

Protecting Data Privacy in the Age of AI-Enabled Ophthalmology

Elyse Tom^{1,*}, Pearse A. Keane^{2,3,*}, Marian Blazes¹, Louis R. Pasquale⁴, Michael F. Chiang⁵, Aaron Y. Lee^{1,*}, and Cecilia S. Lee^{1,*}, and AAO Artificial Intelligence Task Force

¹ Department of Ophthalmology, University of Washington, Seattle, WA, USA

² Medical Retina Service, Moorfields Eye Hospital NHS Foundation Trust, London, UK

³ Institute of Ophthalmology, University College London, London, UK

⁴ Eye and Vision Research Institute, Icahn School of Medicine at Mount Sinai, New York, NY, USA

⁵ Departments of Ophthalmology and Medical Informatics & Clinical Epidemiology, Casey Eye Institute, Oregon Health & Science University, Portland, OR, USA

Correspondence: Aaron Y. Lee,
Department of Ophthalmology,
University of Washington, 325 Ninth
Avenue, Box 359608, Seattle, WA
98104-2499, USA. e-mail:
leeay@uw.edu

Received: April 1, 2020

Accepted: April 2, 2020

Published: July 6, 2020

Keywords: Machine Learning;
Artificial Intelligence; Privacy;
Biomedical Ethics

Citation: Tom E, Keane PA, Blazes M,
Pasquale LR, Chiang MF, Lee AY, Lee
CS. Protecting data privacy in the
age of AI-Enabled ophthalmology.
Trans Vis Sci Tech. 2020;9(2):36,
<https://doi.org/10.1167/tvst.9.2.36>

Introduction

Digital data privacy is a rapidly evolving concept in health care. As electronic records have replaced paper charts, and with the rise of “Big Data” and artificial intelligence (AI), this issue has become increasingly important. Big Data has been defined by the three V’s: volume (large amounts of data), variety (data heterogeneity), and velocity (speed of access and analysis).^{1,2} Analyses of these large datasets have allowed for more powerful assessments of healthcare quality and efficiency with the goal of improving patient care.³ AI is a branch of applied computer science that

uses computer algorithms to perform cognitive tasks that approximate human intelligence, such as clinical decision making.⁴ More specifically, deep learning, a subset of machine learning within the field of AI, has been particularly successful in training powerful algorithms for the classification of medical images and other high-dimensional data.^{5–9} Taken together, these approaches may offer many benefits for patients, including automated screening and triage of disease and treatment optimization. For example, AI-enabled screening of diseases such as diabetic retinopathy, retinopathy of prematurity, and glaucoma could improve early detection and treatment.^{5,10,11} Furthermore, AI has been used for future disease predictions,

in areas ranging from acute kidney injury to age-related macular degeneration and diabetic retinopathy; in the future, such approaches could lead to better preventative strategies.^{12–15} The combination of Big Data and AI also offers many potential benefits for healthcare systems, including increased productivity with decreased costs, as well as reductions in medical error. New data privacy problems have arisen with the use of this technology, however, leading to concerns about the balance between innovation and privacy and the need for better data protection methods that can evolve along with Big Data and AI.

Ethical Considerations

In the United States, the Belmont Report is the most widely recognized ethical framework for health care and the life sciences, and it serves as an essential reference for institutional review boards.¹⁶ The Belmont Report highlights three fundamental principles: respect for persons, beneficence, and justice. Two additional bioethical principles—non-maleficence (often translated as “first, do no harm”) and respect for autonomy—are also considered central to biomedical ethics and AI development.^{17,18} The Belmont Report is not directly applicable to secondary use of de-identified clinical data; in the report, when identifying information has been removed, the use of that data is no longer considered human subjects research. Nonetheless, its core principles are instructive. In particular, with regard to beneficence at a population level, some believe it is unethical to refrain from using clinical data to develop tools that have the potential to benefit others. In contrast, when considering both non-maleficence and respect for autonomy, others weigh the balance of both risks and benefits of such applications for an individual where said individual does not derive benefit.¹⁹ In the United Kingdom, this is recognized in the constitution of its National Health Service (NHS) which pledges “to anonymize the information collected during the course of your treatment and use it to support research and improve care for others.”²⁰ However, with the use of increasingly large clinical datasets, maintaining data privacy and confidentiality—and thus respect for persons—is a challenge.

