

Episode 6: Hidden in the Data: Understanding Bias in Informatics

Transcript

Co-Hosts: Dr. Tiffany Leung (Tiffany), Sarah Takimoto (Sarah)

The DEI Shift Guests: Dr. Sara Murray (Sara), Dr. Michelle Lin (Michelle)

[0:00] Intro

Tiffany: Welcome to The DEI Shift, a podcast focusing on shifting the way we think and talk about diversity, equity, and inclusion in the medical fields. I'm Dr. Tiffany Leung. I'm an internist and clinical informatician, and we're so excited to share with you this episode.

Sarah: I'm Sarah Takimoto, a graduating medical student and soon to be Internal Medicine physician. We'll be your co-hosts for this episode.

Tiffany: Today, we will be talking about how bias can be perpetuated in informatics work—from machine learning, to guiding evidence-based medical care, to algorithms designed to support workflow and the practice or business of medicine. Clinical informatics is the application of informatics and information technology to deliver healthcare services. It's also referred to as applied clinical informatics or operational informatics. So, why would we think of talking about informatics in a podcast about diversity, equity, and inclusion? Well, it's precisely because of this need to be intentional about diversity and informatics that we're covering this topic today.

Sarah: That's right. Human bias exists, and as humans collect data, deploy machines to analyze data, and take actions based on data, our biases are perpetuated. Bias pervades the work of medicine and it's important to remember the use of machine learning, decision support, operations and logistics planning and health systems and various informatic applications to guide clinical practice are not exempt.

Tiffany: So this is the focus of our discussion today with our guests covering the intersection between medicine, health disparities, and informatics. We'll also touch on what rules data has had in the COVID pandemic. Sarah, do you want to introduce our first guest for the podcast today?

Sarah: Yes; we have two fantastic guests who will be helping us learn more about this topic. The first is Dr. Sara Murray. She's an assistant clinical professor in the division of hospital medicine at UCSF and serves as the medical director of clinical informatics for UCSF Health, and her role is leading the advanced analytics and innovation team. She's involved in projects using electronic health record (EHR) data to inform quality and value improvement efforts.

[2:16] Sara: Thank you for having me. It's great to be here.

Tiffany: Our second guest is Dr. Michelle Lin. Dr. Lin is an assistant professor of emergency medicine at Mount Sinai in New York City, where she also serves as director of performance improvement. Her work focuses on measuring and improving health outcomes by integrating new and existing data sources.

Michelle: Thank you; it's great to be here.

Sarah: Thank you both for taking the time to join us on The DEI Shift. We are so excited to have you both here and before we start, we'd love to get to know each of you a little better. We start every episode with 'A Step in Your Shoes' segment where we ask each of our guests to share an element of their background or their culture that has been important in their life.

Transition

[3:06] Sarah: Dr. Lin, would you start us off?

Michelle: Sure, yeah, honestly, the first thing that came to mind was food. I'm Taiwanese American, my parents immigrated from Taiwan in the '70s, and we're in the middle of a COVID epidemic right now and the only foods that I want to eat are the same foods that my mom used to pack for me for lunch, and would make me cringe because inevitably, they would smell, and they were not like everyone else's lunchables or ham and cheese sandwiches. But I find myself wanting those things and having a hard time getting them because I live in New York City without a car. So home cooked Taiwanese comfort food was something I grew up with and will always miss.

[3:59] Tiffany: So Dr Murray, what about you?

Sara: Yeah, so I actually also have a story about food. I grew up in Virginia, and when I was younger, most of the travel that we did was on the east coast, but we took a one trip to San Francisco when I was about eight or nine to visit one of my mom's friends who had moved out here. Over the course of a few days we saw unbelievably beautiful sites but we kept eating in restaurants that serve sourdough and clam chowder, neither of which I liked at the time. So as a child, I actually swore off San Francisco that I'd never returned here because I didn't think there was anything to eat here besides sourdough and clam chowder. And so of course, many years later, a friend convinced me to interview for medical school out here. And now I've lived here for 15 years and it turns out, not only is there a lot to eat here besides clam chowder, but I've actually come to really like sourdough bread. I think the bias that I experienced on my childhood trip is not unlike the biases that we're going to talk about. We worry about in EHR databases, or experiences, or data that we can collect that can really misrepresent reality.

