

In Conversation With... Suchi Saria, PhD

July 1, 2019

In Conversation With... Suchi Saria, PhD. PSNet [internet]. 2019.

<https://psnet.ahrq.gov/perspective/conversation-suchi-saria-phd>

Editor's note: *Dr. Saria is the John C. Malone Assistant Professor of computer science, statistics, and health policy at Johns Hopkins University. Her research focuses on developing next generation diagnostic, surveillance, and treatment planning tools to reduce adverse events and individualize health care for complex diseases. We spoke with her about artificial intelligence in health care.*

Dr. Robert M. Wachter: Tell us a little bit about what got you interested in health care, [artificial intelligence](#) [AI], [machine learning](#), and analytics?

Dr. Suchi Saria: I got interested in AI when I was young. The journey went from robotics into software design and machine learning, which is a way to think about using data to learn and build programs that mimic human intelligence. The further I went in the field, the more exciting it became. I felt like it was all about understanding decision-making. How do humans process gobs of data to figure out what is the right thing to do in any number of situations? We're not so good at combing through lots of data in our head, but it's helpful to see what principles we need to design when we build software to do that. That is how I think of the genesis of machine learning and AI.

RW: What got you jazzed about health care?

SS: Back in 2007, I was at Stanford and I worked on taking large amounts of time-series data and thinking about modeling users over time. I wanted to work on something that could have human impact, and I had an opportunity to visit the neonatal ICU at Stanford. I remember talking to my collaborator, [Dr. Anna Penn](#), who was a neonatologist. As I walked through the floors of the NICU, I said to Anna, "All this data on the monitors looks like time-series data—what are you guys doing with it?" And she said, "What do you mean? We're using it." These are tiny fragile babies. Things like necrotizing enterocolitis or sepsis—if you don't anticipate it and treat it in time, it can have not only devastating consequences almost immediately but also long-term neurodevelopmental outcomes. So, data was my starting point for learning where the use of data in health care was and where I felt we could go bringing this different perspective that I had.

RW: When you saw how stunned she was by your question and began to learn how little use health care makes of its data, did you have theories as to why that was?

SS: No, it took me a lot of learning. Working across disciplines is a very interesting exercise. Because you form a theory, but the problem is you only have a third of the picture in your head. In health care, there are so many stakeholders. I didn't even know if I was asking the right questions in the beginning. I'd ask: "Why aren't you using the data?" But from her point of view, she's hearing: "Of course, we're using the data in clinical care. I'm looking at the monitors." But from my point of view, they are not really using the data, which would be by precisely modeling it, forecasting, and tying it to actions that she might use within delivery. Similarly, I didn't understand at the time that to use the data we needed the right infrastructure. I also didn't understand that we needed the right administrative-level prioritization to know that of the many things they're thinking through, this should be top of mind. It requires thinking through the right infrastructure, the right level of education. Identifying the stakeholders and whether they understand what you're talking about, so they can engage with you and be brought into this process. All of that was learning I had to do.

RW: It sounds like the doctor is not even thinking about the data other than in real time. I look at a monitor; I see it. I look at a lab test on the chart; I see it. They're not thinking about the potential learning from aggregating it and analyzing it in new ways. It also sounds like the consequences of a profound lack of an incentive system to build the infrastructure and prioritize that is different than I imagine you've seen in other industries.

SS: Yes, indeed.

RW: Do other industries resemble health care in a way where they are using data in such robust ways that it could inspire health care, but health care is 5 or 10 years behind?

SS: At the same time as I had this NICU experience, I joined a Stanford company that was building infrastructure when retail was going through this transformation. We used to have a physical bookstore with inventory on hand, but we were shifting toward the paradigm where people were shopping on the web. As a result, we could learn a lot more about a user. Who is our user? Where are they coming from? What do they do? What do they tend to like? Based on that, we could customize their experience to give them what they needed. In this day and age, you cannot imagine a single retail website that doesn't track and monitor users to create a rich user experience. It's amusing, because we go through that whole experience by using thousands of cycles to optimize which shoe to show you next. In health care, we make decisions about which surgery to offer you, and the amount of computation there is not that substantive.

