

WESTERN SCIENCE SPEAKS PODCAST SEASON 4, EPISODE 2

EPISODE TITLE

Artificial Intelligence (AI): Spotting a veiled cry for help on social media

PODCAST SUMMARY

On this episode of Western Science Speaks, Dr. Dan Lizotte and PhD Student Brent Davis discuss how they are using artificial intelligence (AI) to identify people on social media who may be struggling with addiction.

INTERVIEW

You're listening to the Western science speaks podcast. Presented by Henry Standage.

Henry: Hey, welcome to the podcast. Today Western Science Speaks expands its horizons as we have not one, but two guests: [Dan Lizotte](#) who came on to the podcast last year to talk about his research pertaining to artificial intelligence and using it to diagnose and treat physical injuries and we're lucky enough to be joined today by his PhD student, [Brent Davis](#). They're working on a software program that identifies people on social media who appear to be struggling with drug abuse. The whole project, in concept and execution, feels like such a fantastic modern remedy to some of the cries out for help we see online. So I hope you enjoy it. Here's the interview.

Henry: All right. A couple of firsts on the podcast today, our first two-time guest: Dan Lizotte, and our first three- man podcast. Thanks for coming, guys. And we'll start with Brent, why don't you tell us a little bit about the research and what you and Dan do together.

Brent: So the research we're doing is a way to search social media for a topic of interest that you might not know the keywords for. Like if you go on Google and you type in 'cat pictures', you're going to get cat pictures. If you're looking for something a little more relevant to research elsewhere, you might be involving a community of people that you don't really know, so you also don't know their speech patterns. If you're trying to find posts that are related to a topic that is most relevant to that group, it can be very challenging. So, we're trying to design an AI technique that facilitates that kind of research for other people.

Henry: I see. Something that I always see made fun of is: "Parents: how to catch your kids' internet slang." And the advice is always incorrect! And so how do you build a program, that, as this stuff changes fluidly, keeps up with it.

Brent: The kind of new technique that's coming out for that is a technology called 'Word Embedding'. As long as you can find a sample of the kids' speech, their slang, you can develop a technique that will go in and learn a machine-oriented understanding of what that word means. And it'll at least (usually) associate that slang with other slang that can be used interchangeably.

Henry: And where do you move from there? So, you figured out how to embed these words, and then what are the next steps.

Dan: So, once you've got that representation built up of how people are using language online, you can take a lot of different next steps, which is actually kind of a fun thing about this piece of research. One of the next steps is to look for more occurrences of that same kind of discourse - that same kind of way of talking. And, you can move to different platforms. You could learn about a discourse on one platform, like Reddit for example, and you could use what you learned there to maybe search on Twitter for similar kinds of conversations that are happening.

Henry: Right, and you guys are starting off on Reddit, right?

Dan: Yes, we are.

Brent: Yeah, one of the big differences about Reddit is that everything is more likely to be anonymous than Twitter. So the odds of you finding someone that's like put up, this is where I live, this is what I do for a living, usually Reddit accounts don't have that much detail, whereas on Twitter, most people tend to label it with everything; mine has what I do, where I work, a picture of my face. You don't really see that on Reddit, which I think is an advantage for getting this started.

Dan: Because we think that people feel a little bit more free on Reddit to speak their minds and to speak in a way that they know to be understood by the community of people that they want to interact with, we feel like we're getting a better picture of how people are using language.

Henry: Yeah, on Twitter, you could have work colleagues or even family whereas on Reddit you can keep it as low-profile as you want. It's narrow in scope on the subreddits too.

Dan: Exactly. Brent can tell you a little bit more about how we're using that. The other obvious difference, too, is the forum. On Reddit you can post can be as long as you want long as you want, more or less, whereas on Twitter, it's different because you're restricted.

Henry: And you're looking solely at mental health here, right?

Dan: So far.

Brent: Yeah. And mental health and substance abuse in particular. So, one community that is out there, that came up from talking to the health units and seeing what might support them is the opioid users; there's still the opioid epidemic going on. I think it was up to 3000 overdoses in Ontario in 2018 - not sure on the number, I could have gotten that number wrong. But it's been steadily going up. And, they are a hard to reach population; no one's going to go on Twitter, on their work account and be like: "Hey, guys, I'm about to go home and do [whatever their drug of choice is]". But on Reddit, they have a way to talk to each other. And you see some surprising things like safety tips, the overdose kits; it's not what you would expect of people; there's kind of an active community around what they do. And, that's where we're getting the language they use to be able to understand what's going on.

Henry: Can you give me an example of this language? What's something that would pop up on your system as a red flag?

