Trends in Genetics

CellPress OPEN ACCESS

Forum

Patient privacy in AI-driven omics methods

Juexiao Zhou,[1,2] Chao Huang,[3] and Xin Gao[1,2,*]

Artificial intelligence (AI) in omics analysis raises privacy threats to patients. Here, we briefly discuss risk factors to patient privacy in data sharing, model training, and release, as well as methods to safeguard and evaluate patient privacy in AI-driven omics methods.

Sequencing an individual's genome presented a significant challenge three decades ago. Now, many research projects such as The Cancer Genome Atlas (TCGA), the 100,000 Genomes Project, and the Earth BioGenome Project (EBP) have generated a flood of omics data, extracted from millions of individuals through high-throughput sequencing platforms like RNA-seq and single-cell RNA-seq (scRNA-seq). Many AI-driven omics methods took advantage of the omics data explosion and have facilitated significant advances in omics analysis. However, as more data are generated and published, a substantial volume of potentially private genetic information from those omics data may be increasingly exposed. The utilization of AI-driven omics methods further compounds this issue, posing a challenge to patient privacy [1]. Meanwhile, stringent policy regulations such as the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act of 1996 (HIPAA), and the California Consumer Privacy Act have been in place since the 20th century to safeguard patient privacy [2]. In this situation, three interconnected questions are raised to the community as shown in Figure 1. (i) What are the patient privacy issues associated with omics data and AI-driven omics methods? (ii) How can patient privacy be preserved in the development and application of AI-driven omics methods? (iii) How can we evaluate the adequacy of patient privacy protection and assess privacy risks associated with AI-driven omics methods?

Privacy was first defined by Samuel D. Warren and Louis Brandeis as the 'right to be let alone' in 1890. In recent years, the notion of the 'right to be forgotten' (RTBF) has also been integral to defining privacy [2,3]. However, in the realm of omics, privacy assumes a more specific dimension, primarily concerned with protecting patients' genetic data or any sensitive information that could be inferred from omics data. It is important to note that the omics data itself contains the patient privacy. For example, various features extracted from genomics and transcriptomics data, such as SNPs, short tandem repeats (STRs), structural variants (SVs), and gene expression patterns, serve as biomarkers for many human traits, ranging from physical characteristics like eye and hair color to complex attributes like height and disease susceptibility [4]. In certain scenarios, even an individual's protein expression levels in blood plasma from proteomics data analysis could be used to identify individuals [5]. As one of the most advanced sequencing technologies, the analysis of spatial transcriptomics requires the usage of tissue images, which introduces an additional layer of privacy concerns. The aforementioned privacy considerations become especially critical during the sharing and storage of omics data, as these datasets become susceptible to privacy attacks. Methods such as data linkage, DNA phenotyping, pedigree analysis, genotype imputation, and phenotype inference pose potential risks to this process [6]. In addition to the data itself, the advent of AI-driven omics methods introduces privacy concerns. Processes like model training and public model release may cause potential privacy issues. Recent studies underscore the vulnerability of AI-driven omics methods, particularly those employing deep learning (DL) models. It has been demonstrated that, during the training process, it is possible to infer original data by extracting intermediate information like gradients [7]. Even after the model is trained, there remains the potential to extract private information about the training data by attacking the trained AI model [8]. These privacy intricacies necessitate a comprehensive understanding and vigilant management to ensure the responsible and secure advancement of omics and AI-driven methods.

Therefore, recent studies propose various strategies to protect patient privacy across distinct stages of omics data utilization and the development of AI-driven omics methods. One straightforward approach is controlled access, where data is placed behind controlled access barriers. For instance, the National Institutes of Health (NIH) altered its genomic summary results (GSR) data-sharing policies post-2008 to incorporate controlled access measures. Anonymization is another most commonly used strategy, although research indicates its vulnerability. Cryptographic approaches, including homomorphic encryption (HE) [9], trusted execution environments (TEEs), secure multiparty computation (SMPC) [10], federated learning (FL), differential privacy (DP) [1], and blockchain [11] are the most advanced methods [12] being studied in secure omics data sharing, storage, and the development of AI-driven omics methods. However, the majority of studies focus on the application of a singular approach to specific omics tasks. For example, formal privacy guarantees for the participants in research on SNPs, genome-wide association studies (GWASs), and differential gene expression analysis were extensively studied [13]. Limited work examines a

