EPISODE 10

[INTRODUCTION]

[0:00:00.4] JM: Bots on the internet can be malicious, helpful, and everything in between. A bot that responds to all of your tweets might call you a socialist, and that's definitely malicious. Google crawls the web to index Google search. That's definitely helpful. Social media marketing bots schedule 200 Twitter posts to go out throughout the day. That's either a little annoying or a little helpful depending on who you are.

Bots are being used to amplify political viewpoints. An amplified viewpoint can serve as a gravity well for likeminded individuals and help a sparsely supported political cause find its footing. Sometimes, that amplified viewpoint is completely fictional or un-falsifiable. Real people believe that Hilary Clinton is a lizard alien because they have seen that story shared by enough Twitter bots.

So-called fake news is a topic that has been discussed on so many other podcasts. What is not reported is the connection between link-bait and advertising fraud. When a botnet is able to make an article go viral, thousands of people organically click on the link to that article. That organic traffic is used to launder fake clicks. These bots that are spreading "fake news", they might be controlled by conspiratorial Russian hackers, but it doesn't actually have to be that complicated. Anyone who wants to make money from online advertising fraud is incentivized to make salacious media; whether it is real or fake.

Samuel Woolley is the Director of Research at Political Bots. He works with Jigsaw, which is a division of alphabet, formerly known as Google, that seeks to make the internet safer. In today's episode we talk about political bots, advertising fraud and the connection between the two. Also, Software Engineering Daily is having our third meet up, Wednesday, May 3rd at Galvanize in San Francisco. The theme of this meet up is fraud and risk in software. We will have great food, engaging speakers, and a friendly intellectual atmosphere.

© 2017 Software Engineering Daily

If you like this episode about ad fraud — Well, it's related to ad fraud, you'll definitely like the talk that Shalin Dhar is giving at this meet up. To find out more, go to softwareengineeringdaily.com/meetup. Let's get on to this episode.

[SPONSOR MESSAGE]

[0:02:34.4] JM: Spring is a season of growth and change. Have you been thinking you'd be happier at a new job? If you're dreaming about a new job and have been waiting for the right time to make a move, go to hire.com/sedaily today. Hired makes finding work enjoyable. Hired uses an algorithmic job-matching tool in combination with a talent advocate who will walk you through the process of finding a better job.

Maybe you want more flexible hours, or more money, or remote work? Maybe you work at Zillow, or Squarespace, or Postmates, or some of the other top technology companies that are desperately looking for engineers on Hired. You and your skills are in high demand. You listen to a software engineering podcast in your spare time, so you're clearly passionate about technology.

Check out hired.com/sedaily to get a special offer for Software Engineering Daily listeners. A $600 signing bonus from Hired when you find that great job that gives you the respect and the salary that you deserve as a talented engineer. I love Hired because it puts you in charge. Go to hired.com/sedaily, and thanks to Hired for being a continued long-running sponsor of Software Engineering Daily.

[INTERVIEW]

[0:04:01.0] JM: Samuel Woolley is the Direct of Research at Political Bots. Samuel, welcome to Software Engineering Daily.

[0:04:06.3] SW: Thank you for having me.

[0:04:07.2] JM: You are a part of a research team that is investigating the impact of computational propaganda. What is computational propaganda?

[0:04:15.6] SW: Computational propaganda is basically sort of the confluence of automation algorithms and big data used over social media and attempts to manipulate public opinion. A more simple way of putting that is to say that it is the way that different political actors from campaigns to hacking collectives use things like social media bots and attempts to drive up traffic, or to amplify certain perspectives.

[0:04:44.9] JM: Give a few examples.

[0:04:46.7] SW: Sure. During the U.S. Presidential Election, there was several different instances where bots were used surrounding a particular controversy to make different actors look bad. There was a case where a bunch of Twitter bots were launched to tell people to report Ted Cruz to the FCC for using robo-dialing techniques that were outlawed, which is a bots on bots situation.

In the days preceding the election, the U.S. Election, that is, our research team did a research that found that there was nearly — There's hundreds of thousands of bots Twitting in support of Donald Trump that Trump bots outnumbered Clinton bots at a rate of 5 to 1, although there was definitely some Clinton bots in there as well and that a lot of these accounts were being used to send out things like misleading news stories or memes that said, for instance, things like, "You can now vote by text, "trying to get one or the other side to mess up and not vote.

[0:05:56.7] JM: How long has computational propaganda been on the internet?

[0:06:01.2] SW: We have been tracking the use of propaganda over Twitter and Facebook. As long as both of those sides have been around, there's been some attempts to manipulate public opinion using automated techniques or software-driven techniques. That said, the use of bots for various means has existed since the internet was around and it's been around on IRC.

Something that's really important to point out is that the use of automated software to try to get people to do things comes from marketing and from commercial purposes. It kind of started there with spam, and email spam, and other things. Then, the use of bots on social media that actually look like real people is a more new thing. That kind of began around 2010, so building

out profiles that mimic real people or that built to send out specific links, or specific content around a topic, or just to drive up a hashtag's numbers, or to manufacture a hashtag. That's a newer thing.

[0:07:13.8] JM: When you take a hotly contested subject, like who we should elect for president, people have different interpretations of the facts. What's the difference between someone who is advocating an unusual interpretation of the facts and someone who is spreading computational propaganda?

