

Understanding the Promise and Limits of Automated Fact-Checking

Author: Lucas Graves

The last year has seen growing attention among journalists, policymakers, and technology companies to the problem of finding effective, large-scale responses to online misinformation. The furore over so-called 'fake news' has exacerbated long-standing concerns about political lying and online rumours in a fragmented media environment, sharpening calls for technological solutions to what is often seen as a technological problem. This factsheet gives an overview of efforts to automatically police false political claims and misleading content online, highlighting central research challenges in this area as well as current initiatives involving professional fact-checkers, platform companies, and artificial intelligence researchers.

The influence of 'fake news' in different parts of the world remains poorly understood. Initial evidence from the US and Europe suggests that the share of online users who visit false news sites directly is quite limited, and that people exposed to these sites visit mainstream news sources far more (Allcott and Gentzkow 2017; Guess et al. 2018; Fletcher et al. 2018). However, the same studies indicate fabricated news stories may draw disproportionate attention on social networks, outperforming conventional news, and some partisans (e.g. Trump voters in the US) appear to be regular users of false news sites. Little is known about the dynamics by which individual viral stories may influence the opinions and behaviour of specific, targeted audiences around particular events or issues.

In the US and Europe, concern about commercially

or politically motivated misinformation online – in particular about mounting evidence of sophisticated, state-backed campaigns operating from Russia – has fuelled a vigorous debate over policy options. These include a raft of proposals to regulate platform companies like Facebook and Google in new ways, a question under review in the European Commission. Several countries, notably Germany, France, and Ireland, have passed or are considering legislation that penalises the distribution of false information.

These concerns have also drawn new attention to the potential of various automated fact-checking (AFC) technologies to combat false information online. However, deciding the truth of public claims and separating legitimate views from misinformation is difficult and often controversial work (see Graves 2016), challenges that carry over into AFC. Based on a review of current efforts and interviews with both fact-checkers and computer scientists working in this area, this survey of the AFC landscape finds that:

- Much of the terrain covered by human fact-checkers requires a kind of judgement and sensitivity to context that remains far out of reach for fully automated verification.
- Rapid progress is being made in automatic verification of a narrow range of simple factual claims for which authoritative data are available. Even here, though, AFC systems will require human supervision for the foreseeable future.
- Both researchers and practitioners agree that the real promise of AFC technologies for now

lies in tools to assist fact-checkers to identify and investigate claims, and to deliver their conclusions as effectively as possible.

- So far independent, nonprofit fact-checking organizations have led the way in developing and implementing AFC, with little activity from traditional media outlets.
- Some individual AFC tools have been built inexpensively by fact-checking groups. However, advancing capabilities and developing large-scale systems requires continuing support from foundations, universities, and platform companies.

Overview

AFC initiatives and research generally focus on one or more of three overlapping objectives: to spot false or questionable claims circulating online and in other media; to authoritatively verify claims or stories that are in doubt, or to facilitate their verification by journalists and members of the public; and to deliver corrections instantaneously, across different media, to audiences exposed to misinformation. End-to-end systems aim to address all three elements – identification, verification, and correction (see chart).

