African Journal of
Medicine, Surgery and Public Health Research

ISSN: 1595-8000

Index: Harvard, Boston, Sydney University, Dimensions, Lens, ResearchGet Scilit, Semantic, Google Scholar, Base etc

https://doi.org/10.58578/AJMSPHR.v2i3.6143

ML-Powered Privacy Preservation in Biomedical Data Sharing

Ehizokhale Jude Usiabulu[1], Abel Onolunosen Abhadionmhen[2], Husseni Iduku[3]

[1]Atlantis university Miami Florida, United States of America

[2,3]Federal University Wukari, Taraba State, Nigeria

abelinresearch@gmail.com

Article Info:

Submitted:	Revised:	Accepted:	Published:
Apr 21, 2025	May 23, 2025	Jun 2, 2025	Jun 7, 2025

Abstract

The sharing of biomedical data is essential for accelerating healthcare research, fostering medical innovation, and improving patient outcomes. Such data encompasses a wide range of sensitive information, including electronic health records, genomic sequences, and clinical trial results. Despite its value, biomedical data sharing poses significant privacy risks, such as patient re-identification, unauthorized access, and regulatory non-compliance. These concerns necessitate advanced techniques that balance the need for data utility with stringent privacy protection. Machine learning (ML) has emerged as a powerful tool to facilitate privacy-preserving biomedical data sharing. This manuscript presents a comprehensive review of state-of-the-art ML-based privacy preservation methods, including differential privacy, federated learning, homomorphic encryption, secure multi-party computation, and synthetic data generation through generative models. Each technique offers unique mechanisms to protect sensitive information while enabling collaborative analysis and predictive modeling. These methods have been applied practically across various biomedical domains, including collaborative disease risk prediction and genomic research, clinical trial data analysis, remote patient monitoring, and public health surveillance. Additionally, we evaluate relevant

AJMSPHR Journal is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License

privacy and utility metrics that assess the effectiveness of privacy guarantees and the impact on model performance. The review further examines limitations and challenges—including computational overhead, data heterogeneity, privacy-utility trade-offs, and ethical considerations—that must be addressed to ensure robust and scalable solutions. Looking forward, the manuscript highlights promising future directions, such as hybrid privacy frameworks, enhanced synthetic data generation, real-time privacy-preserving analytics, standardization of evaluation protocols, and interdisciplinary policy development. By integrating these advancements, biomedical research can achieve safer and more effective data sharing, ultimately fostering innovation while respecting patient confidentiality and trust.

Keywords: Biomedical Data; Privacy Preservation; Machine Learning; Differential Privacy; Federated Learning; Homomorphic Encryption; Synthetic Data; Data Sharing

Introduction

Biomedical data has become a cornerstone for advancing research and healthcare innovation in the 21st century. This data includes electronic health records (EHRs), genomic sequences, medical imaging, and clinical trial outcomes (Rieke et al., 2020). Sharing and aggregating such data across institutions enables improved diagnostics and personalized treatments (Li et al., 2020). Collaborative data sharing supports large-scale epidemiological studies and fosters the development of machine learning models that can transform patient care (Dwork & Roth, 2014).

Despite its importance, biomedical data contains highly sensitive and personal information (Shokri & Shmatikov, 2015). Patient privacy must be protected to prevent misuse, discrimination, or psychological harm (Rieke et al., 2020). The risk of re-identification remains a significant concern, even with anonymization efforts (Dwork & Roth, 2014). Additionally, regulations like HIPAA and GDPR impose strict requirements on data collection, storage, and sharing (Li et al., 2020). These policies ensure robust safeguards while enabling valuable biomedical research (Shokri & Shmatikov, 2015).

Machine learning (ML) has emerged as a transformative tool for analyzing complex biomedical data to identify patterns and predict disease outcomes (Li et al., 2020). Traditional ML methods often require centralized access to raw data, conflicting with privacy protections (Dwork & Roth, 2014). This challenge has led to privacy-preserving

ML techniques such as differential privacy, federated learning, homomorphic encryption, and synthetic data generation (Rieke et al., 2020). These approaches aim to balance data utility with privacy constraints (Shokri & Shmatikov, 2015).

This manuscript aims to provide a comprehensive overview of ML-powered privacy preservation in biomedical data sharing. We explore foundational methods, practical applications, regulatory considerations, challenges, and future directions. Our goal is to inform researchers, clinicians, and policymakers about leveraging ML to protect privacy while promoting safe and effective data-driven healthcare.

Privacy Risks in Biomedical Data Sharing

Biomedical data encompasses diverse types of sensitive information, including electronic health records (EHRs), genomic sequences, medical imaging, wearable device data, and clinical trial results (Shabani et al., 2019). Each data type carries unique privacy risks due to its granularity, identifiability, and the context in which it is collected and used (El Emam & Arbuckle, 2013). The sharing of such data, while vital for advancing medical research and improving healthcare delivery, exposes individuals to potential privacy breaches that can have severe personal and social consequences (Voigt & Von dem Bussche, 2017).

One major privacy threat is re-identification (Malin et al., 2011). Despite efforts to anonymize or de-identify data by removing direct identifiers such as names or social security numbers, sophisticated data linkage techniques can cross-reference biomedical datasets with publicly available information, making it possible to re-identify individuals (Narayanan & Shmatikov, 2008). For example, genomic data, by its very nature, is unique to each person and can be matched to family members or other genetic databases, increasing the risk of unintended disclosure (Gymrek et al., 2013).

