The human touch in automated fact-checking

How people can help algorithms expand the production of accountability journalism

Bill Adair[†]
Duke Reporters' Lab
Duke University
Durham, NC USA
bill.adair@duke.edu

Mark Stencel
Duke Reporters' Lab
Duke University
Durham, NC USA
mark.stencel@duke.edu

Cathy Clabby
Duke Reporters' Lab
Duke University
Durham, NC USA
catherine.clabby@duke.edu

Chengkai Li
Computer Science & Engineering
University of Texas, Arlington
Arlington, TX USA
cli@cse.uta.edu

ABSTRACT

Work with human fact-checkers during the 2018 midterm elections has proved that algorithms can effectively save journalists considerable time – accelerating their reporting process by automatically sifting through large quantities of online content and identifying political statements whose accuracy deserves scrutiny. However, human intervention is still required.

Based on the experience of Duke Reporters' Lab researchers in 2018, having a journalist selectively edit or annotate algorithmic news choices for a larger group of reporters or editors may both preserve much of the efficiency of this kind of automated reporting while increasing its impact on news coverage.

Bill Adair, Mark Stencel, Cathy Clabby, Chengkai Li. 2018. Fact-checking automation with a human touch. How people can help algorithms expand the production of accountability journalism. ***Computation + Journalism 2019***

Introduction

The digital transformation of the media business has led to job cuts in many newsrooms. Despite stable numbers in radio news and cable TV, and even growth in local television and digital news, the loss of 32,000 newspaper jobs in less than a decade has thinned the ranks of working journalists by nearly a quarter – from 114,000 in 2008 to 88,000 in 2017. [1] With constrained resources, old and new media organizations are increasingly looking for opportunities to use automation for many reporting tasks.

This has created a growing need for automated tools that can perform tasks that are repetitive or time-consuming for journalists. In 2015, a team at the University of Texas at Arlington led by Chengkai Li developed the ClaimBuster algorithm to automate the process of finding factual claims in a political transcript or document. That is a tedious but essential process for the journalists who work in the dozens of news organizations around the world that regularly review the accuracy of statements made by politicians, government officials and other public figures.[2] ClaimBuster analyzes each sentence in the text and assigns a score from 0 to 1 to predict the likelihood that it is a significant claim that a fact-checker would want to examine in a published article. [3]

The Duke Tech & Check Cooperative, a two-year project at the Duke Reporters' Lab to better automate fact-checking, began using ClaimBuster in January 2018 to provide daily alerts of promising claims to check. The initial alerts the lab sent were automated emails to journalists at news organizations such as The Washington Post, the Associated Press, PolitiFact, The New York Times and FactCheck.org.

To find the claims, the Lab's student developers built a system for scraping and parsing content from websites and selected social media feeds. The system submits statements to ClaimBuster's application programming interface (API) [4] for scoring and then automatically sends emails to the fact-checkers listing the top sentences from each source with the highest ClaimBuster scores. Student developers were also able to deliver the results automatically to one national fact-checking team using Slack, a multi-platform communication tool used in many newsrooms.

[*] Bill Adair, Mark Stencel, Cathy Clabby, Chengkai Li. 2018. Fact-checking automation with a human touch. How people can help algorithms expand the production of accountability journalism. Computation + Journalism 2019

[†] Bill Adair, Knight Professor of the Practice of Journalism and Public Policy, Duke Reporters' Lab Duke University, Durham, NC USA bill.adair@duke.edu

During the 2018 midterm elections, the lab incorporated this method to send daily emails to national fact-checkers that included two lists – 25 top-scoring statements from the previous day's CNN program transcripts and another 25 statements from the Twitter feeds of selected candidates and state parties.

In the first three months of the project, fact-checkers at The Washington Post, PolitiFact and FactCheck.org published seven articles on claims they found in the auto-generated emails. But that number dwindled by spring of 2018. Only four more over seven months could be directly attributed to the email alerts, even though the number of alert recipients expanded. This decline in the number of fact-checks published from claims found by ClaimBuster suggested that journalists became less interested in the auto-generated emails over time – an observation confirmed in interviews with several top fact-checkers.

The decline happened to coincide with the launch of a new Reporters' Lab Tech & Check project for North Carolina fact-checkers.[5] This project also relied on ClaimBuster to find claims, but it dispatched a different type of email alert to the Lab's newsroom partners at The News & Observer in Raleigh. Instead of an auto-generated email, this project sent a curated list of suggested claims selected by a journalist who manages the North Carolina project. Some of the claims were recommended by ClaimBuster. Others were identified independently by student researchers. With the claims, the journalist explained why she thought what she was sending might make a good candidate for fact-checking, helped by her understanding of political issues in the state. This change provides us with an opportunity to compare the use of the auto-generated emails with more tailor-made alerts.

