



CSU  
College of Law Library

Journal of Law and Health

---

Volume 35 | Issue 3

Article

---

10-15-2022

## Transcript: Presentation on Artificial Intelligence and Discrimination in Healthcare

Sharona Hoffman

*Case Western Reserve University School of Law*

Follow this and additional works at: <https://engagedscholarship.csuohio.edu/jlh>



Part of the [Health Law and Policy Commons](#), [Law and Race Commons](#), and the [Science and Technology Law Commons](#)

[How does access to this work benefit you? Let us know!](#)

---

### Recommended Citation

Sharona Hoffman, *Transcript: Presentation on Artificial Intelligence and Discrimination in Healthcare*, 35 J.L. & Health 436 (2022)  
*available at* <https://engagedscholarship.csuohio.edu/jlh/vol35/iss3/6>

This Article is brought to you for free and open access by the Journals at EngagedScholarship@CSU. It has been accepted for inclusion in Journal of Law and Health by an authorized editor of EngagedScholarship@CSU. For more information, please contact [library.es@csuohio.edu](mailto:library.es@csuohio.edu).

TRANSCRIPT: PRESENTATION ON ARTIFICIAL INTELLIGENCE AND  
DISCRIMINATION IN HEALTHCARE

Presented By: Sharona Hoffman<sup>1</sup>

**The Journal of Law and Health's  
Digital Health & Technology Symposium**

CLEVELAND-MARSHALL COLLEGE OF LAW  
FRIDAY, APRIL 8, 2022

---

<sup>1</sup> Professor at Case Western Reserve University School of Law

*The following is a transcription from The Digital Health and Technology Symposium presented at Cleveland-Marshall College of Law by The Journal of Law & Health on Friday, April 8, 2022. This transcript has been lightly edited for clarity.*

**Tigan Woolson:**

First today, presenting on Artificial Intelligence (AI) and discrimination in healthcare, we have Professor Sharona Hoffman. Professor Sharona Hoffman is the Edgar A. Hahn Professor of Law, a professor of bioethics, and Co-Director of the Case Western Reserve University Law School's Law and Medicine Center. She has her J.D. from Harvard Law School, an L.L.M. in Health Law from the University of Houston, and an S.J.D. in Health Law from Case Western Reserve University.<sup>2</sup>

In 2017, Professor Hoffman was elected to the prestigious American Law Institute. She is one of the most cited health law scholars in the United States and received the 2021 Case Western Reserve University "Faculty Distinguished Research Award." She has also served twice as a visiting scholar at the Centers for Disease Control and Prevention. Professor Hoffman has published over 60 articles and book chapters on artificial intelligence, health information technology, big data, emergency preparedness, and many other health law topics. She has developed particular expertise and a national reputation in the area of health law and information technology. Her work has appeared in the Georgetown Law Journal, among many other prestigious journals. Professor Hoffman is currently working as a visiting scholar at the National Institute of Health in the Bioethics Department.

**Sharona Hoffman:**

Great! Thank you so much, it is a pleasure to be here. I am going to be talking about AI and discrimination in healthcare, and this is based on an article by the same title that appeared in the Yale Journal of Health Policy, Law, and Ethics a couple of years ago.<sup>3</sup>

I am going to be focusing on concerns, but we do have to acknowledge that AI is actually very promising. It can improve treatment outcomes in a lot of ways. It can help doctors identify patients who are at risk of complications and hospital readmissions. There is a lot of work being done here in Cleveland on images. First, you can use AI to analyze images to help you determine with greater certainty which ones [tumors] are malignant, which ones are aggressive and which ones [tumors] will and will not require chemotherapy.<sup>4</sup> That can spare patients a lot of

---

<sup>2</sup> For more information, see <https://sharonahoffman.com>.

<sup>3</sup> Sharona Hoffman & Andy Podgurski, *Artificial Intelligence and Discrimination in Health Care*, 19 YALE J. HEALTH POL'Y, L., AND ETHICS 1 (2020).

