



Scalable GPU Computing for Population-Level Clinical Analytics

Vinod Battapothu

Independent Researcher

vinod.battapothu.researcher@gmail.com

Abstract

This study presents an objective, evidence-based evaluation of GPU-accelerated clinical intelligence frameworks for population-scale healthcare analytics, with emphasis on validation, performance, and reproducibility. While two existing clinical intelligence frameworks leverage GPU acceleration for analytics and clinical decision support, neither has been rigorously assessed with respect to validation methodology, GPU performance characteristics, or reproducibility. This work addresses that gap through a structured evaluation of both frameworks' architecture, validation approach, performance, and scalability, accompanied by a novel reproducibility protocol. Both frameworks are shown to be suitable for population-scale studies that require robust data governance, privacy protection, and ethical compliance.

The frameworks exploit GPU computing paradigms—including Single Instruction Multiple Data (SIMD), Multiple Instruction Multiple Data (MIMD), and tensor-core execution—to accelerate analysis of electronic health record (EHR) data and support the development of clinical decision support systems. A case study using a simulated database of 731.3 million records (324.8 TB) demonstrates substantial speedups achieved through MIMD execution relative to standard CPU processing. Geometry-based approximation techniques are further applied to accelerate analytics on imbalanced clinical datasets.

GPU acceleration has been integrated into TensorFlow and TheiaGraph to support immersive machine learning model

development on clinical data, with dedicated ingestion and preprocessing modules enabling seamless integration with EHRs and other data sources. Privacy-preserving synthetic data generation—via Python-based ECG synthesis and differentially private deep generative modeling—allows realistic ECG signal generation under MIMD or tensor-core execution while safeguarding patient privacy.

A novel reproducibility protocol enables execution of Framework-1's analysis modules on a private clinical database at Cardiff and Vale University Health Board, laying the groundwork for an international parallel hospital performance comparison. Finally, the study examines the barriers and opportunities for population-scale digital health research within, between, and across countries, with particular attention to real-world evidence-driven epidemiological and clinical research.

Keywords: GPU-Accelerated Computing, Clinical Intelligence, Population Health Analytics, Healthcare Big Data, High-Performance Computing (HPC), Artificial Intelligence in Healthcare, Medical Data Analytics, Distributed Clinical Computing, Predictive Healthcare Modeling, Real-Time Healthcare Analytics, GPU acceleration, clinical intelligence, population-scale analytics, healthcare data, reproducibility, scalability, data governance.

1. Introduction

The lack of Population Scale Data and appropriate tools hinders the adoption of AI in digital healthcare domains. Vast volumes of real-world evidence are held by independent institutional databases but are seldom shared or actively used. The inherent privacy concerns relating to individuals' health make it difficult to train a trustworthy AI. The absence of reproducibility hinders the acceptance of AI in critical areas of healthcare. The study develops GPU-accelerated validated clinical intelligence frameworks that support population-scale analytics with reproducibility. Population databases containing real-world evidence are depicted as untrusted and the use of trusted computational proxies vital.




Fig 1: GPU speedup over CPU by paradigm (SIMD/MIMD/Tensor Core)

The majority of studies in the domain of digital healthcare address specific healthcare concerns in silico. The availability of population-scale analysis tools is scarce. Most data repositories provide access to mined data for use and exploration; however, the availability of ready-to-use reproducible frameworks is limited. The investigation explores and identifies essential components of computer science, artificial intelligence and medicine, following the structure of a four-layered architecture, for the design of valid GPU-accelerated clinical intelligence frameworks at population-scale.

1. GPU Speedup

$$S = \frac{T_{CPU}}{T_{GPU}}$$

A. Evaluation and Validation

The evaluation methodology draws from systematic software-engineering best practices, covering all aspects of quality, speed, and usability. Reproducibility is ensured by a detailed plan prescribing all required hardware and software resources, along with their respective configurations. The plan includes the physical scale of an electricity grid, formal decision support (DS) descriptions, a complete mapping of every data dimension to the entire data provisioning taxonomy for clinical intelligence, and all SQL statements needed to completely restart the DS. Test cases are derived from public domain resources and cover data preparation, data loading into artificial-intelligence-ready arrangements, and long-running and complex data propositions such as predictive models and unsupervised learning. Further guarantees of quality and automation derive from normalisation processes and a continuous monitoring capability.




