What Are the Data Privacy Concerns with AI Medical Imaging?

Rasit Dinc

Rasit Dinc Digital Health & AI Research

Published: November 30, 2020 | AI in Medical Imaging and Diagnostics

DOI: 10.5281/zenodo.17998689

Abstract

The integration of artificial intelligence (AI) into medical imaging has ushered in a new era of diagnostic capabilities, promising to enhance the accuracy a...

What Are the Data Privacy Concerns with AI Medical Imaging?

By Rasit Dinc

The integration of artificial intelligence (AI) into medical imaging has ushered in a new era of diagnostic capabilities, promising to enhance the accuracy and efficiency of healthcare delivery. From detecting cancers to predicting disease progression, AI algorithms are rapidly becoming indispensable tools for clinicians. However, this technological leap forward is not without its challenges. As AI models are trained on vast datasets of medical images, significant concerns have emerged regarding the privacy and security of sensitive patient information. For health professionals, understanding these data privacy issues is paramount to ethically and legally leveraging AI in their practice.

At the heart of the privacy debate is the inherent tension between the accuracy of AI models and the privacy of the data used to train them. Privacy-enhancing technologies (PETs) like **Differential Privacy (DP)** have emerged as a gold standard for mitigating the risk of information leakage from AI systems [1]. DP works by introducing a quantifiable amount of statistical noise to the data, making it difficult to determine whether any single individual's data was used in the training process. However, this comes at a cost. Research has shown that implementing restrictive privacy budgets in DP can lead to a significant degradation in the performance of AI models, particularly when working with smaller or more complex datasets. For instance, one study found that a model's diagnostic accuracy dropped to almost chance level when a very low privacy budget was applied [2]. This creates a difficult trade-off for developers and clinicians: how to balance the need for robust privacy protections with the demand for highly accurate diagnostic tools.

For many years, the primary method for protecting patient data has been de-identification, which involves removing personal identifiers such as names and birthdates from medical records. However, in the age of AI, this approach is proving to be insufficient. Studies have demonstrated that even when direct identifiers are removed, it is possible to re-identify individuals from medical images. For example, the facial contours from an MRI scan can be reconstructed and matched with publicly available photographs, effectively nullifying the de-identification process [1]. This vulnerability underscores the need for more advanced PETs that go beyond simple anonymization and provide a formal guarantee of privacy.

In addition to Differential Privacy, several other PETs are being explored for use in medical imaging. **Federated learning** is a promising approach that allows AI models to be trained on decentralized datasets without the data ever leaving its source institution. This helps to maintain data governance and reduces the risk of a centralized data breach. However, federated learning alone is not a silver bullet. Without the addition of other privacy measures like DP, it can still be vulnerable to data reconstruction attacks. **Cryptographic techniques**, such as homomorphic encryption and secure multi-party computation, offer another layer of protection by allowing computations to be performed on encrypted data. While these methods provide strong privacy guarantees, their practical application can be limited by their computational overhead, especially at the point of inference.

The use of AI in medical imaging is also governed by a complex web of regulatory and ethical considerations. In the United States, the **Health Insurance Portability and Accountability Act (HIPAA)** sets the standard for protecting sensitive patient health information. Similarly, the **General Data Protection Regulation (GDPR)** in the European Union imposes strict rules on the processing of personal data. Health professionals and organizations must ensure that their use of AI complies with these regulations to avoid significant legal and financial penalties. Beyond legal compliance, there are profound ethical questions that must be addressed. These include issues of patient consent, data ownership, and the potential for AI to perpetuate or even amplify existing biases in healthcare data.

Looking to the future, several innovative approaches are being developed to address the data privacy challenges in AI medical imaging. The use of **synthetic data**, where AI models are trained on artificially generated images rather than real patient data, is a particularly promising avenue. This could allow for the development of powerful AI tools without ever putting patient privacy at risk. Another area of active research is the development of **hybrid privacy models** that combine multiple PETs to create a more comprehensive and robust solution. For example, a system might use federated learning to train a model on decentralized data, while also applying DP to provide a formal privacy guarantee.

In conclusion, while AI holds immense promise for the future of medical imaging, the concerns surrounding data privacy are significant and must be addressed proactively. A multi-faceted approach that combines advanced PETs, a strong understanding of the regulatory landscape, and a commitment

to ethical principles is essential. By embracing these measures, health professionals can harness the power of AI to improve patient outcomes while upholding their fundamental duty to protect patient privacy.

References

[1] Ziller, A., et al. (2024). Reconciling privacy and accuracy in AI for medical imaging. *Nature Machine Intelligence*, 6, 764–774.

[2] Suri, A., & Summers, R. M. (2024). Privacy, Please: Safeguarding Medical Data in Imaging AI Using Differential Privacy Techniques. *Radiology: Artificial Intelligence*, 6(1), e230560.
