RegenMed

WHITE PAPER

WEAKNESSES OF AI AND OTHER SYNTHETIC DATA IN HEALTHCARE

September 2025

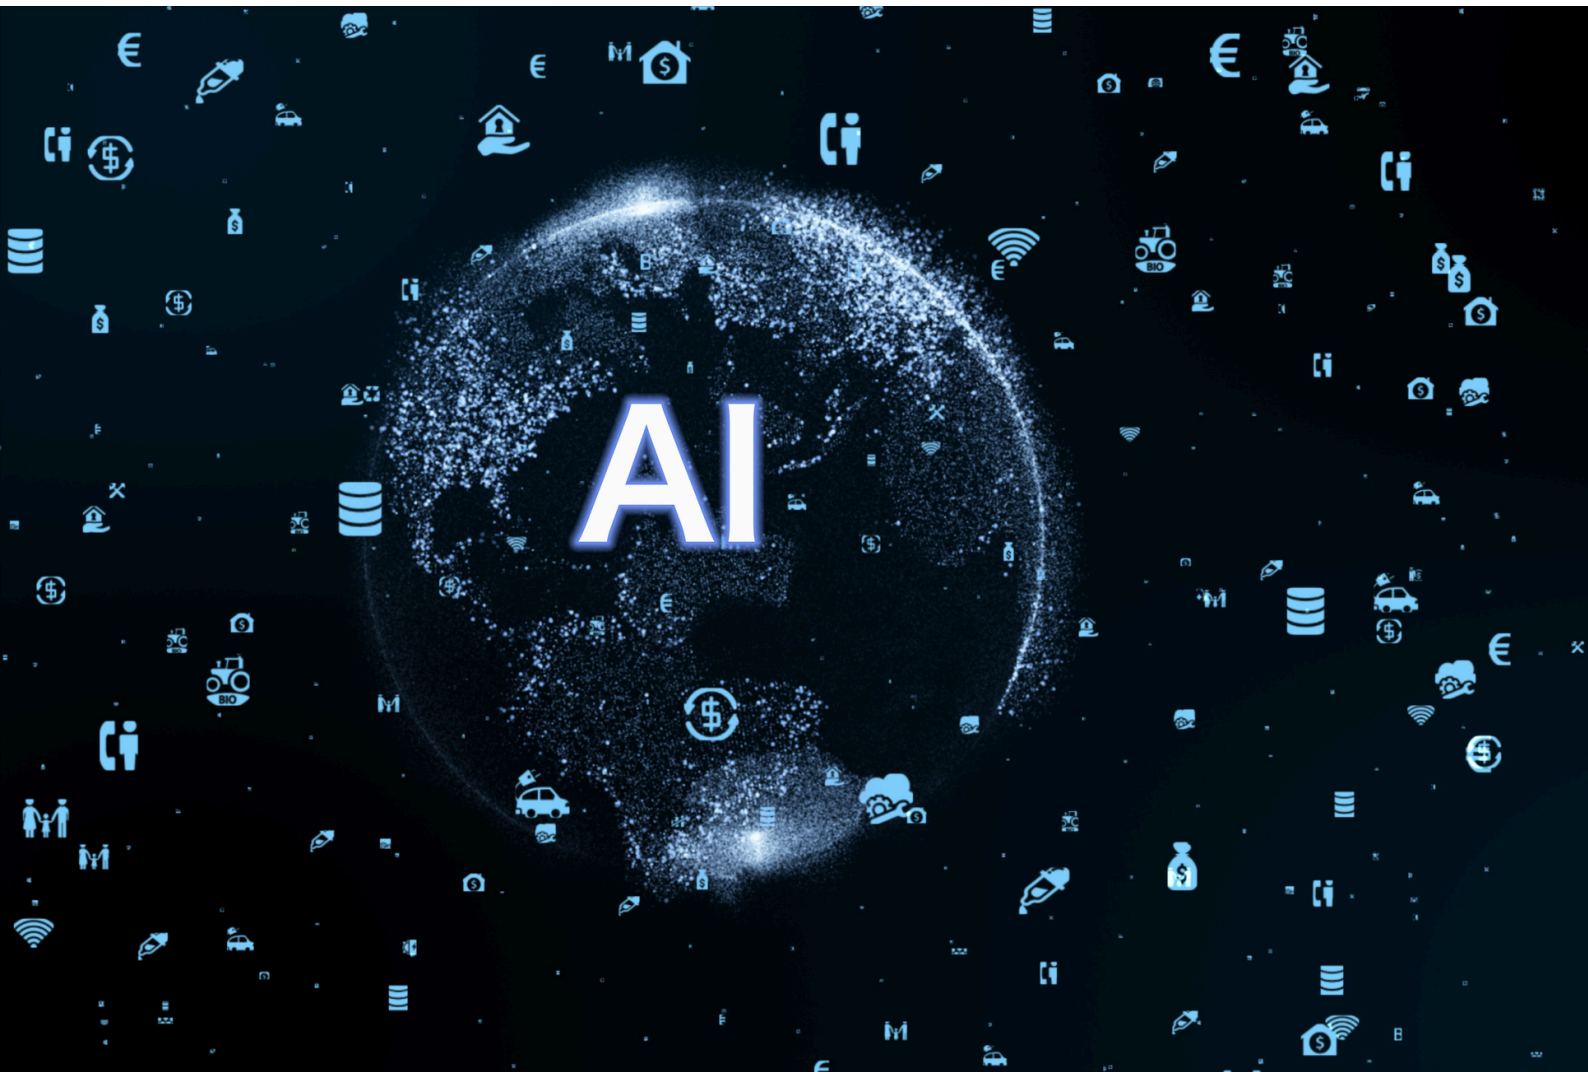

RegenMed

TABLE OF CONTENTS

EXECUTIVE SUMMARY

The rapid rise of artificial intelligence (AI) has positioned AI-generated and synthetic data as a transformative solution for challenges such as data scarcity and privacy compliance. However, an exhaustive analysis of the healthcare landscape reveals that this optimism is often misplaced. While synthetic data can accelerate certain non-critical development processes, its fundamental lack of provenance, realism, and traceability renders it a high-risk tool for any application requiring regulatory or financial defensibility.

This paper establishes that reliance on synthetic data for product validation, regulatory submissions, or reimbursement claims exposes product manufacturers, researchers, payers, and providers to severe vulnerabilities. These risks include the propagation of algorithmic bias, the failure to model critical patient safety scenarios, non-compliance with stringent regulatory requirements, and significant financial penalties. The central finding is that synthetic data, by its very nature, lacks a verifiable link to a real-world patient, making it an indefensible abstraction in the face of payer and regulatory audits.

The analysis concludes that a sustainable and responsible path forward does not lie in a full embrace of synthetic data. Instead, it advocates for a hybrid data strategy that prioritizes the collection and use of auditable Real-World Evidence (RWE). By strategically integrating fully verifiable RWE with carefully applied AI tools, the healthcare ecosystem can build a more trustworthy, transparent, and effective foundation for innovation that withstands the scrutiny of regulators, secures financial integrity, and, most importantly, protects patient safety.

INTRODUCTION: THE DOUBLE-EDGED SWORD OF SYNTHETIC DATA

The healthcare and life sciences industries are navigating a complex intersection of technological innovation, stringent regulation, and deeply sensitive patient information. In this environment, AI-generated and synthetic data have emerged as a purported panacea, promising to unlock new insights while overcoming longstanding barriers to data access. Synthetic data is an artificially created dataset that mimics the statistical properties of real data but does not contain actual verifiable patient information.

It is often generated using sophisticated generative AI models trained on real-world data, enabling the creation of vast datasets that can be used for training algorithms, simulating scenarios, and conducting analytics without the privacy constraints of handling personally identifiable information (PII).

