a claim**, These are sentences that do not fall into any categories and are not claims, e.g., "What do you think?" "Questions to the Prime Minister!"

These categories have proven to broadly cover sentences from political TV shows that Full Fact has encountered over several years. Categories have different levels of occurrence (see Table 1). As previously found [15, 18, 21], “Not a claim” is the most popular category, amounting to about 55% of the annotations.

“Other” is the second largest category with 952 instances, 23% of the whole. It can be broken down into claims that are less well-defined, with formal sub-categories being: “Definitions,” “Voting records,” “Public opinion,” “Trivial claim,” “Support,” “Quote,” “Other other.”⁷ We amalgamate these, because they are likely to overlap, and we wanted our annotators to only select one option. For example, “She said she voted to keep free school meals.” is both a quote and a voting record. Furthermore, high granularity of categories allows one to perpetually think of rarer categories and a high number of categories slows down annotation unnecessarily. To verify if our level of granularity was correct, we split a sample of 160 sentences in “Other” into sub-categories. The vast majority are in the “Other other” category (see Table 1), supporting our chosen level of granularity.

3.3 Crowdsourced Annotation

The annotations were done by 80 volunteers recruited through Full Fact’s newsletter—this meant that volunteers were keen on factchecking. 28,100 annotations were collected for a set of 6,304 sentences extracted from subtitles of four UK political TV shows, 14 episodes in total. TV subtitles were chosen, because 69% of the UK population get their news from TV.⁸

The software used for collecting annotations was Prodigy,^{9,10} a self-hosted annotation platform. It was customised to support multiple annotators with a login and password screen where each user would enter their credentials. Sentences were shown in random order. The preceding two sentences were also shown on the screen to provide context and assist with potential co-references. Once a sentence was annotated five times by different annotators, it was not shown again.

Annotators were encouraged to contact us for any clarifications needed, with thoughtful questions such as: “Where it appears that a claim is dressed up as rhetorical question, should we classify it as a claim? For example, ‘Why should unelected officials in Brussels make rules to stop bananas being sold in bunches of more than 2 or 3?’” To answer this, questions are classified as the claims that they implicitly contain.

3.4 Agreement

At the level of all seven granular categories the inter-annotator agreement is moderate, with a Krippendorff’s alpha [23] of 0.46. However, we attain higher values of alpha of 0.70 and 0.53 when we do the mapping of annotations into the binary claim/non-claim annotation task, following either of the two methods shown in Table 3.

Most of the disagreement was between “Not a claim” and “Other claim.” This showed that it is hard to define the boundary and explicitly list all kinds of claims, as we saw in the process of creating the annotation guidance. The disagreements across all sentence types can be seen in Table 2.

For mapping the seven categories in the schema into the binary classification task distinguishing claims and non-claims, two different mappings were initially proposed (see Table 3). These two methods were proposed by first assuming the two extremes, i.e., Quantity (2) should be deemed a claim, whereas Not a claim (7) belong to non-claim. Subsequently, these two reformulations of the taxonomy were proposed after meeting and brainstorming with factcheckers. While the method in the first row would be reasonable for achieving higher inter-annotator agreement, it would lead to a classification performance prioritising high precision at the

⁷Note the “Other other” category is needed as sentence types previously unseen will inevitably emerge; however, it is only intended for sentences that do not belong to any other category.

⁸“News consumption in the UK:2016” by Ofcom <https://perma.cc/5FDK-BRHD>.

⁹<https://prodigy/>.

¹⁰<https://fullfact.org/blog/2018/feb/how-we-customised-prodigy-ai/>.

CLAIM

And the challenge that we have is this productivity challenge that was laid out in the Budget. S1: Can you use sort of ordinary English? I mean, the productivity challenge is hardly going to get everybody fired up, is it?

S3: What it is is this, it takes people in Germany four days to produce what people in the UK take five days to produce.

