Does Profiling Make Us More Secure?

Shari Lawrence Pfleeger | Dartmouth College
Marc Rogers | Purdue University
Masooda Bashir | University of Illinois
Kelly Caine | Indiana University
Deanna Caputo | MITRE
Michael Losavio | University of Louisville
Sal Stolfo | Columbia University

Profiling means making predictions about likely user behavior based on collected characteristics and activities. Shari Lawrence Pfleeger and Marc Rogers brought together a group of researchers from a variety of disciplines to discuss whether profiling and prediction actually make us secure.

Michael Losavio: "Profiling" is a loaded term and something we should clarify as we go through this discussion.

Sal Stolfo: That's why I like to call it "behavior-based security" or "human user behavior analysis."

Shari Lawrence Pfleeger: There are clear examples of where this kind of profiling—user behavior analysis—is very helpful, examples of where it's sort of creepy, examples of where it might be illegal or unethical, and everything in between. This roundtable addresses what we're doing, what are or should be its limits—legal, ethical, mathematical—and when it's most useful. With all of this in mind, what's the scientific basis for profiling?

Masooda Bashir: It is based on the assumption that we all have stable personality traits. Profiling also assumes that once we hit a certain age, we have a stable personality and characteristics that can be used to predict our later behaviors and therefore past behavior is used to predict future behavior.

Marc Rogers: There are two types of profiling: *clinical*, which is based on personal observations and anecdotal evidence, and *statistical*, which is based on large datasets. So we have to consider the art versus the science.

Losavio: It becomes controversial when profilers use characteristics like race and ethnicity.

Stolfo: There are certain contexts in which profiling is useful, such as credit card protection, which is designed to stop losses. If a transaction is suspect, an alert is issued. It's limited to the owner of the transaction system and the owner of a card. Here, profiling isn't a burden on the person being profiled.

Deanna Caputo: But intended use doesn't always translate into how it's used. If you're being profiled and the profile is incorrect, you end up with a bad profile and the company keeps turning off your credit card.

Bashir: We need to consider whether profiling is used in a criminal investigation or to predict "normal" behavior. Traditionally, profiling occurred after a crime had been committed. The profiling in that context was used to identify the criminal and prevent future crimes. A lot of the profiling happening now tries to predict our behavior, including consumer activity, and therefore has moved beyond criminal behavior.

Kelly Caine: There's another difference: flu trends and global health are population-level predictions, whereas consumer tracking leads to individual prediction. Statistically, we're better at population profiling.

Rogers: There's another distinction: identification of a person versus prediction about the person's behavior. For the latter, we plot trends. What's somebody going to do a year, two years, six months from now? We see a lot of that in the intelligence community, with cyber adversarial predictive analysis trying to predict

1540-7993/12/$31.00 © 2012 IEEE

future behaviors and adapt to them before they actually happen.

Caine: Predicting future behavior and adapting appropriately is tricky on many levels. One concern is that when people aren't even aware of something yet themselves, they certainly can't be aware that they're disclosing that information. How do we reconcile that with protection of privacy?

Rogers: Insurance companies have been collecting actuarial data for decades, predicting whether you're going to die within a few years. They take very large datasets and generate trends; they're very accurate even 50 or 60 years down the line. Most consumers don't realize that their insurance companies have already figured out to within a couple of years of when they think their clients are going to die.

> **"When people aren't even aware of something yet themselves, they certainly can't be aware that they're disclosing that information. —Kelly Caine**

Stolfo: Behavior will be a key authentication technology, so if masqueraders obtain your credentials or steal your passcode, they won't be able to use your machine, because they'll be identified as different from the legitimate user or owner. Recently, DARPA announced a Broad Agency Announcement on advanced authentication to produce *cognitive fingerprints*: predicting what users do typically on their machine and identifying when a change has occurred and whether that change indicates a serious security breach. It involves solving a number of hard technical problems, but those problems aren't insurmountable. I expect in the next few years that such technology will be broadly available.

The privacy implications are important. Protecting privacy depends on how one computes a user model in a large organization that has many users. If algorithms are used that are oriented toward multiclassification, meaning that multiple users' data will be intermixed to compute discriminate classification models, there's a potential privacy breach.