[0:07:38.7] SW: There's a big difference. I don't think that this — It doesn't have to do with opinions about how facts are presented. It has a lot more to do with the rate at which content is being served to people over social media. It's sort of like — I use the phrase amplifying information, or another way that I've had it put to me by people who build these bots is megaphoning information.

While one person might use their Twitter account to send out any variety of a story they consider news, whether it's on the right, or the left, or in between. The difference with computational propaganda is that that person would then go and buy hundreds, if not thousands of proxy accounts on a site, like Twitter, or 10 or 20 accounts on a site like Facebook, and then use those accounts to generate a massive amount of traffic on that same story in an attempt to make it look more grassroots. When in fact, it's actually more astroturf.

[0:08:42.2] JM: We've had cloud computing for a while. We've had good scripting tools, high-level machine learning tools for a while. These are the tools that make bots so easy to scale. Why has it taken so long for these bots to start getting attention?

[0:09:02.8] SW: I think the simple answer is that political campaigns and political actors are about 10 years behind commercial actors and just software engineers in general. Conversations about how Howard Dean and then Barrack Obama made use of digital techniques where it's also kind of funny, because a lot of those techniques had been used in the years preceding those elections. It's the same thing with bots. It's just taken a while for public opinion to swing enough towards concern about information on social media and where it comes from. That, now, we're just starting to see an uptick.

There's something else as well. There's been several events; Brexit, the U.S. President Election. Now, elections coming up in France and Germany, where people are seeing a larger degree of bot-driven information sharing. There's a general fear that the content that is spread by these political bots has specific effects on outcomes. In other words, that there's potential that aspects of Brexit being pushed through, or of Donald Trump getting elected where effected by the spread of computational propaganda. Those big events have also lead to a lot more public awareness and a lot more concern.

[0:10:38.3] JM: Does a hosting provider like AWS police their servers for bots?

[0:10:43.6] SW: I think most savvy hosting providers do police servers for bots. Also, I think, going further, the social media companies themselves police for bots and they're doing a lot of work to try to get rid of malicious content on their platforms, but there is also a real danger for these companies in policing content. They've sold themselves as technology companies, and when they get into policing content, they're judging free speech. In America, free speech, it's an unalienable right. The other thing about it is, is that there's definite legal implications for media companies, for companies like — Sorry. For companies like Facebook and Twitter beginning to act like media companies, there will be a lot of repercussions.

[SPONSOR MESSAGE]

[0:11:46.4] JM: For years, when I started building a new app, I would use MongoDB. Now, I use MongoDB Atlas. MongoDB Atlas is the easiest way to use MongoDB in the cloud. It's never been easier to hit the ground running. MongoDB Atlas is the only database as a service from the engineers who built MongoDB. The dashboard is simple and intuitive, but it provides all the functionality that you need. The customer service is staffed by people who can respond to your technical questions about Mongo.

With continuous backup, VPC peering, monitoring, and security features, MongoDB Atlas gives you everything you need from MongoDB in an easy to use service, and you could forget about needing to patch your Mongo instances and keep it up-to-date, because Atlas automatically updates its version. Check out mongo.db.com/sedaily to get started with MongoDB Atlas and

get $10 credit for free. Even if you're already running MongoDB in the cloud, Atlas makes migrating your deployment from another cloud service provider trivial with its live import feature.

Get started with a free three node replica set, no credit cards required. As an exclusive offer for Software Engineering Daily listeners, use code "sedaily" for $10 credit when you're ready to scale up. Go to mongodb.com/sedaily to check it out, and thanks to MongoDB for being a repeat sponsor of Software Engineering Daily. It means a whole lot to us.

[INTERVIEW CONTINUED]

[0:13:46.4] JM: If I'm hearing what you're saying correctly, what concerns you is not the validity of the information that's being spread through Twitter, or Facebook, or other means. It's the amplification of that information out of proportion with the velocity of true public opinion, of humans. Is that accurate?

[0:14:14.2] SW: Yeah, I think that's accurate. I think that it's a much more subjective practice to police the content and the arguments made in content. Although I do think there are metrics for determining whether or not certain information is vetted properly or comes from empirical evidence.

Our project though is more concerned with one big thing, and the one big thing we're concerned with is the use of automation and spreading information. When we think about political bots, we're thinking about automated accounts that create manipulated information flows.

[0:15:05.8] JM: I remember my second, or third internship doing software engineering, I wrote automated tests to test how web application worked. That work was something that could have been classified as a bot, but the behavior that I was programming into the bot could have been a human. There was no difference between the mouse movements and the clicking and to the logging system, unless I had maybe a key logger or something could tell exactly how my mouse is moving around the screen. Even then, I could have easily recorded my own mouse movement and train a machine learning model to acted just like I do.

How do companies, whether we're talking about Twitter, or Amazon, how do they know who is a bot and who is a human?

[0:16:07.6] SW: That's a really good question. I think that there's a couple of different metrics that companies can look at to figure out that an account is either a bot or a human, and the three broad categories those metrics lie in are temporal signals; how often an account it tweeting, how regularly an account is tweeting. Is it on a very strict time schedule? If we start to parse the bot's behavior, can we tell that it is behaving in a robotic way, in a way that does not match up with human behavior?

Another way is network signals. Looking at who the bot tends to retweet, like, follow, who it has relationships with on social media? A lot of times bots exist on the tangents of a social network. They often are built in what's called a botnet to reproduce content from the other profiles in order to give the illusion of legitimacy, so to make it look as if the account has more followers that are real than it actually does, when in fact, all the followers are just other bots that have been built for that purpose. Network signals are very important.