Sarah: Thank you for sharing that. I love how you were able to tie it into what we're talking about today and I'm also glad that you gave San Francisco and sourdough a chance.

Transition

[5:27] Tiffany: So let's get started. We're talking today because diversity, equity, and inclusion are important to clinical informatics and even maybe sometimes neglected in the fields. What do you think are the important areas of intersection between diversity, equity, and inclusion (DEI), and clinical informatics applications for patient care? Dr. Murray, would you like to comment first?

Sara: Sure, so I think there are a couple issues here. The first and most obvious intersection that I think about is whether the informatics tools that we deploy are adequately accounting for DEI and the patient populations that we're serving. So, if we have a patient portal that our patients are supposed to use to learn about their healthcare conditions or schedule appointments, can our patients with low health literacy or limited English proficiency access the tool in an equitable way? And this is something that we think about a lot, but I think is really challenging to fully address. A slightly deeper and more complex intersection that I think about a fair amount comes into play when we think about using patient data to drive decision support or inform patient care in some way because that data—like all data—can be prone to bias. And so we want to be really conscientious about issues of DEI with respect to that. So finally, I think a related issue is the issue of who our leaders are in informatics. So I'm an assistant program director for our clinical informatics fellowship. And I'll tell you that our first several classes of fellows were not that diverse. I was excited that we were able to recruit our first women last year, but we still have a long way to go, and we know that organizations perform better when there's diversity in leadership. So we really need to make sure that diversity exists in informatics leadership as well.

[7:29] Sarah: I'm so excited to hear about the increasing diversity in the clinical informatics fellowship and the point you brought up about leadership really resonates with me. Dr. Lin, do you also have thoughts about diversity, equity and inclusion in clinical informatics application?

Michelle: Yeah, absolutely. So I identify as a researcher by training and not as an informatician, but clinical and EHR data is one of the data sources I always consider. And we always think about where does data need to come from if we're thinking about something as seemingly straightforward as you know, race for example, right, who's recording that? Who's entering it? Often, if it's an electronic health record, it's a clinician. But we know that patient-reported race, for example, is probably the most accurate. When I look at something like administrative data, which is often used for quality reporting, that's often a registrar, for example, and sometimes they're not directly asking the patient. They may be assuming because in the ER where I work, it's often very busy. The second dimension for diversity, equity, and inclusion I think is with respect to quality measurement, there's a lot that is either incorrectly captured or not captured well. And if we're going to use that data to determine which hospitals and which clinicians are providing high quality care, I think that has a lot of implications because we aren't always accounting for the types of things that influence health and wellbeing. And so, I think we have to think really critically about what's being captured and

how that's being used, and making sure that we're not necessarily penalizing providers and hospitals who choose to care for disproportionately vulnerable groups.

[9:17] Tiffany: That's really great, Dr. Lin, there's something there that you said that caught my attention, in particular, that there are data registrars who are potentially entering that data and making assumptions. Can you expand a little bit more about what you meant by that?

Michelle: You know, when I wear my research hat, I often have to also put on my clinician hat, and you know, when I'm using a variable and a dataset, I often have to consider, okay, well, where in my clinical experience might have that been captured? And when I think back to busy ER shifts and I picture our really amazing registrars who are doing their best in a really stressful environment, particularly in the times of COVID, I see them next to an unstable patient and they're trying to record all the information as quickly as they can to open up a chart because otherwise we can't place orders, and they're doing their best, but, it may or may not capture everything completely or accurately. And those limitations need to be taken into account.

[10:13] Tiffany: Absolutely. That's really a great point because what I think you're raising is that the clinical workflow and that sort of natural prioritization of patient care is absolutely paramount. And yet you're trying to balance that with collecting enough data in a way that's useful. For example, to be able to power quality measurement and other important objectives for the healthcare institution, as well as for the population that you're serving. What examples have each of you seen or studied, where bias and data or how it's used has potentially actually exacerbated health disparities in certain patient populations?