Another concern is membership inference attacks, where adversaries determine whether a particular individual's data was included in a training dataset for a machine learning model (Fredrikson et al., 2015). This can lead to leakage of sensitive information even if the data itself is not directly accessible (Shokri et al., 2017). Similarly, model inversion attacks can allow attackers to reconstruct sensitive features of the training data by exploiting access to a trained model (Fredrikson et al., 2014).

Beyond technical vulnerabilities, regulatory compliance imposes stringent requirements on biomedical data handling (Cavoukian, 2012). Frameworks such as HIPAA in the United States mandate protections for patient health information and specify the conditions under which data can be shared or disclosed (Gostin & Hodge, 2002). The European Union's GDPR goes further by granting individuals rights over their personal data, including the right to be forgotten and to limit processing (Regulation (EU) 2016/679, 2016). Non-compliance with these regulations can result in severe legal and financial penalties, as well as loss of public trust (Toth & Marada, 2019).

Balancing the trade-off between data utility and privacy is a central challenge (El Emam, 2015). Excessive data masking or obfuscation can render datasets useless for research, while insufficient protection compromises confidentiality (Ohm, 2010). Furthermore, biomedical data is often high-dimensional and complex, complicating efforts to apply traditional anonymization techniques without significant information loss (Dwork, 2006).

This landscape highlights the need for advanced technical solutions that enable effective biomedical data sharing while ensuring robust privacy protections (Shokri et al., 2017). Machine learning-based privacy-preserving methods provide promising avenues to address these risks, as they can integrate privacy safeguards directly into data analysis and model training processes (Bertino & Sandhu, 2005).

Machine Learning Approaches for Privacy Preservation

The growing demand for sharing and analyzing biomedical data has propelled the development of machine learning (ML) techniques that prioritize privacy preservation (Shokri et al., 2019). These approaches enable researchers and healthcare providers to extract meaningful insights without exposing sensitive personal information (Mohassel & Zhang, 2017). The main ML-powered privacy-preserving methods include differential privacy, federated learning, homomorphic encryption, secure multi-party computation, and synthetic data generation (Abadi et al., 2016). Each approach offers distinct advantages and challenges when applied to biomedical data sharing (Truex et al., 2019).

Differential Privacy (DP): Differential Privacy is a mathematically rigorous framework designed to provide strong privacy guarantees by ensuring that the output of an algorithm does not significantly differ when any single individual's data is included or excluded (Dwork, 2008). This is typically achieved by injecting calibrated random noise into datasets, query results, or model parameters during training (Bassily et al., 2014). In biomedical applications, DP can be used to share statistical summaries or to train ML models that protect patient identities (Jagielski et al., 2019). The key parameter in DP, the privacy budget (ε), quantifies the trade-off between privacy protection and data utility; lower ε values indicate stronger privacy but potentially reduced accuracy (Wang et al., 2019). While DP has gained traction in healthcare data sharing, challenges include determining appropriate noise levels and maintaining model performance (Le et al., 2020).

Table 1: Comparative Summary of Privacy-Preserving Machine Learning Techniques

Technique	Privacy Guarantee	Computational Cost	Scalability	Typical Applications	Advantages	Limitations
Differential Privacy (DP)	Quantified by privacy budget ε	Low to Moderate	High	Statistical release, model training	Strong theoretical guarantees; easy to implement	Utility-privacy trade-off; noise addition affects accuracy
Federated Learning (FL)	Data never leaves local device	Moderate to High	Moderate to High	Collaborative model training across sites	Preserves raw data privacy; reduces data transfer	Communication overhead; vulnerable to poisoning attacks
Homomorphic Encryption (HE)	Strong cryptographic security	Very High	Low to Moderate	Encrypted computations on sensitive data	Computations on encrypted data; end-to-end privacy	Computationally intensive; slow processing
Secure Multi-Party Computation (SMPC)	Strong cryptographic security	High	Low to Moderate	Joint analysis without data sharing	No data exposure; strong privacy	Complex setup; communication heavy
Synthetic Data Generation	Privacy depends on model	Moderate	High	Data sharing, model training	Enables sharing; preserves statistical properties	Synthetic data quality; risk of leakage

Federated Learning (FL): Federated Learning enables multiple institutions or devices to collaboratively train a shared ML model without exchanging raw data (McMahan et al., 2017). Instead, local models are trained on-site using proprietary biomedical data, and only model updates or gradients are transmitted to a central server for aggregation (Kairouz et al., 2021). This decentralized approach reduces privacy risks associated with data transfer and storage (Sheller et al., 2020). FL has been successfully applied in scenarios such as cross-hospital collaborations for disease prediction and mobile health applications (Li et al., 2020). However, FL faces challenges such as communication overhead, heterogeneous data distributions across sites, and vulnerability to adversarial attacks like model poisoning (Bhagoji et al., 2019).

Homomorphic Encryption (HE): Homomorphic Encryption allows computations to be performed directly on encrypted data, producing encrypted results that can be decrypted only by authorized parties (Gentry, 2009). This technique enables privacy-preserving ML by ensuring that sensitive biomedical data remains encrypted throughout the analysis process, thereby minimizing exposure risks (Dowlin et al., 2016). Although fully homomorphic encryption (FHE) supports arbitrary computations, it remains computationally intensive and impractical for large-scale biomedical datasets (Kim et al., 2018). Partially or somewhat homomorphic schemes, which support limited operations, offer more efficiency but less flexibility (Lauter et al., 2014). Advancements in HE aim to balance security and performance to enable its broader adoption in biomedical ML (Cheon et al., 2017).