The promise of synthetic data is undeniable. It offers a solution to data scarcity, a persistent problem in niche industries or for rare-event scenarios like fraud detection and medical research, where real data is limited, costly, or difficult to obtain. By providing a scalable and privacy-preserving resource, synthetic data can theoretically accelerate AI development and research and development (R&D) efforts.

However, a closer examination reveals that this promise is a double-edged sword. The perceived benefits of synthetic data are fundamentally outweighed by profound weaknesses that compromise its utility in high-stakes healthcare contexts.

This report provides an exhaustive, multi-stakeholder analysis of these critical vulnerabilities, which span technical, legal, financial, and ethical dimensions. It aims to demonstrate that for any application requiring scientific rigor, regulatory defensibility, or financial accountability, the use of purely synthetic data is not a viable strategy.

THE FOUNDATIONAL PITFALLS OF SYNTHETIC DATA

Fidelity and Generalizability: The Inaccuracy of Abstract Models

The core utility of synthetic data rests on its ability to accurately reflect the real-world data from which it was derived. Yet, this is also its most significant technical weakness.

THE FOUNDATIONAL PITFALLS OF SYNTHETIC DATA

The quality of synthetic data is entirely dependent on the quality of its source, adhering to a "garbage in, garbage out" paradigm. If the original data is biased, incomplete, or inaccurate, the synthetic data will inherit and potentially amplify these flaws. A generative model may replicate common patterns and correlations but struggle to capture the complex nuances, subtle temporal dependencies, and, most importantly, rare events and outliers that are present in authentic datasets.

This inability to model edge cases and outliers is not merely a technical limitation; it poses a direct and demonstrable threat to patient safety. A generative model, by its design, prioritizes mimicking the most common statistical patterns to produce a large, realistic-looking dataset.

This inherent focus on commonality means it often fails to accurately represent rare or critical events, such as a patient's unique physiological response to a drug or a subtle, but vital, pattern of disease progression. An AI-enabled medical device or diagnostic tool trained exclusively on this flawed synthetic data will lack the fundamental knowledge to recognize or respond to these critical, but rare, real-world scenarios.

A Virginia Tech study provided empirical validation for this concern, finding that machine learning models trained on patient data alone failed to detect 66% of critical in-hospital deteriorations. This is the process through which a seemingly abstract technical flaw is transformed into a direct patient safety hazard, leading to potential misdiagnosis, delayed treatment, and catastrophic patient outcomes.

Bias and Algorithmic Discrimination: The Reinforcement of Inequity

The use of synthetic data also introduces significant ethical and legal vulnerabilities related to bias. Generative models are often trained on existing datasets that may contain historical or representational biases, which can be propagated and even amplified in the synthetic data they create. This can result in AI systems that produce skewed outputs and discriminatory decisions. For example, a hiring algorithm trained on historical employment data where people of color are underrepresented in high-level positions may perpetuate this inequality. Similarly, a medical device trained on a dataset lacking diversity may be less accurate for certain patient populations, potentially worsening health disparities.

The problem of bias in synthetic data is not a static one; it can create a self-reinforcing feedback loop that entrenches and amplifies systemic inequities over time. The cycle begins with an original real-world dataset that contains historical or representational bias. A generative AI model then replicates and may amplify this bias when creating a synthetic dataset. An AI-enabled tool is subsequently trained on this biased synthetic data, leading to a biased output.

This biased output can then be used to inform real-world decisions — such as patient diagnoses or treatment plans — and the results of those decisions are often collected as new data. If this new data is then used to retrain or update the model, the original bias is reinforced, creating a continuous cycle of skewed and discriminatory outcomes. This dynamic erodes trust in AI systems and has profound ethical and legal consequences, as evidenced by legal frameworks like the EU AI Act, which can impose substantial fines for non-compliance.