Michelle: When we do what's called risk adjustment, which is accounting for some of the things that drive health outcomes and risk for higher or lower utilization in certain populations. The most commonly used approaches include looking at comorbid chronic conditions. Those tend to be hot, captured really well in, for example, primary care settings or in inpatient settings. But if a patient doesn't have a consistently established regular provider, or only comes to the ED, for example, for whatever reason, often because of barriers to access, often those comorbidities don't get captured as accurately or completely. And so when we examine those populations, it might look like they're less sick, but it could actually be that they're more sick. And that's why they can't make it to where they're supposed to be. An example of this would be my former colleague, Dr Obermeyer's work, where he looked at a commonly used commercial insurance algorithm, and it was discriminatory against black patients because they used prior healthcare needs, which were lower because of their barriers to access to predict their future needs, and thus ended up expanding fewer resources on these patients who were actually sicker.

[12:17] Sarah: Thank you for sharing. And Dr. Murray?

Sara: So my favorite example of this is predictive overbooking, which I recently wrote about in the Health Affairs blog. So when patients miss their ambulatory clinic visits, it's a big problem for clinics which rely on the revenue and it's a big problem for other patients who need to see their providers and could have used the spot. And because of this, there's a lot of interest in predicting who might no-show or miss their visit so that you can do something like targeted overbooking where you overbook the patients who are most likely to no-show. The problem is that if both patients arrive, you might have an overworked and rushed clinician, which could negatively impact patient care. So our team unpacked EPIC's first model for predictive overbooking, but what we found was quite concerning. So we found that it included a fair amount of sensitive patient demographic information that we worried might lead to discrimination against vulnerable patient populations. The deeper problem here is that even if you remove all of that information and just base prediction on the patient having had prior no-shows, you still have to think about all of the existing structural inequities that may have influenced that. So, did the patient have adequate access to transportation, childcare, sick-leave, and if not, and now you overbook them, but they managed to make the appointment, they're splitting their time with another patient, while they may be the person who has the greatest need for care.

[14:04] Sarah: You know, as a graduating medical student and a future primary care physician, there's a lot of trust that I place in tools built into the EHR. And when I hear about a no-show prediction model, my first thought is, that sounds great. And after listening to our discussion, I have a lot of fear that during my intern year, I may be using models that perpetuate bias, and I won't even know it. And quite honestly, that's a pretty scary thought.

Tiffany: I think the comment that you're getting at, Sara, is how much do you trust these algorithms and how much do you trust them based on what you know about. And I think this is a really important debate that's coming up in terms of the ethical and social issues around not just artificial intelligence (AI), which is I think the more and more popular, ethical conversation to have, but just generally about any of these sorts of predictive algorithms that may not be as complex. Once you identify these issues with algorithms, what do you do next?

[15:12] Sara: So one of the fundamental questions that you're highlighting is what is it, what does it mean to be trustworthy AI? And this is something that people are thinking about in a lot of industries, not just healthcare. And, in fact, the European Commission published guidelines on this last summer, and there were seven key requirements. So there has to be human agency and oversight. There has to be technical robustness and safety. Algorithms need to encompass privacy and data governance. There needs to be transparency. There needs to be diversity, non-discrimination, and fairness. Algorithms need to consider societal and environmental wellbeing. And finally, there needs to be some accountability and mechanisms for auditing. So I think one of the issues we think a lot about in healthcare is this issue of transparency when we use algorithms. So clinicians are used to algorithms that are really transparent, and we all have these apps on our phone where we'll put in a couple key variables to predict the music score, for example, or if you're trying to calculate a meld or some other score that we know is validated and we know

exactly what the inputs are, we lose that with predictive algorithms when they're deployed in a black box. And that's for obvious reasons; you don't know what the inputs are and you don't know where the predictions are coming from. So there's a lot of thought right now about what it means to be transparent in an algorithm. And companies are working on this. So EPIC is working on this for example, and in many of their clinical algorithms, they will give you the output; the output will not just include the score, but will also include the relative contributions of certain features. So for example, if you're predicting the remaining length of stay, it might tell you that, you know, their length of stay is this, and one of the key features that's contributing is their pain scale or some lab, or something like that.

[17:52] Sarah: It's helpful having you make the distinction between scores and models that we're very familiar with and using some of these newer predictive algorithms that give us the appearance of knowing where they're coming from and what data they're based off of by sharing the key variables and features, but still, are operating and using that data and a black box. Can you tell us a little bit more about why this distinction matters?

