Published in final edited form as:
J Am Coll Radiol. 2020 September ; 17(9): 1159–1162. doi:10.1016/j.jacr.2020.04.007.

Medical Imaging and Privacy in the Era of Artificial Intelligence: Myth, Fallacy, and the Future

Eyal Lotan, MD, PhD, **Charlotte Tschider, PhD**, **Daniel K. Sodickson, MD, PhD**, **Arthur L. Caplan, PhD**, **Mary Bruno, B.Sc. R.T.(R)(MR)**, **Ben Zhang, MSc**, **Yvonne W. Lui, MD**

Eyal Lotan, MD, PhD, Daniel K. Sodickson, MD, PhD, Mary Bruno, B.Sc. R.T.(R)(MR), Ben Zhang, MSc, and Yvonne W. Lui, MD, are from the Department of Radiology, New York University School of Medicine, New York, NY. Charlotte Tschider, PhD, is from the College of Law, University of Nebraska, Lincoln, Nebraska. Daniel K. Sodickson, MD, PhD, is from the Center for Advanced Imaging Innovation and Research, New York, New York. Arthur L. Caplan, PhD, is from the Division of Medical Ethics, New York University, New York, NY.

Twenty years ago, any medical images that actually originated in digital form were printed to film or otherwise rendered analog, and radiographs were housed in jacketed folders in a film library. There was room for privacy breaches but not nearly on the scale possible today. Today, medical images are digitized, stored, sent, and downloaded in a variety of scenarios. With the burgeoning field of computer vision research using artificial intelligence (AI) now being applied in earnest to medical imaging, we face a new era of unique scientific challenges regarding medical image data handling and patient privacy [1,2]. If we are to see societal benefit from technological advances in computer vision applied to medical imaging data (eg, improved diagnostics and patient care), it is imperative that the United States consider new models for identifiability and patient consent.

Across industries and specifically in medicine, there is a growing discussion about individual privacy in the era of big data, because various forms of data may contain identifiable personal information. HIPAA is the major federal law regulating the collection and use of protected health information (PHI) [3]. Medical images, in particular, occupy a unique position in the realm of health information, sharing some characteristics associated with structured, text data in the electronic health record (eg, laboratory values, medication doses, physiologic characteristics), and other characteristics associated with biologic tissue specimens. Structured, text electronic health record data are comparatively easy to aggregate and deidentify. Such data are commonly shared by health organizations with service providers or partner facilities for quality assessment, analytics, and billing. On the other hand, distribution of tissue specimens and the use of deidentified genomic data for secondary research require explicit patient consent. Radiologic images reside in the middle ground between these two forms of medical data. In various ways and on a previously unprecedented scale, we face a major new challenge in determining how best to handle medical images. Their particular combination of ease of transferability, on one hand, and

Yvonne W. Lui, MD: Department of Radiology, New York University School of Medicine, 660 First Avenue, Room 336, New York, NY 10016; yvonne.lui@nyulangone.org.

The authors state that they have no conflict of interest related to the material discussed in this article..

highly personal information content, on the other, reveals weaknesses in our current policy and practice.

THE MYTH OF DEIDENTIFICATION

The US federal law governing health data privacy is HIPAA’s privacy rule, which does not allow the secondary use of PHI without patient authorization or institutional review board waiver [3]. One of HIPAA’s most important strategies for protecting patients from privacy breaches and enabling effective data use and data sharing is data deidentification. The deidentification safe harbor usually permits organizations complying with HIPAA to use, share, and even sell PHI, as long as they remove 18 identifiers, such as name, date of birth, and address. Notably, images are not included on the list of identifiers, although inherently identifiable images, such as facial images, would likely be viewed as protected from sharing. Although the field of data science is rapidly evolving, HIPAA is nearly two decades old. What constitutes appropriate deidentification for medical images used in AI systems needs reassessment.

When considering advanced health care imaging in the AI marketplace, a question arises as to whether complete anonymization (deidentification without the ability to reidentify the data) is even possible. The basic task seems straightforward: selectively remove or codify identifiers in the metadata header content of images. Although nearly all radiologic data use a universal format, DICOM, there are a growing number of exceptions, making it more difficult to standardize processes.

Until the past decade, most medical imaging comprised 2-D images; however, advances in imaging methods have made high-resolution 3-D imaging a reality using state-of-the-art smoothing, interpolation, and super-resolution methods enabling accurate volume rendering. With advances in facial recognition, there is no reason why one would not be able to match images generated from CT or MRI scans to photographs of an individual (Fig. 1). As a result, it is standard practice in medical imaging research to modify images using defacing or skull-stripping algorithms to remove facial features. Unfortunately, such modifications can negatively affect the generalizability of machine-learning models developed using such data (Fig. 2).