☐ Personal experience 1

☐ Quantity in the past or present 2

☐ Correlation and causation (without a quantity) 3

☐ Current laws or rules of operation 4

☐ Prediction 5

☐ Other type of claim 6

☐ Not a claim 7

DATE: 2017-11-23 NAME: Question Time SENTENCE: What it is is this, it takes people in Germany four days to produce what people in the UK take five days to produce. SENTENCE_ID: b09girz4_201 TIME: 00:16:53.520

Fig. 1. Annotation UI in Prodigy.

Table 2. Annotation Disagreements

Qu	12					
Corr	10	10				
Law	2	19	11			
Pred	3	42	27	25		
Other	50	102	129	87	87	
Not	114	58	90	69	103	668
	Pers	Qu	Corr	Law	Pred	Other

The most prominent disagreement is between “Other claim” and “Not a claim.” The labels are shortened versions of those in Figure 1 due to space limitations. (Qu: quantity, Corr: correlation and causation, Pred: predictions, Pers: personal experience.)

Table 3. From Seven Categories to Binary Claim vs. “Not a Claim” Classification

Claim	Non-claim	Omitted	α	N
2	3, 4, 6, 7	1, 5	0.70	6,095
2, 3, 4, 5	1, 6, 7	–	0.53	4,777

N = number of sentences annotated by majority. (1) Personal experience, (2) Quantity in the past or present, (3) Correlation or causation, (4) Current laws or rules of operation, (5) Prediction, (6) Other type of claim, (7) Not a claim.

expense of a lower recall. The method in the second row was ultimately selected for further experimentation, as recall was deemed important by factcheckers and it is in turn more realistic for not omitting any of the seven initial categories.

Hence, we moved on to evaluate our claim detection system on the labels originating from the second row of Table 3. Here, “Other type of claim” is not in the positive class for two reasons. First, there is a lot of disagreement between it and the “Not a claim” class. Second, the kinds of claim in the “Other” section—voting records, quotes, statements about public opinion polls, are less frequently written about by Full Fact.

The agreement of 60% is still low—that is a lot of sentences to throw away if we were only to consider agreement among all the annotators. So, instead, we choose a majority vote where at least three annotators marked the sentence and more than half of them agree. Of the initial 6,304 sentences, this filter selects 4,777 sentences, 3,973 not claims, and 804 claims. This is in line with previous studies where the proportion of claims is 10–30% in political TV [12, 18]. As extra training data, we add 794 claims from the Full Fact database. Out of them 766 are annotated by us as positive, because they fall into our claim categories, for example “The courts have said that the so-called ‘bedroom tax’ is illegal.” The remaining 28 are in the “Other type of claim” category, for example, “The British economy is not only getting better, it is healing.”

If we keep the seven categories in the dataset, instead of mapping them into the two classes, then the same method based on majority votes leads to a slightly smaller dataset with 4,080 sentences, i.e., due to the slightly lower agreement on the broader set of seven categories.

4 METHODS

To capture the diversity of sentences observed during political TV shows, we propose to leverage universal sentence representations. We use InferSent [8] as a method to achieve sentence embeddings. These embeddings are different from averaging word embeddings, because they take word order into account using a recurrent neural network. The method provided by InferSent involves words being converted to their common crawl GloVe implementations before being passed through a bidirectional long-short-term memory (BiLSTM) network [19]. The sentence embeddings were pre-trained on a large dataset of Natural Language Inference tasks.¹¹ Additionally, we also tried concatenating POS and NER information to the embeddings. For each sentence, the POS/NER feature vector was the count of each POS/NER tag in the corpus. We input our sentence representations to a range of supervised classifiers implemented using scikit-learn [32], with the classifiers set to their default parameters. The four classifiers we tested include Logistic Regression, Linear SVM, Gaussian Naïve Bayes and Random Forests, all of which use the default parameters provided with scikit-learn.

We use a number of other features as baselines:

- (1) A number of variants of the state-of-the-art claim detection system by ClaimBuster, using different combinations of TF-IDF, POS and NER features, as in Reference [18].
- (2) Averaging pre-trained word embedding vectors for all words in a sentence. We evaluate:
 - Word2vec [28] via the Gensim implementation [34], using the *GoogleNews* embedding.
 - GloVe [33] trained on Common Crawl, as well as combining them with dimensionality reduction using principal component analysis (PCA).
- (3) TF-IDF representations of sentences with logistic regression. Numbers have a significant role in claims—the “Cardinal Number” part-of-speech tag is the second most discriminating feature in Reference [16], so we try a Spacy NER to replace numbers with “*NUMBER*” during preprocessing.