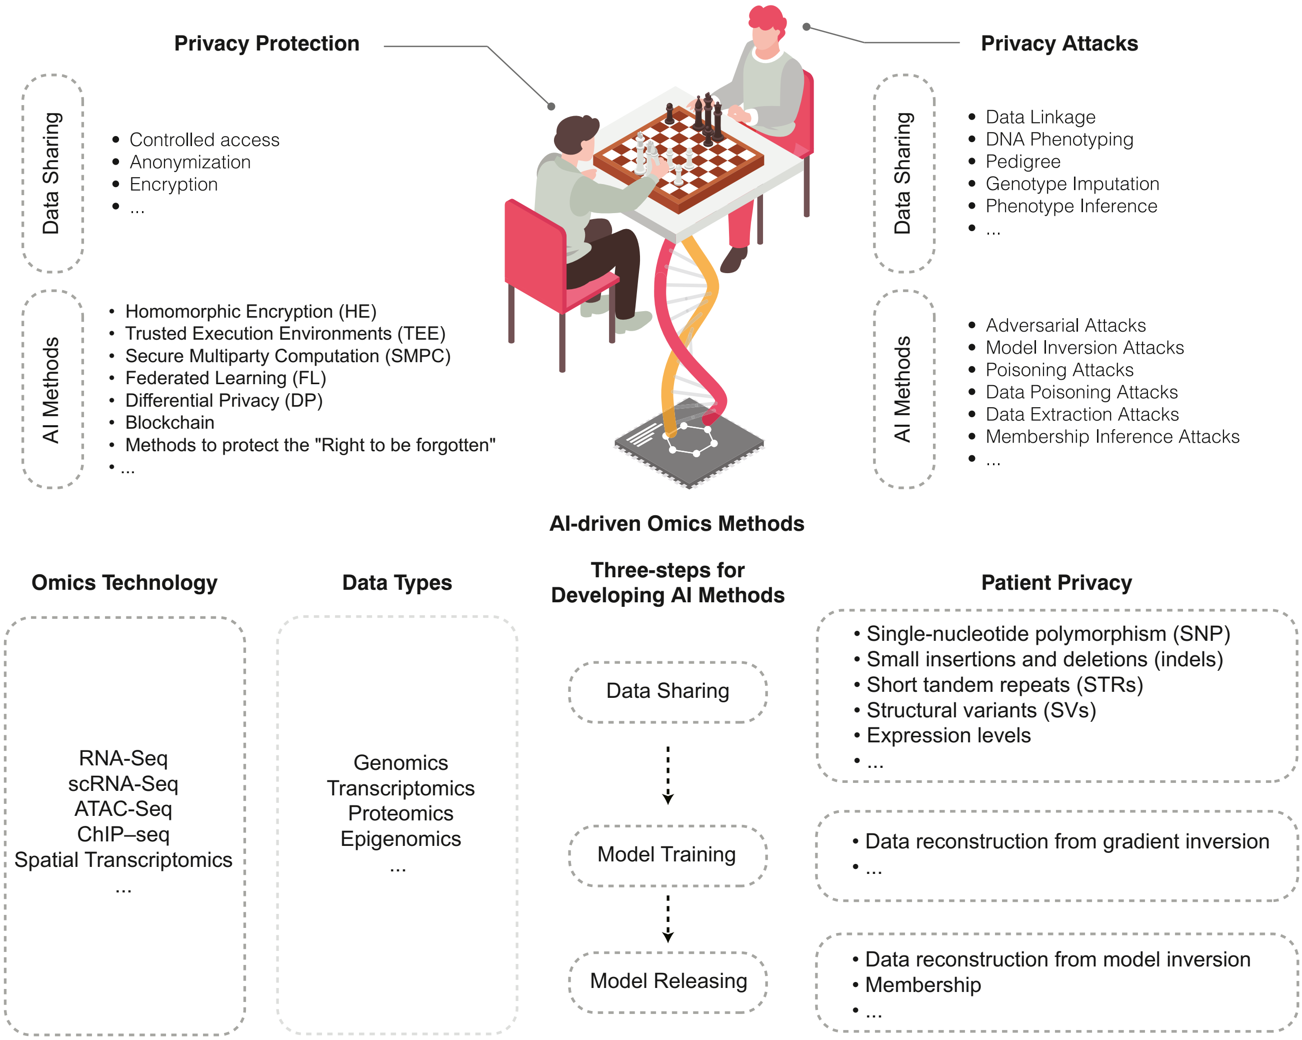

Trends in Genetics

Figure 1. Overview of patient privacy in artificial intelligence (AI)-driven omics methods: risks and solutions.

