AI NOW

New York City Council
Committee on Hospitals, Committee on Technology

Oversight - Electronic Health Records
November 20, 2019

Written Testimony of
Varoon Mathur
Technology Fellow, AI Now Institute, New York University

Good Afternoon Chairpersons Rivera and Holden, and members of the Hospitals and Technology Committees. My name is Varoon Mathur, and I currently serve as a Technology Fellow at the AI Now Institute - an interdisciplinary research institute at NYU, focused on the social implications of artificial intelligence (AI). Thank you for the opportunity to testify today on privacy and security concerns regarding electronic health records (EHRs).

At the AI Now Institute, our research on the use of AI has identified a significant and alarming uptake of AI-based tools and systems within high-stakes domains including criminal justice, education, welfare, employment, and indeed health care.[1] The four key concerns we examine in relation to these systems span areas of bias and inclusion, rights and liberties, labor, and safety and critical infrastructure. This work is of particular importance to the domain of health care, where AI and advanced precision medicine algorithms have been marketed as fulfilling the promise of EHRs, by using "big data" analytics to produce new clinical knowledge, and more precise and tailored diagnostics.[2]

The rapid development and implementation of machine learning (ML) algorithms and data sharing partnerships in the healthcare space brings new challenges around privacy, security, and patient identifiability through EHR data. Most recently, a partnership between Google and Ascension, one of the largest non-profit health systems in this country, became public after a whistleblower working on the project revealed that patient data transferred between Ascension and Google was not “de-identified”.[3] This partnership, in which Google would provide Cloud services to help migrate Ascension’s infrastructure to a Google-managed cloud

[1] Whittaker, Meredith, Kate Crawford, Roel Dobbe, Genevieve Fried, Elizabeth Kaziunas, Varoon Mathur, Sarah Mysers West, Rashida Richardson, Jason Schultz, and Oscar Schwartz. *AI now report 2018*. AI Now Institute at New York University, 2018. https://ainowinstitute.org/aiareport2018.pdf.

[2] MIllard, Mike. 2019. “Machine Learning Will Help EHRs Fulfill Precision Medicine’s Promise.” Healthcare IT News. January 18, 2019. https://www.healthcareitnews.com/news/machine-learning-will-help-ehrs-fulfill-precision-medicines-promise.; “Big Data in Healthcare: Challenges & Promise” n.d. Accessed November 19, 2019. https://catalyst.nejm.org/big-data-healthcare,;“Scientists Outline the Promises and Pitfalls of Machine Learning in Medicine.” n.d. EurekAlert! Accessed November 19, 2019. https://www.eurekalert.org/pub_releases/2019-04/hms-sot040119.php.

[3] “I’m the Google Whistleblower. The Medical Data of Millions of Americans Is at Risk | Anonymous | Opinion | The Guardian.” n.d. Accessed November 19, 2019. https://www.theguardian.com/commentisfree/2019/nov/14/im-the-google-whistleblower-the-medical-data-of-millions-of-americans-is-at-risk

environment, also included Google's development of AI solutions, ostensibly to help support doctors and nurses to improve care in real time.[4]

Google is not the only cloud provider partnering with hospital systems to help migrate patient data and other health information technology (IT) infrastructure to cloud servers owned and managed by large tech firms. Amazon Web Services now provides the ability to subscribe to third party data, enabling healthcare professionals to aggregate data from clinical trials. Microsoft recently announced a partnership with Humana that would provide cloud and AI resources, as it is also helping power Epic Systems' predictive analytics tools for EHRs.[5] In fact, estimates now expect the cloud computing market for healthcare to reach nearly $30 billion by 2026.[6] Meanwhile, recent polls tracking Americans' perception of their experiences with EHRs show that most patients are increasingly concerned with unauthorized access of confidential information.[7]

These new developments raise two key questions regarding the privacy, security, and safety of patient data: 1) how does our definition of protected health information (PHI) change in the age of AI algorithms, given their predictive capabilities which can disclose sensitive information even absent PHI, and 2) how do we assess the utility of EHRs in building more advanced algorithms for better patient care? New research suggests that the rapid deployment of clinical AI tools absent regulatory oversight leaves patients vulnerable to privacy and security breaches. Furthermore, our own research exploring the sociotechnical dynamics of EHRs suggests that these forms of data record limited information, and are unable to capture patients' lived experiences. Thus EHRs do not lend themselves to the development of clinical tools, in spite of the claims made by hospital systems and cloud providers.

Under the Health Insurance Portability and Accountability Act (HIPAA), PHI data is categorized as data that directly and uniquely ties to an individual, with examples including names, birth dates, and email addresses.[8] De-identified data, therefore, would be the removal of such categories from a potential EHR dataset. However, new research shows that it is possible to link two de-identified EHRs of the same patient but from two different data sources accurately using computational methods, so as to create a more complete history of

[4] "Our Partnership with Ascension." n.d. Google Cloud Blog. Accessed November 19, 2019. https://cloud.google.com/blog/topics/inside-google-cloud/our-partnership-with-ascension/.