THE REGULATORY AND COMPLIANCE CONUNDRUM

Payer and Regulatory Audits: The Non-Negotiable Demand for Provenance

In the high-stakes world of healthcare finance and regulation, the lack of provenance is the most significant and insurmountable weakness of synthetic data. Healthcare regulators and payers, including the Centers for Medicare & Medicaid Services (CMS) and private insurers, are intensifying their audit scrutiny. CMS has announced a sweeping expansion of its Risk Adjustment Data Validation (RADV) audits to include all eligible Medicare Advantage plans annually, with sample sizes increased from 35 to up to 200 records per contract.

With billions of dollars in potential clawbacks at stake, this creates a high-stakes environment where compliance is paramount. Private payers are closely mirroring these practices, demanding patient-linked documentation for reimbursement negotiations, value- based care settlements, and utilization management appeals.

The fundamental weakness of synthetic data is not a technical flaw but an existential one: it lacks a verifiable patient of record, making it legally and financially indefensible in any context requiring an audit trail.

An audit requires a clear chain of custody, from a patient encounter to documentation to claim submission. Synthetic data is an abstraction; it has no specific patient, no clinician identifier, no timestamp, and no verifiable EHR origin. Even if it statistically resembles a real dataset, it cannot be tied to a specific beneficiary chart. When an auditor demands a "source of truth" to support a diagnosis or claim, synthetic data cannot provide it. This lack of provenance and traceability means that synthetic data is automatically rejected as valid evidence in payer and regulatory audits, leading directly to financial clawbacks, penalties, and reputational harm.

THE REGULATORY AND COMPLIANCE CONUNDRUM

FDA's Cautious Stance: Navigating the Total Product Life Cycle

The U.S. Food and Drug Administration (FDA) is actively engaged in developing a nuanced regulatory framework for AI, signaling a cautious approach that views AI as a tool to be managed, not a black box to be blindly trusted. The agency has issued comprehensive draft guidance for developers of AI-enabled medical devices and AI use in drug development, emphasizing the need for transparency, bias mitigation, and robust validation throughout a product's life cycle.

The FDA explicitly requires a high degree of accountability from developers. When synthetic data is used in a regulatory submission, the agency asks for a "comprehensive explanation of how the data were generated and why they are fit-for-purpose". This is a high bar for accountability that is often challenging for complex generative models.

This cautious stance is complemented by the FDA's increasing receptivity to Real-World Evidence (RWE), which is derived from authentic patient data sources such as electronic health records (EHRs) and claims databases. The FDA encourages sponsors to provide an assessment of RWE to evaluate its relevance and reliability, as it can provide a more comprehensive understanding of how a new drug or therapy performs in the "real world" beyond the controlled environment of a traditional randomized controlled trial (RCT). High quality (verifiable, fit-for-purpose) RWE is valued for its ability to provide insights into how treatments perform in specific patient subgroups, with co-morbidities, and over a patient's lifetime. This is a key distinction: regulators view validatable RWE as the preferred tool for a holistic view of a product's safety and efficacy, placing synthetic data in a secondary, supplementary role for specific, non-critical applications.

Table One summarizes key FDA guidance documents on AI and data usage, illustrating the agency's focus on transparency and a risk-based approach.

WEAKNESSES BY STAKEHOLDER

For Product Manufacturers & Developers: From Lab to Market Failure

For product manufacturers and developers, the allure of synthetic data is its potential to accelerate R&D and clinical trial simulation. However, this can create a false sense of security. While synthetic data can be used to test models internally, its inability to accurately capture the complexity and unpredictability of real-world conditions means a model trained exclusively on such data may fail unexpectedly when released into the market.

This can lead to costly late-stage failures[26] and, more critically, to product recalls.

A significant, documented risk for manufacturers is the high likelihood of recalls for AI- enabled devices that have not undergone rigorous clinical validation. The FDA's 510(k) pathway, through which many AI devices are cleared, does not require prospective human testing or clinical trials.