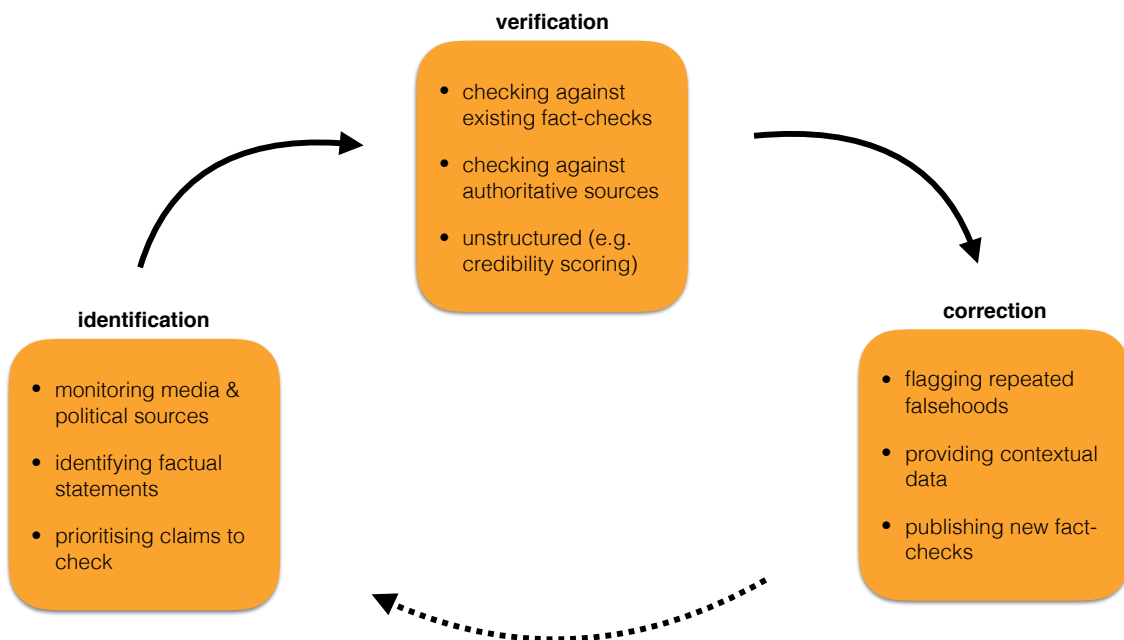
The first proposals to automate online fact-checking appeared nearly a decade ago. Over the last several years a growing research literature has embraced AFC

as an interesting problem in artificial intelligence, intersecting with practical experiments by fact-checkers.¹ Two recent programming competitions, the ‘Fast & Furious Fact Check Challenge’ and the ‘Fake News Challenge’, allowed research teams from around the world to test different AFC techniques on common problem sets.² Dr Andreas Vlachos, a lecturer at University of Sheffield, remarks on the increased attention:

We published our first paper in 2014. To us, apart from our interest in politics, we thought it was a great challenge for artificial intelligence to actually work on this problem. [But] for better or worse, Trump’s election increased the interest.

Meanwhile, real-world AFC initiatives have enjoyed a wave of additional funding in the last two years. Full Fact, a London-based fact-checking charity, began developing AFC tools in 2016 with a €50,000 grant from Google and recently announced £500,000 additional funding from the Omidyar Foundation and the Open Society Foundations. The Duke Reporters Lab, based at Duke University, received \$1.2m in late 2017 to launch the Tech & Check Cooperative, a hub for AFC projects, from the Knight Foundation, the Facebook Journalism Project, and the Craig Newmark Foundation. In January, Factmata, a London-based startup developing an AFC platform, announced \$1m in seed funding.

CORE ELEMENTS OF AUTOMATED FACT-CHECKING



¹ Useful research overviews are in Cohen et al. 2011; Hassan et al. 2017; Vlachos and Riedel 2014.

² See <https://www.herox.com/factcheck/guidelines>; <http://www.fakenewschallenge.org>

Approaches to AFC

Real-world AFC efforts begin with systems to monitor various forms of public discourse – speeches, debates, commentary, news reports, and so on – online and in traditional media. This is a difficult problem that may involve scraping transcripts and other material from media or political pages, monitoring live subtitle feeds, or using automatic transcription.³

Once monitoring is in place, the central research and design challenge revolves around the closely linked problems of identifying and verifying factual claims, explored below. A tension exists in that success in the first complicates the second, widening the range of claims that must be verified. In practice, AFC implementations constrain the problem by drawing on the work of human fact-checkers and/or by sharply limiting the kinds of claims being checked.

IDENTIFYING CLAIMS

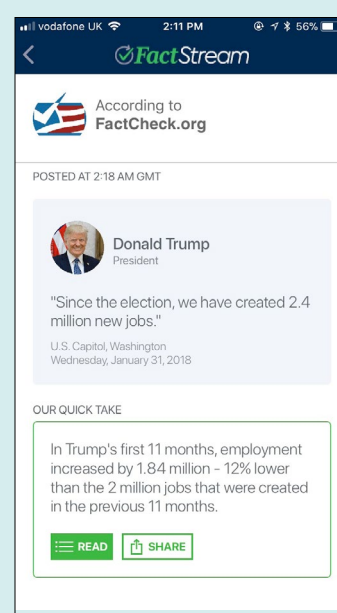
The greatest success in AFC research has come in the area of extracting discrete factual claims from a text such as a speech or an article. The most common approach relies on a combination of natural language processing and machine learning to identify and prioritise claims to be checked. For instance, ClaimBuster, an AFC platform developed at the University of Texas-Arlington, at a cost of roughly \$150,000 so far, trained on about 20,000 sentences

from past US presidential debates, classified by paid human coders, to learn to distinguish ‘check-worthy’ factual claims from opinions and boring statements (Hassan et al. 2017). In a test during a US primary debate in 2016, more than 70% of actual claims checked by fact-checkers at PolitiFact and CNN were among the top fifth of statements identified by ClaimBuster.⁴