This comparison is based on a relatively short time period (the first North Carolina alerts were sent in early September – about seven weeks before this paper was written) and the analysis is affected by some additional personal contact between the Lab staff and the fact-checkers. To better explore the connections, we interviewed fact-checkers at FactCheck.org, PolitiFact and The Washington Post about the Tech & Check Alerts.

The usage numbers and the interviews indicate that while fact-checkers benefit from automated journalism, such automation will be most effective if it is first curated. As PolitiFact Editor Angie Holan said in an interview, "You know a tip from a human being is going to cut through all the clutter and be worth sending to us."

Automating ClaimBuster

Researchers at the University of Texas at Arlington began developing the ClaimBuster algorithm in December 2014, with advice and input from collaborators at Duke University, Google, and Stanford University.

The project's development was partially supported by NSF grant #1408928. Over time the project has received several other grants, including a Knight Prototype Fund award from the Knight Foundation (for prototype development), NSF grant #1565699 (for advancing research toward commercialization), NSF grant #1719054 (for research on end-to-end automation of fact-checking), and a subaward of the grant to the Tech & Check Cooperative from the Knight Foundation, Facebook and Craig Newmark (for developing various tools for professional fact-checkers).

The algorithm was derived from transcripts of U.S. presidential debates from 1960 to 2012. Students and other volunteers marked each sentence of the transcripts that seemed checkable, providing input for the machine-learning algorithm to derive models that can score sentences for their promise as check-worthy factual claims.

The Texas team's work generated some press attention, but little use by U.S. fact-checkers. A website allowed the press and public to submit text for scoring and an API provided the means to automatically submit content and receive ClaimBuster scores. A Slack[Bot?] app was developed to integrate ClaimBuster API with Slack. As a demonstration project, a ClaimBuster Twitter feed (https://twitter.com/claimbustertm) generated a stream of high-scoring tweets from a list of several hundred politically active accounts. But no fact-checkers incorporated this freely available tool into their workflow.

Researchers at the Reporters' Lab still suspected ClaimBuster could accelerate the fact-checkers' reporting process, in part based on informal experiments that compared ClaimBuster scores of political speeches with lists of statements that national fact-checkers wrote about. A student developer also was able to demonstrate that he could automatically produce meaningful results by scraping TV transcripts posted on the web by national media and submitted the text to the ClaimBuster API.

Usage of Tech & Check Alerts

With financial support from John S. and James L. Knight Foundation, the Facebook Journalism Project and the Craig Newmark Foundation, the Lab launched its Tech & Check Cooperative in 2017. Among its goals: deploying ClaimBuster as a daily tool for political fact-checkers across the United States as a stepping stone to a process that could automatically fact-check some political statements in real time. As noted above, the automatically generated daily email alerts included 25 statements taken from a day's worth of CNN transcripts and another 25 statements culled from roughly 1,200 to 1,300 tweets over the past day from selected, politically focused Twitter feeds. (Behind the scenes, student research developers worked on adding new feeds to the alert, including statements from the Congressional Record, Facebook posts and ads from political figures and organizations, and transcripts from other TV news programs. Each involved technical challenges we did not overcome in time to share with national fact-checkers, though we made enough headway to begin deploying those tools early in the 2020 election cycle – which unofficially commenced as this paper was being written in October 2018.)

Early on, the Lab had a clear sense that these daily CNN and Twitter alerts were being read. On at least four occasions between

January and June, Washington Post Fact Checker Glenn Kessler was the the first to notify us about a glitch in our parsing mechanism or email process. His average response time: 15 minutes. [6]

From January to August 2018, claims dispatched in these alerts led to 11 national fact checks. While a promising start, more than half of these fact checks – six of eleven – were published in the first several weeks of the program.

So what happened? Why did the pace slow? From more recent experiences, we have a clue. As it turns out, the robots may have needed a little human touch.

Un-Checked

Early on, emails sent to fact-checkers every morning included identical introductory language. ("Good morning, fact-checkers! This edition of our Tech & Check Alerts features claims that the Duke Reporters' Lab automatically scooped up and then prioritized using the ClaimBuster algorithm developed by our computer science partners at the University of Texas, Arlington.") The alerts did not try to disguise the fact they were automatically generated. A tagline said, "This automated email is part of an experimental alert service developed by the Tech & Check Cooperative at the Duke Reporters' Lab. It is strictly intended to help fact-checkers spot potentially newsworthy statements and claims. It is not for public distribution. No humans on the Tech & Check team reviewed or verified these statements or their attribution before this alert was sent."