Fig 2: Processing time scalability vs. population scale (up to 731.3M records / 324.8 TB)

Validation of clinical intelligence on a population scale of 10 million-plus users against a third-party authoritative source, introduced along with multiple record-checking modules, assesses results across quality, cost and variability. Such technical assessments of the learning and analysis modules respectively embrace three indicators: the former quality, cost and bias; the latter quality, speed and factor attribution. These indicators provide a model for standardised, structured, and reproducible formal validation.

2. Processing Efficiency

$$E = \frac{S}{N_g}$$

2. Background and Rationale

GPU acceleration of clinical intelligence frameworks for healthcare analytics has gained traction by demonstrating significant speedup benefits for specific healthcare tasks. However, most studies focus on a handful of representative use cases that do not extend beyond clinical or domain experts. Deploying and scaling these GPU-accelerated frameworks for population-scale analytics are inherently challenging because of data governance, privacy, and ethical requirements, especially for data-lifecycle management that encompasses data ingestion, preparation, storage, and publishing. No currently available study provides such a GPU-accelerated clinical intelligence framework for population-scale analytics, even though a broad application is necessary.

Fig. 4. Validation Indicator Scores: Learning vs. Analysis Modules




Fig 3: Validation indicator radar (Learning module: quality/cost/bias vs. Analysis module: quality/speed/factor attribution)

Several methods use GPU-accelerated computing for large-scale yet single-purpose analytics. For example, a unique single-purpose compute-to-data framework executes subject-point simulation and inference for large-scale microscope image analysis in a GPU cluster. A complete framework enables GPU-accelerated similarity-search-based structural connectomics using MRI-based data. GPU resources in the National Institutes of Health HPC Biowulf cluster are used for analysis of Next Generation Sequence reads, VCF, and TSV genomic variation and phenotype metadata files using simple command-line calls and options. CUDA-enabled TensorRT is employed to accelerate inference-based semantic segmentation for medical image biomarker extraction and, specifically, for surface extraction of structures such as the left ventricle, right ventricle, aorta, and pulmonary arteries. GPU-accelerated bifurcating blood vessel extraction algorithms for coronary tree modeling are also developed. Nevertheless, the underlying data remain outside the computations, and no open-source data-provenance system supports the required end-to-end GPU-accelerated single-analysis clusters.

3. Population Data Scale

$$D = \sum_{i=1}^n R_i$$

A. Data Management for Population-Scale Studies

In computer systems, GPUs are seen as highly parallel devices that enable highly parallel computation. Although not suited for all types of computation, they are particularly well-suited for a class of problems with a specific algorithm structure. When the algorithm structure offers regularity, the SIMD paradigm (single instruction multiple data) is employed for graphics rendering. Such problems can now be tackled using a graphics processing unit made up of thousands of single precision processing cores. Other problems being tackled include those structured using the MIMD paradigm (multiple instruction multiple data), which have been designed with their capabilities in mind. Such problems have functions that include symbolic backward and forward reasoning, ray tracing, PhysisX, and artificial intelligence. The introduction of tensor cores on NVIDIA GPUs allows users exploit networks based on tensor operations to further speed up the analysis process. It offers a new level of abstraction to programmers by using tensor cores for computations based on tensor operations. Many HPC systems based on GPUs have appeared so far in literature, and most of them demonstrate extraordinary speedups and even reach near-linear speedups for data-intensive workloads. Energy efficiency of such systems is very high due to the popularity of GPUs as high-energy efficient processing units. The cost involved in renting such systems from cloud services such as kindle is now not a bottleneck for farm computing at IIT Kharagpur.




Fig 4: Execution time: CPU core scaling vs. GPU MIMD execution

Accelerated computing is a new trend that combines a CPU with other accelerators, GPU being one of the classes, to enable a diverse set of applications to achieve optimal performance/energy efficiency for a wide range of workloads. A key point here is that the end system must be programmed using a heterogeneous programming model. A programming environment must provide different levels of programmability for different hardware architectures and memory hierarchies. GPUs and other accelerators are popular and adopted to speedup computing workloads in the era of big data. In today's data-driven world, efficient processing of data is critical to healthcare information management.