A number of fact-checking outlets around the world have begun relying on software to help spot claims to check. In the US, for instance, the Duke Reporters Lab recently deployed a tool that uses ClaimBuster to deliver potentially interesting claims to fact-checkers at PolitiFact, FactCheck.org, the *Washington Post*, and the Associated Press (see the box). However, so far these systems can only identify simple declarative statements, missing implied claims or claims embedded in complex sentences which humans recognise easily. This is a particular challenge with conversational sources, like discussion programmes, in which people often use pronouns and refer back to earlier points.

It is also important to note that the ‘ground truth’ established by training algorithms on human work is neither universal nor permanent. For instance, ClaimBuster has been optimised to detect debate claims and does somewhat less well harvesting statements on Twitter. More broadly, the meaning and

One hub for automated fact-checking projects is the Duke Reporters Lab at Duke University. **Tech & Check Alerts**, in beta testing since early 2018, automatically generates a daily email newsletter neatly listing 15 promising political claims harvested from transcripts of CNN programming. The programme uses the ClaimBuster API but identifies the speaker and strips out statements by journalists; modules are being developed to pull claims from the Congressional Record, the California legislature, and the Facebook feeds of candidates in contested congressional races. Today the email goes out at 10 a.m. EST to PolitiFact, FactCheck.org, the *Washington Post*, and the Associated Press. Another new project, **FactStream**, offers live, ‘second-screen’ fact-checking of major political events via a mobile app. Fact-checkers use the platform to respond to dubious claims in real time, issuing either a capsule summary from a previous fact-check, or, for new claims, a ‘quick take’ adding context. Its first public test came during the 2018 State of the Union address, fed by alerts from PolitiFact, FactCheck.org, and the *Washington Post*; reportedly more than 3,000 people used the app at some point during the speech. (See: <https://www.poynter.org/news/review-live-fact-checking-takes-center-stage-state-union> and [http://www.niemanlab.org/2018/01/factstream-debuted-live-fact-checking-with-last-nights-sotu-howd-it-go.](http://www.niemanlab.org/2018/01/factstream-debuted-live-fact-checking-with-last-nights-sotu-howd-it-go/))



³ A seminal discussion of monitoring and other core AFC challenges which informs this report is in Babakar and Moy 2016.

⁴ See <https://www.poynter.org/news/holy-grail-computational-fact-checking-and-what-we-can-do-meantime>. A longer term comparison is reported in Hassan et al. 2017.

the importance of a particular statement may shift depending on historical or political context. Will Moy, director of Full Fact, gives the example of claims about the EU – polls show UK residents cared very little about the issue until the Brexit campaign brought it into the headlines.

Mevan Babakar, the groups' digital product manager, highlights the difference between knowing a factual statement has been made and understanding what is being claimed, a vital step in determining the importance of a question:

Identifying a factual statement is not easy but it is consistently possible. If you show me a sentence I can probably tell you if it's a claim. Understanding the meaning of a claim is hard – you need to understand the geography, what years it's referring to, and so on. Understanding how important a claim is is even harder, because it changes depending on who's doing the asking, and it changes depending on the political context, and that's something that's shifting all the time.

VERIFYING CLAIMS

The conclusions reached by professional fact-checking organizations often require the ability to understand context, exercise judgement, and synthesise evidence from multiple sources. Many claims don't lend themselves to simple true-or-false verdicts. But even seemingly straightforward statements that can be debunked by people – for instance, the now-infamous Brexit campaign claim that the UK would save £350m per week by leaving the European Union – present a thorny challenge for automated verification. Despite some progress no AFC system performs this reliably today. Echoing a widespread view among researchers in this area, Vlachos argues that expectations should remain modest:

The kind of fact-checking that PolitiFact does, or Full Fact, they do much more advanced things than the kinds of things I'm able to do today, or that I'm really able to do in the next 5 or 10 years at least. ...