RW: When you thought about how little we mine of the user's experience in health care versus in retail, did you have a sense of whether the issue was mining patient data or mining provider data or both? In both parties, the predicament is that things could be learned about the work from the data that we've not been taking advantage of.

SS: Instead of what's the data and whose data is it, we could almost start from the problem. I see a potential for AI to be deeply transformative in health care. It feels like it's grossly underutilized. But as I talk to senior leadership and senior stakeholders, I realize that we're talking about the solution not the problem. There is a lot of hype and noise around it, and instead of building positive momentum for the field by talking about solutions, it's just creating a lot of confusion about what is and is not possible. I'd like to change that.

There are problems everywhere, and then working backward to see what's the problem we're solving? What data would we need in an ideal world to do it well? Then, thinking through what an ideal solution would look like and putting it together.

RW: If you watched the "Jeopardy" match with [Watson](#), you sort of figure this is not that hard, and AI is going to take over. Yet as you mentioned, there has been a fair amount of hype. In health care, what are some hazards you think we need to make sure we're thinking deeply about as we implement more of these systems?

SS: There are several example issues to think through. Let's start with [alert fatigue](#). I see many systems where they say they're taking advantage of the electronic health record (EHR) to implement systems for detecting at-risk patients. In a way it's great, because it's building organizational muscle for implementing electronic systems to do things. If we do it not very thoughtfully, you could easily imagine scenarios where you have systems alerting all over the place. For instance, I lost my nephew to sepsis—a scenario close to my heart. I've worked on it extensively and tracked what others are doing. I know systems where they have alert systems and get over 250 alerts a day in a single 250-bed hospital. If only 1 in 10 alerts are right, that is a huge amount of burden on our physicians and on the whole care team. Instead of making people more productive, more efficient, or more effective, we're actually harming patients.

RW: When you approach a problem like that, how do you even begin to unravel it?

SS: When I talk to health care systems, they say they're taking advantage of the EHR and they're detecting these patients using "like AI." The challenge is they are often using a one-size-fits-all model. Patient symptoms are often very heterogeneous. What looks normal for a patient who's immunocompromised is not at all the same as someone who's a young athlete. These one-size-fits-all systems often generate a ton of alerts, which then lead to fatigue. Done correctly, we would be thinking very hard about the patient's context and taking that into account. We must fundamentally solve for these kinds of heterogeneity as we're generating these kinds of alerts. Tuning into what the patient's context is and in that context what's normal, and then using that to identify who looks like they're deviating and going downhill. If done correctly, it can dramatically improve accuracy and reduce false alerting.

RW: And that is done through the analysis of many patients to say that, in this context, this alert always turns out to be a false-positive so don't fire an alert? Or is it done through a rules-based mechanism that says the literature says that in these kinds of patients, the alerts are always meaningless? Or is it a combination of those two things?

SS: The good part about the rules-based approach is if you know exactly what you're looking for, you can program it in and then you're good to go. The challenge with the rules-based approach is that your patient population is highly heterogeneous. So, your rules-based approach forces you into often a very coarse rule set, which goes back into the one-size-fits-all. You're not tackling patient context to the level of detail as one should. The view now is that we just take millions and millions of records and put it into a magical algorithm, and this magical algorithm—you might call it deep learning—is going to figure out what the answer is and now we're good to go. To me, this has several issues: not only is it unlikely that it will be accurate because often the data is cruddy and contains things that you need to correct for, but also, you're

not leveraging any of the domain, clinical, or institutional knowledge that you've developed over time. In practice, what I've found is that the right solutions are often in the middle. It means you want to take advantage of what you know in terms of heterogeneity in patient population, the kinds of contexts to account for, the ways to correct for biases. But at the same time, what machine learning, or AI, is good at is learning patterns from data. Once you've titrated it to a very small homogeneous population in that context, it can surface ideas precisely for what other kinds of signs that start showing up. What are the patterns we're seeing in terms of signals going awry as this person is heading toward an adverse event?

RW: In looking at the alert fatigue problem, people also talked about physiologic plausibility. For example, in the ICU you have the heart rate and the blood pressure monitors and each one has a fixed parameter. But if your heart rate immediately goes from 60 to 150 and your blood pressure doesn't drop, that cannot be physiologically real. Yet today, an alert fires because each channel is independent. How important is that as a problem?