Data Protection and Privacy

One of the main limitations of machine learning and deep learning approaches is their requirement for large datasets for development and testing—datasets that are typically an order of magnitude or

even greater than those collected in most prospective clinical trials. Compared to other medical specialties (e.g., obstetrics), ophthalmology has benefited from the widespread availability of large, well-curated imaging datasets and thus is often seen as being at the forefront of AI-enabled health care.²¹ Although the availability of anonymized datasets has been a boon for technological advancement, it also represents a significant risk. The principle of beneficence requires that healthcare professionals “do no harm”; yet, breaches of patient privacy can cause major harms and can also have unintended consequences. These could potentially impact one’s employment or insurance coverage² and may even allow computer hackers to obtain Social Security numbers and personal financial information.²²

Removal of all potentially identifiable information from large datasets can be a daunting task. In fact, it is now clear that, even with the most rigorous efforts, there will always remain at least a theoretical risk of re-identification.²³ This is not an issue unique to ophthalmology, as it is now conceivable to apply facial recognition software to three-dimensional reconstructions of computed tomography of the head. In addition, features from the periocular region have been used to identify the age of patients using machine learning algorithms.²⁴ Gender, age, and cardiovascular risk factors have been identified from fundus photographs.¹³ Even for datasets not involving medical images, and even without the use of advanced or future technologies, it may be possible to identify individuals by linkage with other datasets. This is particularly the case as patient information generally accumulates over time.²⁵

Data Sharing

Another problem related to privacy and AI is managing the exchange of data in an ethically acceptable way. AI typically requires specialized technical expertise and powerful computer resources. In the case of a rare disease, for example, consolidation of data from multiple institutions would be required. As a result, datasets must be shared outside of the institution in which they were generated. If executed poorly, such data sharing may increase the risk of data breaches.

Data sharing that involves major multinational corporations in the pharmaceutical and technology sectors is also of great concern. Monetization of clinical data is a trending topic lately, as evidenced by the oft-repeated phrase “data is the new oil.”²⁶ These

increasing relationships between healthcare companies and academic research data can heighten the risk of malicious privacy violations. Although detailed discussion of this issue is outside the scope of this article, the use of exclusive contracts or licenses that prohibit sharing of routinely collected clinical data is another cause for concern. Thus, exclusive arrangements that restrict or preclude the widest possible patient benefit for clinical data could undermine the Belmont principle of justice.

Models for Consent

In many countries, research ethics authorities do not require individual consent for retrospective research on de-identified datasets. This is generally well accepted in ophthalmology, and the majority of ophthalmic clinical research including the AAO IRIS (Intelligent Research in Sight) Registry takes place using this model.²⁷ However, this practice is sometimes questioned in the context of machine learning, where the clinical data themselves are used to develop algorithms. Many patients are supportive of the use of their data to improve health care and research but feel that they should be asked to give permission first.²⁸ At first glance, this would appear to be the most ethically sound approach—it is certainly an appropriate model for interventional research studies such as clinical trials. However, this approach is cumbersome or not feasible for large, historical datasets of routinely collected data. There are also challenges in prospectively deploying such a consent model, particularly when seeking permissions from patients for unforeseen future uses of their de-identified data. It would not be true informed consent if patients are asked to sign up to extensive terms and conditions before each episode of care or to agree to future uses of their data about which they have not yet been “informed.”

The use of opt-in models has also been proposed as a lighter touch approach; however, this means that only the most engaged patients, who actively take steps to get involved, will be included. Of course, the act of de-identifying itself may require consent from the covered entity. There is an increasing awareness of the potential for significant ethical risks in this regard with the use of AI, particularly concerning racial bias.²⁹ For these practical and ethical reasons, an opt-out model is often preferred. As part of a review into the security and use of NHS data in the United Kingdom in 2016, the National Data Guardian recommended that a national opt-out model should be introduced, rather than one based on an opt-in consent.³⁰ In 2020, clinical AI researchers from Stanford University proposed the