But a typical fact-check has been reported to take a day. So if we're able to save time by automating some of the simpler aspects, that's where I see the role of automation here. I don't see it as a way of replacing humans, it's more like increasing productivity because we don't have enough fact-checking at the moment.

Two primary approaches to automatic verification are matching statements to previous fact-checks or consulting authoritative sources. A third family of techniques infers credibility from secondary signals.

CHECKING AGAINST PREVIOUS FACT CHECKS

The most effective approach to automatic verification today is to match statements against a library of claims already checked by one or more fact-checking organizations. This leaves difficult questions of judgement to human researchers, using automation to boost their reach and responsiveness when false claims resurface.

A number of fact-checking outlets are beginning to use this internally as a way to flag repeat offenders. For instance, Full Fact's in-house AFC platform constantly monitors an array of media outlets, as well as Prime Minister's Questions, for claims the group has already checked (see box page 6). Similarly, the Duke Reporters Lab expects to test a system within months which will match statements chosen by ClaimBuster against the libraries of FactCheck.org, PolitiFact, and other fact-checkers who use Share the Facts, a common tagging system that now covers more than 10,000 fact-checks.⁵ In this way, the software will be able to identify an interesting claim and point to related fact-checks, which may yield a 'quick hit' story, explains lab co-director Mark Stencel:

Our goals are to accelerate the reporting process but also accelerate the production of new fact-checks. ... This is our whole model, which is not to try to conquer all of the big problems of automated fact-checking all at once, but to break down the assorted challenges into solvable tasks that over time will add up to automated instantaneous fact-checking, at least in some instances.

Besides its limited scope, this method faces two obstacles. First, while NLP algorithms can reliably capture close variants of a statement, paraphrasing remains a substantial challenge. As a result, a trade-off exists between 'recall' and 'precision': matching more instances of a claim always comes at the expense of accuracy, potentially leading to false positives. To optimise the balance, Full Fact writes custom search queries for each claim it monitors (but is experimenting with machine learning to improve the process). Pablo Fernandez, in charge of AFC efforts at Argentina's Chequeado, argues that human gatekeeping will be required for the foreseeable future:

⁵ Share the Facts implements the ClaimReview schema, an open standard for coding the different components of a fact check, such as the claim and the verdict, in a machine-readable way. See <https://www.poynter.org/news/google-now-highlighting-fact-checks-search>

Chequeado, a fact-checking nonprofit based in Buenos Aires, began using a beta version of its own AFC platform in January 2018. Called **Chequeabot**, the system was developed in-house over the last year (at a cost of roughly \$20,000 so far) in part because Spanish-language NLP algorithms remain poorly developed. In its current version, the programme monitors presidential speeches and about 30 media outlets across Argentina to find claims to check; the group has published a handful of articles based on suggestions by Chequeabot, and constant feedback from the fact-checkers trains the algorithm to focus on statements that are both 'checkable' and interesting.

Another Chequeabot feature now in development matches statements against previous fact-checks, and in some cases against official statistics, in order to automatically generate a draft of a tweet or a post for a human fact-checker to review. The platform will be shared with other fact-checking organisations in South America, and with news organisations interested in political fact-checking.

Right now what we are trying to do is enhance the way fact-checkers work, because natural language processing is not that accurate, [especially] with things where you know there are a lot of grey areas. ... Right now we think we have to have a man in the middle.

Second, even subtle changes in the wording, timing, or context of a claim can make it more or less reasonable. A good example can be seen in the fine distinctions fact-checkers had to draw between various versions of the £350m Brexit claim, which were 'inaccurate to different degrees' depending on the precise wording.⁶ Even a perfectly accurate statistic can misinform without the proper context; Babakar offers the example of the UK murder rate, which appears to spike in 2003 because killings by a notorious serial killer were officially recorded that year.

CHECKING AGAINST AN AUTHORITATIVE SOURCE

A steeper challenge at the centre of current research is to verify claims against the same kinds of original information sources relied on by human fact-checkers. In theory, this has the potential to vastly expand the range of statements which can be checked automatically. But it requires that, having identified a discrete claim to check, the AFC system can recognise the kind of data called for, and that the data are available from an authoritative source in a form the software can use.