[5] "AWS Data Exchange | Amazon Web Services." n.d. Amazon Web Services, Inc. Accessed November 19, 2019. https://aws.amazon.com/data-exchange/.; Thorne, James. n.d. "Microsoft Lands Another Healthcare Partnership, This Time with Humana to Take Care of Aging Seniors – GeekWire." Accessed November 19, 2019. https://www.geekwire.com/2019/microsoft-lands-another-healthcare-partnership-time-humana-take-care-aging-seniors/.; "Ochsner Health System Adopts Epic's Machine Learning Platform Powered by Microsoft Azure." n.d. Accessed November 19, 2019. https://azure.microsoft.com/en-us/resources/videos/ochsner-health-system/.

[6] Lagasse, Jeff. n.d. "Healthcare Cloud Computing Growth Due in Part to Curbing Infrastructure Costs | Healthcare Finance News." Accessed November 19, 2019. https://www.healthcarefinancenews.com/news/healthcare-cloud-computing-growth-due-part-curbing-infrastructure-costs.

[7] Muñana, Cailey, Ashley Kirzinger, and Mollyann Brodie. 2019. "Data Note: Public's Experiences With Electronic Health Records." The Henry J. Kaiser Family Foundation (blog). March 18, 2019. https://www.kff.org/other/poll-finding/data-note-publics-experiences-with-electronic-health-records/.

[8] "What Is Considered Protected Health Information Under HIPAA?" 2018. *HIPAA Journal* (blog). April 2, 2018. https://www.hipaajournal.com/what-is-considered-protected-health-information-under-hipaa/.

a patient without using any PHI of the patient in question.[9] Similarly, last month a New York Times article reported new research that showed it is possible to create a reconstruction of patients' faces using de-identified MRI images, that could then be identified using facial recognition systems.[10] These examples demonstrate how vulnerabilities within large technology infrastructure present serious security and privacy challenges for the collection and use of EHR data, and that these may be beyond the reach of HIPAA protections. Such concerns are echoed in a recent class action complaint filed in response to the partnership between the University of Chicago Medical Center and Google, which states that Google is "uniquely able to determine the identity of almost every medical record the university released" due to its expertise and resources in AI development.[11]

Trading the privacy and security of individual patients in order to leverage precision clinical care incorrectly assumes that EHR data and infrastructure are inherently viable for training of machine learning models. Yet research demonstrates that this premise is misguided because it fails to consider two key challenges: (1) EHR infrastructure was originally constructed for billing and other administrative tasks, rather than clinical care; and (2) EHR data is both incomplete and flawed because it is missing important data for a variety of populations and is incapable of capturing all of the data necessary for precision clinical care. For example, a Michigan State University study showed that EHR tend to function more for administrative record keeping rather than a tool for clinical care. This is because EHR are structured to reflect the interests of political and corporate stakeholders, recording what is important to them, and not necessarily what matters to patients.[12] Though EHR infrastructure has evolved over time, it is still riddled by the structural flaws and presumptions that motivated its initial development. Moreover, research conducted by Dr. Elizabeth Kaziunas, a Postdoctoral Fellow at AI Now, demonstrated the ways in which the social construction of health data (how it is shaped by the interests of institutions and corporate stakeholders), along with the design limitations of our current health information systems, like EHRs, result in a failure to capture important types of health information. Specifically, gaps in the EHR can result from health disparities within communities, and can inadvertently exclude certain patient populations, as well as the under-reporting of chronic illnesses by individual patients due to associated stigmas.[13] The significant limitations of EHRs mean that machine learning tools informed and trained by such data are likely to be highly biased. And this suggests the urgent

[9] Hejblum, Boris P., Griffin M. Weber, Katherine P. Liao, Nathan P. Palmer, Susanne Churchill, Nancy A. Shadick, Peter Szolovits, Shawn N. Murphy, Isaac S. Kohane, and Tianxi Cai. "Probabilistic record linkage of de-identified research datasets with discrepancies using diagnosis codes." *Scientific data* 6 (2019): 180298.

[10] Kolata, Gina. n.d. "You Got a Brain Scan at the Hospital. Someday a Computer May Use It to Identify You. - The New York Times." Accessed November 19, 2019. https://www.nytimes.com/2019/10/23/health/brain-scans-personal-identity.html.

[11] Wakabayashi, Daisuke. n.d. "Google and the University of Chicago Are Sued Over Data Sharing - The New York Times." Accessed November 19, 2019. https://www.nytimes.com/2019/06/26/technology/google-university-chicago-data-sharing-lawsuit.html.