Sara: I think the really tricky thing here is this question of showing the features or the variables to a clinician who's like rapidly working through their job, or to front desk staff who are doing all the scheduling work for a clinic, is showing those variables enough? And I would argue "no" because you have to sit there and think about all of the relationships between these different variables that we're including in the model and whether missing data could have informed these variables in some systematic way, or whether there's confounding or bias, to truly interpret the output of any given model. So that's one of the reasons I think that we need to think about what our broader regulatory framework should look like so that we don't put all of the responsibility on our frontline staff or clinicians.

[19:29] Tiffany: You know, I think your point is really well taken in terms of how feasible it is for the human brain to process all the variables that are needed to be able to come up with a particular output or recommendation that a computer can compute so much faster and more easily than we can. So I think part of it, too, then is of course, circle back to this sort of evaluation, and as you mentioned, the regulatory framework. It sounds like the regulatory framework needs that level of transparency and some way to be able to investigate and make sure that the algorithm—or whatever that technology is—that it's not systematically introducing or perpetuating biases towards certain patient populations then.

[20:17] Sarah: And you bring up another great point in your examples that we haven't yet established. What is the regulatory process for ensuring that trustworthy AI is being built into our electronic health record?

Sara: You know, so the FDA for example, has been rapidly evolving their regulatory framework for these types of tools. The FDA framework right now asks two questions. So first, can the provider

independently review the basis for the software recommendation? And that gets into some of the transparency issues I discussed earlier. And then, you know, is machine learning or the logics and inputs explained and available? I think again, the problem with the current regulatory framework is this is very much a gray zone because what does it really mean for a provider to understand the recommendation of an algorithm. And, I think we need to shift how we're thinking a bit. So when EPIC releases a tool, that tool can come to a major academic medical center like UCSF where we have our advanced analytics and innovation team dedicated to validating these tools internally before we deploy anything. Or, that tool can go to a community hospital that doesn't have those resources, doesn't have data scientists on staff. And so we can't expect our health systems to do that work because that's not how it's being marketed. It's not being marketed that you have to do this extensive validation work before you take it. It's being marketed as a tool that is meant to be deployed at the point of care. And there are very good reasons for that, but I think that we need higher level oversight of some of these tools, particularly if they can be turned on at the point of care quite easily because they're built into an electronic health record.

Transition

[22:46] Tiffany: So that's a nice lead into our next question. What are some of the solutions that you think the medical and informatics communities need in order to ensure that information technology advances aren't widening the digital divide?

[23:00] Michelle: I think we have to think about the ways that different populations do access the internet, and understand that the vast majority of lower SES individuals have internet access, it's just on their phone. So instead of relying on something that's an app based platform, for example, something that's readily accessed via a web browser that's not going to take hours to load. When we think about engaging people in healthcare through technology, I think it needs to be accessible in that way. That's certainly the first thing that comes to mind.

Sara: I think the solutions are a bit different. Whether we're talking about patient facing technologies or deployment of artificial intelligence, we need to advocate, as an informatics community, for more regulatory oversight, particularly for these commercial algorithms that are built into the electronic health record. And that oversight needs to include diverse representatives, patient advocates, ethicists, and others who understand how to consider the potential biases and algorithms and how they might be used.

[24:23] Sarah: We've seen the power of real time data and the impact it can have in guiding public health policy and medical decisions. Would you be able to talk about how clinical informatics has played a role in responding to COVID-19, specifically in relation to our current topic of conversation?

[24:38] Sara: So we think a lot about health equity at UCSF and we have a health equity council that leads that work. My group's been working with them to build standard ways of capturing race and ethnicity, language, etc., which sounds really simple, but it isn't with the HR data because of the multitude of ways that we capture this data. So we're looking at that data for COVID across the spectrum right now: who's accessing care, who's getting tested, who's testing positive, who's being hospitalized and having complications. And what did those breakdowns look, you know, according to race, ethnicity, limited English proficiency, payor mix, etc. And we're very fortunate not to have the surge here in San Francisco that they're seeing in New York City, but we're monitoring that data fairly proactively for any trends in disparities that we could intervene upon. The other tool that we need to think about with regards to DEI in COVID and supporting our vulnerable patient populations are some of the patients facing self-diagnostic tools that are deployed. And so we've done that at UCSF. We have tools that are meant to evaluate symptoms and help us make decisions with patients whether they need to come in for care or not. And I think we need ongoing thoughtfulness around whether these tools are adequately serving the entire patient population.