combination of approaches to further enhance privacy protection for the development of AI-driven omics methods. Recent studies, like swarm learning [11] and PPML-Omics [1], explore combinations like FL with blockchain technology and decentralized differential private FL algorithms, respectively, to bolster the private development of AI-driven omics methods. Despite their potential, these methods are not without trade-offs. Methods like HE and blockchain-based strategies, as seen in swarm learning, encounter a trade-off between privacy protection and computational overhead. By contrast, SMPC and FL grapple with the trade-off between communication overhead and privacy protection. Methods in the track of DP, such as PPML-Omics, face a trade-off between privacy protection and the utility of AI-driven omics methods. Therefore, reducing the cost and achieving better privacy protection is one of the main tasks in studies of the patient privacy of AI-driven omics methods. An additional noteworthy challenge pertains to safeguarding patients' RTBF in the context of published AI-driven omics methods. This is particularly crucial for pretrained DL models, given recent findings suggesting the potential retention of substantial private information by AI-driven omics methods [2]. Addressing these challenges is vital to ensure the responsible and ethical advancement of AI-driven methods in omics research.

To ascertain the efficacy of patient privacy protection in the development of AI-driven omics methods and safeguarding of the RTBF against released AI-driven omics methods, suitable privacy risk assessment methods are necessary. The delicate balance between privacy risks and potential benefits must be thoroughly understood to identify the most appropriate privacy methods or policies to secure patient privacy in the context of AI-driven omics methods. However, modeling privacy risk is challenging and often contingent on available information and the capabilities of potential adversaries, such that few studies have discussed this topic in biology [6]. Recognizing that privacy protection and privacy attacks are integral components of the same landscape, privacy attacks could be used to conduct a privacy risk assessment to a certain extent. While this approach has been extensively explored within the AI community in recent years, particularly in the context of adversarial attacks, model inversion attacks, poisoning attacks, data poisoning attacks, data extraction attacks, and membership inference attacks [14], its application to the domain of biology, especially in the context of patient omics data [4,15], deserves more studies, as shown in Table 1. Delving into the study of privacy protection for patient's omics data and assessing the resilience of AI-driven omics methods under various attack scenarios provides a promising avenue for future research. Understanding how these methods fare under the scrutiny of potential privacy attacks enables the refinement and enhancement of privacy measures, ensuring a more robust defense against adversarial attempts to compromise patient privacy in the evolving landscape of AI-driven omics research.

Table 1. Detailed scenarios and solutions using privacy technologies across three stages of AI-driven omics method development to protect patient privacy

Stage of development of AI-driven omics methods	Attack	Privacy risk	Privacy solution
Data sharing	Man-in-the-middle (MitM) attack	Intercepting communication between data centers to steal or manipulate omics data	Encryption protocols or digital signatures
	Data breaches	Unauthorized access to sensitive omics data in databases	Controlled access
	Insider threats	Malicious actions by authorized individuals leading to omics data leakage or misuse	Restriction of access to sensitive data on a need-to-know basis
	Data linkage attack and re-identification	Combining seemingly anonymized omics datasets enabling re-identification of individuals	Anonymization, DP, or k-anonymity techniques to prevent re-identification from linked datasets
	DNA phenotyping	Prediction of physical traits, such as facial features, hair color, and skin color, from DNA samples	Controlled access
	Pedigree analysis	Disclosure of sensitive family health information or potential identification of individuals through familial connections	Anonymization, pseudonymization, or encryption techniques to safeguard privacy
	Genotype imputation	Prediction of missing genetic variations in a dataset and increased risk of re-identification due to expanded genetic data	Controlled access, encryption, or use of DP techniques to mitigate the disclosure risk associated with imputed genotypes
Model training	Data poisoning attacks	Injection of malicious data into omics datasets leading to biased model training or exposure of sensitive genetic information	Encryption, SMPC, or FL approaches to protect raw data during model training
	Gradient construction attack	Exploitation of gradients to reverse engineer or reconstruct sensitive input omics data	Employment of SMPC or FL for secure gradient aggregation or DP for gradient perturbation
Model release	Membership inference attacks	Determination of whether an individual's genomic data was used in the model's training process, leading to privacy breaches	Application of DP mechanisms to prevent inference of membership in the training dataset
	Model inversion attacks	Extraction of sensitive genomic information from trained models, compromising individual privacy	Utilization of DP techniques to limit the leakage of sensitive information from model outputs
	Adversarial attacks	Manipulation of model outputs to induce misclassifications or extract sensitive genomic information	Training of robust models with adversarial training techniques to resist adversarial attacks
	Inference data leakage	Risk of input omics data being leaked when using online inference services such as Machine Learning as a Service (MLaaS)	Implementation of encrypted online inference techniques