High-resolution rendering of the face from medical images may ultimately be less controversial than imaging of other body parts, because it is standard practice in medical imaging research to modify images to remove facial features using defacing or skull-stripping algorithms. However, increasingly powerful methods of part-to-part comparison render nearly all anatomic areas potentially distinct [4]. Recently, investigators have shown that patterns of brain activation and metabolic signatures reveal characteristics unique to an individual [5]. Such studies raise further questions down the line regarding the possibility of revealing not only anatomic information but information reflecting a person’s thoughts [6].

THE FALLACY OF CONSENT

Although HIPAA permits the secondary use of deidentified patient information, in today’s AI world, we have seen how HIPAA-defined deidentification is not effective in preventing privacy harms and balancing scientific interests. The richness of medical imaging data combined with personally identifiable image content complicates our previous notions about how best to protect patient privacy. Where deidentification is not possible, organizations regulated by HIPAA must solicit patient authorization, including explicit consent, which restricts use to details shared in that authorization and limits use to a specified time period or purpose.

In the European Union, data privacy standards were promulgated in the General Data Protection Regulation, which will frequently require both (1) notice of AI used to make decisions and (2) explicit consent for collection and use of sensitive personal data, including medical data. Without both notice and consent, collection and use of medical data, exempting scientific research purposes, requires complete anonymization.

In theory, patient consent models should reinforce established clinical and research practices for informed consent. However, consent models vary in specificity of data use terms, and it is as yet unclear which (if any) type of consent may be the most appropriate for medical imaging applications, specifically for AI. Aggregated health data can be used to understand population health patterns, enhance research, and improve AI algorithms, providing real benefit to society. Unfortunately, if consent is used as a curative step for irresponsible data-handling practices, it is unlikely to meaningfully enhance patient choice [2]. Furthermore, restricting data use to limited, disclosed purposes, common in notice and consent models such as HIPAA authorization, is incompatible with the populating of large, high-quality training data sets and the ongoing learning process that are essential to modern AI systems.

THE FUTURE OF MEDICAL IMAGING DATA USE

As the amount of data composed of or including medical images increases, our capacity to extract private information from such data also advances. What was once considered innocuous can no long be treated as innocent. For AI in medical imaging, a consistent approach to patient information and education, including policies and consent models, as well as a modernized version of deidentification, is needed to appropriately consider patient privacy interests without hindering technological development. Indeed, patients are willing to share their data to assist developing AI tools, as long as it cannot be traced back to them as individuals [7]. As a field, radiologists and data scientists working in the medical imaging space should seek active involvement from the AMA, patient advocacy organizations, and global standard-bearing organizations, such as the International Standards Organization, to establish standards to guide regulatory bodies charged with delivering such directives.

A balanced solution likely involves (1) making information about data collection and use available, succinct, and understandable for patients; (2) building on existing relationships of trust between institutions and their patients; and (3) simultaneously pursuing more effective deidentification models that reduce identifiability rather than aiming for the false goal of complete anonymization. Current best practices for deidentification in radiology include working with manufacturers to avoid placement of identifiable data in proprietary DICOM fields, optimizing protocols to minimize data risks, using validated and tested protocols for deidentification, and investigating safer means of data sharing such as containerization and blockchain. There are calls for formation of advisory committees to periodically review how best to define the concept of identifiability with advances in imaging technology and how to best implement the needed safe-guards [8]. It will not be enough to expand HIPAA to include more and different fields, nor does it make sense to require greater and more complex consent which patients cannot understand. Devising appropriately innovative and dynamic solutions will require dynamic conversations between scientific investigators, policymakers, and the public.

REFERENCES

1. Price WN, Cohen IG. Privacy in the age of medical big data. Nat Med 2019;25:37–43. [PubMed: 30617331]
2. Tschider C The consent myth: improving choice for patients of the future. Wash L Rev 2018;96:1505–36.
3. The Health Insurance Portability and Accountability Act of 1996 (HIPAA). Pub L No 104–191, August 21, 1996; Security and Privacy 45 CFR Part 160.103 (2002).
4. Decker SJ, Ford JM. Forensic personal identification utilizing part-to-part comparison of CT-derived 3D lumbar models. Forensic Sci Int 2019;294:21–6. [PubMed: 30445251]
5. Finn ES, Shen X, Scheinost D, et al. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. Nat Neurosci 2015;18:1664–71. [PubMed: 26457551]
6. Bauer AJ, Just MA. Brain reading and behavioral methods provide complementary perspectives on the representation of concepts. Neuroimage 2019;186:794–805. [PubMed: 30458304]
7. Adams SJ, Tang R, Babyn P. Patient perspectives and priorities regarding artificial intelligence in radiology: opportunities for patient-centered radiology. J Am Coll Radiol. In press.
8. Geis JR, Brady AP, Wu CC, et al. Ethics of artificial intelligence in radiology: summary of the joint European and North American multisociety statement. J Am Coll Radiol 2019;16:1516–21. [PubMed: 31585696]