[26:27] Tiffany: And I think that's maybe a good segway into what we, and I'm sure our listeners have heard more and more lately about emerging data, about how COVID is disproportionately affecting certain minority populations in the US—in particular, black and Latinx populations. Dr. Lin, do you want to comment on that based on your experiences in the emergency rooms in New York?

Michelle: Yeah, certainly from my anecdotal experience and from preliminary data from the CDC and the New York City Department of Health, you know, I think that we've seen disproportionately higher rates of both infection and mortality in persons of color, often because they are still essential workers, whether they're in healthcare, food service, etc. Also because they are living with many more chronic comorbid conditions that predispose people to poor outcomes with COVID. And so I think those are some of the things that we are looking to dive into more. You know, we're also actively examining the way that we collect race and ethnicity data, both in our direct clinical care to use down the road for research. And then also with some of our telehealth technology, which has expanded really, really quickly, but the vast majority of that is video-based. Most providers are still English-speaking and while we do have, for example, translator services, I think it's taken a lot of learning on the fly to integrate that to make sure that we're reaching all the populations who need those services.

[28:20] Sara: So we've actually had interesting findings here with regards to the deployment of telehealth. So at UCSF, like many other institutions, we've had a massive uptake in telehealth. But one of the things that we've found that has been surprising is we've had increased uptake in our limited-English proficiency patient population. And I think one of the reasons that that might be happening is because we have a really good system with Zoom, where we automatically bring translators on. And so the experience I think actually might be a bit more seamless than the experience of going to in-person visit. So that's been a surprising finding for us right now and we're still trying to understand why that's the case.

Sarah: I think that's really exciting to hear about. And it kind of hints at a possibility and a hope that there's a potential for telehealth to help us close disparities. I think we focus a lot about how in light of the COVID pandemic, we're seeing a greater digital divide, but if we thoughtfully employ existing resources, we can actually help close it.

[29:35] Michelle: We have seen differences in, for example, the payor mix. So, you know, I think previously our telehealth was almost exclusively like certain subsets of commercial patients with contracts, but because of the expansion of telehealth services, we've seen a lot of Medicaid patients and Medicare, so we see a lot of older patients using the platform sometimes with help from a relative, often referred by their PCP. So I do think it's changing and perhaps widening some gaps while also closing others.

Sara: Yeah, I think you can start to imagine the possibilities here, right? You know, childcare may be challenging for some people, and now they can have their visit at home while their kids are in the next room. Taking time off from work to travel is challenging for people, and now, their visit is kind of sneaking into the corner for a couple of minutes for a quick call with their doctor. And so I think there's a lot of opportunity here to use telehealth to better serve more vulnerable patient populations.

Sarah: And this kind of goes back to that first example that we talked about, in terms of developing the no-show prediction model, and how we talked about how it sort of penalized patients using past utilization data, but that if we are thoughtful about how we're using the data and acknowledging that bias in the data exists, that the implementation or use of that can actually be really positive. And I know one of the examples, Dr. Murray, that you've talked about in the past is instead using this model to identify patients who may need a rideshare voucher or may need additional support in terms of reminding and exploring the reasons why they're not able to make it to that appointment. Bias in the data and these disparities exist, and have been brought to light by the use of machine learning and COVID-19 pandemic, but existed previously, and so, knowing that these exist, how can we then be more thoughtful about what we implement?

[31:55] Sara: I think that's exactly right. So when we learned about the bias and the predictive no-show model, and we started sharing that with our senior leadership across the organization, our recommendation wasn't to not use the model. Our recommendation was to be incredibly thoughtful about what we use the model for. And so to use the model to overbook patients, or do something that could be punitive to individuals, or on a larger level of vulnerable patient populations was clearly not the right thing to do. But the model could be useful for outreach and targeted outreach or support that we could give to patients to help them get to their appointments. So I think that even more important than whether underlying bias exists in our model is what we end up using that model for.