Brent: One of the big ones is the word 'fentanyl'. Right, you just call it 'fent'. So, if you didn't know that, and you're trying to do a keyword search, you search for fentanyl you may not find much. It is very common that short forms get used, informal names get used.

Dan: Oh, yeah, a lot of language around buying and selling. And I can't think of any of the specific words right now. But a lot of that language around buying and selling is different than it is in other communities. And, I should say, you could just read a whole bunch on the subject. and eventually you would get it right, or you would identify some of these things. But yeah, I think one of the things that we're providing is a way to use the computer to highlight the words that are most distinctive of the discourse in that subreddit without having to read through it all manually.

Brent: Who is this information for? So, you're making you're kind of making a shortcut for people so that they don't have to go on a full method investigation for a year to figure out that 'fent' means fentanyl. So, who are you supplying this to?

Dan: This project was really instigated by an epidemiologist with Ottawa Public Health - Cam McDermott, he's a great guy. And we were talking about different ways in which AI techniques might be able to help people working in public health. He really expressed his frustration over how hard it is, for folks working in public health who want to be helping and delivering services to these vulnerable populations, to actually get in touch with them. It's really hard. And so that kind of piqued our interest. And so, they've always been our kind of number one client, not just Cam, but public health organizations in general. So, they're definitely at the top of our minds.

Henry: And they've got two major universities there, as well as Algonquin college. I don't know what else is there, but have they used it toward student life there at all yet?

Dan: No, they haven't but we are collaborating with folks at Carleton University in the Geography Department because they're really interested in how the place where you are interacting correlate with the vulnerabilities you have. What do you have access to? How does it impact what kind of services you can get? The folks at Carleton have a lot of expertise in working with these populations and really helping them tell their own stories about what their experiences are like, which we really like. Working with the social scientists at Carleton, I think, is going to give us a way to really bring this beyond helping the health units to help the people on social media tell their stories; help them tell the rest of us what they're really going through.

Henry: Right- without the veil of anonymity. Yeah. So problem I've had is that Reddit doesn't like it when I promote the podcast on their subreddit'. Apparently, I don't add to the community, I just plug the podcast. But I imagine for you guys, with the Ottawa people, it's probably pretty rare that you find someone who talks about their location, being in Ottawa, and then kind of being the ideal person for who you're working with. So how are you going to be able to find lots of people in that demographic?

Dan: Yeah.

Brent: The kind of bewildering part is there are many more addicts than you would expect. There are people on there that are not thinking that someone's going to come and find their comments. And, they do they say "I'm in Ottawa, and I bought this"; not exactly in those terms, but spread out over 20-30 posts, you can sometimes find something that explicit. The other part of what we're doing, involving public health news, requires the expertise of the geography folks. They want to look in an area. And like

you said, there are subreddits dedicated to particular areas; Western is in area, Ontario has its own subreddit, Canada, London, we've got Western and London, Ontario. So those are all spots where you can go and find the same usernames used in different areas, because usually, someone doesn't do all of their posting on London, Ontario. You can identify some people that likely have a tie to London, Ontario pulling usernames from there. And then you can retrieve their history. And that's where we started to bring in the machine learning apparatus to help us out

Henry: Right, so they're prominent on Reddit, Ottawa Senators fans or something and then yeah, you have a pretty big hint there or something like that.

Dan: Yeah, exactly. And with that was one of the things we started thinking about at the beginning of this: what would be those flags [hints]? And it's like, since fans, and maybe government stuff, maybe [tags like] Ottawa Blues Fest, like there's a lot of things that folks in that region, for example.

Henry: Right, definitely better to start off with a major city.

Dan: Yeah, I think so. But we're not just interested in identifying people so we can go to their house, and I don't think we'd ever be able to do that, which is just as well. But the cool thing about this is that it's a two-way street. Through social media, the health units can reach out to these individuals without exposing them, without having to know exactly where they're at, without having to know their name. Even just with things like, "did you know, we have this program? Did you know, safe injection sites are here, here and here?" Even stuff that basic is stuff that they can't target right now, because they're not sure who really needs it. So that's really the goal of the project in terms of what we hope we can deliver to them.

Henry: Have you thought about how you would ever possibly expand to other social media sites?

Brent: The advantage of how we build the AI technique is that as long as it's from a social media site that has text, it can be used. Nothing that we are doing so far as specific to like Reddit score, karma votes, it's entirely based just on the words they use. So it actually translates to Twitter very easily, and potentially other social media, too.

Henry: Yeah, weirdly, I see a lot of heavy stuff on Facebook these days, which you would assume wouldn't be the case. But Facebook is like if I'm feeling like particularly strong that day. I'll check out Facebook and see what old people from high school are doing and maybe stuff hasn't panned out perfectly.