Then, the last one is semantic signals. Thinking through that, it's what the content is. If the content is really rote, if it appears as the same content over and over and over again, or if the account is only sharing links, or if the account is only sharing content on one issue, or if when the account is engaged in conversation with a human user, if the conversation doesn't hold up for more than a few lines, then you can tell that, more than likely, that counts a bot.

Moreover, I think that when you look at these three signals lined up and aspects of these three signals, there's so many other different things that I haven't mentioned. When you line them all up, you can really start to identify accounts, because maybe they'll be able to get around the network signal, but they won't be able to get around the temporal, or the semantic signal. There's definite ways of catching bots.

[0:18:38.0] JM: Let' take all three of those. You've got the temporal, the network structure, and the semantic. The temporal; there are social media marketing managers who schedule 200 Twitter posts to go out throughout the day using buffer, or now you can use the functionality built in to Twitter. Temporal doesn't seem like a reliable way of doing it. I think more and more people

are getting to where they like to schedule lots of Twitter posts, because they're like, "Okay. I need to be hygienic on Twitter. I can't just post 200 things at once, so I'm just going to schedule it." I think the tools are headed towards the direction of democratizing that.

The network structure, if you've got bots that are just interacting with each other and they're not colliding with the real network as much, those bots wouldn't concern me as much, because who cares about those? They're just interacting with each other and maybe they're amplifying the signal among each other. Maybe they're confusing Twitter algorithms, but Twitter should be able to detect that they're on the margins of the network. Those probably don't matter. Yet, we are engaging with these bots. I'm not sure if the network structure — Basically, if the bots are interacting with humans, that is what is concerning.

You could imagine a scenario where you've got a bot that is interacting with a bunch of — I don't know. I'm not sure. The network structure one seems tough to me. I'll give you that. Maybe that's one plausible.

The semantic one, I've got people I know from high school where I see the information, or — Yeah, people I know from high school. Let's go with that. That's anonymzing. The stuff that they post is as idiotic and fact — Not fact-holding. It looks like stuff that a bot would post. I'm like, "This is so idiotic, and yet I know you're a person, because I went to high school with you."

The semantic analysis, I'm not sure I buy that either, because — I did a show with somebody — Robbie, from Automated Insights who — That company, they generate content from — You give them a stock stick, or history, or the numerical history of a baseball game, like, "This person hit a homerun at this point in the game," and the bot will write a detailed news article based on that numerical data. They've surveyed people, the people can't tell one from the other.

I guess where I'm going with this is none of those signals that you just described to me are plausible enough to me that you could actually use them to reliably detect bots on automated basis.

[0:21:42.0] SW: Right. I think what you've done is break them apart into singular things on their own. Again, I think that when combined all three together, tend to catch a lot more accounts

than any one on its own. Your analysis of talking through each one, you have fair points. What I'm saying here is that when all three are combined, there is a greater chance of finding automated accounts.

That said, it's an imperfect method. There are other ways that you can look into whether an account is a bot, or is software-driven or not. When people time their tweets, they're still using software to get tweets out there. The difference is that a lot of times, the content that they're generating is self-generated. It's not generated using a database. That's a little different.

A lot of the times, if someone is using if, this, then, that, or during the election, a lot of people are using Patriot Journalist Network, the content is created by means of the bots. The bots searches for the content and then provides it for the account. You can see using a few different techniques who the client is that — Or where the client is Tweeting from. The question is; are they tweeting from Twitter? Are they tweeting from Twitter for business? Are they tweeting from the API, or are the using another client to tweet an to automate their presence?

One of the things that we've done is unpack that and been able to say, "Well, we can see most of the time, this account is tweeting from if, this, then, that, so it's done 30,000 tweets from if, this, then, that. One time, in the last few days, it signed on to Twitter and done a few human tweets." That's one way that we've also looked at this. There's a few other things that people can do. I think that they get at points that you made about whether or not the interaction is the most important thing or not. I think I'd actually disagree with that. I wouldn't say that bots interacting with people is the most important aspect of what a bot can do.

I think that social media companies are particularly concerned with basically manufactured ideas or manufactured thought streams. Bots like that, there's a question of means and ends. You have the bot, which is your means, but then what are you attempting to have the bot do? Are you attempting the bot to have the bot manipulate public opinion directly? Talk to people and share articles amongst a community and attempt to get them to change the way they think? Are you doing something a bit more subtle with the goal of manipulating hashtag, creating hashtag? Basically, getting a hashtag to trend, or getting something to show up in the newsfeed at a higher rate, or something which is completely different. Are you trying to create confusion?

Is the goal of building your botnet to create confusion rather than to change people's perspective of facts?

I think that that latter thing; creating confusion, is something that we are seeing on a large scale in Eastern European countries, like Ukraine and Poland. We're seeing bots being used to basically cause confusion so that someone can fill that vacuum with some kind of solid action. We also saw that during the U.S. Election. We continue to see it right now in U.S. politics.

[0:25:16.4] JM: You described the ensembling as being important. You're not taking any of these three strategies for identifying bots versus humans atomically. You are ensembling them together. That is what is so valuable about Facebook, and Google, and perhaps Twitter as identity platforms, because you do so much through your Google accounts and through your Facebook account that even if you took the most sophisticated replay attack, which is you — Let's say somebody had a piece of malware on my computer that was recording all of my Google actions, all of my Facebook actions, everything. It was the Panopticon into my browser and they were copying everything that I could do.