<sup>4</sup> *Id.* at 9.

uncomfortable and problematic treatments. It also can help identify candidates for clinical studies, and of course, it's been used a lot with COVID-19.<sup>5</sup> There were studies that used AI to try to predict patients' disease courses to predict which patients are going to get sick or very sick and need hospitalization, and which are not, and it was also used to analyze patient records to identify effective treatments.<sup>6</sup>

Now first looking at the privacy concerns we have. When you have AI, you are looking through a lot of patients' records and you are handling a lot of information, so you have to be concerned about whether it's going to be stored and handled securely and if privacy will be maintained. There are questions about the physician-patient relationship; will physicians become redundant and unnecessary, eventually? Will we have medicine by computer? You'll go in, you'll describe your symptoms to the computer, and it will spit out a diagnosis and treatment course. Will that be the new reality?

Second, the other very real problem is whether doctors are going to trust AI suggestions if those suggestions are not what they would naturally be inclined to do. For example, if the AI says "no, this patient doesn't need chemotherapy," is that going to be problematic for the doctor? Is the doctor not going to want to follow that suggestion because of concerns about liability? Definitely something important to consider.

However, I am going to spend most of my time on this issue of discrimination. We know that AI has led to discrimination in a variety of other areas. In criminal justice, and this has gotten a lot of media attention, they use AI for race-based recidivism predictions to determine who is most likely to reoffend and to tailor punishments accordingly. Not surprisingly, the software has often mislabeled Black offenders as more likely to commit crimes again and mislabeled white offenders as unlikely to commit crimes again, and in reality, this is not even the pattern these offenders have fallen into if/when recommitting criminal acts.<sup>7</sup>

There is a well-known hiring case where Amazon used an algorithm in order to try to pick out resumes from the pile and determine who should be interviewed.<sup>8</sup>

---

<sup>5</sup> *Id.* at 11.

<sup>6</sup> *Id.* at 11; see also Wim Naudé, *Artificial Intelligence vs. COVID-19: Limitations, Constraints and Pitfalls*, 35 *AI & Soc'y* 761, 761-62 (2020).

<sup>7</sup> *Id.* at 5; see also Julia Angwin ET AL, *Machine Bias*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>; Melissa Hamilton, *Debating Algorithmic Fairness*, 52 *UC DAVIS L. REV.* 261, 264 (2019) (reporting that the risk tool's corporate owner denied the allegation and stated that its reanalysis of the data led it to conclude that "the tool was unbiased as blacks and whites had similar positive predictive values for recidivism"); Sandra G. Mayson, *Bias In, Bias Out*, 128 *YALE L.J.* 2218, 2221-22 (2019) (discussing algorithmic risk assessment in the criminal justice system and its racial impact).

<sup>8</sup> *Id.*; see also MICHAEL KEARNS & AARON ROTH, *THE ETHICAL ALGORITHM* 60-61 (2020); Katherine Maher, Opinion, *Without Humans, A.I. Can Wreak Havoc*, *N.Y. TIMES* (Mar. 12, 2019), <https://www.nytimes.com/2019/03/12/opinion/artificial-intelligence-wikipedia.html>.

And disproportionately, it rejected women's applications, and it picked men's applications.<sup>9</sup> Why? Because the training data that was used to train the algorithm came largely from men, so the algorithm learned that men are good candidates and women are not. Similar problems exist in the area of housing. For example, Facebook kept minorities from seeing housing ads in areas that were supposed to be “all white.”<sup>10</sup> So, I was curious to see if the same kind of thing happens in healthcare, and of course, it does. There are a lot of anecdotes, a lot of stories about discrimination or bias in AI algorithms.

Here is another well-known story: there was a high-risk management algorithm program that was used to refer patients with chronic illnesses to case management programs, and these were very beneficial. They gave support. They gave resources. The problem: the program was not referring minority patients for these resources. Why? Because the AI incorrectly used past medical expenditures as a proxy for medical need. So, if you had spent a lot of money in the past on medical care, it assumed that you were very sick and needed this high-risk management program, and if you did not, it assumed you were healthy and did not need that extra support. Of course, this was very problematic because there are systemic barriers that minorities face: for example, not spending money on healthcare because they do not have it to spend or because there are other access barriers.