Dan: I agree. And I wonder if that's something that's changing over time. We have some great expertise at Western too who are thinking about the impacts of social media on different aspects of people's lives. [Jackie Burkell](#) in FIMS is a great example. She's a great researcher and she looks at different aspects of social media. One of them I think, being the idea that you don't always present your best self on social media; you pretend that things are always super sunny, even when they're not. But I wonder if we're seeing that change over time. That's the other thing about this project. And the thing about social media in general, that everything is evolving so fast that we're kind of watching it happen in front of us as we develop these techniques. And it'll be interesting to see where it goes and how we must adapt the work.

Henry: Yeah, and Reddit as kind of an intellectual community, in a way. You guys said they have all these safety tips and there's actually a real discussion going on. If I were to do an opioid, that would be the

place that I would trust to talk about it. With Facebook, I think you have people more just shouting out into the void, if that makes sense.

Dan: Yeah, I agree with that. And I think, like what Brent was saying, but the opioids community is really remarkable; there are people on that subreddit who are public health psychologists and people who care about harm reduction. So, professionals are already using that subreddit, for example. So, we're seeing people with a lot of expertise and a huge amount of lived experience having these conversations. It's an amazing thing that's being captured right now, as people are going through this.

Henry: When you came on last summer, we talked about physical AI diagnosing stuff? Have you noticed any similarities or major contrasts between the two projects?

Dan: I feel like there's one big difference and one big similarity. In the diagnosis setting, you usually have measurements of what we call health outcomes -what happened with somebody- and you're trying to figure out the causes that led to whatever that outcome was, whether it was good or bad. Then you're trying to figure out what you could have done differently. We can sort of do that with this project, but not nearly as well. And it's a lot more about just saying, "look, there's this whole big bunch of data and we're just trying to understand what the patterns are in it. We're trying to relate that to what's happening in people's actual experience." So this project, in a lot of ways, has been a lot more open ended, which has been exciting and interesting, but really challenging. It's a lot more difficult if you don't have a problem, like, here's your specific problem, and I'm going to do this, this, this, and then I'm going to solve it. And then it's going to pick a treatment for somebody. This has been different; it's really tools for exploring the data and trying to make use of it when you can't necessarily get that kind of clear goal. So it's been cool to be doing both of these kinds of research. We're lucky working in AI to have different aspects to explore. But I think in both cases, you always have to be so careful; to do your best to understand where the data are coming from, how they're being produced, and what they really mean. And, whether the data comes from Reddit, or electronic health records, if you're going to do a really good job, you really need to take the time not just to understand the AI methods and stuff, although that's key, but to appreciate what that domain of knowledge is and, how can you add to that.

Henry: How much data are you getting on a daily basis?

Brent: We're not doing the data collection, but the repository that hosts this, per day, there's probably two or three gigs per day going up; gigabytes of just pure text going up on Reddit every day. I think someone was saying there's something like 2000 or 3000 posts every second. Not all of them are going to be relevant. And that's part of the challenge of working in social media, dealing with the huge sea of information that's out there. I think it was something like there was 3 billion Reddit post last time we did a full count of everything we had. So, there's, there's a huge wealth of information. And that's part of why we build the search technique. Because really, until you go and see what's on there, in a lot of ways you have no idea what you are going to find. There's definitely a lot of communities on there that posts all kinds of things.

Henry: I guess. So is it levels of filters, and you kind of narrow it down slowly?

Brent: I think that's a good way to describe it. And then the last stage of the research that I find the most interesting is when we build a Smart Filter that's customized to seek what we want to look at. Because there's a lot you can find by just grabbing a subreddit, grabbing a list of usernames that are

associated with what you want. But when you can find relevant information with something that has no idea or understanding about any of these topics, well, that's the fascinating part to me.

Henry: I think this is super interesting research. And I think you guys are kind of getting in on the ground floor, but it's probably going to be one of the most prominent forms of research over the next two or three decades. It's important to be able to figure out how to identify people on social media, because I think one thing we've seen from recent tragedies where people are overwhelmed by stuff on social media is that there are always warning signs when you look back.

Dan: Thanks. That was a lot of fun. We have got our website under construction. It is at philab.uwo.ca. and we're slowly getting research projects up. So, if people are interested to know more, go check it out.

Henry: That concludes this episode of the Western Science Speaks podcast. I think with a research area like this, you're going to see it blow up where people try to figure out how can we deal with some of the mental health issues we see from people in my generation. And Dan and Brent are at least the first I've heard properly, addressing part of the problem. So, definitely a research project to keep your eyes on. I'm Henry Standage, signing out. Thanks for listening.