A Comprehensive Review of Digital Twin Models for Stroke Prediction: Towards Software-Driven and Agent-Based Approaches

Sheetal Shinde[1] and . Akanksha Pal [2]

[1] PhD Research Scholar, Oriental University Indore, MP, India
[2] Assistant Professor, Oriental University Indore, MP, India
sheetalshankarshinde@gmail.com

Abstract. Stroke has remained among the top contributors to mortality and chronic disability over the past few years. The prediction of the likelihood of patients having strokes has remained a challenge despite the development of risk scores, statistical models, and machine learning algorithms, due to the limitations associated with the dynamic interactions between various genetic, physiological, behavior-related, and environmental factors. In addition, most healthcare-related digital twins employ a vast array of IoT sensors that not only increase healthcare costs but also lack the necessary scalability. This review considers a software-based digital twin solution that combines machine learning and agent-based modeling for the prediction of stroke risk within a dynamic, personalized approach. The solution is more effective since it does not depend on hardware-based monitoring systems, but instead on data gathered from various sources including clinical data, patient behavior, and environmental factors. Through the ABM component, interactions are simulated among the various risk factors, which are then processed by the machine learning component into dynamic risk scores that change based on the patient situation. Through a systematic review of various models used in predicting the risk of stroke, including risk scores and more advanced deep learning and hybrid models, the various advantages and disadvantages are highlighted. The proposed software-based digital twin solution is more scalable and inclusive, and can work within various healthcare systems without the need for expensive IoT infrastructure.

Keywords: Digital Twin; Stroke Prediction; Agent-Based Modeling; Machine Learning in Healthcare; Software-Driven Models; Personalized Medicine; Preventive Healthcare..

1. Introduction

Stroke continues to be one of the most threatening health issues faced by the world-wide population, ranked as the second-leading cause of death and a significant source of permanent disability worldwide. According to the World Health Organization, this condition affects the larger population in lower and middle-class countries predominantly [1]. The root cause lies in the obstruction or hemorrhage of blood supply reaching the brain, causing major health, economic, and social consequences [2]. Unpredictability related to factors such as high blood pressure, diabetes, and genetic predisposition has made stroke both common and difficult to forecast [3].

Traditional models for predicting stroke rely on statistical methods or machine learning alone with a small set of clinical features. Although these models give reasonably

M. Kaushik et al. (eds.), *Proceedings of the International Conference on Responsible, Risk-aware, and Regulated AI (RRRAI 2026)*, Advances in Intelligent Systems Research 210,
https://doi.org/10.2991/978-94-6239-723-1_30

good predictions, they are not adaptable to individual patient situations. Recent developments in digital twin technology provide a promising approach for modeling patient health behaviors in a virtual setting.

Currently, most digital twinning solutions for healthcare depend on hardware or IoT components, where continuous sensor data streams require advanced infrastructural support. This creates a pressing need for software-driven digital twinning solutions. This study fills this knowledge gap by developing a hybrid digital twin framework integrating ML predictions and ABM behaviors for stroke risk assessment.

however, risk-scoring systems have a major drawback in being overly reliant on data from particular demographics, neglecting key individual-specific factors. Machine learning algorithms have increased prediction accuracy but tend to be computationally intensive and inapt as adaptive models due to their dependence on large quantities of quality data. Recently, the notion of a Digital Twin has appeared as a revolutionary strategy in medicine, seeking to build a virtual replica for each unique individual [11]. Such a replica could integrate data from various sources, simulate disease evolution, and predict future trajectories, enabling physicians to analyze risks on a dynamic basis and provide more personalized treatment options [12]. Exploratory analyses have indicated that digital twins could be applied in cardiology, oncology, and rehabilitation, profoundly influencing stroke treatment and prevention [13].

2. Bibliometric Analysis of Digital Twin and Stroke Prediction Research

A bibliometric study was performed on publications regarding Digital Twin applications in healthcare, with a focus on stroke prediction. Fig. 1 shows the annual trend of publications from 2013 to 2024. There is a low growth rate between 2014 and 2019, with fewer than 50 documents in some years. A sharp rise is noticed post-2020, with a remarkable year-by-year increase. After 2021, the number of publications exceeds 100, then doubles in 2022, reaching a peak of almost 600 publications in 2024, demonstrating a tremendous increase in research interest.