I'm not sure that they would be able to train something to operate convincingly like me, or convincingly like a human in a way that Google would not be able to detect. Maybe they would. I'm not 100% sure of that, but I —

[0:26:25.9] SW: At least it would take a lot more time and effort to do something like that than it would do to the more simple stuff that we're talking about.

[0:26:34.6] JM: That's right. This is one of the things where I'm like, "God, I love Facebook and Google," and people who are afraid of surveillance, for example, "Okay. Yeah, sure. Surveillance is an issue." It's also an issue that I need to have some sort of way of verifying who I am in the digital realm and verifying others. If I'm engaging with somebody on Twitter, if somebody says, "Hey, I really didn't like your last podcast episode." I would prefer to know if that's a human, or a bot. Maybe someday I'll have bots that are actually listening and studying what I'm doing and they want to know the material.

Talk about some of the techniques that you're using to study political bots.

[0:27:26.9] SW: That's a really great question. We're using mixed research methods to study bots. We're not just using quantitative techniques, which is what we've been covering. We're also qualitative techniques. At the moment, we're doing a series of 10K studies in collaboration with Jigsaw. It's both the Oxford Internet Institute which is where I'm based and where the computational propaganda project is based. Also, getting some help from Jigsaw and Google to do research on 10 different countries and run a diagnostic on what the state of computational propaganda writ large looks like in those countries.

We're looking at Russia, China, the United States. Again, we're also looking at the Ukraine, and Poland. We're looking at Germany, because of the upcoming election. Then, we're also looking at Brazil, and Rwanda, and a couple of others. The goal is to basically say, "Here is what has been said," so to do content analysis of the media in each of those countries and look through all of the media reports that have said or alleged that there was automated attacks, or that there was computational spread of misleading information, then to do a series of interviews, at least 10, with experts on the ground. Also, and more importantly, I think, with the people who are building the bots.

We go into a county like Poland, or a country like the Ukraine, and we use the context that we have built over the last four years to get in contact with the people who are actually building bots on behalf of political campaigns and talk to them about what their tactics are. I think that human-centric way of doing research, of telling the stories of the actual software engineers and of what they've done, why they've done it, who they've done it for, and all of that, is actually a really important thing that has not been gotten that by a lot of computer science research.

That social science, that real time social science is pretty important. Then, the other thing that we do is collect a bunch of data around particular issues. Say, for instance, the German election, we just collected a bunch of data for the federal nomination process. The actual election is taking place later on. The parties basically put their candidates up and we are able to collect a ton of data surrounding that and analyze it and create an analysis of what's sorts of information were appearing, how likely it was the accounts were automated. Then, use all our qualitative and quantitative tactics for doing research to produce a holistic document. That's really the approach of our project at Oxford. That's how we do things. We don't only do the

quantitative, or the big data work. We don't only do social science work. We do a combination of both.

[0:30:22.4] JM: What is so terrifying about these bots is that they can muddy the waters of truth and, like you said, create confusion. Basically, the thesis that you described earlier was the goal of these bots is to create confusion so that maybe a political strong man or somebody can come in with hard action, because everybody is so confused, "What is the truth?"

How has your — I'm sure as — I totally love the approach of going in, boots on the ground, talking to these programs, "Hey, why are you making a bunch of bots?" "Hey, why are you writing a selenium script to generate thousands and thousands and thousands of tweets from a Twitter account which Hilary Clinton with fire in her eyes and red lighting all over the — Just these horrifying Twitter accounts. What is making you do this?"

As you're talking to these people, you must be getting such a nuanced view about what is going on. I'm sure your views are so much more nuanced than what we see in the — I hate to use the term, mainstream news articles. Tell me some of the nuance. Tell me some of the nuance. What are you finding out talking to people? Are the mainstream news article getting it right?

[0:31:56.5] SW: I hope that you're right, that we are a bit more nuanced. The goal of this work is to provide more context, more nuanced, and more — Honestly, more empirical evidence that this is going on and about how this is going on.

One of the main things, I would say, is that there's a real misconception that the user of bots to spread political propaganda only occurs when political campaigns, or politicians, or states use this methods. In fact, on sites like Twitter, a lot of the time, we see lone actors or groups of people spreading this information, people with some computer coding skills, or even people with very little computer coding skills, but who have enough knowledge to use one of the automated platforms for sending out their messages, or sending out different things.

I think that that's democratization of the way the information sharing goes on, but also the democratization of the ability to use automation to magnify a view is actually a really important thing that isn't often caught in media article about this.

Another thing is that a lot of people who build bots on behalf of political campaigns do so purely, because they are an employee of a subcontractor who works for that political campaign. Some of them may be don't have a bias, or an interest in the issue. Often times, because this is a transnational phenomenon, we see that the people who are building or launching the bot, or building or launching the information campaigns that go along with them are in other countries. They'll be in India, or they'll be in the Philippines, or Russia. Then, they have complete disassociation with the information that's being spread.

This lines up with the stories that BuzzFeed did about the Moldovan teenagers who spread fake news. The way that Twitter and Facebook exist, and Redit, and Kik, and Periscope and all these other things is that it's as international platforms for communication. Any one person with a vested interest in trying to manipulate the way that things go on can manipulate the flows.