[33:26] Tiffany: I think that was really well said. And it reminds me also of this event that I went to a couple of months ago at my university for women in data science. One of the key messages that I took away from her talk really was that all technology, I mean, information technology included in medicine needs to be ethical by design, starting from the foundational prototype and development phases of really any technology that's going to be applied. And I think Dr. Murray, you touched on this before as well, and you did as well, Dr. Lin, was about the multidisciplinary and having the representativeness and diversity that's needed in the workforce and leadership to make sure that we're really taking into account the diversity of our patients as well as we're designing and implementing these tools to improve our practices.

Sara: So before you build anything, you should think about what problem you're trying to solve, how you'd like to solve it equitably and ethically. And then, if artificial intelligence or informatics tools are part of your toolbox, those need to be built and evaluated by the same principles. And we think about this a lot on our team, and I think it's particularly important in medicine where the health of individuals and populations of individuals are what's at risk here.

[34:23] Michelle: To echo Dr Murray's point about how the information is deployed and integrated into policy. I think the example that always comes to mind is readmissions and how those are used. And I think hospitals for example, with excess readmissions, often get financial penalties. And what we've found is that often those hospitals tend to be safety net institutions. Thinking not only about where the data comes from, but also what the implications are...does that necessarily mean that we shouldn't report the data on safety net institutions? Maybe not necessarily, right? Because I think quality is important for those institutions and for the populations that they serve. But maybe that means that we shouldn't financially penalize them because they may have fewer resources to begin with.

Transition

[35:21] Tiffany: Dr. Murray, Dr. Lin, do you have any final comments or thoughts as we bring our episode to a close?

[35:29] Michelle: I haven't had the privilege of being in a conversation that was four or five women talking about data and algorithms, and there should be more of this.

Tiffany: I'm totally with you.

Sara: Yeah, I am also totally with you. So I lead our health system's advanced analytics and innovation team, but for the COVID-19 analytics needs, we actually took an unconventional approach and pulled together a group of six data and analytics leaders from across the organization, ranging from IT to quality, who all happen to be women. And we built a data analytics and visualization collaborative that we called the DAV collaborative, and it's a cross-functional team that

we are using to drive the analytics for our COVID-19 response and recovery. What we found is that we've been incredibly effective when we fostered collaboration and teamwork, where people have historically worked, and we built consensus. And while there has been a fairly amount written about the success of women leaders in the COVID-19 response, I think in health IT and informatics, this is more broadly an example of how critical it is to ensure diversity in leadership so that different perspectives are represented that may ultimately change the approach and lead an organization to success.

[36:53] Sarah: Well, thank you all. This was a wonderful discussion about the power of machine learning, the need to recognize perpetuated bias, and the power of women in medicine—and specifically, clinical informatics.

Sara: Thank you for having us.

Michelle: Thank you.

Tiffany: Absolutely, and thank you both so much for sharing your expertise and bringing light to these important issues. It's really so needed to set a benchmark for how we best adopt technologies equitably for our patients' best outcomes, so thank you so much!

[37:22] Sarah: And thank you to our listeners for joining us on another enlightening episode of The DEI Shift, you can find more resources including our show notes and guestwork on our website at thedeishift.com. You can also follow us on Instagram and Twitter @thedeishift—that's @deishift for the latest information—and share with us your thoughts and feedback on today's episode content. Thank you for joining us, and tune in next time.

[37:47] Outro

Disclaimer: The DEI Shift podcast and its guest provide general information and entertainment, but not medical advice. Before making any changes to your medical treatment or execution of your treatment plan, please consult with your doctor or personal medical team. Reference to any specific product or entity does not constitute an endorsement or recommendation by The DEI Shift. The views expressed by guests are their own, and their appearance on the podcast does not imply an endorsement of them or any entity they represent. Views and opinions expressed by The DEI Shift team are those of each individual, and do not necessarily reflect the views or opinions of The DEI Shift team and its guests, employers, sponsors, or organizations we are affiliated with.

Season 1 of The DEI Shift podcast is proudly sponsored by the [American Medical Association's Joan F. Giambalvo Fund for the Advancement of Women](#) and the [American College of Physicians Southern California Region III Chapter](#).

The DEI Shift theme music is by Chris Dingman. Learn more at www.chrisdingman.com.

Contact us: thedeishift@gmail.com, [@thedeishift](https://www.instagram.com/thedeishift), thedeishift.com