For AI researchers, the central problem is to parse statements in terms that make sense to a database. Vlachos says his own efforts do reasonably well with claims that directly name an entity, a property and a numerical value – say, 'Lesotho has a population of 2 million.' But AFC algorithms struggle with even straightforward 'single-predicate' claims that relate

multiple elements, like 'Lesotho is the smallest country in Africa.'

In practice, fully automatic verification today remains limited to experiments focused on a very narrow universe of mostly statistical claims. For instance, both Argentina's Chequeado and the UK's Full Fact are developing purpose-built AFC modules designed to match claims about specific public statistics, such as the unemployment or inflation rate, against the official figures. Both groups have campaigned to make more official statistics available as structured data which are friendlier to developers. It is worth noting that access to data tends to be more limited where fact-checking is needed most, in authoritarian environments with few independent media outlets (Graves and Cherubini 2016).

Similarly, the ClaimBuster platform includes a module, still in the early stages of development, which reformulates claims as a question for Wolfram Alpha, a general-interest structured knowledge base. This widens the set of available facts, but in practice only a tiny fraction of statements harvested from real political discourse can be tested this way. Chengkai Li, a professor at UT-Arlington and one of the creators of ClaimBuster, agrees that the most important bottleneck is caused by data:

The big challenge is the lack of data sources. Understanding the claim and formulating the query and sending the query to the source, that's one challenge. But another challenge is the lack of authoritative and comprehensive data. It's not just about the technical solutions, it's about the lack of data quality.

However, Li also notes that taking advantage

⁶ <https://fullfact.org/europe/foreign-secretary-and-uk-statistics-authority-350-million-explained>

of structured data sources will require greater sophistication in understanding claims. One approach being tested by his lab is to build up a taxonomy of different kinds of claims, with input from professional fact-checkers, to guide how statements are parsed. Li gives the example of the statement that the United States has 'the highest rate of childhood poverty of any major country on Earth'. Defining a class of claims about 'ranking' would alert the AFC algorithm to look for specific elements like the basis of comparison (child poverty rates), the comparison set (major countries), and so on.

UNSTRUCTURED AND NETWORK APPROACHES

Another avenue of research involves less structured or "non-reference" approaches to verification (Babakar and Moy 2016). Rather than looking up a specific authoritative reference, these methods search more widely and may rely on a variety of content- or network-related signals to make inferences about the likely truthfulness of a claim.⁷

For instance, Vlachos explains, a way to test the claim that 'Lesotho is the smallest country in Africa' without logically interpreting it is to search for similar language across a large textual source, or across the entire Web. In experiments using Wikipedia as a trusted source and a dataset of 125,000 claims, for example, a team led by one of his students can predict correctly whether

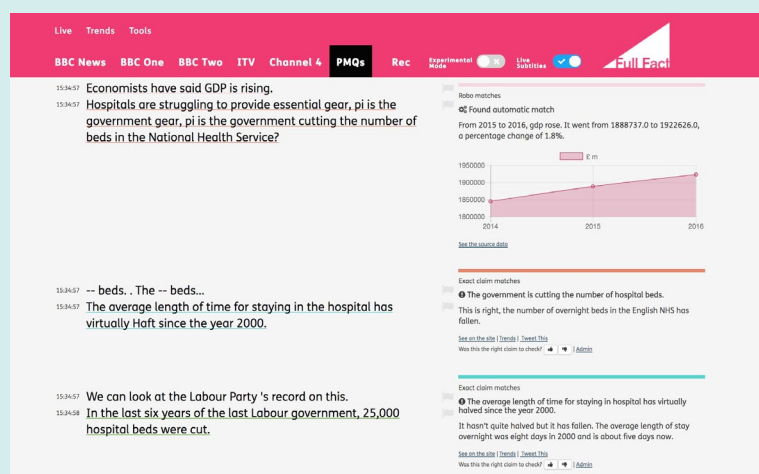
a single-predicate claim is supported or refuted (or whether there is not enough evidence) about 25% of the time (Thorne et al. 2018).