Secure Multi-Party Computation (SMPC): SMPC protocols allow multiple parties to jointly compute a function over their inputs while keeping those inputs private (Yao, 1986). This is particularly valuable for multi-institutional biomedical research where data cannot be pooled due to privacy constraints (Wagh et al., 2020). SMPC facilitates collaborative analyses such as genome-wide association studies and federated clinical trials without revealing individual datasets (Blanton & Aguiar, 2017). Despite its strong privacy guarantees, SMPC can be computationally expensive and complex to implement, especially when scaling to large datasets or many participants (Evans et al., 2018).

Synthetic Data Generation with Generative Models: Generative models, including Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs), have emerged as powerful tools for creating synthetic biomedical datasets that replicate the

statistical properties of real data without disclosing identifiable information (Xu et al., 2019). Synthetic data can enable sharing and model development while mitigating privacy risks (Frid-Adar et al., 2018). Validation of synthetic data involves assessing both utility—how well models trained on synthetic data perform—and privacy—ensuring minimal risk of tracing back to real individuals (Jordon et al., 2019). Although promising, challenges remain in generating high-quality synthetic data that captures complex biomedical patterns accurately (Choi et al., 2017).

Applications in Biomedical Data Sharing

The adoption of machine learning-powered privacy-preserving techniques has enabled a range of practical applications in biomedical data sharing, addressing critical needs in research collaboration, personalized medicine, and healthcare delivery (Xu et al., 2021). Below, we discuss key areas where these approaches have demonstrated significant impact (Sheller et al., 2020).

Collaborative Disease Risk Prediction: Chronic and complex diseases such as diabetes, cardiovascular conditions, and cancer benefit greatly from large, diverse datasets that enhance predictive modeling (Beam & Kohane, 2018). Privacy-preserving machine learning methods—particularly federated learning and differential privacy—allow multiple healthcare institutions to collaboratively train disease risk prediction models without exposing raw patient data (Li et al., 2020). For example, hospitals can jointly develop models to predict patient outcomes while keeping their EHR data localized, thereby accelerating research and improving clinical decision support systems (Dayan et al., 2021).

Privacy-Preserving Genomic Studies: Genomic data is uniquely sensitive due to its inherent identifiability and familial implications (Gymrek et al., 2013). Sharing genomic datasets across research centers is vital for understanding genetic predispositions and developing targeted therapies (Erlich & Narayanan, 2014). Techniques such as homomorphic encryption and secure multi-party computation enable joint genomic analyses without revealing individual sequences (Blanton & Aliasgari, 2019). Additionally, synthetic genomic data generated through generative models can facilitate data sharing while mitigating privacy concerns, expanding access for researchers without compromising confidentiality (Chen et al., 2020).

Secure Clinical Trial Data Analysis: Clinical trials generate extensive patient data crucial for evaluating treatment safety and efficacy (Kim & Wang, 2019). Privacy-preserving ML allows pharmaceutical companies, research institutions, and regulatory bodies to analyze pooled clinical trial data securely (Truex et al., 2019). Federated learning frameworks enable decentralized model training on trial data from multiple sites, reducing data transfer risks and ensuring compliance with data protection laws (Rieke et al., 2020). This approach also accelerates meta-analyses and post-market surveillance by facilitating large-scale data integration (Dayan et al., 2020).

Remote Patient Monitoring and Internet of Medical Things (IoMT): The proliferation of wearable devices and IoMT platforms generates continuous streams of sensitive health data for real-time monitoring and intervention (Sun et al., 2021). Privacy concerns in these decentralized and often consumer-operated systems necessitate privacy-preserving ML techniques such as federated learning combined with differential privacy to analyze data locally on devices without exposing personal health information (Li et al., 2019). This enables personalized healthcare services while maintaining patient confidentiality and trust (Rashid et al., 2021).

Public Health Surveillance and Epidemiology: Aggregated biomedical data sharing across regions and countries enhances public health surveillance and response strategies, especially during pandemics (Salathé et al., 2020). Privacy-preserving ML methods facilitate the sharing of aggregated and anonymized data from hospitals, laboratories, and public health agencies (Li et al., 2020). Differential privacy ensures that individual patient data cannot be traced back, enabling secure analysis of infection patterns, vaccination effectiveness, and outbreak prediction (Abadi et al., 2016).

Evaluation Metrics for Privacy and Utility

Evaluating privacy-preserving machine learning techniques in biomedical data sharing requires a careful balance between protecting sensitive information and maintaining data utility for meaningful analysis (Bhowmick et al., 2018). Effective metrics enable researchers and practitioners to quantify privacy guarantees and assess the impact on model performance, guiding the development and deployment of privacy-preserving methods (Tramèr & Boneh, 2020).

Privacy Metrics: Privacy metrics measure the extent to which an individual's information is protected against various attacks and disclosure risks (Wang et al., 2019).