For our implementation of the ClaimBuster system, we use the Watson Natural Language Understanding API, the updated version of the Alchemy API used by the original authors. All other features were implemented as outlined in Reference [18].

¹¹<https://nlp.stanford.edu/projects/snli/>.

The ClaimRank [12] model was harder to re-implement. To maintain impartiality Full Fact cannot use the data on sentence speakers when selecting which claims to check. This special care is also encouraged in the computer science research for systems that are integrated into the infrastructure of society by the ACM code of ethics [1]. Additionally, our dataset did not have data on applause, laughing, or speaker crossover. Even though we unfortunately could not use one third of ClaimRank’s features,¹² we trained the FNN, SVM and logistic regression classifiers on the remaining features in our dataset [12].¹³

4.1 Experiment Settings

The dataset consists of 5,571 sentences (4,777 from annotations and 794 from the Full Fact database of claims), of which 1,570 are claims and 4,001 are not claims, which gives a 30/70 class imbalance. We use stratified fivefold cross-validation to train and test our models. We use precision, recall and F_1 -score measures to assess classifier performance.

We show the best-performing classifier for any given feature set. We also show 95% confidence interval for the precision and recall using binomial distributions. This demonstrates possible overlap in results between different models. The interval is wide for recall due to the small number of positive examples. The next section will present the results of applying these methods.

5 CLAIM DETECTION

Here, we present results for the binary classification, using the class grouping shown in the second row of Table 3.

5.1 Analysis of Results

Table 4 shows the results of our experiments. Note that in this task recall is especially important as factcheckers do not want to miss out important claims, and hence our priority is to maximise recall while also keeping a good balance of precision and recall as measured by the F_1 score. Interestingly, the simple approach of TF-IDF achieves high precision but low recall. We call our new model “CNC,” which stands for “Claim/No Claim.” It achieves a better balance of precision and recall; logistic regression classifier gives the highest overall F_1 score of 0.83, outperforming all other techniques. In the interest of space and clarity, we do not include the results of the other classifiers (i.e., Linear SVM, Gaussian Naïve Bayes, and Random Forests). The use of POS and NER features in our model has no effect on the performance. GloVe embeddings achieve performance close to our method with F_1 scores 2% lower and substantially lower recall scores. Despite the overlap in precision scores between GloVe and our method, the overlap is minimal in terms of recall. Our CNC model also clearly outperforms the state-of-the-art method by ClaimBuster at 0.79 F_1 ; hence our method yields F_1 scores that improve ClaimBuster by over 5% in relative terms. ClaimBuster performs similarly to CNC in terms of precision, albeit with substantially lower recall scores. ClaimRank has the best precision scores across the board, but with the lower recall scores. CNC achieves a 6% relative improvement in F_1 -score over ClaimRank.

6 MULTI-CLASS CLASSIFICATION

We further test our classifier on the broader set of seven categories, using the subset of 4,080 sentences where there was enough agreement at this level of granularity. We train a multinomial logistic regression on the features from CNC, the best performing binary classifier.

¹²<https://github.com/pgencheva/claim-rank>.

¹³Note that ClaimRank is originally a ranking system that determines the check-worthiness of claims, so here, instead, we implemented several binary classifiers to make it work as a binary claim detection system.