A crucial element in strengthening such approaches, and one which can also be used to assist human fact-checkers, is stance detection: determining whether a particular document supports the claim in question (see Ferreira and Vlachos 2016). The 'Fake News Challenge' concluded in late 2017 let computer scientists compare stance detection algorithms using a common library of real-world rumours and news reports from a rumour-tracking project run by journalists. A challenge scheduled for October 2018 will test these methods against more structured AFC techniques in delivering final verdicts about claims.⁸

Other research has focused on interpreting a variety of signals related to content or social context that may speak to credibility. These range from stylistic features, like the kind of language used in a social media post or a supposed news report, to clues based on the network position of a source (the sort of information Google uses to rank search results) or the way a particular claim or link propagates across the internet. (A useful overview is in Shu et al. 2017.)

Such probabilistic approaches draw on adjacent areas of AI research, like rumour detection, which shift the

The UK-based fact-checking charity Full Fact plans to release two AFC tools in 2018, developed with more than £500,000 in foundation support and meant for adoption by fact-checkers around the world. Both tools rely on a monitoring infrastructure that constantly scans major newspapers as well as broadcast news and parliamentary sources, using subtitle feeds as well as text to speech conversion. Designed to help the fact-checkers respond quickly to misinformation, **Live** uses human-tailored searches to scan these sources in real time for statements relating to claims which Full Fact has already investigated. An exact match automatically pulls up the relevant conclusion; in other cases, Live tries to display contextual information such as official statistics. **Trends** is based on the same engine but takes a longer view, offering claim-by-claim reports on the career of individual falsehoods. The purpose is to show fact-checkers whether their efforts have made a difference and help them to target their interventions more strategically.



⁷ One recent paper in this area argued, 'In other words, the important and complex human task of fact checking can be effectively reduced to a simple network analysis problem, which is easy to solve computationally' (Ciampaglia et al. 2015).

⁸ See <https://sheffieldnlp.github.io/fever>

problem from determining veracity to scoring reliability. This can resemble the kinds of inferences platform companies make in surfacing promising material and ‘down ranking’ sites or posts associated with problematic sources. In fact, some efforts piggyback on the complex language- and network-analysis capabilities of Google and Bing, using them as inputs to other AFC algorithms (see e.g. Karadzhev et al. 2017).

However, both researchers and practitioners argue that source credibility cannot be a substitute for assessing the factual accuracy of individual statements. One problem is that reliable sources make mistakes. As Vlachos indicates,

The most dangerous misinformation for each of us comes from the sources we trust. Philosophically, I don't want my model's decisions to be affected by the source, even though the source matters. I'm not saying one should never look at it, but we should also have models that ignore that part. Because everybody says incorrect things.

This points to a wider tension in the push for effective large-scale measures to counteract the spread of online disinformation: The impulse to promote trusted institutional sources can threaten pluralism and diversity in online discourse. Babakar, of Full Fact, notes that a mistake from an organization like the Office of National Statistics can do a lot of damage precisely because it is so trusted. She continues,

By upgrading certain sources we are implicitly downgrading others. ... There are cases where a minority publication may be more credible than a national newspaper, for example. My main question with credibility scores is who might you be unintentionally silencing and are their voices actually vital to the debate?

Discussion

This factsheet has offered an overview of the landscape of automated fact-checking initiatives and research. It documents rapidly growing activity in this area from both academic researchers and professional fact-checking organizations, as well as the consensus within both groups that fully automated fact-checking remains a distant goal. The most promising developments today are in AFC tools that help fact-checkers to respond more quickly and effectively to political lies, online rumours, and other forms of misinformation.

Real-world AFC tools are developing rapidly. Supported by foundations, platform companies, and other charitable sources, a handful of fact-checking organizations on different continents have emerged as hubs for developing and implementing automation technologies for the wider global community of political fact-checkers based in news outlets, universities, and civil-society groups. Several outlets are now using automation in a supporting role to help find interesting and important political claims to check. Progress is also being made in matching some claims against previous work; despite limits, this year will see the official launch of new tools to track where false claims are being repeated and to automatically bring up related fact-checks or other relevant information to help fact-checkers intervene quickly.

However, the potential for automated responses to online misinformation that work at scale and don't require human supervision remains sharply limited today. Researchers are exploring both more and less structured approaches to automated verification, reflecting wider divisions in the AI landscape. Despite progress, AFC techniques which emulate humans in comprehending the elements of a claim and checking them against authoritative references are constrained by both the current science and by a lack of data; researchers suggest one path forward is to build up recognition of different kinds of claims in a gradual and ad hoc way. Another family of techniques assesses the quality of information based on a complex array of network signals, making judgements about a message or its source in ways that may be opaque to humans. It is unclear how effective various unstructured approaches will prove in responding to different kinds of misinformation, for instance, false claims from political elites as opposed to viral online rumours. These approaches may also be vulnerable to mistakes from reputable sources, and raise difficult questions about protecting open and diverse political expression online.

