UC SANTA BARBARA

THE Current

July 29, 2019

Sonia Fernandez

Modeling Misinformation

We live in an era of too much information — an endless stream of status updates, memes, reposts, infographics, quotes and hashtags roll daily through our social media feeds, meant to express viewpoints, drum up solidarity, provide information, change minds or cause controversy.

Problem is, the average online browser/social media user doesn't have the time or wherewithal to investigate the legitimacy or provenance of everything that shows up on their feeds. And it is this vulnerability that less scrupulous content generators exploit to spread misinformation, with results that can range from a little egg-on-face embarrassment to downright life-changing or potentially deadly consequences.

For UC Santa Barbara computer scientist William Wang, this chaotic morass is fertile grounds for exploration. Wang believes that deep learning techniques, when deployed on the text and hyperlink network of online posts and news articles, can help us with some of the critical thinking heavy lifting. This concept lies at the heart of his three-year project "Dynamo: Dynamic Multichannel Modeling of Misinformation."

"So the question is, given a post, how would you be able to understand whether this is specifically misleading or if this is a genuine post," Wang said, "and, given the structure of the network, can you identify the spread of misinformation and how it is going to be different compared to standard or nonstandard articles?"

It’s a tall order, especially in the social media arena, which has leveled the playing field between legitimate, established news websites and questionable sites that do their best to look official, or appeal to a user’s emotions before they can step back and question the source of their information.

However, thanks to natural language processing — Wang’s specialty — the text in these posts and articles can be used to reveal information about their creators and propagators, such as their affiliations, ideologies and incentives for posting, as well as who their intended audience may be. The algorithm crawls through millions of news articles posted by anonymized users on platforms such as Twitter and Reddit and examines the articles’ titles, content and links. The purpose is to get a sense not only of the entities behind them, but also of their patterns of dissemination across the network.

“A lot of us take websites for granted and casually retweet or repost misinformation and that’s how it gets propagated, cascades and spreads virally,” Wang said. “Some of the most important questions we’re asking are: What are the patterns? What are the incentives?”

To find out, he and his team proposed a learning mechanism that susses out why certain stories get reposted or retweeted in addition to whether the content itself is true or false. Along the way, Wang said, they could figure out who is involved in the spread of the misinformation and what patterns might emerge in that process. Images will also become part of the dataset, he added.

Later on, the researchers plan to integrate other aspects of their work with misinformation, such as clickbait, which uses catchy, often sensational titles to lure readers into clicking a link that at best sends them to a dodgy website, or at worst, steals their information.

“Clickbait mainly is low-quality articles which can indeed contain a lot of misinformation and false information because they have to exaggerate,” Wang said. Together with computer science Ph.D. student Jiawei Wu, the team developed a method called “reinforced co-training,” which employs an efficient system of labeling a few hundred articles that are then used to train a machine learning classifier to label what it thinks may be clickbait in an enormous, million-story dataset.

“Then we take these newly labeled instances and retrain the classifier,” Wang said. “This iterative process allows us to collect more label data over time,” he added, which refines the accuracy of the tool.

Using artificial intelligence to understand and find patterns in the tidal wave of text we send each other every day would give us insight on how we, intentionally or unwittingly, propagate misinformation.

“That’s really the beauty of natural language processing and machine learning,” Wang said. “We have a huge amount of data in different formats, and the question is: How do you turn unstructured data into structured knowledge? That’s one of the goals of deep learning and of data science.”

Tags
Artificial Intelligence

About UC Santa Barbara

The University of California, Santa Barbara is a leading research institution that also provides a comprehensive liberal arts learning experience. Our academic community of faculty, students, and staff is characterized by a culture of interdisciplinary collaboration that is responsive to the needs of our multicultural and global society. All of this takes place within a living and learning environment like no other, as we draw inspiration from the beauty and resources of our extraordinary location at the edge of the Pacific Ocean.