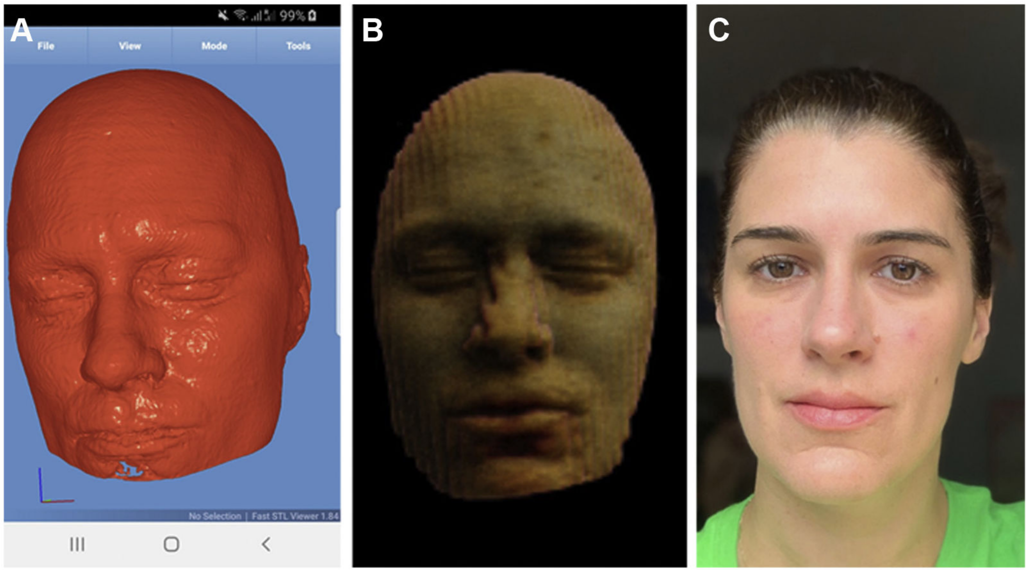

Fig 1.

Whereas in the past, dedicated medical imaging software was required, surface rendering of high-resolution (1-mm) medical imaging data is now easily achievable using freely available software such as Fast STL Viewer, an app intended for 3-D printing, shown here in use on an Android mobile device (A). Advances in image postprocessing methods such as slice interpolation allow the generation of a facial likeness even from 2-D, thick-section images (5 mm), albeit with noticeable stairstep artifacts (B). A digital photograph of the same individual (C) shows that identifying people using reconstructed images alone may not be trivial; however, advances in facial recognition algorithms are likely to resolve current shortcomings. Images used with permission.

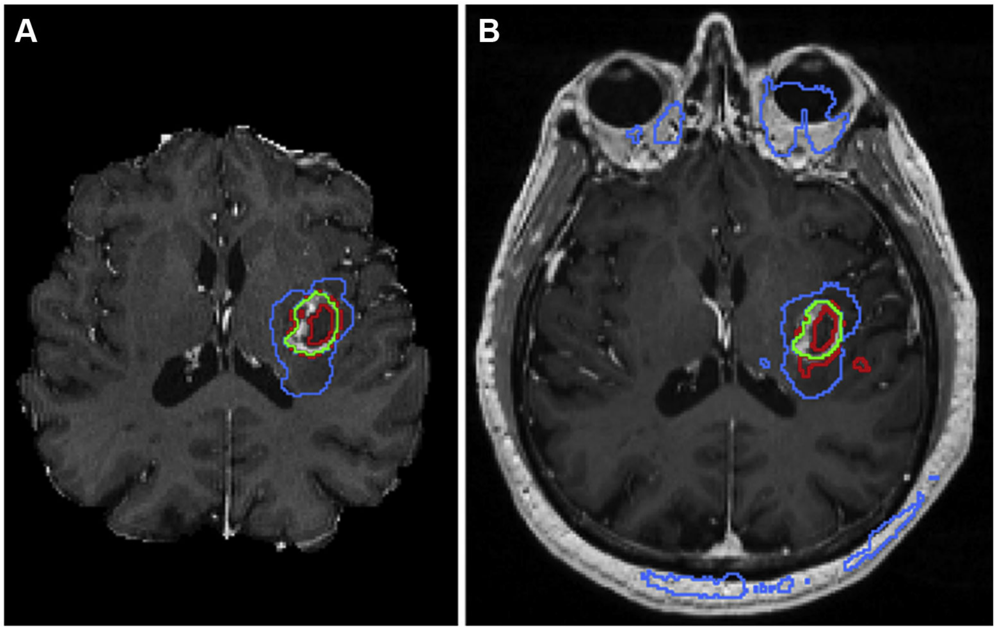

Fig 2.

Implementation of one of the top-performing segmentation algorithms from the 2017 Medical Image Computing and Computer Assisted Intervention Brain Tumor Segmentation Challenge shows excellent results (A) when data are preprocessed using skull-stripping and registration algorithms before inputting into a segmentation neural network; however, when the algorithm is deployed on a clinical MRI data set, the resulting segmentation contains many errors (B). Here, image segments identified as enhancing tumor are outlined in red, regions of enhancing tumor with any central nonenhancing components are shown in green, and regions of tumor with surrounding edema are shown in blue.