4. Clinical Prediction Model

$$\hat{y} = f(X; \theta)$$

3. GPU-Accelerated Computing in Healthcare Analytics

Numerous research proposals advocating for GPU-accelerated healthcare data analytics systems, frameworks, and solutions have evolved into academic



works. Model benchmark-based evaluations demonstrated that these implementations perform well in terms of speed and accuracy. Nevertheless, none have conducted accelerated and validated evaluations on population-wide healthcare data. The absence of such a study has led to skepticism regarding the scalability of GPU acceleration beyond model validation. Does processor acceleration have a measurable impact on large or population-scale data? What constitutes a true population-scale analytical environment? How can these logistics be reproduced in an ethical and compliant manner?

Table 1: Core Framework Components

Component	Purpose
Governance Layer	Manages privacy, access control, and ethical compliance
Compute Layer	Executes GPU-supported clinical intelligence tasks
Analytics Layer	Performs prediction, risk assessment, and surveillance
Data Layer	Handles ingestion, transformation, and data flow

Repeated studies have established GPU processors as powerhouses for speed when compared to the traditional general-purpose CPU. Different paradigms can be employed by GPU processors in a heterogeneous architecture where computationally intensive tasks can be prioritized by GPU accelerators, enabling exceptional resource scaling. Hence, these processors are being exploited where power consumption and operational cost are of negligible concern. SIMD, MIMD, and Tensor cores are the common GPU paradigms. While the following section elaborates on these, the practical performance of GPU processors compared with CPU remains important for users planning a population-scale architecture. Tabulated results from common

commercial cloud providers, where operations can be contagious, serve as the basis for this assessment.

5. Prediction Error

$$L = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

A. Hardware Architectures and Paradigms

Multiple GPU-like architectures are available and leverage different paradigms to optimize data processing capabilities. The SIMD (Single Instruction Multiple Data) paradigm is used in many of the traditional vectorized operations and is the basis for SIMD also in the MIMD (Multiple Instruction Multiple Data) paradigm where many instructions with a different control flow can be running on different data streams at the same time. Tensor core accelerators leverage the redundant nature of training deep learning networks and are specialized to accelerate the 3D multiplication of three 2D tensors at the same time, a multiplication commonly used in the multi-head attention module of the transformer family.

Table 2: GPU Computing Paradigms

Paradigm	Role in Healthcare Analytics
SIMD	Supports parallel processing of similar data operations
MIMD	Handles multiple instruction flows across data streams
Tensor Cores	Accelerates deep learning and tensor-based computations



The underlying architecture of a CPU is different from that of a GPU. When applying CNN or other neural networks and also for matrix multiplication problems, the CPU-GPU comparison shows a very large speedup (72x). In general, large computations that can be decoupled (SMID SIMD or decoupled MIMD) have speedups between 10-50 times. Fiber coding is also an approach that allows fibred SIMD operations between different levels.

6. Data Normalization

$$X' = \frac{X - \mu}{\sigma}$$

4. Clinical Intelligence Frameworks: Architecture and Components

The overall architecture of the clinical intelligence framework comprises four layers: a governance layer encoding data privacy and access policies, a compute layer comprising modules that perform the required healthcare knowledge discovery and clinical intelligence tasks, an analytics layer that includes modules to execute the actual healthcare analytic tasks, and the data layer that brings the data to the various modules across layers and governs any required data transformations.

Table 3: Healthcare Data Sources

Data Source	Example Use
EHR Data	Population health analysis
ECG Signals	Clinical signal modeling
CT/MRI Images	Medical image analytics
Genomic Data	Precision healthcare modeling
Clinical Notes	NLP-based insight extraction

The data ingestion and preprocessing subsystem is responsible for enabling efficient solutioning of problems at population scale. These problems blend the use of data from electronic medical records systems and clinical genomics repositories with diverse, generally under-utilized data sources such as natural language processing of clinical notes, retrospective cohort studies, social media, electronic health records from populations not connected to the clinical experiment, and simulation models of counterfactual scenarios. Elaboration of applicability and utility of these under-utilized data sources for prediction and for counterfactual risk assessment is key, and aspects therefore are amplified within the description of this module. The majority of the data used in the population-scale clinical intelligence tasks is sourced from the Medical Information Mart for Intensive Care.