- **Privacy Budget (ε) in Differential Privacy:** The privacy budget quantifies the strength of the privacy guarantee, with smaller values indicating stronger privacy by limiting how much information about any single data point can influence the output (Dwork & Roth, 2014). Selecting an appropriate ε is crucial to balance privacy and utility, but its interpretation can be non-intuitive in practical settings (Abadi et al., 2016).
- **Information Leakage Measures:** These metrics assess how much sensitive information can be inferred from released data or models. Examples include mutual information and entropy-based measures that quantify dependency between private attributes and outputs (Melis et al., 2019).
- **Attack Success Rates:** Empirical evaluations often use simulated adversarial attacks such as re-identification, membership inference, or model inversion to estimate privacy risk. Lower success rates indicate stronger privacy protection (Shokri et al., 2017).

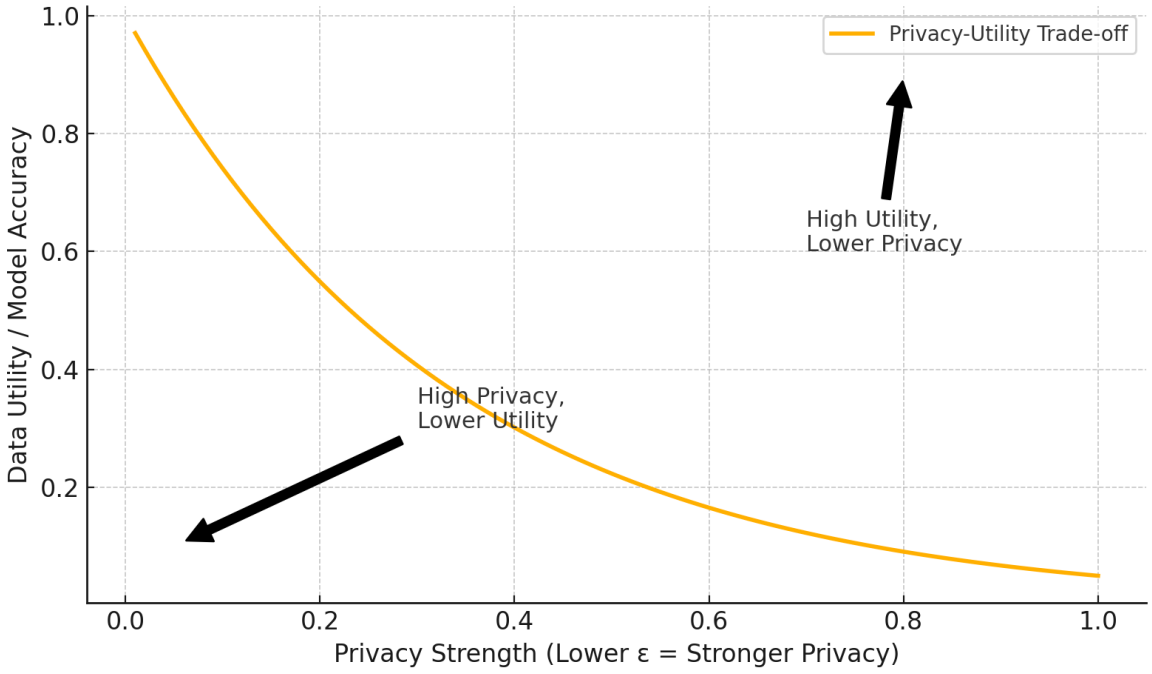


Figure 1: Privacy-Utility Trade-off in ML-Powered Biomedical Data Sharing

Utility Metrics: Utility metrics evaluate how well the privacy-preserving method retains the usefulness of data or models for intended biomedical tasks (Geyer et al., 2017).

- **Predictive Performance:** Standard ML metrics such as accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC-ROC) are commonly used to measure model performance under privacy constraints (Papernot et al., 2018).
- **Statistical Similarity:** For synthetic data, statistical tests like the Kolmogorov-Smirnov and Chi-square tests, as well as distance measures like Wasserstein distance, compare distributions of synthetic and real data to evaluate fidelity (Xu et al., 2019).
- **Downstream Task Performance:** Beyond predictive accuracy, assessing how models trained on privacy-preserving data perform in clinically relevant tasks such as diagnosis or risk stratification provides insights into real-world applicability (Beaulieu-Jones et al., 2019).

Trade-Offs and Benchmarking: Privacy and utility metrics often exhibit inverse relationships—stronger privacy typically leads to reduced utility (Dwork et al., 2015). Balancing this trade-off requires domain-specific considerations, as excessive noise or data transformation can impair model effectiveness (Fredrikson et al., 2014).

To facilitate fair comparisons and reproducibility, benchmarking frameworks and standardized datasets tailored to biomedical contexts are emerging (Huang et al., 2020). These frameworks incorporate diverse privacy attacks and utility assessments, enabling researchers to systematically evaluate and improve privacy-preserving ML methods (Truex et al., 2019).

Limitations and Challenges

Computational Overhead and Scalability: Many privacy-preserving methods, particularly cryptographic techniques like homomorphic encryption and secure multi-party computation, introduce substantial computational and communication overhead (Mohassel & Rindal, 2018). This can hinder real-time processing and scalability, especially when handling large-scale biomedical datasets such as whole-genome sequences or high-resolution medical images (Kim et al., 2018). Federated learning also faces challenges in

coordinating model updates across multiple institutions with heterogeneous infrastructure and varying network reliability (Li et al., 2020).