use of a similar model for the development of AI in radiology.¹⁹

Real-World Case Study

With ophthalmology at the forefront of AI-enabled health care and potentially acting as an exemplar for other medical specialties, the specialty has had to engage with these issues. For example, the collaboration between Moorfields Eye Hospital in the United Kingdom (led by P.A.K.) and DeepMind, an AI company, adopted a multipronged approach. First, it addressed an area with clear patient benefit, the development of a triage tool for macular diseases using optical coherence tomography (OCT) images.³¹ Second, all OCT scans used were de-identified to the standards described by the UK Information Commissioner’s Office Anonymisation Code of Practice,³² as well as according to *BMJ* guidance for the sharing of clinical data.³³ Third, contractual safeguards were put in place for non-exclusive data sharing that prohibits linkage with other datasets or attempts at re-identification. Finally, and most importantly, an active program of patient and public engagement was undertaken, with the aim of ensuring public transparency. This included early communication with the major eye disease charities and the Royal College of Ophthalmologists, as well as with the NHS Health Research Authority,³⁴ as well as also providing information for those patients who preferred to opt out of the research, either at a local or national level.³⁰

Additional Legal Considerations

Many of the privacy concerns associated with Big Data and AI exist due to gaps in existing laws and regulations regarding traditional medical data. The Health Insurance Portability and Accountability Act (HIPAA) was enacted in 1996, prior to the rise of Big Data and AI. HIPAA regulates the use of protected health information (PHI) in the United States and requires de-identification of data by two mechanisms: (1) expert determination (expert risk assessment for a particular use) and (2) safe harbor (the removal of 18 prespecified identifiers).³⁵ Whereas HIPAA is meant to ensure data protection on the part of healthcare providers and healthcare systems, these regulations may be inadequate for managing the ever-larger amounts of data associated with medical care today. For example, PHI can be shared without consent for treatment, payment, and operational purposes,³⁶ and

HIPAA does not cover data generated outside of health entities, such as patients, providers, and insurers, covered by the act.² Examples of unregulated data include data from smart watches, mobile health applications, internet search engines, social media, and consumer-initiated health tests, such as genetic testing, all of which can be triangulated to re-identify individuals.² In Europe, patients are protected under the General Data Protection Regulation (GDPR), which was implemented by the European Union in 2018 to regulate personal data protection. Under the GDPR, all health data are considered personal data, but there are exceptions under which health data may be used without consent such as for research purposes if safeguards are instituted.³⁷ The GDPR introduced a weaker version of de-identification known as pseudonymization, which is the removal of only directly identifying data.³⁸

New Approaches

Some of the most promising approaches for protecting data privacy in the era of Big Data and AI are those that take advantage of the technology itself. One strategy, differential privacy, involves describing patterns of groups in the dataset rather than individuals.³⁹ Federated learning (or collaborative learning) and distributed models are machine learning techniques that can be used to protect data by training algorithms across multiple servers using separate data samples.^{40–42} In this method, training code and models are brought into each data silo and trained in situ while the data remain in place. The combined model parameters, trained across many locations, would effectively have been trained on all available data without risking data breaches from allowing outsider use.

Similarly, training local generative adversarial networks (GANs)⁴³ and then sharing GANs instead of data may mitigate re-identification risks.^{44,45} Each deep learning model would be trained to recapitulate the statistical distribution of the training set and would generate synthetic image examples that are different from the original images. This method would require each hospital to train an AI model to synthetically generate examples, and the resulting models would be transferred outside of the protected environments to generate synthetic images while still capturing disease-relevant imaging features. However, GANs have not been fully explored and must be applied with caution. An important metric for this approach is to ensure that the generated images are sufficiently different from the original images to preserve privacy. For example, the de-identification of a color fundus photograph may require that the resulting synthetic image cannot be re-

identified by retinal vessel configuration^{46,47} or that the membership to the training set cannot be established. An important tradeoff is performance difference in the real world when models are trained with the original data compared with generated data that preserve privacy.⁴⁵ If the generated synthetic datasets are too different from the original images, then the performance of the models trained with synthetic data would suffer and risk safety and efficacy when deployed.

Even when the data are not shared, there are other AI-related privacy issues that must be examined. Although individual patient data and imaging are safeguarded, traditionally the trained parameters or weights of AI models are not considered private or at risk for privacy breaches. However, large-parameter deep learning models, when trained with relatively few examples, can overfit and “memorize” these examples. AI models have been created to perform model inversion, where one AI model will attempt to reconstruct images with which another AI model was trained, which could potentially expose private data.⁴⁸ Recently, even federated learning schemes with differential privacy have been overcome using GANs.^{49,50} Clearly, the tools that are developed to protect data privacy will have to adapt quickly as AI technology evolves. It is also clear that such tools cannot be utilized in isolation; careful consideration of ethical and legal frameworks, with the adoption of appropriate safeguards, will also be necessary. Furthermore, it is important to also consider potential unintended ethical consequences; for example, the use of GANs and federated learning could lock in incumbents who have the resources to develop such systems, thus inhibiting wider dissemination of clinical data for patient benefit.