This creates a business incentive for companies, particularly large, publicly traded ones, to bypass time-consuming clinical validation in favor of faster, more cost-effective methods, such as training on synthetic data. A study published in the *Journal of the American Medical Association (JAMA)* found a direct correlation: devices with no clinical validation were more likely to be recalled. The study also found that publicly traded companies, which had a lower rate of clinical validation, accounted for over 90% of recall events.

This reveals a critical failure in the product development pipeline, where the public is, in essence, becoming the "test" subject for products not rigorously validated with real-world human data. Nearly half of the AI device recalls in the study occurred within the first year of approval, suggesting that reliance on insufficient data can lead to dangerous performance gaps in real-world applications.

THE EVOLVING LEGAL AND LITIGATION LANDSCAPE

The Rise of State-Level Regulation: A Patchwork of Laws

The lack of a single, comprehensive federal framework governing the use of AI in healthcare has created a fragmented, state-by-state regulatory landscape that poses significant compliance challenges for national organizations. In this void, states are taking action to legislate these tools in response to concerns about misrepresentation, accuracy, and bias. Each state is approaching the issue differently, creating a complex and costly compliance challenge.

This patchwork means that a product or process that is legal in one state may be illegal in another. For example, a Nevada law bans the deployment of AI systems that provide mental or behavioral health care services, while a pending bill in Illinois would prohibit AI from making independent therapeutic decisions. This forces companies to either restrict their offerings to specific states or invest heavily in a complex, state-specific compliance strategy to navigate these varying requirements.

Table Two provides a summary of key state-level AI regulations.

THE EVOLVING LEGAL AND LITIGATION LANDSCAPE

Beyond Regulation: Liability and Accountability

The risks of synthetic data extend beyond regulatory non-compliance to the realm of litigation and broader legal liability. If a misdiagnosis or adverse patient event is directly linked to an AI tool trained on inaccurate or biased data, the manufacturer, the healthcare provider, and even the creator of the data could be held liable for negligence or malpractice.

Furthermore, international and ethical frameworks, such as the EU's General Data Protection Regulation (GDPR) and the EU AI Act, are influencing the conversation around synthetic data. These frameworks emphasize principles like fairness, transparency, and accountability, which are directly challenged by the inherent opacities of synthetic data.

The concept of "group harms" is particularly relevant, where the aggregate statistics from a synthetic dataset could reveal sensitive information about groups to which an individual belongs, leading to unfair decisions — such as an insurer erroneously raising rates based on a flawed risk profile. This introduces a new dimension of liability that goes beyond individual privacy to collective injury.

STRATEGIC SOLUTIONS AND RECOMMENDATIONS

The analysis presented in this paper indicates that in the high-stakes environment of healthcare, the purported benefits of AI-generated and synthetic data are eclipsed by its profound weaknesses. The responsible path forward does not lie in a purely synthetic data strategy but in a sophisticated, hybrid approach that combines the best of auditable, real- world data with carefully applied AI tools.

Prioritizing Real-World Evidence (RWE): The Foundation of a Robust Data Strategy

The strategic alternative to synthetic data is a focus on high-quality Real-World Evidence. Such RWE is derived from real-world data (RWD) verifiable all the way to primary sources and which is routinely collected during clinical practice. Unlike synthetic data, RWE is grounded in authentic patient consent, contains verifiable provenance, and is fully auditable by regulators and payers.

STRATEGIC SOLUTIONS AND RECOMMENDATIONS

The unique value of such RWE lies in its ability to provide a comprehensive, patient-centric understanding of a drug's use, safety, and efficacy in the complex, uncontrolled environment of routine clinical care. This makes it the superior foundation for regulatory submissions and financial defensibility. A strategic shift from data generation to data curation is recommended, focusing on creating auditable, patient-consented RWE from validated sources.