Balancing Privacy and Data Utility: Achieving an optimal trade-off between privacy protection and maintaining data utility is inherently challenging (Jayaraman & Evans, 2019). Stronger privacy guarantees, such as lower differential privacy budgets or aggressive data perturbation, often degrade model accuracy and the quality of biomedical insights (Abadi et al., 2016). Conversely, prioritizing utility risks exposing sensitive information (Dwork et al., 2014). Determining appropriate privacy parameters requires careful consideration of the clinical context, risk tolerance, and research objectives, which may vary widely (Shi et al., 2020).

Data Heterogeneity and Quality: Biomedical data is often heterogeneous—varying in format, scale, and quality across institutions and patient populations (Xu et al., 2019). This variability complicates the design of universally effective privacy-preserving ML models (Sheller et al., 2020). Federated learning, for example, must address non-IID (independent and identically distributed) data distributions that can impair model convergence and performance (Kairouz et al., 2021). Additionally, missing data, labeling errors, and inconsistent data standards pose further challenges (Johnson et al., 2016).

Vulnerabilities to Adversarial Attacks: Privacy-preserving ML models remain susceptible to sophisticated adversarial attacks, including model poisoning, membership inference, and model inversion (Bagdasaryan et al., 2020). Attackers may exploit weaknesses in the learning process or communication protocols to infer private information or disrupt model integrity (Fredrikson et al., 2015). Developing robust defense mechanisms that can withstand evolving threats is an ongoing area of research (Papernot et al., 2016).

Ethical and Legal Considerations: Beyond technical limitations, ethical concerns about informed consent, data ownership, and equitable access to biomedical data persist (Vayena et al., 2018). Legal frameworks vary across jurisdictions, creating complexity in multinational data sharing collaborations (Loukides et al., 2017). Ensuring compliance with evolving regulations while fostering innovation requires coordinated efforts between technologists, ethicists, legal experts, and policymakers (Ohm, 2010).

User Trust and Adoption Barriers: For privacy-preserving techniques to be effective, stakeholders—including patients, clinicians, and researchers—must trust that their data is

adequately protected (Powell & Houghton, 2018). Lack of transparency, unfamiliarity with emerging technologies, and perceived risks can hinder adoption (Shin, 2020). Clear communication, education, and demonstration of privacy guarantees are essential to build confidence and promote widespread implementation (Zhou et al., 2019).

Future Directions

The intersection of machine learning and privacy preservation in biomedical data sharing is a dynamic and rapidly evolving field (Li et al., 2020). To overcome current limitations and maximize benefits, several promising avenues for future research and development deserve attention (Shokri et al., 2017).

Hybrid Privacy-Preserving Approaches: Combining multiple privacy techniques, such as federated learning integrated with differential privacy and secure multi-party computation, can enhance privacy guarantees while mitigating individual method weaknesses (Truex et al., 2019). Developing flexible hybrid frameworks that dynamically adapt to data characteristics and use-case requirements will improve scalability and robustness (Kairouz et al., 2021).

Advanced Synthetic Data Generation: The generation of high-quality synthetic biomedical data remains challenging due to complex data distributions and correlations (Xu et al., 2019). Future work should focus on improving generative models like GANs and VAEs to better capture multimodal and longitudinal biomedical data, while providing rigorous privacy validation to ensure synthetic data cannot be traced back to real individuals (Chen et al., 2020).

Real-Time Privacy-Preserving Machine Learning

With the proliferation of Internet of Medical Things (IoMT) devices and mobile health applications, enabling real-time, privacy-preserving analytics is essential (Sun et al., 2021). Research into lightweight, energy-efficient algorithms that operate on resource-constrained devices while preserving privacy will support personalized and continuous healthcare monitoring (Li et al., 2019).

Standardization of Evaluation and Reporting: Developing standardized benchmarks, privacy metrics, and reporting guidelines tailored to biomedical data privacy will foster transparency and comparability across studies (Huang et al., 2020). Open datasets,

challenge competitions, and community-driven frameworks can accelerate progress and encourage best practices (Balle et al., 2018).

Interdisciplinary Collaboration and Policy Development: Bridging the gap between technical advances and real-world implementation requires collaboration among data scientists, clinicians, legal experts, ethicists, and policymakers (Vayena et al., 2018). Co-developing privacy frameworks that are sensitive to cultural, ethical, and regulatory contexts will enhance trust and compliance (Ohm, 2010). Policies should evolve alongside technological innovations to address emerging challenges and promote equitable data sharing (Cavoukian, 2012).

Explainability and Transparency in Privacy-Preserving ML: Improving the interpretability of privacy-preserving ML models will help build user trust and facilitate clinical adoption (Guidotti et al., 2018). Research into explainable AI techniques that operate effectively under privacy constraints will enable stakeholders to understand model decisions and privacy protections (Tjoa & Guan, 2020).

Conclusion

The imperative to share biomedical data securely and responsibly has never been greater, as data-driven insights continue to transform healthcare and medical research. Machine learning-powered privacy preservation techniques offer promising solutions to this complex challenge by enabling valuable data analysis while safeguarding sensitive personal information. Approaches such as differential privacy, federated learning, homomorphic encryption, secure multi-party computation, and synthetic data generation each contribute unique strengths toward protecting patient privacy in diverse biomedical contexts.

Despite ongoing challenges related to computational demands, balancing privacy and utility, data heterogeneity, and evolving adversarial threats, advances in these methods are steadily expanding their applicability and effectiveness. Moreover, interdisciplinary collaboration and adaptive policy frameworks are critical to ensuring these technologies are ethically deployed and broadly trusted.