All of that being said, powerful political actors, whether it'd be governments, or actors within governments, or whether it's intelligent services, or militaries, they definitely have a foot up on the regular Joe Schmo on the street, because they have a few different things. They have a lot more money, so they can scale a lot more. They have a lot more in the areas of expertise and other resources that have to do with intelligence. Then, they also have the ability to use it half the time to do this stuff.

Their resource-base is much bigger. When a subcontractor that's working on behalf of the Russian government does something like this, as in the company that Adrian Chen reported on in his piece, The Agency in the New York Times Magazine, you start to see more sophisticated attacks. You start to see more complex and much more confusing attacks.

[0:35:31.5] JM: I believe all of that — I find all of that to be tremendously plausible. There is another dimension to this that I want to discuss with you that I never hear people talking about. It's one of the these things where I don't know if I have a tinfoil hat on my head or if I'm right about this. Software Engineering Daily has done numerous shows about advertising fraud. Anybody's who's a regular listener to the show knows that I believe advertising fraud is a huge problem and that nobody talks about it, because it's within nobody's best interest to talk about it, because it would hurt the internet's income.

I can scrape the internet for a bunch of content, and I can make a site that looks plausible. I can scrape the internet for recipes that are made with beef. I can set up beefrecipes.com, it's a WordPress site. It's got some beef recipes that I scraped from the internet. I can set up automated ads on that site using an ad network, and then I can pay for bot traffic to come to that site. As long as I make more money from the ads, then I lose from the bot traffic that I'm paying for. I have an arbitrage, and this is a huge arbitrage opportunity. It's really easy to make money from it. I can pay for infinite bot traffic and make much as I want until I get caught.

From my analysis, there is a deep connection between advertising fraud and political bots, and the reason is that advertising fraud works better if you have some human traffic coming to your website, because then you can launder the fake clicks using the protection of the real clicks. There are fake clicks going to every site. We all know that. The question is; what are the sites that have some human traffic going to them? The more human traffic you get, the more fake clicks you can launder.

Sites are incentivized to actually bring in some human traffic just to disguise an avalanche of fake traffic. In order to bring some human traffic, you need some plausible original material. You can't just get beef recipes, because that's really obvious that this is a fake site. It was harvested from information that sitting around from the internet. You need some plausible [inaudible 0:38:04.6] Macedonian teenagers writing crazy stuff about Hilary Clinton, turning babies into hamburger meat or something.

Because that plausible material turns out to be what some people categorize as computational propaganda, or "fake news", and that's where this circle becomes so profitable, and so proliferate. Yes, we can talk about people subcontracting to Indians, Russians, or whatever, to just make propaganda to shift public opinion. It doesn't even have to be that malicious. You can just be looking to make money. You can just ben an ad fraudster. It's almost like the automated trading business. The automated trading business makes money as long as they're variants. They don't care if it's a huge uptick, or a huge downtick. You're hedged on the upside. You're hedged on the downside. All you need is variants, the more variants, the better.

© 2017 Software Engineering Daily

[0:39:09.4] SW: That variants, you're spot on there, and you're also really on to something with the connection between ad fraud and what's going on with the political bots. There's a deep connection there, and a lot of the use of political bots has grown out of the desire to create income, or to pull over the eyes of the powers that be.

I think there're a couple of things there. One is that there's definitely people that have used bots to spread fake news stories, but then also had human interaction around those news stories. The Denver Guardian, the whole thing that blew up during the election. There is a guy that went on 60 Minutes recently on their fake new segment about how, yeah, he had done all of these to make a dime. He was making money. In fact, he was feeling remorseful about the ways that this happened.

The other thing, which is building off of what you said about creating variants and traffic I think is something really important. I think that all of the fancy machine learning aside and all of the innovations happening in an AI aside, if you just look at the ways in which the more sophisticated political bot, or computational propaganda campaigns happen, then you start to see really quickly that it's never just bots alone that create the more successful campaigns. It's a combination of bots and humans together that drive a really successful campaign.

It's important to note that nothing stops someone from signing into Twitter on the bots account, and generating some real tweets. That does a couple of different things. One; it gives the situation, the account in question more legitimacy. Also, two, it has the function of getting around the algorithms that search for automated behaviors. It creates that variants that's needed in order to have people interacting with the content.

[0:41:20.0] JM: How much power does a Facebook or a Google have to stop computational propaganda?

[0:41:28.2] SW: I think that social media companies and also search-oriented companies have quite a bit of power. I think that if Twitter has the information about whether or not lots and lots of Russian accounts were used to spread information during the U.S. Presidential Campaign, for instance, that they could share this and it could have a big effect on what we know about certain

things. Whereas like, if researchers are attempting to track IP addresses, their work is spotty at best because of the use of VPNs and all sorts of other things.

I think, also, that companies can do something else, which is they can have a degree of transparency about which accounts display a high degree of automation. I know there're lots of good accounts out there. I know that the New York Times, and the Washington Post and all of these other groups use automation to get their tweets going.

I think that there's a difference when it comes to looking at a typology of bots, or looking at all the different types of bots. If you look at the bots that are used to spread, or what's being called fake content, or misinformation, then you're able to say, "Wow! These accounts are only interacting around one topic. We should probably be flagging them as automated accounts just so people know."