[12] Hunt, Linda M., Hannah S. Bell, Allison M. Baker, and Heather A. Howard. "Electronic health records and the disappearing patient." *Medical anthropology quarterly* 31, no. 3 (2017): 403-42. https://doi.org/10.1111/maq.12375.

[13] G. M. Weber, W. G. Adams, E. V. Bernstam, J. P. Bickel, K. P. Fox, K. Marsolo, V. A. Raghavan, A. Turchin, X. Zhou, S. N. Murphy, and K. D. Mandl, "Biases introduced by filtering electronic health records for patients with "complete data"," *Journal of the American Medical Informatics Association : JAMIA*, vol. 24, pp. 1134–1141, Nov. 2017.;Elizabeth Kaziunas, Michael S. Klinkman, and Mark S. Ackerman. 2019. Precarious Interventions: Designing for Ecologies of Care. Proc. ACM Hum.-Comput. Interact. 3, *CSCW,* Article 113 (November 2019), 27 pages. DOI: https://doi.org/10.1145/3359215; Tiffany C Veinot, Hannah Mitchell, Jessica S Ancker, Good intentions are not enough: how informatics interventions can worsen inequality, *Journal of the American Medical Informatics Association*, Volume 25, Issue 8, August 2018, Pages 1080–1088, https://doi.org/10.1093/jamia/ocy052.

need for more regulatory oversight over algorithms developed within hospital systems and deployed in partnership with cloud technology companies.

Given the large number of world-class health systems in New York City that will continue to utilize more cloud services for EHR storage and integration, and continue to pursue AI development, this Committee has a unique opportunity to spearhead city-wide legislative efforts that can address the current challenges. We provide three forward-looking policy recommendations that this council should pursue.

Policy Recommendations for New York City Council

1. **Require New York City health systems procuring AI/ML solutions, alongside Cloud server solutions, to conduct Algorithmic Impact Assessments as part of notifying and obtaining consent from patients.[14]**

   In 2018, AI Now published the Algorithmic Impact Assessment (AIA) framework, which offers a means for assessing algorithmic systems, while also providing the public with meaningful opportunities to evaluate the potential impacts if such a system would be adopted, before an agency has committed to its use. This process fosters transparency and trust between agencies and the communities they serve, and is especially important to ensure that patients are aware of how their health records are being used, and have the opportunity to consent before their records are used for training AI models. Such measures would also ensure clear reporting on what types of data are being shared by health systems with cloud service providers.

2. **Require New York City health systems to publicly state whether social-media data is combined with EHR data for patient surveillance or monitoring of patient well-being.**

   Public health agency use of social media data to identify disease outbreaks and predict epidemics before they occur raises significant concerns around surveillance, especially since such predictions are usually made without consent from patients whose data they rely on.[15] Such tools also raise issues regarding accuracy: there is mounting evidence that algorithms predicting health outcomes using social media data are inaccurate, and prone to significant bias.[16] These specific problems are compounded in the context of EHR data, and therefore clear justification must be made available through public disclosures.

[14] Reisman, Dillon, Jason Schultz, Kate Crawford, and Meredith Whittaker. "Algorithmic impact assessments: A practical framework for public agency accountability." *AI Now Institute* (2018). https://ainowinstitute.org/aiareport2018.pdf.

[15] Graham Dodge, "Using Social Media as a Public Health Surveillance Tool," Becker's Hospital Review, March 2, 2017, https://www.beckershospitalreview.com/population-health/using-social-media-as-a-public-health-surveillance- tool.html.; Ebele Mogo, "Social Media As A Public Health Surveillance Tool: Evidence And Prospects," Sickweather, https://enterprise.sickweather.com/downloads/SW-SocialMedia_WhitePaper.pdf.

[16] Shirin Ghaffary, "The Algorithms That Detect Hate Speech Online Are Biased against Black People," Vox, August 15, 2019, https://www.vox.com/recode/2019/8/15/20806384/social-media-hate-speech-bias-black-african-american- facebook-twitter.

3. **Conduct city-wide disparate impact evaluations around the current uses of EHRs in order to identify potential socioeconomic disparities arising from the use of AI/ML health solutions.**

   A recent study found that an algorithm trained on patient data and used to screen for patients in need of "high-risk care management" was substantially biased against black patients.[17] This was due to the fact that the algorithm used health care costs as a proxy for health needs, but failed to account for the fact that disparities exist between patients and thus their ability to access care, which results specifically in black patients having fewer health care dollars spent on them. Such examples detail how EHRs and similar patient data do not fully capture the sociotechnical context of their use, and can lead to further inequity within health care systems. It also shows why disparate impact analysis must be a central component of any assessments conducted around algorithmic tools procured within the city and that is used to inform decisions around health care resource allocation.

[17] Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. "Dissecting racial bias in an algorithm used to manage the health of populations." *Science* 366, no. 6464 (2019): 447-453.