Conclusions

AI and Big Data have introduced privacy concerns that require solutions and updated regulations. New regulations governing privacy must be created to protect against inappropriate use of data, accidental disclosures, and weaknesses in de-identification techniques.⁴ However, we must also acknowledge that the overprotection of data may be detrimental to the data-driven innovation that ultimately improves our overall healthcare system.² Ophthalmology has been at the forefront of AI development, but there is also much to learn from other medical specialties that have adopted AI and are also confronting these issues. A successful balance is possible as thoughtful solutions that can adapt to evolving technology are implemented. Education of our patients and the public, alongside transparency about usage and sharing, will become vital as this field rapidly matures.

Acknowledgments

This research has been funded by National Institutes of Health Grants NIH/NEI K23EY029246, R01AG060942, R01EY015473, R01EY19474, P30EY10572, and K12EY27720 and by an unrestricted grant from Research to Prevent Blindness. The sponsors or funding organizations had no role in the design or conduct of this research.

AAO Artificial Intelligence Task Force Members

Michael F. Chiang, MD (Chair), Departments of Ophthalmology and Medical Informatics & Clinical Epidemiology, Casey Eye Institute, Oregon Health & Science University, Portland, OR, USA

Michael D. Abramoff, MD, PhD, Retina Service, Departments of Ophthalmology and Visual Sciences, Electrical and Computer Engineering, and Biomedical Engineering, University of Iowa, Iowa City, IA, USA; IDx, Coralville, IA, USA

J. Peter Campbell, MD, MPH, Department of Ophthalmology, Oregon Health & Science University, Portland, OR, USA

Pearse A. Keane, MD, FRCOphth, Institute of Ophthalmology, University College London, UK; Medical Retina Service, Moorfields Eye Hospital NHS Foundation Trust, London, UK

Aaron Y. Lee, MD, MSCI, Department of Ophthalmology, University of Washington, Seattle, WA, USA

Flora C. Lum, MD, American Academy of Ophthalmology, San Francisco, CA, USA

Louis R. Pasquale, MD, Eye and Vision Research Institute, Icahn School of Medicine at Mount Sinai, New York, NY, USA

Michael X. Repka, MD, MBA, Wilmer Eye Institute, Johns Hopkins University, Baltimore, MD, USA

Rishi P. Singh, MD, Cole Eye Institute, Cleveland Clinic, Cleveland, OH, USA

Daniel Ting, MD, PhD, Singapore National Eye Center, Duke-NUS Medical School, Singapore, Singapore

Disclosure: **E. Tom**, None; **P.A. Keane**, DeepMind Technologies (C), Roche (C), Novartis (C), Apellis (C), Bayer (F), Allergan (F), Topcon (F), Heidelberg

Engineering (F); **M. Blazes**, None; **L.R. Pasquale**, Verily (C), Eyenovia (C), Nicox (C), Bausch + Lomb (C), Emerald Bioscience (C); **M.F. Chiang**, Novartis (C), InTeleretina, LLC (I); **A.Y. Lee**, U.S. Food and Drug Administration (E), Genentech (C), Topcon (C), Verana Health (C), Santen (F), Novartis (F), Carl Zeiss Meditec (F); **C.S. Lee**, None; **M.D. Abramoff**, IDx (I, F, E, P, S), Alimera (F); **J.P. Campbell**, Genentech (F); **R. Singh**, Genentech (C), Novartis (C), Apellis (F), Bayer (C), Carl Zeiss Meditec (C), Aerie (C), Graybug (F), Regeneron (C); **D. Ting**, EyRIS (I, P), Novartis (C), Ocutrx (I, C), Optomed (C)

* ET and PAK contributed equally to the work presented here and therefore should be considered equivalent co-first authors. AYL and CSL contributed equally to the work presented here and therefore should be considered equivalent co-last authors.