Table 4. Results for the Claim/No Claim Experiments

Features	Classifier	P	R	F ₁	P-interval	R-interval
TF-IDF	LogReg	0.90	0.59	0.70	0.89–0.91	0.56–0.61
TF-IDF+num. preproc.	LogReg	0.91	0.59	0.70	0.90–0.92	0.56–0.61
Word2Vec	SVM	0.85	0.75	0.78	0.84–0.86	0.73–0.77
GloVe	LogReg	0.89	0.76	0.81	0.88–0.90	0.74–0.78
GloVe+PCA	LogReg	0.89	0.75	0.81	0.88–0.90	0.73–0.77
CB	LogReg	0.90	0.59	0.70	0.89–0.91	0.56–0.61
CB +POS	LogReg	0.88	0.68	0.76	0.86–0.89	0.66–0.71
CB +NER	LogReg	0.88	0.60	0.71	0.87–0.89	0.58–0.63
CB +POS+NER	LogReg	0.87	0.71	0.78	0.86–0.88	0.68–0.73
CB	SVM	0.84	0.70	0.76	0.83–0.85	0.69–0.73
CB +POS	SVM	0.86	0.74	0.79	0.85–0.87	0.72–0.76
CB +NER	SVM	0.84	0.71	0.77	0.83–0.85	0.69–0.73
CB +POS+NER	SVM	0.86	0.75	0.79	0.85–0.87	0.73–0.77
ClaimRank	LogReg	0.93	0.65	0.77	0.92–0.94	0.63–0.67
ClaimRank	SVM	0.93	0.53	0.67	0.92–0.94	0.51–0.55
ClaimRank	FNN	0.89	0.61	0.72	0.87–0.91	0.58–0.62
CNC	LogReg	0.88	0.80	0.83	0.87–0.89	0.78–0.82
CNC+POS	LogReg	0.88	0.80	0.83	0.87–0.89	0.78–0.82
CNC+NER	LogReg	0.88	0.80	0.83	0.87–0.89	0.78–0.82
CNC+POS+NER	LogReg	0.88	0.80	0.83	0.87–0.89	0.78–0.82

CB: ClaimBuster. The 95% confidence intervals are from binomial distribution adapting [22].

Table 5. Multi-class Classifier Performance CNC Model

Class	P	R	F1	N
Not a claim	0.77	0.90	0.83	2,235
Other type of claim	0.59	0.55	0.57	952
Quantity (past/present)	0.80	0.79	0.79	403
Prediction	0.60	0.27	0.37	181
Personal experience	0.72	0.39	0.50	124
Correlation/causation	0.50	0.13	0.21	107
Current laws/rules	0.27	0.04	0.07	78
microavg / total	0.71	0.73	0.70	4,080
macroavg / total	0.61	0.44	0.48	4,080

CNC model.

6.1 Analysis of Results

Table 5 shows the results for the multi-class classification experiments. These results reaffirm our expectations that, beyond the binary classification of claims and non-claims, classification at a finer granularity becomes more challenging. This is especially true for the categories with the smallest number of instances, such as “Current laws” or “Correlation or causation.” We achieve low F1 score for these categories, however the small number of instances may have a significant impact on this. Looking at bigger classes, “Quantity” (relatively easy to

Fig. 2. Confusion matrix for the multi-class classifier.

identify by looking for numbers and quantitative words) and “Not a claim” (the most popular category) yield the best F_1 -scores. The results for “Other type of claim” are also good, which only tend to be confused with “Not a claim” – this is in line with human annotator disagreement in Figure 1. Overall results are reasonably good when we measure with a microaveraged F1 score of 0.70, however it shows significant room for improvement when we measure it by macroaveraged F1 score of 0.48. We aim to expand our dataset in the near future to circumvent these issues.

Figure 2 shows a confusion matrix of the predictions made by the multi-class classifier. This again corroborates that “Quantity” is the easiest category to predict, with very few mistakes deviating to that category. Categories where more mistakes occur include “Not a claim” and “Other type of claim,” which the classifier has a tendency to mispredict when the actual label is a different one. Collection of more annotated sentences is expected to yield improved performance on these categories, however the binary classifier is for now more accurate and realistically usable.

7 DEPLOYMENT AND IMPACT

Model deployment has impacted Full Fact and the crowdsourcing exercise has educated the volunteer community.