Winter labs in Toronto, Canada developed a learning tool that tried to identify people with Alzheimer's disease using speech samples.<sup>11</sup> However, there was an issue: the training data came from native English speakers who had a particular accent, so it worked well for them, but it misdiagnosed people with Alzheimer's disease if they had a different accent, if they were not native English speakers, or if they had some hesitation in their speech because it had not been trained to analyze certain kinds of speech samples. Just yesterday, I saw a story in the New York Times titled “Can AI-driven Voice Analysis Help Identify Mental Disorders,” with the subtitle reading “Early tests have been promising but issues involving bias, privacy, and mistrust of black box algorithms are possible pitfalls.”<sup>12</sup> So here you go again, AI might be trying, based on voice alone, to analyze and diagnose mental disorders.

---

<sup>9</sup> *Id.* at 13-16.

<sup>10</sup> *Id.* at 5; see also Margot E. Kaminski & Andrew D. Selbst, Opinion, *The Legislation That Targets the Racist Impacts of Tech*, N.Y. TIMES (May 7, 2019), <https://www.nytimes.com/2019/05/07/opinion/tech-racism-algorithms.html>.

<sup>11</sup> Sharona Hoffman & Andy Podgurski, *Artificial Intelligence and Discrimination in Health Care*, 19 YALE J. HEALTH POL'Y, L., AND ETHICS 1, 17 (2020). See also KATHLEEN C. FRASER ET AL., *Linguistic Features Identify Alzheimer's Disease in Narrative Speech*, 49 J. ALZHEIMER'S DISEASE 407, 407 (2016).

<sup>12</sup> Ingrid K. Williams, *Can A.I.-Driven Voice Analysis Help Identify Mental Disorders?*, N.Y. TIMES (Apr. 5, 2022), <https://www.nytimes.com/2022/04/05/technology/ai-voice-analysis-mental-health.html>.

Then you have algorithms that explicitly adjust for race, so you get more points or fewer points based on whether you are white or black. There are some studies that show that certain populations are more or less likely to have certain diseases.<sup>13</sup> These algorithms take it a step too far. If you walk into the emergency room with particular symptoms if you are black, you get a certain number of points, and if you are white, you get a certain number of points. Some of the points are just for your race. And so, Black patients routinely get diagnosed as less likely to have heart failure and kidney stones, even if they have clear symptoms of them, because of this race adjustment. There is an algorithm that is supposed to recommend whether someone who has had a prior cesarean delivery should be allowed to try for a natural birth or should immediately be sent for another C-section.<sup>14</sup> Again, this adjusts for race. So, if you are Hispanic or Black it will say no natural birth, and you have to go for a C-section. If you are white, the algorithm is more likely to say you can try for a natural birth. This results in minority women getting unnecessary surgeries.

This can all lead to health disparities in a variety of ways. First, because of race adjustments or bias, minorities may suffer worse medical outcomes because of the AI. On the contrary, we also said that AI can help a lot. AI can promote more accurate diagnoses and more accurate treatment recommendations. However, if this is an expensive resource that a lot of facilities are not able to have because they are in resource-poor areas, this too can exacerbate health disparities. Therefore, you can have the benefit of AI at the Cleveland Clinic but not at a poor rural hospital. This is another health disparity issue. If these algorithms, or if AI, focus on race, adjust for race, or work better for one group but not another, this will raise stigmatization concerns and concerns about technology emphasizing racial differences inappropriately.

So, what are the reasons for algorithmic bias? First, there are often deficiencies in training data. If the algorithm is not trained on data from women, is not trained on data from people who are non-native English speakers, or is not trained on data from minorities, the algorithm will work well for the groups it was trained to analyze. However, it very well may not work appropriately for groups that it wasn't trained for. Second, there is feedback loop bias.<sup>15</sup> Bias can actually be baked into AI. For example, women are often misdiagnosed when they have cardiac symptoms. Doctors do not catch heart attacks or cardiac problems because women's symptoms are often different from men's symptoms. So, if the algorithm is trained on data that misdiagnoses women because it uses male benchmarks, algorithms will continue to misdiagnose women until they are updated. Third, the issue of poor study design. We saw that with the example of the study that

---

<sup>13</sup> Sharona Hoffman & Andy Podgurski, *Artificial Intelligence and Discrimination in Health Care*, 19 YALE J. HEALTH POL'Y, L., AND ETHICS 1, 20-23 (2020).

<sup>14</sup> *Id.*

<sup>15</sup> *Id.* at 15-16.

inappropriately used past medical expenditures as a proxy for medical need, and we have a lot of studies that have this kind of a problem.