References

1. Mooney SJ, Westreich DJ, El-Sayed AM. Commentary. *Epidemiology*. 2015;26:390–394.
2. Price WN, Nicholson Price W, Glenn Cohen I. Privacy in the age of medical big data. *Nat Med*. 2019;25:37–43.
3. Hoffman S. *Electronic Health Records and Medical Big Data: Law and Policy*. Cambridge, UK: Cambridge University Press; 2016.
4. He J, Baxter SL, Xu J, Xu J, Zhou X, Zhang K. The practical implementation of artificial intelligence technologies in medicine. *Nat Med*. 2019;25:30–36.
5. Abramoff MD, Lavin PT, Birch M, Shah N, Folk JC. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *NPJ Digit Med*. 2018;1:39.
6. Ting DSW, Cheung CY-L, Lim G, et al. Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes. *JAMA*. 2017;318:2211–2223.
7. Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*. 2016;316:2402–2410.
8. Lee CS, Tyring AJ, Wu Y, et al. Generating retinal flow maps from structural optical coherence tomography with artificial intelligence. *Sci Rep*. 2019;9:5694.
9. Kihara Y, Heeren TFC, Lee CS, et al. Estimating retinal sensitivity using optical coherence

- tomography with deep-learning algorithms in macular telangiectasia type 2. *JAMA Netw Open*. 2019;2:e188029.
10. Brown JM, Peter Campbell J, Beers A, et al. Automated diagnosis of plus disease in retinopathy of prematurity using deep convolutional neural networks. *JAMA Ophthalmol*. 2018;136:803–810.
 11. Liu H, Li L, Wormstone M, et al. Development and validation of a deep learning system to detect glaucomatous optic neuropathy using fundus photographs. *JAMA Ophthalmol*. 2019;137:1353–1360.
 12. Yan Q, Weeks DE, Xin H, et al. Deep-learning-based prediction of late age-related macular degeneration progression. *Nat Mach Intell*. 2020;2:141–150.
 13. Poplin R, Varadarajan AV, Blumer K, et al. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nat Biomed Eng*. 2018;2:158–164.
 14. Wen JC, Lee CS, Keane PA, et al. Forecasting future Humphrey visual fields using deep learning. *PLoS One*. 2019;14:e0214875.
 15. Arcadu F, Benmansour F, Maunz A, Willis J, Haskova Z, Prunotto M. Deep learning algorithm predicts diabetic retinopathy progression in individual patients. *NPJ Digit Med*. 2019;2:92.
 16. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research*. Washington, DC: U.S. Government; 1978.
 17. Beauchamp TL, Childress JF. *Principles of Biomedical Ethics*. 8th ed. New York: Oxford University Press; 2019.
 18. Abramoff MD, Tobey D, Char DS. Lessons learnt about autonomous AI: finding a safe, efficacious and ethical path through the development process. *Am J Ophthalmol*. 2020;214:134–142.
 19. Larson DB, Magnus DC, Lungren MP, Shah NH, Langlotz CP. Ethics of using and sharing clinical imaging data for artificial intelligence: a proposed framework. *Radiology*. 2020;295:675–682.
 20. Department of Health & Social Care. Guidance: The NHS Constitution for England. Available at: <https://www.gov.uk/government/publications/the-nhs-constitution-for-england/the-nhs-constitution-for-england>. Accessed March 31, 2020.
 21. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25:44–56.
 22. Shi M, Jiang R, Hu X, Shang J. A privacy protection method for health care big data management based on risk access control. *Health Care Manag Sci*. 2019;23:1–16.
 23. Rocher L, Hendrickx JM, de Montjoye Y-A. Estimating the success of re-identifications in incomplete datasets using generative models. *Nat Commun*. 2019;10:3069.
 24. Kamarajugadda KK, Polipalli TR. Extract features from periocular region to identify the age using machine learning algorithms. *J Med Syst*. 2019;43:196.
 25. Manrique de Lara A, Peláez-Ballestas I. Big data and data processing in rheumatology: bioethical perspectives. *Clin Rheumatol*. 2020;39:1007–1014.
 26. Martínez AG. No, data is not the new oil. Available at: <https://www.wired.com/story/no-data-is-not-the-new-oil/>. Accessed March 31, 2020.
 27. Chiang MF, Sommer A, Rich WL, Lum F, 2nd Parke DW. The 2016 American Academy of Ophthalmology IRIS Registry (Intelligent Research in Sight) database: characteristics and methods. *Ophthalmology*. 2018;125:1143–1148.
 28. Wellcome Trust Ltd. Why an opt-out rather than an opt-in or consent? Available at: <https://understandingpatientdata.org.uk/news/why-an-opt-out>. Accessed March 31, 2020.
 29. Noor P. Can we trust AI not to further embed racial bias and prejudice? *BMJ*. 2020;368:m363.
 30. National Data Guardian for Health and Care. Review of data security, consent and opt-outs. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/535024/data-security-review.PDF. Accessed March 31, 2020.
 31. De Fauw J, Ledsam JR, Romera-Paredes B, et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nat Med*. 2018;24:1342–1350.
 32. Information Commissioner's Office. Anonymisation: managing data protection risk code of practice. Available at: <https://ico.org.uk/media/for-organisations/documents/1061/anonymisation-code.pdf>. Accessed March 31, 2020.
 33. Hrynaskiewicz I, Norton ML, Vickers AJ, Altman DG. Preparing raw clinical data for publication: guidance for journal editors, authors, and peer reviewers. *BMJ*. 2010;340:c181.
 34. NHS Health Research Authority. How we're supporting data-driven technology. Available at: <https://www.hra.nhs.uk/planning-and-improving-research/research-planning/>