AFC has been an area of unusually close collaboration between researchers and practitioners. Further progress will depend mainly on two factors: continued financial support for both basic research and real-world experiments, and progress by government and civil society groups in establishing open data standards. Traditional news organizations, whose fact-checking initiatives have larger reach and greater scale, also have much to contribute — and potentially to gain — by becoming more active in this arena.

References

- Allcott, Hunt, and Matthew Gentzkow (2017). 'Social Media and Fake News in the 2016 Election', *Journal of Economic Perspectives*, 31(2) (May): 211–36. <https://doi.org/10.1257/jep.31.2.211>.
- Babakar, Mevan, and Will Moy (2016). *The State of Automated Factchecking*. London: Full Fact, Aug. <https://fullfact.org/blog/2016/aug/automated-factchecking/>.
- Ciampaglia, Giovanni Luca, Prashant Shiralkar, Luis M. Rocha, Johan Bollen, Filippo Menczer, and Alessandro Flammini (2015). 'Computational Fact Checking from Knowledge Networks', *PLOS ONE*, 10(6) (17 June): e0128193. <https://doi.org/10.1371/journal.pone.0128193>.
- Cohen, Sarah, Chengkai Li, Jun Yang, and Cong Yu (2011). 'Computational Journalism: A Call to Arms to Database Researchers', in *Proceedings of CIDR 2011: 5th Biennial Conference on Innovative Data Systems Research*. Asilomar, CA: CIDR, 148–51. <http://web.eecs.umich.edu/~congy/work/cidr11.pdf>.
- Ferreira, W., and A. 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*. San Diego, CA: ACL, 1163–1168. <http://www.aclweb.org/anthology/N16-1138>.
- Fletcher, Richard, Alessio Cornia, Lucas Graves, and Rasmus Kleis Nielsen (2018). *Measuring the Reach of 'Fake News' and Online Disinformation in Europe*. Oxford: Reuters Institute for the Study of Journalism, Feb. <https://reutersinstitute.politics.ox.ac.uk/our-research/measuring-reach-fake-news-and-online-disinformation-europe>.
- Graves, Lucas (2016). *Deciding What's True: The Rise of Political Fact-Checking in American Journalism*. New York: Columbia University Press.
- Graves, Lucas and Federica Cherubini (2016). *The Rise of Fact-Checking Sites in Europe*. Oxford: Reuters Institute for the Study of Journalism, Nov. <https://reutersinstitute.politics.ox.ac.uk/our-research/rise-fact-checking-sites-europe>.
- Guess, Andrew, Brendan Nyhan, and Jason Reifler (2018). *Selective Exposure to Misinformation: Evidence from the Consumption of Fake News during the 2016 US Presidential Campaign*. <http://www.dartmouth.edu/~nyhan/fake-news-2016.pdf>.
- Hassan, Naeemul, 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*. Halifax, Canada: ACM, 1803–12. <https://dl.acm.org/citation.cfm?id=3098131>.
- Karadzhov, Georgi, Preslav Nakov, Lluís Marquez, Alberto Barron-Cedeno, and Ivan Koychev (2017). *Fully Automated Fact Checking Using External Sources*. 1 Oct. <http://arxiv.org/abs/1710.00341>.
- Shu, Kai, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu (2017). 'Fake News Detection on Social Media: A Data Mining Perspective', *SIGKDD Explorations Newsletter*, 19(1) (Sept.): 22–36. <https://doi.org/10.1145/3137597.3137600>.
- Thorne, James, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal (2018). 'FEVER: a large-scale dataset for Fact Extraction and VERification', in *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. New Orleans: ACM.
- Vlachos, Andreas, 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*. Baltimore, MD: ACL, 18–22. <http://www.aclweb.org/anthology/W14-2508>.

Acknowledgments

Lauren Jackson contributed valuable research assistance to this factsheet. This research was supported by Google as part of the Digital News Initiative (CTR00220).

ABOUT THE AUTHOR

Lucas Graves is a Senior Research Fellow at the Reuters Institute for the Study of Journalism at the University of Oxford.