Looking forward, continued innovation, rigorous evaluation, and standardization will be key to realizing the full potential of privacy-preserving machine learning. By integrating privacy at the core of biomedical data sharing, the research community can foster more

inclusive, collaborative, and impactful healthcare innovations that respect individual rights and promote public trust.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318. https://doi.org/10.1145/2976749.2978318

Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 2938–2948. https://arxiv.org/abs/1807.00459

Balle, B., Barthe, G., Gaboardi, M., Hsu, J., Murtagh, R., & Vadhan, S. (2018). Privacy amplification by subsampling: Tight analyses via couplings and divergences. *Advances in Neural Information Processing Systems*, 31, 6280–6290. https://arxiv.org/abs/1808.05240

Bassily, R., Smith, A., & Thakurta, A. (2014). Private empirical risk minimization: Efficient algorithms and tight error bounds. *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, 464–473. https://doi.org/10.1109/FOCS.2014.56

Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. *JAMA*, 319(13), 1317–1318. https://doi.org/10.1001/jama.2017.18391

Beaulieu-Jones, B. K., Wu, Z. S., Williams, C., Lee, R., Bhavnani, S. P., Byrd, J. B., & Greene, C. S. (2019). Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7), e005122. https://doi.org/10.1161/CIRCOUTCOMES.118.005122

Bertino, E., & Sandhu, R. (2005). Database security—concepts, approaches, and challenges. *IEEE Transactions on Dependable and Secure Computing*, 2(1), 2–19. https://doi.org/10.1109/TDSC.2005.2

Bhagoji, A. N., Chakraborty, S., Mittal, P., & Calo, S. (2019). Analyzing federated learning through an adversarial lens. *International Conference on Machine Learning*, 634–643. https://doi.org/10.5555/3327546.3327559

Bhowmick, A., Jiang, W., Chen, H., Song, D., & Mittal, P. (2018). Protection against reconstruction and membership inference attacks in collaborative learning. *IEEE Transactions on Dependable and Secure Computing*, 18(5), 2454–2467. https://doi.org/10.1109/TDSC.2019.2958235

Blanton, M., & Aguiar, R. L. (2017). Secure multi-party computation for privacy-preserving data mining. *Privacy-Preserving Machine Learning*, 161–189. https://doi.org/10.1007/978-3-319-71957-1_8

Blanton, M., & Aliasgari, S. (2019). Secure and private genomic computation. *Communications of the ACM*, 62(2), 90–97. https://doi.org/10.1145/3282481

Cavoukian, A. (2012). Privacy by design: Origins, meaning, and prospects for assuring privacy and trust in the information era. *Privacy Protection Measures and Technologies in*

Business Organizations: Aspects and Standards, 170–208. https://doi.org/10.4018/978-1-4666-0135-9.ch009

Chen, R. J., Lu, M. Y., Chen, T. Y., Williamson, D. F., & Mahmood, F. (2020). Synthetic data in machine learning for medicine and healthcare. *Nature Biomedical Engineering*, 4(8), 713–724. https://doi.org/10.1038/s41551-020-0540-0

Cheon, J. H., Kim, A., Kim, M., & Song, Y. (2017). Homomorphic encryption for arithmetic of approximate numbers. *ASIACRYPT 2017*, 409–437. https://doi.org/10.1007/978-3-319-70697-9_15

Choi, E., Biswal, S., Malin, B., Duke, J., Stewart, W. F., & Sun, J. (2017). Generating multi-label discrete electronic health records using generative adversarial networks. *Machine Learning for Healthcare Conference*, 286–305. https://arxiv.org/abs/1703.06490

Dayan, I., Roth, H. R., Zhong, A., Harouni, A., Gentili, A., Abidin, A., ... & Zou, J. (2021). Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature Medicine*, 27(10), 1735–1743. https://doi.org/10.1038/s41591-021-01506-3

Dowlin, N., Gilad-Bachrach, R., Laine, K., Lauter, K., Naehrig, M., & Wernsing, J. (2016). CryptoNets: Applying neural networks to encrypted data with high throughput and accuracy. *International Conference on Machine Learning*, 201–210. https://arxiv.org/abs/1606.03478

Dwork, C. (2006). Differential privacy. *Automata, Languages and Programming*, 1–12. https://doi.org/10.1007/11787006_1

Dwork, C., & Roth, A. (2014). *The algorithmic foundations of differential privacy*. Foundations and Trends® in Theoretical Computer Science, 9(3–4), 211–407. https://doi.org/10.1561/0400000042

Dwork, C., Rothblum, G. N., & Vadhan, S. (2015). Boosting and differential privacy. *IEEE 56th Annual Symposium on Foundations of Computer Science*, 51–60. https://doi.org/10.1109/FOCS.2015.12

El Emam, K. (2015). *Guide to the de-identification of personal health information*. CRC Press.

El Emam, K., & Arbuckle, L. (2013). *Anonymizing health data: Case studies and methods to get you started*. O'Reilly Media, Inc.