- [how-were-supporting-data-driven-technology/](#). Accessed March 31, 2020.
35. U.S. Department of Health & Human Services. Guidance regarding methods for de-identification of protected health information in accordance with the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule. Available at: <https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html>. Accessed March 23, 2020.
 36. U.S. Department of Health & Human Services. The HIPAA Privacy Rule. Available at: <https://www.hhs.gov/hipaa/for-professionals/privacy/index.html>. Accessed March 14, 2020.
 37. Lee D, Park M, Chang S, Ko H. Protecting and utilizing health and medical big data: policy perspectives from Korea. *Healthc Inform Res*. 2019;25:239–247.
 38. Berger B, Cho H. Emerging technologies towards enhancing privacy in genomic data sharing. *Genome Biol*. 2019;20:128.
 39. Abadi M, Chu A, Goodfellow I, et al. Deep learning with differential privacy. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. New York, NY: Association for Computing Machinery; 2016:308–318.
 40. Bonawitz K, Eichner H, Grieskamp W, et al. Towards federated learning at scale: system design. Available at: <https://arxiv.org/pdf/1902.01046.pdf>. Accessed June 22, 2020.
 41. Konečný J, Brendan McMahan H, Yu FX, Richtárik P, Suresh AT, Bacon D. Federated learning: strategies for improving communication efficiency. Available at: <https://arxiv.org/pdf/1610.05492.pdf>. Accessed June 22, 2020.
 42. Geyer RC, Klein T, Nabi M. Differentially private federated learning: a client level perspective. Available at: <https://arxiv.org/pdf/1712.07557.pdf>. Accessed June 22, 2020.
 43. Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative adversarial nets. In: Ghahramani Z, Welling M, Cortes C, Lawrence ND, Weinberger KQ, eds. *Advances in Neural Information Processing Systems 27*. Red Hook, NY: Curran Associates; 2014:2672–2680.
 44. Choi E, Biswal S, Malin B, Duke J, Stewart WF, Sun J. Generating multi-label discrete patient records using generative adversarial networks. Available at: <http://arxiv.org/abs/1703.06490>. Accessed March 20, 2020.
 45. Mukherjee S, Xu Y, Trivedi A, Ferres JL. Protecting GANs against privacy attacks by preventing overfitting. *ArXiv*, 2020;abs/2001.00071.
 46. Farzin H, Abrishami-Moghaddam H, Moin M-S. A novel retinal identification system. *EURASIP J Adv Signal Process*. 2008;2008:280635.
 47. Mariño C, Penedo MG, Penas M, Carreira MJ, Gonzalez F. Personal authentication using digital retinal images. *Pattern Anal Appl*. 2006;9:21.
 48. Fredrikson M, Jha S, Ristenpart T. Model inversion attacks that exploit confidence information and basic countermeasures. In: *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. CCS '15. New York, NY: Association for Computing Machinery; 2015:1322–1333.
 49. Hitaj B, Ateniese G, Perez-Cruz F. Deep models under the GAN: information leakage from collaborative deep learning. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. New York, NY: Association for Computing Machinery; 2017:603–618.
 50. Bagdasaryan E, Veit A, Hua Y, Estrin D, Shmatikov V. How to backdoor federated learning. Available at: <https://arxiv.org/pdf/1807.00459.pdf>. Accessed June 22, 2020.