But in other cases, like anomaly detection and algorithms, which aren't classifiers in the traditional sense, we can compute a model based entirely on the user's own data. This has a better chance of keeping private what users actually do on their machines. I think it's feasible to expect that our behavior on our own machines could help us protect them from being misused by either malware that's a proxy for an external hacker or an insider who gains access and masquerades as you on your own machine.

Let's not forget that profiling is a very useful area of security analysis. It's very targeted but very useful.

Bashir: Well, authentication is now primarily password or credential based. If behavior is used as a key authentication technology, would a change in behavior be a serious security breach?

Pfleeger: What are the sensitivities and specificities? That is, how bad are the false positives and false negatives? This will tell you how much you should worry about using profiling.

Losavio: You should also ask, "What are the consequences to the individual from a particular form of use?" Targeted advertising might be an irritation, but it's certainly much less intrusive than an arrest warrant. We're just now seeing statistics-based cases being made for search warrants in the criminal justice complex, to see a person's computer, to search his home. How far can we go with that?

Rogers: We must think carefully when we're talking about rates here. In certain circumstances, if you make a false positive that results in somebody's going to jail, that's a lot more serious than making a false positive that results in someone's trying to log in to a computer system again. We should concentrate on how accurate it needs to be versus how useful it needs to be.

Caine: But even the advertising examples aren't harmless. They can affect real social interactions. When advertising is repeatedly directed at them, people become aware of faults, things that they "should be taking care of." Maybe you could frame that as being good, but in a lot of cases, there could be a serious impact on relationships that we haven't even realized yet.

Bashir: If you don't know what's being collected and why it's collected as well as how it is going to be used, you can't make informed choices about protecting your privacy. We need to consider the rules and guidelines of how data is collected from human subjects in social science research studies and how informed consent is an important aspect of any data collection and analysis.

Pfleeger: Right after Google announced that it was unifying its privacy policy, I heard a woman on a talk show saying that she was appalled that after she had done an online search about diabetes, she started getting emails tailored to people with diabetes. So every day, she now goes online and searches for a different disease, just to try to throw everything off track. What should we do about user behaviors that try to avoid the very kinds of purportedly beneficial things profiling offers and that could make predictions worse instead of better?

Stolfo: We should do "fog computing" as opposed to "cloud computing." The notion of having avatars that represent various versions of your personality would help protect you. Many people have just given up so much personal information, it's lost forever and everybody knows it forever—except if you suddenly start populating your profile with bogus information. You know the facts, your friends know the facts, but others don't. Decoys or other personas could be a very useful way of protecting people's privacy by simply making their public data no longer certain.

Caputo: But there are consequences if they're trying to find a job and the prospective employer finds the bogus persona.

Stolfo: But we don't want folks to rely on social media sites to make decisions about people.

Caputo: But then who believes the truth? It's just like when someone tells a lie on the stand. The jury no longer believes the rest of what he has to say.

Rogers: We have the same problem in the criminal world: it's called staging behaviors. Traditionally, the criminals who weren't stupid figured out that the reason why the others got caught was because they left clues behind that pointed to their personality. The cleverer criminals combated this by staging their crime scene, trying to make it look like somebody that they're not. It's not new.

Caine: It's much broader than even criminal justice. It's basic psychology. We present different selves to different groups of people. That's just a behavior that humans engage in. People engage in avoidance, modification, and alleviatory behaviors to preserve their privacy. It doesn't mean that they're necessarily trying to hide something nefarious. They just don't want a corporation knowing everything about them, or they don't want all of their friends knowing every single thing about them. And those behaviors happen offline, too. What's exacerbated online is that people don't always know all the information being collected about them, and they don't have the analytical capability that corporations and governments have to engage in predictive modeling about how they'll behave in the future. I want to reiterate that this could be a paradigmatic shift. It's now information that I haven't even come to realize about myself that potentially puts me at risk, the future consequences of which I have no idea.

Rogers: Once everybody starts creating different personas, it's hard for them to not have at least a little bit of a common thread. Given enough data, time, and effort, I think the identification will still happen. That is, in the short term, personas might help confuse the system, but once multiple personas become a habit that everybody uses, they lose their effectiveness.