Leveraging Hybrid Data Models: A Path to Comprehensive Insight

The most effective data strategy overcomes the limitations of any single data source by leveraging a hybrid approach. Real-world data often exists in fragmented silos. Claims data provides a comprehensive longitudinal view of patient encounters but lacks clinical depth. EHR data is often fragmented and difficult to aggregate at scale, and again difficult if not impossible properly to verify.

A hybrid data ecosystem links these disparate datasets at the patient level, providing a more complete and comprehensive view of a patient's journey.[31] This approach combines the "what" (claims data), the "why" (EHR notes) and clinically-specific, fully verifiable and longitudinal Circles Datasets. It maintains the crucial elements of provenance and auditability that synthetic data lacks. This strategy streamlines analytics, provides a richer foundation for decision-making, and can lead to significant cost and time savings.

Adopting Hybrid Clinical Trials: Blending Rigor with Reality

Similarly, the future of clinical research is moving toward a hybrid model that combines the scientific rigor of traditional in-person Randomized Controlled Trials (RCTs) with the patient-centric benefits of Decentralized Clinical Trials (DCTs). This model allows researchers to collect both traditional, controlled RCT data and continuous RWE from remote sources, ensuring a more comprehensive understanding of a drug's performance in both controlled and real-world settings. This approach can enhance patient recruitment, improve data quality, and accelerate development timelines while gaining increasing regulatory acceptance.

The allure of synthetic data's speed and cost-savings is often a mirage, collapsing under the weight of regulatory, financial, and scientific scrutiny. In the high-stakes world of healthcare, the fundamental weaknesses of synthetic data — its lack of provenance, its inability to model critical nuances and outliers, and its potential to propagate dangerous biases — render it an unacceptable tool for any application that impacts patient care or requires regulatory and financial defensibility.

For **product manufacturers**, the risk of product recall and market failure linked to insufficient validation is a documented reality. The strategic priority must be to conduct rigorous real-world clinical validation for all patient-facing AI tools. Synthetic data should be used for non-critical internal model testing, but not for final product validation or regulatory submissions.

For **researchers**, while synthetic data can accelerate the research pipeline, it is an unreliable tool for the study of complex, nuanced, or rare conditions. Final tools must be validated and fine-tuned on real-world data to ensure scientific rigor and generalizability.

For **payers and providers**, the use of unverified, synthetic data exposes organizations to significant financial clawbacks and legal liability. A rigorous data governance framework must be implemented to ensure that all data used for reimbursement or clinical decision-making is verifiable and auditable, and third-party vendors must be held to this same standard.

For all stakeholders, engaging with regulators and policymakers is crucial to help shape a more consistent and responsible regulatory framework. In the high-stakes world of healthcare, the future is not purely synthetic, but a sophisticated hybrid — a blend of robust, auditable RWE and carefully applied AI tools to create a more trustworthy and effective healthcare ecosystem.

TABLE ONE

Guidance Document	Publication Date	Key Recommendations and Focus
Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations (Draft Guidance) (16)	Jan 2025	Provides comprehensive recommendations for AI-enabled devices across the Total Product Life Cycle (TPLC). Recommends strategies to address transparency and bias, with suggestions for product design, development, and documentation. Encourages sponsors to provide detailed information about device uses, inputs, outputs, architecture, and validation datasets in marketing submissions.
Considerations for the Use of Artificial Intelligence to Support Regulatory Decision Making for Drug and Biological Products (Draft Guidance) (18)	Jan 2025	Provides recommendations for the use of AI to produce data intended to support regulatory decision-making for drugs and biologics. Proposes a risk-based credibility assessment framework to evaluate an AI model's credibility for a specific use. It highlights challenges such as bias, variability in training datasets, and the opaque nature of AI models.
Use of Real-World Evidence to Support Regulatory Decision-Making for Medical Devices (Guidance) (21)	N/A	Acknowledges the use of Real- World Evidence (RWE) in regulatory submissions. Recommends that sponsors provide an assessment of "fit-for- purpose" data for RWE sources, evaluating both their relevance and reliability. For synthetic data, the guidance recommends a "comprehensive explanation of how the data were generated and why they are fit-for-purpose".(21)
Artificial Intelligence and Machine Learning Software as a Medical Device Action Plan (17)	Jan 2021	Outlines the FDA's approach to regulating AI/ML-based software as a medical device (SaMD). It signals a shift from the traditional regulatory paradigm to one that accommodates the adaptive nature of AI/ML technologies, while emphasizing the need for ongoing safety and effectiveness reviews.