National fact-checkers told us that they paid less attention to the automated daily alerts as the pace of their midterm election coverage increased. Eugene Kiely, director of FactCheck.org, said he reviews the alerts "pretty much every day" to see "who's being quoted and whether it's something we'd even be interested in." But he also told us he looked at the alerts most closely "when it's slow and we don't have anything particularly in the works." Glenn Kessler of The Washington Post Fact Checker gave a similar answer: "I have to admit that, because things have been so busy, I have not looked at it every day as I have in the past."

One factor that limited the impact was the signal-to-noise ratio in the daily alerts. On any given day, a few statements from either the CNN or Twitter lists would stand out as good fact-checking material. But fact-checkers had to read through 45 or so other statements to find them. Said PolitiFact editor Angie Drobnic Holan, "When we're busy, I tend not to open [the Tech & Check emails]."

Also, the CNN feed in the daily alerts was continuously polluted with statements made by the network's journalists. Reporters' Lab tools were unable to consistently screen out the journalists' names and statements, mostly because of transcription errors and outdated CNN staff lists on the cable network's website. "There are some people that are TV commentators or ex-party officials who are acting as political pundits" Kiely said. "We're not going to fact check them, because we focus on the president, the top officials in the administration, and congressional leaders..."

As for Twitter, the Lab's focus on gathering claims from state parties and selected candidates in close congressional races during the 2018 midterm elections made that part of the daily alerts less valuable to the fact-checkers. "I find that tweets have not been particularly that helpful, just because it's been a little too granular," Kessler told us. "It's a lot of local issues. It has to be a big deal for us to do an uber-local fact-check."

The fact-checkers we spoke to describe the impact of these accumulated distractions as an effect much like "banner blindness." That online phenomenon was first documented with that label in a 1998 usability study by Jan Panero Benway and David M. Lane of Rice University. They found that people, over time, were likely to ignore promotional messages on a web page, even when the content was customized and tailored to grab their attention. [7] Rather than an "uncanny valley," in which technology all-too-creepily imitates life, our valley was too canned. It was recognizably robotic – and easy to dismiss.

Holan, the PolitiFact editor, said that even though the alerts did some useful culling of possible claims, most in the email were still not usable. "You know when you click open the report that three-fourths of it is not going to be of interest." Her conclusion about the automated service: It's a helpful first screen of possible claims, but that "identifying good fact-checks for human audiences is still best done by humans."

The Human Touch

What's different about the human-curated NC Fact-Checking alerts?

Instead of simply forwarding all the potential fact-checking claims ClaimBuster identified to reporters with the same here-you-go greeting, the Reporters' Lab manager does some editing.

Each day, she reads ClaimBuster's 50 top claims ripped from the Twitter accounts of North Carolina candidates for US Congress, the state legislature and and other significant political players in the state. She deletes irrelevant claims, often artifacts of ClaimBuster's imperfect aim at finding newsy political claims. And she keeps bullseye claims that her experience as a journalists tells her might appeal to a fact-checking journalist looking for the next story. Some days she sends claims collected by students and the Reporters' Lab bots. Other days it's all ClaimBuster.

One mixed-content alert composed in September, for example, read like this:

"As Election Day nears, politicians are talking a lot about $$, specifically about how their opponents misuse it due to bad ideas or plain old avarice. The examples below were collected by our students and student-made bots."

The ClaimBuster content in that alert included this:

From student-made bots and ClaimBuster:

Rep. Ted Budd, R	U.S. House 13	ICYMI: It's true: Kathy Manning gave nearly $1 million to liberals like Nancy Pelosi and Hillary Clinton – but she certainly won't tell you that. #Retweet this ad if you stand with Ted. #tedbudd #nc13 #pelosiinsider #ncpol #NCGOP https://t.co/qrkI5UPpZF	Link
Steven Buccini, D	State House 59	Curious why my opponent didn't mention that this mailer was paid for by the tens of thousands of dollars he's taken from the greedy insurance, pharmaceutical, and hospital lobbies.	Link

Not every offering in this format, including the two above, gets selected for a fact-check. But some do. These two were dispatched one October morning and selected by a fact-checking reporter the same day:

"Hello fact-checkers,

ClaimBuster offers these new political claims on the amendments:

One NC Democrat says the voter ID amendment might require all of us to get a new voter ID after every election. Really?