I think there's something akin to what Wikipedia does with articles that haven't been written property or need to be expanded upon, that social media companies could do at the very least. I don't think it's enough for social companies to be — Or for companies writ large to be putting this upon the shoulders of simple society actors. To be saying that it's a job of the AP, or the job of CNN, or of Snopes to be doing real time fact checking. I think that that's not very fair, because the reality is that those simple society actors and even researchers like me do not have access to all of the information that the companies have access to.

I think the real crime and the real travesty lies in not using that data in some ways for the benefits of democracy, because if you're like me, I want to believe that social media and all forms of new media can be used to benefit democracy. There're still aspects of them that are very empowering for freedom, if companies don't have some buy-in, then very little will happen.

[0:44:10.0] JM: You work with Google on their Jigsaw project. What is Jigsaw?

[0:44:15.4] SW: Jigsaw is Google's human rights-centered sort of think-tank and incubator. What Jigsaw is build tools that are informed by research for communities that are facing things like censorship, or attacks online. Jigsaw is working to protect people, and it's a really

interesting organization. It's actually not part of Google. It's actually a part of Alphabet, the large Google parent company. It has a degree of autonomy there.

Jigsaw just released something called Perspective, which, for instance, Perspective is a tool that detects hate speech in the comment section. For that, that tool, Jigsaw, paired up with the New York Times and the Wikimedia Foundation and used a bunch of data to build a machine learning algorithm that pretty successfully detects hate speech in comment sections.

[0:45:15.1] JM: When I'm reporting on this advertising fraud stuff, it's such an uphill battle, because nobody will talk about anything. Honestly, I think if people at Google could stop advertising fraud, they would. I think it's just a problem that's too hard to stop and I don't think anybody is willing to admit that, "Yeah, a giant CPG company. We actually can't tell you how many bots saw those ads that you ran on our ad network." It's like really hard for anybody to talk about. Obviously, the media companies don't want to talk about it, because they already have enough problems shifting to the digital age.

In an ideal world, the people in Jigsaw would be, like you said, segregated enough from the Google cash cow to actually research advertising fraud to talk about it some. Have you had any conversations with people at Jigsaw about ad fraud?

[0:46:19.3] SW: I haven't. My work has been more focused on sort of the things that we've talked about. Obviously, I was brought on as a fellow there to think about bots and automation. I'm sure that there're folks at Jigsaw that are thinking about this and they have a fantastic team of some really smart people. My purview is quite narrow there, and so I can't really speak beyond what goes on with my own research.

[0:46:44.0] JM: I understand. What about ISIS? ISIS is harnessing the internet, or something. I don't know exactly what they're doing. Are they using political bots, or are they just using — Making YouTube videos and communicating with people over Signal or something.

[0:47:03.6] SW: No. I think that terrorists groups have certainly made use of bots and of political bots. I think that there's been a big onus upon social media companies and on states to detect and get rid of content related to things like ISIS, or groups or like ISIS.

We've seen a lot more success with both the companies and with governments in getting rid of terrorist-related information. I think that we're just starting to see a pick up on the broader question of automation being used to send out other types of political information. Yes, ISIS has used political bots, but I think that they're being heavily, heavily targeted and broader to control in many ways.

[0:48:03.1] JM: You're saying the ISIS bots are being heavily targeted.

[0:48:05.7] SW: Yeah. I think that up until really recently, it's been — The last couple of years has been — There's been a focus from the social media companies specifically on terrorist-related speech, or extremist speech. Now, there's starting to be a pivot that says, "Of course, extremists' speech is really important, and it's important if terrorists are using automation to scale up their efforts."

For instance, ISIS uses this app a while back that since been gotten rid of, called The Dawn of Glad Tidings app. They basically took control of people's social media accounts that were sympathetic to ISIS and used them to create the illusion of solidarity, or to tweet out things that ISIS wanted pirated.

Now, the shift has been to actually discussing computational propaganda more broadly and the ways in which people manufacture consensus or manipulate public opinion using bots at an electoral level, or during security crisis and all sorts of other things.

[SPONSOR MESSAGE]

[0:49:13.4] JM: For more than 30 years, DNS has been one of the fundamental protocols of the internet. Yet, despite its accepted importance, it has never quite gotten the due that it deserves. Today's dynamic applications, hybrid clouds and volatile internet, demand that you rethink the strategic value and importance of your DNS choices.

Oracle Dyn provides DNS that is as dynamic and intelligent as your applications. Dyn DNS gets your users to the right cloud service, the right CDN, or the right datacenter using intelligent

response to steer traffic based on business policies as well as real time internet conditions, like the security and the performance of the network path.

Dyn maps all internet pathways every 24 seconds via more than 500 million traceroutes. This is the equivalent of seven light years of distance, or 1.7 billion times around the circumference of the earth. With over 10 years of experience supporting the likes of Netflix, Twitter, Zappos, Etsy, and Salesforce, Dyn can scale to meet the demand of the largest web applications.

Get started with a free 30-day trial for your application by going to dyn.com/sedaily. After the free trial, Dyn's developer plans start at just $7 a month for world-class DNS. Rethink DNS, go to dyn.com/sedaily to learn more and get your free trial of Dyn DNS.

