




MOTHERBOARD



LIAR

Teaching Machines to Detect Fake News Is Really Hard



DANIEL OBERHAUS

May 7 2017, 9:00am



Image: Gage Skidmore/Flickr

Researchers combined meta-data with over 10,000 examples of IRL fake news in a dataset to train machine learning algorithms how to



MOTHERBOARD

One of the most difficult parts of limiting the spread of fake news is that humans have a hard time filtering what is legit from what is bogus. But if humans can't tell the difference between what's real and fake, could machines do any better?

In an attempt to answer this question, UC Santa Barbara computer scientist William Wang created LIAR, the largest ever database of fake news in an effort to train machines to automatically detect deception.

Comprised of 12,836 examples of statements culled from a decade's worth of short statements from the Pulitzer prize-winning politifact.com, LIAR is an order of magnitude larger than any other fake news database that has been created in response to the 2016 election.

To help a machine learning algorithm understand the statements, each of the 12,836 quotes was tagged with information about the truthfulness, subject, context, speaker, state, party, and prior history of inaccurate statements, in addition to a "lengthy analysis report" for the statement. These tags then were integrated into the database as metadata for each quote.

Next it was time to train neural networks to see whether or not they could be taught to identify fake news. To train this neural network, the researchers used about 10,000 examples from the collected statements so the machine could learn what fake news looks like—that is, which words are subject might be more predictive of fake news.

After training the neural network on this large corpus of statements with metadata, it was then given another 1,000 statements from the quote corpus. But this time it didn't know in advance where the statement fell on the six point truth metric used on Politifact, where statements are judged as either 'pants-fire,' 'false,' 'barely-true,' 'half-true,' 'mostly-true,' and 'true.'

According to Wang, teaching a machine to judge on such a statement with this level of granularity is significantly more difficult than just having it judge on a binary of true or false. Still, their neural net was able to do pretty well after being trained on this corpus of about 10,000 statements.



MOTHERBOARD

deep learning on text data, our system can improve the accuracy performance to .27, which is actually a 35 percent improvement compared to random guessing."

While a 27 percent accuracy rate for a deception detector isn't going to singlehandedly put a stop to the proliferation of fake news any time soon, it underscores the difficulty of teaching a machine to recognize deception.

"I believe that this is just a preliminary study, there is tons of work that has to be done for accurate detection of fake news," he told me via email. "It's very difficult to identify fake news if a machine does not fact-check an external database. In the future, my student and I will be trying to design fact-check algorithms to look for external evidence in addition to context and metadata."

M

- MACHINES
- AI
- PULITZER
- FAKE NEWS
- MACHINE LEARNING
- NEURAL NETWORK


Watch This Next



A Smarter Gun



MOTHERBOARD



VICE