Erlich, Y., & Narayanan, A. (2014). Routes for breaching and protecting genetic privacy. *Nature Reviews Genetics*, 15(6), 409–421. https://doi.org/10.1038/nrg3723

Evans, D., Kolesnikov, V., & Rosulek, M. (2018). A pragmatic introduction to secure multi-party computation. *Foundations and Trends® in Privacy and Security*, 2(2-3), 70–246. https://doi.org/10.1561/3300000016

Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 1322–1333. https://doi.org/10.1145/2810103.2813677

Fredrikson, M., Lantz, E., Jha, S., Lin, S., Page, D., & Ristenpart, T. (2014). Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing. *2014 IEEE Symposium on Security and Privacy*, 81–95. https://doi.org/10.1109/SP.2014.16

Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). Synthetic data augmentation using GAN for improved liver lesion classification. *2018 IEEE*

15th International Symposium on Biomedical Imaging (ISBI 2018), 289–293. https://doi.org/10.1109/ISBI.2018.8363571

Gentry, C. (2009). A fully homomorphic encryption scheme. *Stanford University*. https://crypto.stanford.edu/craig/

Geyer, R. C., Klein, T., & Nabi, M. (2017). Differentially private federated learning: A client level perspective. *arXiv preprint* arXiv:1712.07557. https://arxiv.org/abs/1712.07557

Gostin, L. O., & Hodge, J. G. (2002). Personal privacy and common goods: A framework for balancing under the national health information infrastructure. *Maryland Law Review*, 62, 215–260.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 1–42. https://doi.org/10.1145/3236009

Gymrek, M., McGuire, A. L., Golan, D., Halperin, E., & Erlich, Y. (2013). Identifying personal genomes by surname inference. *Science*, 339(6117), 321–324. https://doi.org/10.1126/science.1229566

Huang, Z., Wang, T., Liu, Y., & Song, D. (2020). Benchmarking and analyzing privacy preserving machine learning: A case study of membership inference attack. *Proceedings of the 28th ACM International Conference on Multimedia*, 1748–1757. https://doi.org/10.1145/3394171.3413658

Jagielski, M., Oprea, A., Biggio, B., Liu, C., Nita-Rotaru, C., & Li, B. (2019). Differentially private learning needs better features (or much more data). *Advances in Neural Information Processing Systems*, 32, 9368–9378. https://arxiv.org/abs/1904.02744

Jayaraman, B., & Evans, D. (2019). Evaluating differentially private machine learning in practice. *arXiv preprint* arXiv:1902.08990. https://arxiv.org/abs/1902.08990

Johnson, A. E., Pollard, T. J., Shen, L., Lehman, L. W., Feng, M., Ghassemi, M., ... & Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3, 160035. https://doi.org/10.1038/sdata.2016.35

Jordon, J., Yoon, J., & van der Schaar, M. (2019). PATE-GAN: Generating synthetic data with differential privacy guarantees. *International Conference on Learning Representations*. https://openreview.net/forum?id=HJGs0sC5Ym

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2), 1–210. https://doi.org/10.1561/2200000073

Kim, H., & Wang, S. (2019). Privacy-preserving clinical trial analytics using secure multi-party computation. *Journal of Biomedical Informatics*, 98, 103275. https://doi.org/10.1016/j.jbi.2019.103275

Kim, M., Song, Y., & Cheon, J. H. (2018). Homomorphic encryption for arithmetic of approximate numbers. *ASIACRYPT 2017*, 409–437. https://doi.org/10.1007/978-3-319-70697-9_15

Lauter, K., Naehrig, M., & Vaikuntanathan, V. (2014). Can homomorphic encryption be practical? *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*, 113–124. https://doi.org/10.1145/2517470.2517477

Le, T., Kairouz, P., Nissim, K., & Murtagh, F. (2020). On the utility of differential privacy for health data. *Journal of Privacy and Confidentiality*, 10(1). https://doi.org/10.29012/jpc.v10i1.1280

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

Li, X., Huang, W., Yang, W., Wang, S., & Zhang, Z. (2019). Multi-user data privacy protection based on federated learning. *IEEE Access*, 7, 160180–160191. https://doi.org/10.1109/ACCESS.2019.2954423

Malin, B., Karp, D., & Scheuermann, R. H. (2011). Technical and policy approaches to balancing patient privacy and data sharing in clinical and translational research. *Journal of Investigative Medicine*, 59(5), 759–765. https://doi.org/10.231/JIM.0b013e3182183e53

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. *Artificial Intelligence and Statistics*, 1273–1282. https://arxiv.org/abs/1602.05629

Melis, L., Song, C., De Cristofaro, E., & Shmatikov, V. (2019). Exploiting unintended feature leakage in collaborative learning. *2019 IEEE Symposium on Security and Privacy (SP)*, 691–706. https://doi.org/10.1109/SP.2019.00058

Mohassel, P., & Rindal, P. (2018). ABY3: A mixed protocol framework for machine learning. *2018 IEEE Symposium on Security and Privacy (SP)*, 35–52. https://doi.org/10.1109/SP.2018.00013

Mohassel, P., & Zhang, Y. (2017). SecureML: A system for scalable privacy-preserving machine learning. *2017 IEEE Symposium on Security and Privacy (SP)*, 19–38. https://doi.org/10.1109/SP.2017.13

Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. *2008 IEEE Symposium on Security and Privacy*, 111–125. https://doi.org/10.1109/SP.2008.33

Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, 57, 1701–1777.