[INTERVIEW CONTINUED]

[0:51:07.4] JM: ISIS seems like a useful thing to study, because if you look at ISIS as the base case for a kind of political bot that we just want to completely silence, completely obliterate from the internet, if we can, you can study the techniques that are successful in obliterating it and then think about what are the ways to use those in more subtle fashion to subdue these more — These bots — I would say Twitter is doing — They're doing an earnest job trying to subdue these bots without subduing the free flow of information. I'll applaud Twitter. Does that sound like a reasonable strategy?

You're obviously studying this stuff right now and you're keeping an open mind. I'm sure you are starting to think of strategies, or maybe you've already thought of some strategies for what are the best ways to subdue these forces. I guess can you use ISIS as the base case that you can induct upon and figure out more complex strategies for subduing these problematic bots?

[0:52:28.8] SW: I think that that's a really interesting idea. I think that using case studies whether it's be ISIS or whether it'd be — You mentioned ad fraud earlier as a basis for looking for the types of bots that are in action, or the ways of policing accounts. I think that there's something that's definitely true there.

I think that what needs to happen though is that people — The research has to be done in depths about the ways in which ISIS-driven accounts, or ISIS-driven content differs greatly from other content. I think there's something else, which is this kind of a larger point is that this is contextual to political events in different countries. That's not to say that you can't build a tool that polices content or that you can develop some software that gets rid of, or identifies bots on a platform, but it is to say that the ways in which attacks happen, or with the way in which computational propaganda happens in the United States might look a little bit different from the way that it happens in Indonesia, or the way that it happens in Australia.

We need to be really cognizant of the generalizations that we do make when we develop these tools and we need to be quite secure in the typologies [0:54:01.3] build the bots, because when we build a typology and we say, "These are the types of bots that are used to spread propaganda, or that are used to spread hate speech," that they do have some crossover from country to country. I think that that definitely is the case. I think that is the case that accounts that have been automated, that have been used by ISIS do mirror a lot of the tactics that are used by accounts. They come out of Russia, or accounts that come from actors in America. There is certainly transnational information flow on how these bots get built. Yeah, I think that research nuance, that's sort of the place — That's the place where I hope that my work in research can help.

[0:54:45.2] JM: People often talk about Twitter bots. They don't talk as much about Facebook bots. My theory is that it's not that Facebook is necessarily better at policing bots than Twitter. It's that Facebook allows you to curate your network and build a more insular worldview more easily. You don't see those bots, but they are on Facebook. Do you find that theory plausible, or is there some bigger mote that Facebook has against bots?

[0:55:16.9] SW: I think the mote that Facebook has is the network structure. I think that Facebook and Twitter look completely different. Twitter has, in some ways, done something that's fairly palatable and creating a fairly open and democratic platform where the APIs are open and where anyone can really build software and launch it within the platform if they stay within certain parameters.

Also, Twitter, you follow people you don't know to put it simply, and you follow sorts of different people. On Facebook, you only really follow people you know or you know through a friend. Then, also Facebook has done more to limit the ways in which software developers can interact with the platform.

[0:56:08.1] JM: I had Brad Stone on the show recently. He's, I think, senior editor at Bloomberg. He's the head of Bloomberg Technology reporting. We were talking about how there are more new sources than ever before. Power is no longer centralized in the esteemed institutions. You can get your news by just reading the Washington Post, or the New York Times, or Fox, or Wall Street Journal, or you can scroll through your Twitter feed, you can look at Facebook, you can look at Hacker News, you can synthesize a viewpoint from a wide array of individuals and news sources.

The news has become decentralized, and I said that to Brad. I was like, "I love it. I love being able to tip my toes into the cesspool of bots and angry people, that is Twitter, and to look at Google news and see articles from the Washington Post, and Wall Street Journal, and I love to dip into Facebook and see somebody I know who has shared something from BuzzFeed. I would never go to buzzfeed.com, I don't think. Maybe I should go to buzzfeed.com. It's probably fine. It's probably great."

Brad was describing this as completely exhausting. By the way, this is the ecosystem that has allowed me to build a podcast by myself and actually get traction. It seems great, but do you find it exhausting? Are things better than ever before because there is so many new sources, or is it now worse? Is it now just completely confusing and there's no centralized view of the truth, or are we in some fluctuative point that is going towards an area of more stability? I don't know. What are your thoughts on this?

[0:58:05.8] SW: I think there were influx. I hate to be the academic that qualifies what you just said, but I'm going to qualify it. I think that there are certain benefits to having a really diverse information ecosystem and for people to be able to access all sorts of different types of information. I think that the reality is a bit more nuanced. I think that people tend to only really engage a lot of the time with the same websites over and over and over again.

I know that people I know very well, people in my own family, when they get on the internet, they access the same 5 or 10 sites. This isn't to say that the filter bubble, as it's called, or homophily is what academics call sort of birds of a feather flock together. They go to the same websites in this context. It's not to say that it is all encompassing. I don't think that that's true, and the research that I've seen doesn't bear that up.

I think that there are certainly democratic benefits to having lots and lots of different source of information, but I think that we're going through some growing pains here. The internet was scaled so fast, and this happened — The spread of the internet has happened in an exponential rate far beyond what we've seen with other technologies that are similar, which you can't really point to anything. If you're talking media technologies, it's been an exponential growth far beyond even the media technologies that we've had that have been developed in the last 50 years, like the television.

I think that, because the information landscape is getting more and more complex, I think that the way that people confront it and deal with it is also getting more complex. I think that there's a move towards simplifying this going on. I think that things like fact checking are really important. I also think that people being, in a sense, accustomed to the confusion that different sources of information generate is something else entirely.