Stolfo: I disagree. I tend to think of this as a traditional cat-and-mouse game between attackers and defenders. Deception has been an effective tool since the time of the pharaohs.

Losavio: Yes, but with enough power to compute those same patterns on your behalf as a service, you can create other patterns to confuse your real pattern: a useful technology to protect people. You as a defender have the same technical basis as the adversary to protect yourself.

Caine: Marc's right. If we end up with a high false-positive rate, we can probably figure out things statistically. But then we're also throwing a lot of people in a category.

Stolfo: As a scientist, I tend to worry also about false negatives. You can create a folk theorem that says, "Any machine-learning algorithm will have false positives; most people think that way." But with hard work and some practical implementations in a limited context, you can manage false positives, which aren't evil, per se; they're informative. And depending on your mitigation or response, false positives might not even be that dangerous. Folk theorems are easy for people to cite and state. But without doing an actual experiment with real systems and environments, you can't make an assertion saying something is bad or good based on a rate that you have no idea about. You have to do the science.

Rogers: But we also have to remember that we're dealing with people who aren't as good at making things up or doing things as a computer program might do. There will be a common filter, a behavioral characteristic that will leak through no matter what personas are put there. That's just human nature. We know it from 50 years of research—we're just not as clever as we think we are.

Caputo: A lot of good social science demonstrates that people leak when they aren't trying to—through their actions.

Stolfo: But don't you think that that knowledge is useful for people building a defensive system as a service? They can identify when leakage is occurring and use algorithms to find patterns that can confuse whatever information is leaked. In other words, even though people themselves will make mistakes, automated systems and services can in fact assist humans in not making as many of them.

Caputo: Automatic systems don't anticipate consequences, so putting that new profile out there to confuse is risky.

Stolfo: What we have now is a mess: nothing but risk. Any adversary has unfettered access to everybody's data and information.

Losavio: Obfuscation testing is probably a good idea. But consider the woman who researched other diseases to throw off any tracking. The danger is that somebody collecting that data might say, "This is a seriously ill woman. We don't want to provide her with health insurance." If we test this type of obfuscation, we've got to comply with an internal review board's ethical requirement for human subject research. By contrast, retailers have practically no ethical limitations whatsoever; neither do life or medical insurance companies.

Stolfo: Regulators will see that this is happening, and there would be obvious ways of dealing with falsehoods. If you think one step forward, if those methods of deception do cloud a fact for an adversary, then they cloud that fact for everybody—in which case, they can't be depended on. Business people aren't stupid. If they make bad decisions on bad data, their business will suffer, and they won't use information if they suspect it's wrong.

Losavio: Yes, if in the aggregate it affects profitability. But if businesses can cut off people who might be high-risk, high-expense individuals in their insurance pool, they'll do so.

Stolfo: Exactly—if everybody is high risk because their profiles are fog, there is no business. Businesses need alternative ways to get to the facts, for example, through doctors' offices.

Pfleeger: Certain decisions are being made on the basis of collected data, but many people don't know that these decisions are based on profiling—for example, you can't get on a plane, you can't get a mortgage, or you can't get a credit card. You're usually not given a reason why. If you find out that the profiling made a mistake, you might be able to go through some sort of process to fix it. But if you're a business person trying to get on a plane and you're denied a boarding pass because of something found in a background check or a search term, you don't necessarily have good options. How do we manage the uncertainty in the process, and where do we find the right balance between telling people what's going on and not telling them? There's a difference between the security of the individual and the security of the society or group in which the individual is functioning.

Rogers: That is such a delicate balance. By letting people know that you're actually collecting information and observing them, and that you're using it to make a system analysis or profile, it could cause them to slightly alter what they're doing. It's a trade-off between how accurate you need to have your models versus how much information you can give them. Potentially, people become deceptive.

Bashir: Initially, people might change behavior a little bit just to meet that social desirability or expectation. But if it goes on long enough, people tend to return to their normal behavior. They should know how much information is being collected and why.