TABLE TWO

State	Law(s)	Effective Date(s)	Key Provisions		
Utah	HB 452, SB 226	May 7, 2025	**Mental Health Chatbots:** Requires clear and conspicuous disclosure that a chatbot is an AI and not a human at the beginning of an interaction and upon user prompting. Prohibits the sale or sharing of user data. (29)	**General Chatbots:** Prohibits chatbots that mislead users into believing they are human.(29)	
Nevada	AB 406	Jul 1, 2025	**Mental and Behavioral Health Care:** Bans AI systems from providing services that would, if provided by a human, constitute mental or behavioral health care.29 Prohibits providers from using AI for direct care, though administrative use is permitted.(29)		
Texas	HB 149, HB 1445, SB 1974	Sept. 2025-Jan. 2026	**Incitement/D iscrimination:** Prohibits the development or deployment of AI systems that would "incite or encourage" physical self- harm or criminal activity, or that discriminate.(29)	**Provider Disclosure:** Requires providers using AI for health services to disclose this fact to the patient at the time of service, with an emergency exception.29	**Payor Use:** Prohibits automated systems from making adverse determinations on medical necessity or prior authorization, requiring a physician or peer review.29

TABLE TWO

State	Law(s)	Effective Date(s)	Key Provisions
Illinois	HB 1806 (Pending)	N/A	A pending bill would prohibit AI systems from making independent therapeutic decisions, directly interacting with clients in therapeutic communication , or generating treatment plans without a licensed professional's review and approval. This could substantially impair the use of AI in mental health services.(29)

REFERENCES

1. Synthetic Data in AI: Solving Data Scarcity and Privacy Challenges | by Hybrid Minds, accessed September 7, 2025, https://medium.com/@hybrid.minds/synthetic-data-in-ai- solving-data-scarcity-and-privacy-challenges-824d8a380c10
2. Synthetic data generation: a privacy-preserving approach to accelerate rare disease research - PMC, accessed September 7, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC11958975/
3. Synthetic Data: A Real Fix for Clinical Trials?, accessed September 7, 2025, https://www.clinicaltrialvanguard.com/article/synthetic-data-a-real-fix-for-clinical-trials/
4. Synthetic Data vs Real Data: Benefits, Challenges - Research AIMultiple, accessed September 7, 2025, https://research.aimultiple.com/synthetic-data-vs-real-data/
5. The benefits and limitations of generating synthetic data - Syntheticus, accessed September 7, 2025, https://syntheticus.ai/blog/the-benefits-and-limitations-of-generating-synthetic- data
6. What can I not do with synthetic data? - BlueGen AI, accessed September 7, 2025, https://bluegen.ai/what-can-i-not-do-with-synthetic-data/
7. Synthetic Data - what, why and how? - Royal Society, accessed September 7, 2025, https://royalsociety.org/-/media/policy/projects/privacy-enhancing technologies/Synthetic_Data_Survey-24.pdf
8. Machine learning models fail to detect key health deteriorations, Virginia Tech research shows, accessed September 7, 2025, https://news.vt.edu/articles/2025/02/virginia-tech- study-published-in-communications-medicine-.html
9. Bias Mitigation via Synthetic Data Generation: A Review - MDPI, accessed September 7, 2025, https://www.mdpi.com/2079-9292/13/19/3909
10. What is Data Bias? | IBM, accessed September 7, 2025, https://www.ibm.com/think/topics/data-bias
11. Synthetic Health Data: Real Ethical Promise and Peril - PMC, accessed September 7, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC11555762/
12. The Black Side Of Synthetic Data For Medical Sector - Labelvisor, accessed September 7, 2025, https://www.labelvisor.com/the-black-side-of-synthetic-data-for-medical-sector/
13. Study links AI medical device recalls to gaps in FDA regulatory process, accessed September 7, 2025, https://www.medicaldesignandoutsourcing.com/study-links-ai-medical- device-recalls-to-gaps-in-fda-regulatory-process/
14. Synthetic data in medicine: Legal and ethical considerations for patient profiling - PMC, accessed September 7, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC12166703/
15. FDA Issues Comprehensive Draft Guidance for Developers of Artificial Intelligence- Enabled Medical Devices, accessed September 7, 2025, https://www.fda.gov/news- events/press-announcements/fda-issues-comprehensive-draft-guidance-developers- artificial-intelligence-enabled-medical-devices
16. Artificial Intelligence in Software as a Medical Device - FDA, accessed September 7, 2025, https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence- software-medical-device