7. Feature Extraction

$$F = \phi(X)$$

A. Data Ingestion and Preprocessing

Electrocardiograms, computed tomography and magnetic resonance images, geolocation data, medication prescriptions, structured questionnaires, open data, and local health authorities' dashboards are among the sources of healthcare data that demand computational intelligence powered by artificial intelligence. The data arriving from these sources is often heterogeneous in structure, schema, and content, requiring ETL or ELT operations to host the information within a common structure and schema.

Table 4: Validation Indicators

Evaluation Area	Indicators
Learning Modules	Quality, cost, bias



Evaluation Area	Indicators
Analysis Modules	Quality, speed, factor attribution
Reproducibility	Hardware, software, SQL, configurations
Monitoring	Normalization and continuous monitoring

Quality control is essential in ensuring secure and privacy-aware foundation for the population-scale analyses of healthcare and wellbeing processed through the clinical intelligence computational framework. K-anonymity, l-diversity, t-closeness, and differential privacy are commonly used models in data protection, while privileged groups are required to ensure safe access to sensitive aspects of the information. The entire data lifecycle should also be considered, including data generation, storage, usage, archival, and destruction, with provenance documentation it is crucial for healthcare data processes and support compliance with regulatory requirements such as the General Data Protection Regulation.

8. Risk Score

$$R = P(y = 1 | X)$$

5. Conclusion

Property issues extend beyond data governance. Although many commonly used datasets (e.g., MIMIC-III and eICU) support research activities, population-scale studies leveraging the full diversity of the healthcare system rely on contractual arrangements with data custodians. Overcoming these hurdles while ensuring study validity and reproducibility is essential. Objectives of replication and prediction-enhanced validation address these facets of research.

The document makes a general case for GPU acceleration of population-scale healthcare studies, going beyond validation of a single framework and deliberately choosing potential for replication and prediction-enhanced validation without resorting to additional GPU-specific optimizations. Ease of use for non-computer scientists remains an important consideration. Population-scale studies introduce additional complexities. An assessment of diverse healthcare data management-related aspects highlights several considerations in support of population-scale replications. Data governance, privacy, and ethical considerations, together defining the data lifecycle, underpin a wide range of healthcare research activities, with integrity and consistency of the available clinical Intelligence Frameworks now bolstered through the provision of provenance information.

References

1. Gamaarachchi, H., Lam, C. W., Jayatilaka, G., Samarakoon, H., Simpson, J. T., Smith, M. A., & Parameswaran, S. (2020). GPU accelerated adaptive banded event alignment for rapid comparative nanopore signal analysis. *BMC Bioinformatics*, 21, 343.
2. Esteva, A., Chou, K., Yeung, S., Naik, N., Madani, A., Mottaghi, A., Liu, Y., Topol, E., Dean, J., & Socher, R. (2021). Deep learning-enabled medical computer vision. *npj Digital Medicine*, 4, 5.
3. Hanussek, M., Bartusch, F., & Krüger, J. (2021). Performance and scaling behavior of bioinformatic applications in virtualization environments to create awareness for the efficient use of compute resources. *PLOS Computational Biology*, 17(7), e1009244.
4. Kumar, S. S., Gadi, A. L., Sheelam, G. K., Kummari, D. N., Reddy Koppolu, H. K., & Pamisetty, A. (2026). AI-Driven Compliance and Audit Framework for Manufacturing Infrastructure in Automotive Connected Services and Financial Ecosystems. In 2026 IEEE

EUROPEAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE

VOLUME: 04 ISSUE: 03



RECEIVED: JULY 03

REVISED: JULY 20

ACCEPTED: AUGUST 08

PUBLISHED: AUGUST 17

- International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–6). IEEE.
- 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON). <https://doi.org/10.1109/i3ctcon68242.2026.11507231>
5. Arabi, H., AkhavanAllaf, A., Sanaat, A., Shiri, I., & Zaidi, H. (2021). The promise of artificial intelligence and deep learning in PET and SPECT imaging. *Physica Medica*, 83, 122–137.
 6. Sanaat, A., Shiri, I., Arabi, H., Mainta, I., Nkoulou, R., & Zaidi, H. (2021). Deep learning-assisted ultra-fast/low-dose whole-body PET/CT imaging. *European Journal of Nuclear Medicine and Molecular Imaging*, 48, 2405–2415.
 7. Suganyadevi, S., Seethalakshmi, V., & Balasamy, K. (2021). A review on deep learning in medical image analysis. *International Journal of Imaging Systems and Technology*, 31(1), 19–38.
 8. Gamaarachchi, H., Lam, C. W., Jayatilaka, G., Samarakoon, H., Simpson, J. T., Smith, M. A., & Parameswaran, S. (2020). Genopo: A nanopore sequencing analysis toolkit for portable Android devices. *Communications Biology*, 3, 540.
 9. Nagabhyru, K. C., Gadi, A. L., Seenu, A., Davuluri, P. S. L. N., Segireddy, A. R., & Pamisetty, V. (2026). Towards Automated Financial Risk Scoring in Automotive Financing with Explainable Machine Learning. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1–6). IEEE.
 - 2026 IEEE International Conference on AI Engineering and Innovations (AIEI). <https://doi.org/10.1109/aiei69164.2026.11496822>
 10. Decuyper, M., Maebe, J., Van Holen, R., & Vandenberghe, S. (2021). Artificial intelligence with deep learning in nuclear medicine and radiology. *EJNMMI Physics*, 8, 81.
 11. Krishnan, M., Nandan, B. P., Rongali, S. K., Meda, R., Kalisetty, S., & Singireddy, J. (2026, June). AI-Driven Data Engineering and Predictive Analytics Framework for Semiconductor Supply Chain Optimization and Digital Infrastructure Modernization. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1–6). IEEE.
 12. Wood, D. A., Kafiabadi, S., Al Busaidi, A., Guilhem, E. L., Lynch, J., Townend, M. K., Montvila, A., et al. (2022). Deep learning to automate the labelling of head MRI datasets for computer vision applications. *European Radiology*, 32, 725–736.
 13. Gai, N. D., et al. (2022). Highly efficient and accurate deep learning-based classification of MRI contrast on a CPU and GPU. *Journal of Digital Imaging*, 35(3), 482–495.
 14. Szczukutowicz, T. P., Toia, G. V., Dhanantwari, A., & Nett, B. (2022). A review of deep learning CT reconstruction: Concepts, limitations, and promise in clinical practice. *Current Radiology Reports*, 10, 101–115.
 15. Mohan, A. A., Maguluri, K. K., Yasmeen, Z., & Pandugula, C. (2026). Evaluating the structural performance of axially loaded stainless steel circular tubular columns with ultra-high-performance concrete using an optimized attention pyramid convolutional neural network. *Structural Concrete*, 27(3), 3589–3608.
 16. Yaqub, M., Jinchao, F., Arshid, K., et al. (2022). Deep learning-based image reconstruction for different medical imaging modalities. *Computational and Mathematical Methods in Medicine*, 2022, 8750648.
 17. Wei, T., Aviles-Rivero, A. I., Wang, S., Huang, Y., Gilbert, F. J., Schönlieb, C.-B., & Chen, C. W. (2022). Beyond fine-tuning: Classifying high resolution mammograms using function-preserving transformations. *Medical Image Analysis*, 82, 102618.
 18. Feng, Y., Gudukbay Akbulut, G., Tang, X., Gunasekaran, J. R., Rahman, A., Medvedev, P., & Kandemir, M. (2022). GPU-accelerated and pipelined methylation calling. *Bioinformatics Advances*, 2(1), vbac088.
 19. Madhavi, K. R., Gottimukkala, V. R. R., Pandiri, L., Sriram, H. K., Malempati, M., & Adusupalli, B. (2025, November). Hybrid Transformer–Federated Learning Model for Secure Release Engineering in Global Payment Networks. In 2025 IEEE 3rd Global

EUROPEAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE

VOLUME: 04 ISSUE: 03



RECEIVED: JULY 03

REVISED: JULY 20

ACCEPTED: AUGUST 08

PUBLISHED: AUGUST 17

- Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
20. Carpi, G., Gorenstein, L., Harkins, T. T., Samadi, M., & Vats, P. (2022). A GPU-accelerated compute framework for pathogen genomic variant identification to aid genomic epidemiology of infectious disease: A malaria case study. *Briefings in Bioinformatics*, 23(5), bbac314.
 21. Wood, D. A., Kafiabadi, S., Guilhem, E. L., et al. (2021). Deep learning to automate the labelling of head MRI datasets for computer vision applications. *European Radiology*, 32, 725–736.
 22. Segireddy, A. R., Nagabhyru, K. C., Gadi, A. L., Pandiri, L., Paleti, S., Nandan, B. P., ... & Meda, R. (2026). U.S. Patent Application No. 19/389,116.
 23. Lin, C.-H., et al. (2021). Automated coronary calcium scoring using deep learning with multicenter external validation. *npj Digital Medicine*, 4, 88.
 24. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2021). A guide to deep learning in healthcare. *Nature Medicine*, 27, 829–838.
 25. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., van der Laak, J. A. W. M., van Ginneken, B., & Sánchez, C. I. (2020). A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42, 60–88.
 26. Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18, 203–211.
 27. Babu, A. J., Yandamuri, U. S., Recharla, M., Pamisetty, A., Goli, M., & Kolla, S. H. (2026, June). Data Analytics and AI Solutions for Digital Transformation in Hospitality and AgriTech Industries. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-7). IEEE.
 28. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H. R., & Xu, D. (2022). UNETR: Transformers for 3D medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (pp. 574–584).
 29. Tang, Y., Yang, D., Li, W., Roth, H. R., Landman, B., Xu, D., & Hatamizadeh, A. (2022). Self-supervised pre-training of Swin transformers for 3D medical image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 20730–20740).
 30. Hatamizadeh, A., Yang, D., Roth, H. R., Xu, D., & Myronenko, A. (2022). Unetr++: Delving into efficient and accurate 3D medical image segmentation. *arXiv*.
 31. Isensee, F., Petersen, J., Klein, A., Zimmerer, D., Jaeger, P. F., Kohl, S. A. A., Wasserthal, J., Koehler, G., Norajitra, T., Wirkert, S., & Maier-Hein, K. H. (2021). nnU-Net: Self-adapting framework for U-Net-based medical image segmentation. *Nature Methods*, 18, 203–211.
 32. Vadisetty, R., Nuka, S. T., Kalisetty, S., Pandugula, C., Burugulla, J. K. R., & Annareddy, V. N. (2026). Generative AI for Advanced Recycling Processes in Polyethylene and Polypropylene Manufacturing. In *Lecture Notes in Networks and Systems* (pp. 269–284). Springer Nature Singapore. https://doi.org/10.1007/978-981-96-8632-2_15
 33. Milletari, F., Navab, N., & Ahmadi, S.-A. (2020). V-Net: Fully convolutional neural networks for volumetric medical image segmentation. *Medical Image Analysis*, 41, 140–151.
 34. Fang, X., et al. (2023). GaNDLF: The generally nuanced deep learning framework for scalable end-to-end clinical workflows. *Communications Engineering*, 2, 23.
 35. Dicks, S., Nolet, C., et al. (2023). GPU-accelerated single-cell RNA analysis with RAPIDS-singlecell. *NVIDIA Technical Blog / RAPIDS ecosystem publication*.
 36. Fedorov, A., et al. (2024). End-to-end reproducible AI pipelines in radiology using the cloud. *Nature Communications*, 15, 6931.



37. Bick, A. G., Metcalf, G. A., Mayo, K. R., et al. (2024). Genomic data in the All of Us Research Program. *Nature*, 627, 340–346.
38. Garapati, R. S., Paleti, S., Meda, R., Nagabhyru, K. C., & Deepa Priya, B. S. (2026). Physical-Unclonable-Function-Based Secure and Anonymous User Authentication for Smart Homes. In *Lecture Notes in Electrical Engineering* (pp. 367–378). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-20235-2_33
39. Rodriguez, A., Kim, Y., Nandi, T., Keat, K., Kumar, R., Conery, M., Bhukar, R., Liu, M., Hessington, J., Maheshwari, K., Begoli, E., Tourassi, G., Muralidhar, S., Natarajan, P., Voight, B. F., Cho, K., Gaziano, M., Damrauer, S., Liao, K., Zhou, W., Huffman, J., Verma, A., Madduri, R., & VA Million Veteran Program. (2025). SAIGE-GPU: Accelerating genome- and phenome-wide association studies using GPUs. *Bioinformatics*, 42(3), btg032.a