SS: Absolutely. Imagine if you had to make a decision, but each signal was being shown to you one at a time and you had to make a very complex multifactorial decision about whether this patient is headed toward sepsis. You'd probably struggle with it as a physician. That just makes no sense at all, this idea of rules implemented one signal at a time. We should take as much of the patient context into account, as much of the longitudinal history into account, and as much of historical records we have of other patients into account to learn in that very specific context what the right thing to do is.

RW: Diagnosis is an area that people think a lot about and might be one of the harder problems in artificial intelligence. How hopeful are you that the computer will either take over for physicians or meaningfully augment what they do?

SS: This will probably depend a lot on the specific type of problem we're talking about. My belief is that in the next 10 years or so, 80% of the problem, if not more, is going to be of the type where you can think of it like a human-machine team. Essentially, we're leveraging our data to surface the right kinds of intuitions just in time to say here's a patient who looks like they're deteriorating. Here's a patient who looks like their care needs to be escalated, and here's why. Now it gives something for the caregiver to engage with and react to. Previously, they wouldn't be standing by the bedside watching the patient continuously. But here is an example where just having that prompt, and of course not prompting all the time, but if the prompt was accurate and the prompt could be trusted, then that would be a very productive relationship. There is complementarity here in terms of skill sets. When I say complementarity—we're not good at processing gobs of data in our heads. But we're very good at thoughtfully considering context because we're right in front of the patient in ways that the machine doesn't have context at all. But the machine can look at gobs and gobs of data to understand where something subtle might be developing that you might have otherwise missed.

RW: What is your sense of how AI is doing in radiology, dermatology, pathology, and fields like that, and what risks flow from that kind of work?

SS: We're seeing a lot of activity and productive use in radiology, especially ways in which the images can be processed automatically to prioritize what important tasks need to be looked at, which images need to

be looked at first, or problematic areas within an image. Then there are also what I would call augmentation opportunities, where you're learning patterns from these images. So scanning retinal images to identify risk of diabetic retinopathy is a good example where people are using AI to augment and scale and increase access. But overall, I see opportunities in a lot of places. There is confusion at scale about productive areas to work in. Radiology has a clear business model, and companies are making a lot of progress. We're starting to see more progress in clinical operational areas as well, but relatively early in terms of successful scalable projects that people are discussing.

RW: Do you worry at all about the problem that has surfaced in aviation, where you still have a human operator, but the human operator may not be paying much attention because the machine is right virtually all the time. Then when the machine fails, the human operator is not particularly good at the task anymore because they've never had to do it in real-world emergency circumstances.

SS: This is where education matters so much and how we think about educating our stakeholders as we roll out any kind of technology or any solution. In my own personal life, I had my nephew who went undetected until it was too late and we lost him to sepsis. That sort of harm is happening every day. There is that statistic about a [jumbo jet](#) crashing on a daily basis and we're doing very little to prevent it. There is harm on both sides, which is us not doing anything is in its own right very harmful. If we thoughtlessly implement technology for technology's sake and rush through things, it can also be harmful. So, what is that middle ground? It has to be a combination of rigorously validated careful risk analysis and deep understanding of what's being put out there. But also careful education and pairing so that it's clear that the individuals who are using it understand how to not become overreliant in the wrong ways.

RW: I like that answer, because sometimes you can just be pointing to all of the potential harms from technology and forget that the harms that happen every single day in a nontechnologically robust health care system, they happen all over the place. My stepdaughter is a first-year medical student. Will she have a job 15 years from now?

SS: Absolutely. To me, the health care field is a lot about the human touch. It matters—you're going through emotional times, often tough times. Things are happening; you're fearful. You want to have clarity on what's going on and what needs to get done. Is everything being taken care of correctly? This is going to sound funny because so much of people's experience of technology is through the EHR, which they have very strong impressions of. But if done right, I see the use of AI and the use of technologies that can really tune into context as a way to make it possible to bring the human side of health care front and center again, and making it possible to deliver care—that you're focused on your patient and trying to do the right thing and being enabled to know what's the right thing to do.