Caputo: There's research on youths who don't understand that the stuff they post has consequences. The data has shown that education leads to awareness and understanding that persist for a short period of time. But in the long run, it doesn't change their behavior unless they've had immediate consequences, such as identity theft. Simple awareness doesn't seem to be breeding change.

Caine: Some of that research is flawed, because there's usually no option where the collection isn't happening. You either use the system and information is collected about you, or you don't use the system. We haven't had a lot of research yet that gets at the question of whether people would prefer to use a system that simply doesn't collect the information but still gives the same benefit.

More generally, I think people should be aware of what information is being collected about them and what consequences come with that. That's the only way to strike this balance between society and individual. Otherwise, when consequences come—you show up at the airport, something happens to you, you have no idea why, or what information led to that—you aren't able to modify your behavior so that it doesn't happen again. Instead, you begin to live in an incredibly stressful, uncertain world, and we have no evidence about what kind of societal consequences that breeds. To say that we just end up with a better society because we think that's probably what will happen—no, we need a lot more research to determine that.

Stolfo: But if the commercial world finds that folks are distressed because of profiling, I believe strongly it'll improve in time.

Caputo: The counterintelligence and counterterrorism communities will tell you that revealing what we're looking for gives our adversaries an advantage: all the data they need to game those profiles. We have to weigh the stress against the protection offered by the profiling. So you don't go through the security line and hear, "Okay, that's a false positive. You're upset that you missed your vacation, and I'm sorry you're really stressed out." But neither do you hear that profiling kept somebody else from going through that same security line with malicious intent. The system doesn't reveal what it deterred or prevented, because it gives away its collection capabilities to the adversary. It's a tough balancing act.

> **What are the true consequences of an action? How does it upset people? How does it protect us from our adversaries? —Deanna Caputo**

Rogers: London probably has the most intrusive video system in the world. The number of video cameras is just incredible. And over the long term, researchers found—guess what—that the cameras had no impact on decreasing the crime rate whatsoever. So the claim that, by monitoring everybody, you're going to make it safer hasn't held up. A lot of what we claim is very anecdotal; there hasn't been a lot of empirical research to support either side yet.

Bashir: Very few empirical studies can guide us in how we're supposed to go forward. How can we research some of these topics, so we can get real data and then make informed decisions?

Caputo: We have to combine researchers with operational folks. Some of the research is done in a vacuum. What are the true consequences of an action? How does it upset people? How does it protect us from our adversaries?

Caine: It would be fantastic to get the data out in aggregate about how many actions have been averted using these kinds of different methods.

Pfleeger: But that's the problem in cybersecurity in general, isn't it? The absence of attack doesn't really tell us anything. We don't know whether it's because of something we did, or something we haven't done. Having good measures of the effectiveness of prediction and profiling is difficult.

Rogers: And you have to be very careful about mistaking correlation for causation.

Pfleeger: Which naturally leads to the next question: What are the legal and ethical issues? How do we deal with the false positives and negatives? What are legal and ethical ways to use profiling? For example, what if the President, based on prediction or profiling, thinks that part of the critical infrastructure is in jeopardy and wants to shut it down? Should that decision be based on some sort of prediction? The consequences could be much more severe than not being able to get on a plane.

Losavio: There is something of a data security regime in place under the US Privacy Act. It has certain basic fairness provisions in it, one of which is revealing what you're doing with data as a federal agency, and if you're going to act on data based on some analysis you've conducted, you must give notice to the people who will be affected—and they have a chance to challenge it. However, one of the problems is that you can't challenge their decision. Suppose, for instance, you receive federal benefits, and the government says, "We've got an analysis under our match program, and we think you're no longer eligible for these benefits." You're told you have 30 days to file a response if you think the judgment is incorrect. You say, "You've got me wrong, and I want to see the data." You think the government has associated you with another individual with the same name. But it's different for private data holders in the US. It's nothing at all like what they have in Europe. If you were to try to implement equivalent options in the private sector through legislation, that might be a way to protect people, by requiring big data brokers to do something correctly or face liability for it.