In recent years, there has been significant progress in making AI-driven omics methods more private and secure. Alongside these technological advancements, it is equally important to improve the rules and guidelines that govern how we use these technologies. By tackling both the technical and the regulatory challenges related to privacy together, we can empower people to actively participate in scientific research, ultimately improving our understanding of omics and benefiting medical research.

Acknowledgments

We apologize to all those scientists whose important work could not be included in this Forum due to constraints on space and number of references. This work was supported by King Abdullah University of Science and Technology (KAUST). J.Z. and X.G. were supported in part by grants from the Office of Research Administration (ORA) at King Abdullah University of Science and Technology(KAUST) under award number FCC/1/1976-44-01, FCC/1/1976-45-01, REI/1/5202-01-01, REI/1/5234-01-01, REI/1/4940-01-01, RGC/3/4816-01-01, REI/1/0018-01-01, REI/1/5414-01-01, REI/1/5289-01-01, and REI/1/5404-01-01.

Declaration of interests

The authors declare that they have no competing interests.

[1]Computer Science Program, Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Thuwal, 23955-6900, Kingdom of Saudi Arabia
[2]Computational Bioscience Research Center, Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Thuwal, 23955-6900, Kingdom of Saudi Arabia
[3]Ningbo Institute of Information Technology Application, Chinese Academy of Sciences (CAS), Ningbo, China

*Correspondence:
xin.gao@kaust.edu.sa (X. Gao).

https://doi.org/10.1016/j.tig.2024.03.004

References

1. Zhou, J. *et al.* (2024) PPML-Omics: a privacy-preserving federated machine learning method protects patients' privacy in omic data. *Sci. Adv.* 10, eadh8601
2. Zhou, J. *et al.* (2023) A unified method to revoke the private data of patients in intelligent healthcare with audit to forget. *Nat. Commun.* 14, 6255
3. Rosen, J. (2011) The right to be forgotten. *Stan. L. Rev. Online* 64, 88
4. Gürsoy, G. *et al.* (2022) Functional genomics data: privacy risk assessment and technological mitigation. *Nat. Rev. Genet.* 23, 245–258
5. Geyer, P.E. *et al.* (2021) Plasma proteomes can be reidentifiable and potentially contain personally sensitive and incidental findings. *Mol. Cell. Proteomics* 20, 100035
6. Bonomi, L. *et al.* (2020) Privacy challenges and research opportunities for genomic data sharing. *Nat. Genet.* 52, 646–654
7. Zhao, B. *et al.* (2020) idlg: improved deep leakage from gradients. *arXiv*, Published online January 8, 2020. https://doi.org/10.48550/arXiv.2001.02610
8. Fredrikson, M. *et al.* (2015) Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, Association for Computing Machinery
9. Kim, M. *et al.* (2021) Ultrafast homomorphic encryption models enable secure outsourcing of genotype imputation. *Cell Syst.* 12, 1108–1120.e1104
10. Froelicher, D. *et al.* (2021) Truly privacy-preserving federated analytics for precision medicine with multiparty homomorphic encryption. *Nat. Commun.* 12, 5910
11. Warnat-Herresthal, S. *et al.* (2021) Swarm learning for decentralized and confidential clinical machine learning. *Nature* 594, 265–270
12. Kaissis, G.A. *et al.* (2020) Secure, privacy-preserving and federated machine learning in medical imaging. *Nat. Mach. Intell.* 2, 305–311
13. Zolotareva, O. *et al.* (2021) Flimma: a federated and privacy-aware tool for differential gene expression analysis. *Genome Biol.* 22, 338
14. Rigaki, M. and Garcia, S. (2023) A survey of privacy attacks in machine learning. *ACM Comput. Surv.* 56, 101
15. Wan, Z. *et al.* (2022) Sociotechnical safeguards for genomic data privacy. *Nat. Rev. Genet.* 23, 429–445