REFERENCES

17. Artificial Intelligence for Drug Development - FDA, accessed September 7, 2025, https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial- intelligence-drug-development
18. FDA Issues Draft Guidance Documents on Artificial Intelligence for Medical Devices, Drugs, and Biological Products | Sterne Kessler, accessed September 7, 2025, https://www.sternekessler.com/news-insights/client-alerts/fda-issues-draft-guidance- documents-on-artificial-intelligence-for-medical-devices-drugs-and-biological-products/
19. Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations | FDA, accessed September 7, 2025, https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial- intelligence-enabled-device-software-functions-lifecycle-management-and-marketing
20. Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations - FDA, accessed September 7, 2025, https://www.fda.gov/media/184856/download
21. Real-World Evidence in Healthcare: A Complete Guide | Veradigm, accessed September 7, 2025, https://veradigm.com/real-world-evidence/
22. Real-World Evidence and Big Data Group - ISOP, accessed September 7, 2025, https://isoponline.org/special-interest-groups/real-world-evidence-and-big-data-group/
23. Addressing the Limitations of Medical Data in AI - FDA, accessed September 7, 2025, https://www.fda.gov/medical-devices/medical-device-regulatory-science-research- programs-conducted-osel/addressing-limitations-medical-data-ai
24. Top 20+ Synthetic Data Use Cases - Research AIMultiple, accessed September 7, 2025, https://research.aimultiple.com/synthetic-data-use-cases/
25. Linrodostat | Advanced Drug Monograph - MedPath, accessed September 7, 2025, https://trial.medpath.com/drug/report/cc4c008e96387b02
26. AI devices with no clinical validation tied to more recalls, study finds - Healthcare Dive, accessed September 7, 2025, https://www.healthcaredive.com/news/ai-medical-devices-no- validation-more-recalls-jama/758998/
27. Why unsupported software is a risk to healthcare organizations - Paubox, accessed September 7, 2025, https://www.paubox.com/blog/why-unsupported-software-is-a-risk- to-healthcare-organizations
28. Health AI Policy Tracker - Manatt, Phelps & Phillips, LLP, accessed September 7, 2025, https://www.manatt.com/insights/newsletters/health-highlights/manatt-health-health-ai- policy-tracker
29. Auditing and Generating Synthetic Data with Controllable Trust Trade-offs - arXiv, accessed September 7, 2025, https://arxiv.org/html/2304.10819v3
30. Driving Healthcare Innovation Through Hybrid Data - Forian Blog, accessed September 7, 2025, https://blog.forian.com/driving-healthcare-innovation-hybrid-data

REFERENCES

31. What is a Hybrid Clinical Trial? - RxE2 | Pharmacists Unbound, accessed September 7, 2025, https://rxe2.com/what-is-a-hybrid-clinical-trial/