Full Fact’s factcheckers are provided with a UI that shows a live feed of transcripts from television in a tool called “Live,” which aids live factchecking (see Figure 3). When the claim detection model identifies a claim in

Fig. 3. Claims highlighted in bold in the “Live” factchecking tool during Prime Minister’s Questions on 12 Sept 2018. The transcript errors are from the closed captions broadcast.

these sentences, it highlights it in bold. In addition to this, the factcheckers have the ability to manually highlight (in yellow) any claim they believe to be of interest. The integration of claim detection model has provided several benefits. It has saved time—factcheckers can quickly skim through the text looking for claims, instead of reading the entire text. (Though, to ensure 100% recall there are still designated people who watch the entire programme or read the entire transcript.) As a consequence, some factcheckers have started skimming just the claims in the transcripts in areas that are outside of their immediate domain. Going through the entire transcript was not viable for them prior to this automation.

We analysed four live factchecking sessions and noted that all claims manually highlighted in yellow had also been detected by the model, demonstrating the high recall. The precision however is harder to measure during user deployment. The model does detect claims that are not highlighted in yellow, but this is to be expected. The reason for this is that factcheckers rely on their domain expertise and awareness of current affairs to decide which claims are of specific interest to them—as opposed to highlighting all claims that they notice. These considerations are impossible to encode in a model as discussed in our claim definition section.

A factchecker gave the following feedback on the system: “*Claim detection is very useful after I have finished live factchecking a show and reviewing it to decide what to write a longer piece about. I no longer have to read the whole transcript, just the highlighted bits.*”

Outside of the Full Fact factcheckers, there has also been anecdotal evidence of impact on the volunteer community. The annotation exercise has been educational for them. They became more scrupulous media consumers. This was also a notable side product in other crowdsourced factchecking initiatives such as TruthSquad.¹⁴

8 CONCLUSION

Through leveraging the factcheckers at Full Fact, and through academia-industry collaboration, we have developed the first annotation schema for claim detection informed by experts. This has enabled us to create an annotated dataset made of sentences extracted from transcripts of political TV shows. We have introduced and

¹⁴“Crowdsourced Fact-Checking? What We Learned from Truthsquad” 2010 <https://perma.cc/J8AS-YU8E>.

tested a classifier that leverages universal sentence representations that, with an F1 score of 0.83, outperforms a range of baseline classifiers, including a well-known method by ClaimBuster, by over 5% in relative terms. While we achieve competitive results for binary claim classification, there is room for improvement when we need finer granularity of classification into the seven categories in the annotation schema.

Our plans for future work include expanding our dataset to other languages and collaborating with other factchecking organisations. We also wish to collect more data using the same annotation schema from other news sources, such as social media, digital and print outlets. Another small potential improvement is using the counts of first-/second-person pronouns to detect “Personal Experience” category, as proven useful in Reference [30].

ACKNOWLEDGMENTS

We thank the 80 volunteers who participated in the annotation task, in particular, Andreas Sampson Geroski. Also, thanks to Joseph O’Leary, Amy Sippitt, Will Moy, Ed Ingold, Vigdis Hanto, Mats Tostrup, Heri Ramampiaro, Jari Bakken, Benj Petitt, Michal Lopuszynski, Ines Montani, Gael Varoquaux, and David Corney.