And then, of course, there is a problem with the entire concept of race. “Race is a social construct.” What does it really mean? Why would an algorithm adjust or give someone extra points or fewer points if they look black or white? Often when we talk about race, we are really talking about genetic variations which may be more common in certain populations, but they're never exclusive to a population. And so, if you are trying to create AI, you should be focusing on genetic variations and not on somebody's skin color.

Sometimes outcomes or health status is influenced by social determinants of health.<sup>16</sup> Does a person have too much stress? Do they have good nutrition? Do they have opportunities to exercise? And so on. And so again, the AI should focus on the social determinants of health and on the health status, rather than on skin color. Finally, of course, we have millions of people with mixed ancestry today, so, if for purposes of AI, you have to label somebody White or Black and that is it, you are going to make a lot of mistakes.

I do remember that this is a law talk, so let's talk about some gaps in legal protection. First, let's discuss disparate impact. Hopefully, and presumably, when healthcare providers use AI, they are not intending to hurt minority groups, or to hurt women, or to hurt anybody else. So, the harm is unintentional. It is a case of disparate impact, where a facially neutral practice that was supposed to work well for everyone has a disparate impact on minorities or women or some other group. Now, in the areas of employment, or housing, private parties can sue for disparate impact. If there is an employment practice that harms Black people or somebody else unintentionally, you can sue for that. The same is not true in healthcare. You cannot sue for disparate impact if you are a private party, no matter what harm you suffered because of unintentional discrimination.

Also, there is an issue with the Americans with Disabilities Act (ADA) and other anti-discrimination acts. The Americans with Disabilities Act prohibits discrimination based on current or past disabilities.<sup>17</sup> A lot of these algorithms are designed to predict whether someone will develop a problem in the future. Will they develop cognitive decline? Will they develop heart problems? Will they develop diabetes? The Americans with Disabilities Act does not prohibit discrimination that is based on predictions of future harm. And so, if an employer, for example, wants to say, “well, I have this AI result that says you will develop cognitive decline in five years, so I do not want you in my workforce,” they can. That is a big problem, and I think this is an important legislative gap that, in this era of AI and predictions, we need to address.

---

<sup>16</sup> *Id.* at 21 n.148.

<sup>17</sup> *Id.* at 31 n.241.

Finally, how is AI regulated? We have an unclear FDA regulatory approach at this time. The FDA only regulates some forms of AI.<sup>18</sup> They have some proposals to regulate more, but they really have not fleshed them out. We are not sure how the FDA will approach regulation of AI in the future. Will it be comprehensive, or sort of piecemeal? Interestingly, in 2019, there was a proposal for the Algorithmic Accountability Act.<sup>19</sup> This bill was sponsored by Cory Booker and some others. One of the things the bill would have done was require companies and users to study the algorithms they use, identify bias, and correct the problems. This would have placed some responsibility on users. However, like most legislation these days, the bill did not pass. It did get reintroduced in February of 2022, so there could be some developments in the future.<sup>20</sup>

So, very, very briefly, what are my recommendations? I think that algorithmic fairness should be a key element in designing, validating, and implementing all AI. It should be something that everyone thinks about: people who are developing it, people who are using it, and people who are regulating it. We have to keep in mind that fairness is a very complex concept, and you probably cannot be fair to everyone. You probably are going to have to make choices. For example, maybe you can achieve group fairness, but not individual fairness, not fairness to every single individual within the group. Or maybe you have to choose between process fairness and outcome fairness. So, it would be naive to say, “well let's just have algorithmic fairness.”; it is much more nuanced than that.

As I already suggested, we need some law reform. I think we do need to make disparate impact available as a tool for litigants, for parties that are harmed. I also think it is time to revise the ADA in order to extend protection to people who are currently healthy but have suffered discrimination because of predictions of future illness or future disability. I have written a lot about this. Nobody listens, but I keep optimistically repeating this point anyway.

---

<sup>18</sup> *Id.* at 37-38.

<sup>19</sup> Algorithmic Accountability Act of 2019, H.R. 2231, 116th Cong. (2019).

<sup>20</sup> Algorithmic Accountability Act of 2022, H.R. 6580, 117th Cong. (2022).