Joseph Fowler III, D	State House 76	Food for thought. If we pass the voter ID amendment the NCGA could require everyone to get a special voter ID that would expire after each election. Be careful what you wish for. #gotv #ncga #ncpol https://t.co/uuVTwTM0fx	Link

The same candidate is trying to convince people who hunt and fish that the GOP doesn't have their interests at heart after all. True?

Joseph Fowler III, D	State House 76	It is once more time for hunters and fishermen to stand up for what we love. If this hunting and fishing amendment is adopted it will enable the NCGOP to control your NC Wildlife Endowment Fund. Vote against all 6 amendments. #ncgotv #ncwf #nixallsix https://t.co/es9cpyGr1U	Link

It's not surprising that every emailed pitch does not inspire a story. That never happens in newsrooms.

But a combination of human touch and human tools appears to improve engagement with these alerts. The key is a genuine voice with a changing message. The tools range from simple copy editing ability to a tad of analysis, abilities that reporters bring to their desks every day.

Conclusion

The use of ClaimBuster as part of the Duke Tech & Check Cooperative and the North Carolina Fact-Checking Project demonstrates that automated reporting tools can handle important journalism tasks that reduce editorial workloads. ClaimBuster has shown that its findings can be the basis of significant journalism. However, a human touch is needed to increase the effectiveness of algorithms that alert editors of possible political falsehoods.

That requires more experimentation, especially in the case of the daily Tech & Check Alerts that the Reporters' Lab automatically sends to national fact-checking partners. That might mean adapting a format that is more like the conversational, highly curated approach we use with the alerts sent to the PolitiFact North Carolina reporters who work at The News & Observer. Another possibility is a less stylized approach that involves having an editor who reviews and annotates ClaimBuster's daily suggestions, removing or heavily annotating statements that otherwise would not be of any interest to fact-checkers. The second approach would address what we heard from some journalists, who told us they would still prefer to see large numbers of claims each day – just without irrelevant and distracting clutter that ClaimBuster often flags now. As Kiely of FactCheck.org put it, "The more claims the better, always."

REFERENCES

[1] Based on Bureau of Labor Statistics data analyzed by the Pew Research Center. Elizabeth Grieco. 2018. Pew Research Center's FactTank blog. Newsroom employment dropped nearly a quarter in less than 10 years, with greatest decline at newspapers. Retrieved Oct. 31, 2018 from http://www.pewresearch.org/fact-tank/2018/07/30/newsroom-employment-dropped-nearly-a-quarter-in-less-than-10-years-with-greatest-decline-at-newspapers/

[2] The Duke Reporters' Lab has continuously tracked fact-checking projects around the world since 2015: https://reporterslab.org/fact-checking/ The International Fact Checking Network based at the Poynter Institute for Media Studies in St. Petersburg, Florida, also maintains a more selective list of more than 50 projects that comply with that organization's Code of Principles based on reviews submitted independent evaluators: https://ifcncodeofprinciples.poynter.org/

[3] The technical details of ClaimBuster algorithms and system are described in Hassan et al., "Toward Automated Fact-Checking: Detecting Check-worthy Factual Claims by ClaimBuster", KDD 2017, Hassan et al., "ClaimBuster:The First-ever End-to-end Fact-checking System", PVLDB 2017, and Jimenez et al., "An Empirical Study on Identifying Sentences with Salient Factual Statements", IJCNN 2018; The first article about ClaimBuster in the media was W. Gardner Selby, "Will robots take over fact-checking the information overload?" Austin American-Statesman, Sept. 23, 2016: https://www.statesman.com/news/20160923/will-robots-take-over-fact-checking-the-information-overload (accessed Oct. 22, 2018)

[4] ClaimBuster: Automated Live Fact-checking Retrieved Oct. 29, 2018 from https://idir.uta.edu/claimbuster

[5] Catherine Clabby. 2018. Duke Reporters' Lab: Reporters' Lab joins N&O, UNC Reese News Lab on major fact-checking project. Retrieved Oct. 31 from https://reporterslab.org/reporters-lab-joins-no-unc-reese-news-lab-on-major-fact-checking-project/

[6] Based on emails Kessler sent to the Tech & Check Cooperative team during the 2018 campaign on Jan. 19, (9 minutes); Jan 21 (15 minutes); March 9 (3 minutes); and June 8 (33 minutes). The half-hour delay on March 9 also may reflect Kessler's observation to us that the alerts were a more essential part of his daily workflow earlier in the year, when election news was less intense.

[7] JP Benway, DM Lane. Banner blindness: Web searchers often miss "obvious" links Itg Newsletter 1 (3), 1-22 Retrieved Oct. 31, 2018 from http://www.ruf.rice.edu/~lane/papers/banner_blindness.pdf