REFERENCES

- [1] ACM. 2018. ACM Code of Ethics. Retrieved from <https://www.acm.org/code-of-ethics>.
- [2] Sarah A. Alkhodair, Steven H. H. Ding, Benjamin C. M. Fung, and Junqiang Liu. 2020. Detecting breaking news rumors of emerging topics in social media. *Info. Process. Manage.* 57, 2 (2020), 102018.
- [3] Mevan Babakar and Will Moy. 2016. *The State of Automated Factchecking*. Technical Report. Full Fact, London, UK.
- [4] Alberto Barrón-Cedeno, Tamer Elsayed, Reem Suwaileh, Lluís Màrquez, Pepa Atanasova, Wajdi Zaghouani, Spas Kyuchukov, Giovanni Da San Martino, and Preslav Nakov. 2018. Overview of the CLEF-2018 CheckThat! Lab on automatic identification and verification of political claims. Task 2: Factuality. In *Proceedings of the Conference and Labs of the Evaluation Forum (CLEF’18)*.
- [5] Yixuan Cao, Hongwei Li, Ping Luo, and Jiaquan Yao. 2018. Toward automatic numerical cross-checking: Extracting formulas from text. In *Proceedings of the World Wide Web Conference*. International World Wide Web Conferences Steering Committee, 1795–1804.
- [6] Giovanni Luca Ciampaglia. 2018. The digital misinformation pipeline. In *Positive Learning in the Age of Information*. Springer, 413–421.
- [7] Giovanni Luca Ciampaglia, Prashant Shiralkar, Luis M. Rocha, Johan Bollen, Filippo Menczer, and Alessandro Flammini. 2015. Computational fact checking from knowledge networks. *PLoS One* 10, 6 (2015), e0128193.
- [8] Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 670–680.
- [9] D. J. Flynn, Brendan Nyhan, and Jason Reifler. 2017. The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics. *Polit. Psychol.* 38, S1 (2017), 127–150.
- [10] Diane Francis and Mevan Babakar. 2016. HeroX Factchecking challenge. Retrieved from <https://www.herox.com/factcheck/5-practise-claims>.
- [11] Kim Fridkin, Patrick J. Kenney, and Amanda Wintersieck. 2015. Liar, liar, pants on fire: How fact-checking influences citizens’ reactions to negative advertising. *Polit. Commun.* 32, 1 (2015), 127–151.
- [12] Pepa Gencheva, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, and Ivan Koychev. 2017. A context-aware approach for detecting worth-checking claims in political debates. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*. INCOMA Ltd., 267–276.
- [13] Lucas Graves. 2018. *FACTSHEET: Understanding the Promise and Limits of Automated Fact-Checking*. Technical Report. Reuters Institute for the Study of Journalism, University of Oxford, Oxford, UK.
- [14] Lucas Graves, Brendan Nyhan, and Jason Reifler. 2016. Understanding innovations in journalistic practice: A field experiment examining motivations for fact-checking. *J. Commun.* 66, 1 (2016), 102–138.
- [15] Vigdis Hanto and Mats Tostrup. 2018. *Toward Automated Fake News Classification—On Building Collections for Claim Analysis Research*. Master’s thesis. NTNU.
- [16] Naemul 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*. ACM, 1803–1812.
- [17] Naemul Hassan, Chengkai Li, and Mark Tremayne. 2015. Detecting check-worthy factual claims in presidential debates. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*. ACM, 1835–1838.
- [18] Naemul Hassan, Gensheng Zhang, Fatma Arslan, Josue Caraballo, Damian Jimenez, Siddhant Gawsane, Shohedul Hasan, Minumol Joseph, Aaditya Kulkarni, Anil Kumar Nayak, et al. 2017. ClaimBuster: The first-ever end-to-end fact-checking system. *Proc. VLDB Endow.* 10, 12 (2017), 1945–1948.