I think that other people, for instance, Danah Boyd just wrote an article. Danah Boyd of Data & Society wrote an article about media learning and talking about the need for critical thinking and media-base thinking; getting people to actually think in depth about the media which they consume.

I think there're a couple of things that play that are both solutions for addressing this problem. There is definitely software that we can build that can do fact checking, or that can identify bots. There're also all sorts of social things that need to happen as well, which involve teaching people about the way that this media happens or integrating critical thinking into education in an earlier age. Sure, I think when people bring up media education that they kind of get slapped down and said, "Well, that education system is bad as it is," but we have to start somewhere.

© 2017 Software Engineering Daily

I think that coming at this problem, the problem of computational propaganda from multiple angles is probably the best way to deal with it. I don't think that a technology-based solution is going to be the main fix on its own. I don't think that a social solution on its own is going to be main fix. I think that both of these things have to come together in order to create some kind of real change in the way that people approach this and the way that they consume this information.

[1:01:45.4] JM: Did you mention just now that the filter bubble has not been born out in your research? You've looked at it and you've said, "No. There's not a filter bubble."

[1:01:55.5] SW: No. I'm saying there's different opinions on the ways in which the filter bubble plays out. I'm saying that, "Sure. There're plenty of research out there that shows that people do consume the same news over and over and over again. Often times, their consumption of that news results in them seeing the same stuff on social media sites."

However, there's research from groups like the Social Media & Politics Lab at New York University that shows that the ways that we think about things like the filter bubble should be more nuanced. Actually, the ways in which these information flows go are more complex. There's a guy at that — What they call the SMAP lab at NYU named Joshua Tucker who's actually done some amazing recent work on this question of information flows and the ways in which people share information and kind of challenges the state of thought that says, "Yes, unequivocally, people only look at the same information over and over and over again."

That said, there's also a research by folks like Shanto Iyengar at Stanford's Communication Department that does show that there is a massive amount of echo chamber going on online and that we should be concerned.

I think that the reality lies in the middle, that there's some degree of it, but that it's not always the case. I think that, again, it's contextual to the social situation and the social group in question, and that that's something that often times gets gauzed over and the need to create a fairly black and white picture.

[1:03:39.6] JM: Let's close off with just — Do you have any crazy stuff that you've seen recently? When I look in deeply into this bot space, I'm like — I just see the weirdest things, whether it's like somebody — I don't know. Pizzagate is the first thing that comes to mind. There are some many weird things. I remember reading this article recently, it was like, "For Chan is deeply involved in this mess of bot content." I don't know. Tell me something weird, or unexpectedly you've seen recently studying this stuff.

[1:04:14.5] SW: Yeah, I think that looking at the Chans, 4chan, but also HN is definitely something where you'll find surprising information. There's work that's being done at the Data & Society Research Institute, but also by our own project looking at the ways that other groups online, like groups on HN been fundamental in the spread of different memes that have really gone viral online that are political. There's a sort of a show I think that [inaudible 1:04:51.8]. It was called something like *The Trolls Among Us* that got really deep into this stuff.

The other thing that I think is really kind of interesting, though not necessarily crazy, or super surprising, but something that people should really focus in on if they're going to research or look into this more deeply, is sort of the human element.

There's always a person behind the bot. It's really for us to sort of disassociate the fact that people are using bots as a tool. They are software tool that allow these people to spread their own ideas and that I think the future, by that, I mean the next year or two is not going to — The future of computational propaganda, that is, in the next year or two is not going to be dictated necessarily by machine learning or sophisticated way of doing individual political marketing. I think that it's going to be a lot more focused upon the use of sort of cyborg technologies. Using a combination of humans and automation to attempt to reach voters, or attempt to manipulate opinion.

I think that that's sort of where things are going that you have a bot that's running, but you also have a team of people that are around to really moderate what the bot does and says and also to engage with other human beings once a bot makes contact.

[1:06:15.5] JM: We'll conclude our episode of science fiction engineering. Seriously, I can't imagine somebody can write science fiction that is more compelling than what's going on in our actual world today.

[1:06:33.9] SW: I know.

[1:06:34.0] JM: You just read the news and it's more entertaining than any science fiction book that can be written.

[1:06:38.9] SW: Yeah. I actually just read a new article out in Vanity Fair by Mike Mariani, and I chatted with him about the article before he wrote it. He actually wrote that there's a propaganda book written by a Russian fellow with a pseudonym about how the new world war is going to be a confusing one, and it turns out that the guy that wrote the book is actually an advisor to Putin. Maybe we need to go there and read that book to figure out what the propaganda of the future looks like.

[1:07:07.2] JM: Okay. Well, if you ever come out with another — Or a new report, or if you have some kind of capstone on your recent research or something, you're going to come back on because this I really an interesting conversation.

[1:07:18.6] SW: Yeah, thank you for having me. It's been great.

[END OF INTERVIEW]

[1:07:21.7] JM: Thanks to Symphono for sponsoring Software Engineering Daily. Symphono is a custom engineering shop where senior engineers tackle big tech challenges while learning from each other. Check it out at symphono.com/sedaily. Thanks again Symphono for being a sponsor of Software Engineering Daily for almost a year now. Your continued support allows us to deliver this content to the listeners on a regular basis.

[END]