Papernot, N., Abadi, M., Erlingsson, Ú., Goodfellow, I., & Talwar, K. (2018). Semi-supervised knowledge transfer for deep learning from private training data. *International Conference on Learning Representations*. https://openreview.net/pdf?id=rJzIBfZAb

Papernot, N., McDaniel, P., Goodfellow, I., Jha, S., Celik, Z. B., & Swami, A. (2016). Practical black-box attacks against machine learning. *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, 506–519. https://doi.org/10.1145/3052973.3053009

Powell, T., & Houghton, L. (2018). Patient perspectives on data sharing and privacy in biomedical research. *Journal of Medical Ethics*, 44(12), 809–813. https://doi.org/10.1136/medethics-2018-104872

Rashid, F., Al-Qurishi, M., Al-Salman, A., & Qureshi, K. N. (2021). Privacy-preserving frameworks for smart healthcare applications: A comprehensive review. *IEEE Access*, 9, 17114–17136. https://doi.org/10.1109/ACCESS.2021.3050982

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation). (2016). *Official Journal of the European Union*, L119, 1–88.

Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3, Article 119. https://doi.org/10.1038/s41746-020-00323-1

Salathé, M., Althaus, C. L., Neher, R., Stringhini, S., Hodcroft, E., Fellay, J., & Wilder-Smith, A. (2020). COVID-19 epidemic in Switzerland: On the importance of testing, contact tracing and isolation. *Swiss Medical Weekly*, 150, w20225. https://doi.org/10.4414/smw.2020.20225

Shabani, M., Borry, P., & Gordo, M. (2019). Challenges of genomic data sharing: Ethical, legal and social considerations. *Human Genomics*, 13(1), 1–12. https://doi.org/10.1186/s40246-019-0198-2

Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., ... & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10(1), 12598. https://doi.org/10.1038/s41598-020-69250-1

Sheller, M. J., Reina, G. A., Edwards, B., Martin, J., & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10(1), 12598. https://doi.org/10.1038/s41598-020-69250-1

Shi, Y., Cao, Y., Zhang, Q., Li, Y., & Xu, L. D. (2020). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646. https://doi.org/10.1109/JIOT.2016.2579198

Shin, D. (2020). Demystifying "hackers": The social context of cybersecurity vulnerability. *Computers in Human Behavior*, 103, 233–243. https://doi.org/10.1016/j.chb.2019.09.010

Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (pp. 1310–1321). https://doi.org/10.1145/2810103.2813687

Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. *2017 IEEE Symposium on Security and Privacy (SP)*, 3–18. https://doi.org/10.1109/SP.2017.41

Sun, L., Yang, X., Wang, L., & Meng, W. (2021). Privacy-preserving medical data sharing and analytics in IoMT systems. *IEEE Internet of Things Journal*, 8(3), 1684–1695. https://doi.org/10.1109/JIOT.2020.3027589

Tjoa, E., & Guan, C. (2020). A survey on explainable artificial intelligence (XAI): Toward medical XAI. *IEEE Transactions on Neural Networks and Learning Systems*, 32(11), 4793–4813. https://doi.org/10.1109/TNNLS.2020.3027314

Toth, R., & Marada, M. (2019). GDPR and the role of privacy by design. *Journal of Cyber Security Technology*, 3(4), 195–205. https://doi.org/10.1080/23742917.2019.1679983

Tramèr, F., & Boneh, D. (2020). Slalom: Fast, verifiable and private execution of neural networks in trusted hardware. *Advances in Neural Information Processing Systems*, 33, 16560–16572. https://arxiv.org/abs/1806.03287

Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019). A hybrid approach to privacy-preserving federated learning. *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*, 1–11. https://doi.org/10.1145/3356401.3380292

Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS Medicine*, 15(11), e1002689. https://doi.org/10.1371/journal.pmed.1002689

Vayena, E., Salathé, M., Madoff, L. C., & Brownstein, J. S. (2018). Ethical challenges of big data in public health. *PLoS Computational Biology*, 14(2), e1006280. https://doi.org/10.1371/journal.pcbi.1006280

Voigt, P., & Von dem Bussche, A. (2017). *The EU General Data Protection Regulation (GDPR)*: A Practical Guide, 1st Ed. Springer. https://doi.org/10.1007/978-3-319-57959-7

Wagh, S., Setty, S., & Jha, S. (2020). Secure multi-party analytics over vertically partitioned data. *Proceedings on Privacy Enhancing Technologies*, 2020(1), 202–222. https://doi.org/10.2478/popets-2020-0009

Wang, Y. X., Balle, B., & Kasiviswanathan, S. P. (2019). Subsampled Rényi differential privacy and analytical moments accountant. *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, 1226–1235. https://arxiv.org/abs/1808.00087

Wang, Y., Yurochkin, M., Sun, S., Papailiopoulos, D., & Khazaeni, Y. (2019). Federated learning with matched averaging. *International Conference on Learning Representations*. https://arxiv.org/abs/2002.06440

Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2019). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1), 1–19. https://doi.org/10.1007/s41666-020-00080-8

Xu, L., Skoularidou, M., Cuesta-Infante, A., & Veeramachaneni, K. (2019). Modeling tabular data using conditional GAN. *Advances in Neural Information Processing Systems*, 32, 7335–7345. https://arxiv.org/abs/1907.00503

Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), Article 12. https://doi.org/10.1145/3298981

Zhou, J., Wu, J., Ding, Z., & Lin, Y. (2019). Understanding user trust in healthcare artificial intelligence applications. *International Journal of Medical Informatics*, 129, 185–193. https://doi.org/10.1016/j.ijmedinf.2019.06.011