- [19] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.* 9, 8 (1997), 1735–1780.
- [20] Viet-Phi Huynh and Paolo Papotti. 2018. Toward a benchmark for fact checking with knowledge bases. In *Companion of the Web Conference 2018*. International World Wide Web Conferences Steering Committee, 1595–1598.
- [21] Israa Jaradat, Pepa Gencheva, Alberto Barrón-Cedeño, Lluís Màrquez, and Preslav Nakov. 2018. ClaimRank: Detecting check-worthy claims in Arabic and English. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- [22] Ron Kohavi. 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Montreal, Canada, 1137–1145.
- [23] Klaus Krippendorff. 1980. *Reliability*. Wiley Online Library.
- [24] Sejeong Kwon, Meeyoung Cha, and Kyomin Jung. 2017. Rumor detection over varying time windows. *PloS ONE* 12, 1 (2017), e0168344.
- [25] Ran Levy, Yonatan Bilu, Daniel Hershcovich, Ehud Aharoni, and Noam Slonim. 2014. Context dependent claim detection. In *Proceedings of the 25th International Conference on Computational Linguistics: Technical Papers (COLING'14)*. 1489–1500.
- [26] Yunfei Long, Qin Lu, Rong Xiang, Minglei Li, and Chu-Ren Huang. 2017. Fake news detection through multi-perspective speaker profiles. In *Proceedings of the 8th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, Vol. 2. 252–256.
- [27] Jing Ma, Wei Gao, Zhongyu Wei, Yueming Lu, and Kam-Fai Wong. 2015. Detect rumors using time series of social context information on microblogging websites. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. ACM, 1751–1754.
- [28] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. Retrieved from <https://arXiv:1301.3781>.
- [29] Preslav Nakov, Alberto Barrón-Cedeño, Tamer Elsayed, Reem Suwaileh, Lluís Màrquez, Wajdi Zaghouani, Pepa Gencheva, Spas Kyuchukov, and Giovanni Da San Martino. 2018. CLEF-2018 lab on automatic identification and verification of claims in political debates. In *Proceedings of the Conference and Labs of the Evaluation Forum (CLEF'18)*.
- [30] Joonsuk Park and Claire Cardie. 2014. Identifying appropriate support for propositions in online user comments. In *Proceedings of the 1st Workshop on Argumentation Mining*. 29–38.
- [31] Ayush Patwari, Dan Goldwasser, and Saurabh Bagchi. 2017. TATHYA: A multi-classifier system for detecting check-worthy statements in political debates. In *Proceedings of the ACM Conference on Information and Knowledge Management*. ACM, 2259–2262.
- [32] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.* 12 (2011), 2825–2830.
- [33] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'14)*. 1532–1543.
- [34] Radim Řehůřek and Petr Sojka. 2010. Software framework for topic modelling with large corpora. In *Proceedings of the LREC Workshop on New Challenges for NLP Frameworks*. ELRA, Valletta, Malta, 45–50. Retrieved from <http://is.muni.cz/publication/884893/en>.
- [35] Baoxu Shi and Tim Weninger. 2016. Discriminative predicate path mining for fact checking in knowledge graphs. *Knowl.-Based Syst.* 104 (2016), 123–133.
- [36] Prashant Shiralkar. 2017. *Computational Fact Checking by Mining Knowledge Graphs*. Ph.D. Dissertation. Indiana University.
- [37] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explor. Newslett.* 19, 1 (2017), 22–36.
- [38] James Thorne and Andreas Vlachos. 2018. Automated fact checking: Task formulations, methods and future directions. In *Proceedings of the 27th International Conference on Computational Linguistics*. 3346–3359.
- [39] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: A large-scale dataset for fact extraction and verification. Retrieved from <https://arXiv:1803.05355>.
- [40] Andreas Vlachos and Sebastian Riedel. 2014. Fact checking: Task definition and dataset construction. In *Proceedings of the ACL Workshop on Language Technologies and Computational Social Science*. 18–22.
- [41] 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*, Vol. 2. 422–426.
- [42] You Wu, Pankaj K. Agarwal, Chengkai Li, Jun Yang, and Cong Yu. 2014. Toward computational fact-checking. *Proc. VLDB Endow.* 7, 7 (2014), 589–600.
- [43] Fan Yang, Yang Liu, Xiaohui Yu, and Min Yang. 2012. Automatic detection of rumor on sina weibo. In *Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics*. ACM, 13.
- [44] Xichen Zhang and Ali A. Ghorbani. 2020. An overview of online fake news: Characterization, detection, and discussion. *Info. Process. Manage.* 57, 2 (2020), 102025.

- [45] Arkaitz Zubiaga. 2018. Learning class-specific word representations for early detection of hoaxes in social media. Retrieved from <https://arXiv:1801.07311>.
- [46] Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018. Detection and resolution of rumours in social media: A survey. *ACM Comput. Surveys* 51, 2 (2018), 32.
- [47] Chaoyuan Zuo, Ayla Ida Karakas, and Ritwik Banerjee. 2018. A hybrid recognition system for check-worthy claims using heuristics and supervised learning. In *Proceedings of the Conference and Labs of the Evaluation Forum (CLEF'18)*, Vol. 8995. CEUR-WS.org, 171–175.

Received October 2019; revised May 2020; accepted July 2020