Caine: I'd be very careful in intervening in a case like that when we have a prediction that an individual might be more likely to commit a crime. It's one thing to predict that people might engage in a behavior, and it's another thing to actually encourage them and change their behavior so that it conforms to what your algorithm believes they might eventually do. Confirmation bias already exists in our current datasets, and algorithms can further exacerbate those kinds of biases.

Pfleeger: When is profiling most useful, least useful, and how should it fit into cybersecurity tools?

Rogers: It's least useful when you don't have any empirical data to support the conclusions you're making, and they're instead based on clinical or anecdotal evidence.

Stolfo: In the very limited case of user authentication, it could be remarkably useful. It gets more complicated in environments typified by change. But in an environment where you don't download and load applications every day, and you have essentially stable systems for some period of time, it could be a very effective tool in identifying masqueraders.

Losavio: Given the size of these datasets and the questions you're dealing with, there may be no other way to practically deal with this. How do you filter against what you use it for and the injuries that it might potentially cause people if you're incorrect, whether it's a privacy violation or a decision on employment, or health insurance, or whatever? How do you protect the liberty of citizens to be left alone and not have the government intrude on their actions? There's an interesting parallel with drug courier profiling. If you look at US law regarding what it calls the "drug courier profile," the courts make a distinction between "grounds to direct an investigation toward an individual" and "grounds to take legal action," like detention of an individual. The US Supreme Court has been very careful to avoid actually resolving that legal issue, because it's not quite sure how to deal with it. The resolution involves juggling public safety with efficiency and liberty.

Rogers: And if you move away from the individual, it can be incredibly useful for target hardening. How do you make potential targets less attractive to those people who want to attack them? We refer to it as *victimology* in the social sciences. How do we build more robust, more secure operating systems, computer networks, databases? This is another practical application that's less of a privacy issue, because you aren't necessarily identifying an individual.

Losavio: Profiling tells you what to target, where to focus. The courts are willing to give the police leeway to do some things that are a limited invasion of a citizen's freedom of movement. They can say, "Excuse me, sir. Can I talk to you for a moment?" rather than, "Stop! You're under arrest," based on this correlation of characteristics.

Pfleeger: But because we look at people differentially, we currently prosecute people differentially. The bias in the baseline dataset results in bias in our actions.

Losavio: The police always say, "Race wasn't a factor in our stopping this individual," even though African-American females are nine times as likely to be given a detailed search at an airport as Caucasian females. Just a coincidence.

Caine: Randomization is a huge counter to this issue, as are the set of fair information practices. If we're trying to maintain the same level of privacy that we had 40 years ago, before extensive information technology systems existed, we need to embrace those practices.

Rogers: We also have a sliding target here. What is privacy now versus what privacy was considered to be 20 years ago? How will we define privacy 10 years down the road? The time scale has interesting implications. People give up privacy for a perceived societal benefit.

Pfleeger: But there has to be a demonstrable, near-term yet persistent societal benefit. For example, companies don't report data breaches because there are no long-term consequences. But airlines don't punish those reporting problems, so there has been good compliance with reporting.

Caine: There are potentially bad outcomes if we think we're doing this well but in fact we're doing it really badly.

Pfleeger: Clearly, there are far more questions about profiling than answers. But equally clearly, we need to find answers, because the impact of profiling can be profoundly positive or negative. The issues raised here are good topics for future roundtables. ■

Shari Lawrence Pfleeger is the director of research for the Institute for Information Infrastructure Protection at Dartmouth College. Contact her at shari.l.pfleeger@dartmouth.edu.

Marc Rogers is a professor and university faculty scholar at Purdue University. Contact him at rogersmk@purdue.edu.

Masooda Bashir is assistant director for social trust initiatives at the University of Illinois's Information Trust Institute. Contact her at mnb@illinois.edu.

Kelly Caine is a principal research scientist in the School of Informatics and Computing at Indiana University. Contact her at caine@indiana.edu.

Deanna Caputo is a lead behavioral psychologist at MITRE. Contact her at dcaputo@mitre.org.

Michael Losavio teaches criminal procedure and law and computer forensics at University of Louisville. Contact him at mmlosa01@louisville.edu.

Sal Stolfo is a professor of computer science at Columbia University. Contact him at sal@cs.columbia.edu.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*