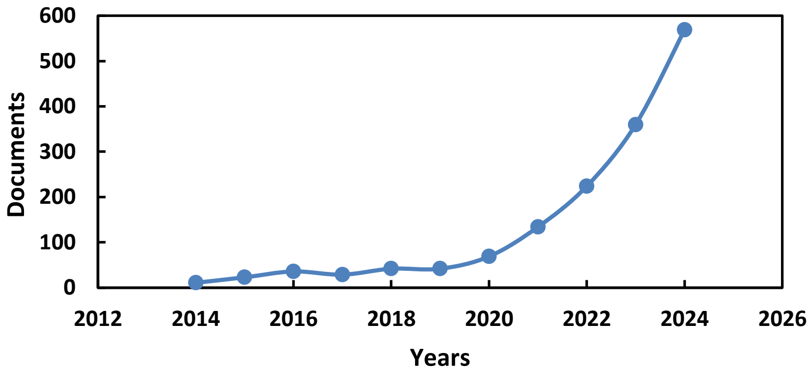

Fig. 1. Annual growth of publications related to Digital Twin and stroke prediction (2013–2024).

Fig. 2 illustrates the number of published documents according to author contributions. The top contributing authors include Theotokatos G., Lip G.Y.H., Ding S., Covache-Busuioc R.A., and Tadros M., reflecting the multidisciplinary nature of this research field spanning engineering, cardiology, and neurology

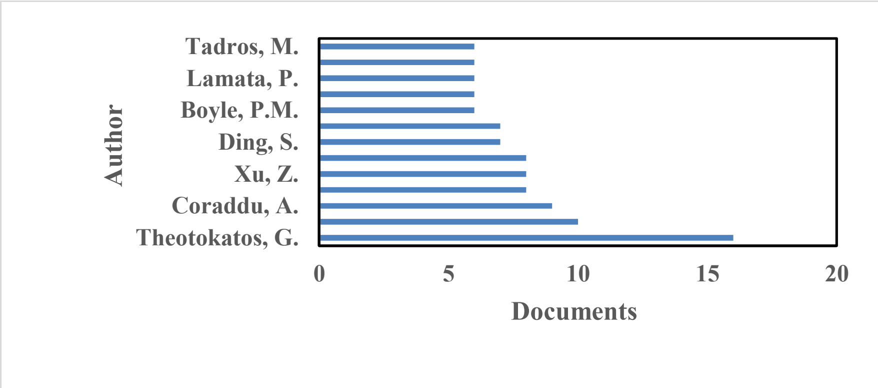

Fig. 3 depicts the number of published documents according to author affiliation from top worldwide institutions. Harvard Medical School leads with the highest number of publications (approximately 47 documents), followed by the University of Cambridge (~32), Massachusetts General Hospital, King's College London, and University College London (~28–30 each). The Ministry of Education of the People's Republic of China, the University of Edinburgh, and the University of Washington contribute an average of 25–27 publications, while Inserm, University of Strathclyde, and the University of Oxford report around 23–24 documents each.

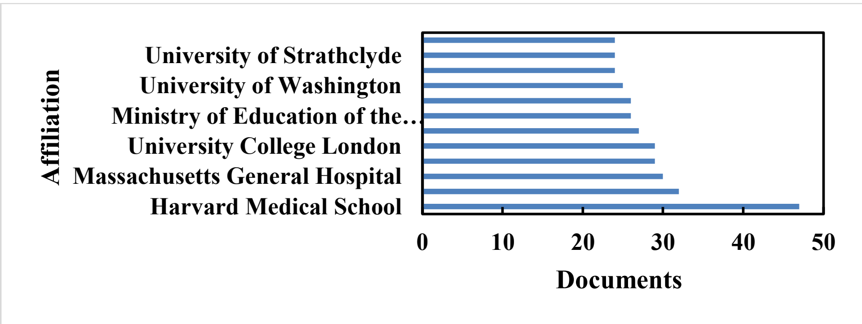

Fig. 3. Leading institutions contributing to research on Digital Twin and stroke prediction.

Fig. 4 illustrates the country-wise distribution of published documents. The United States leads with over 570 documents, followed by China (~450) and the United Kingdom (~300). India contributes approximately 220 documents. European nations such as Italy and Germany contribute 140–160 documents each, while Australia and Canada each report just over 120. The Netherlands, Spain, and France contribute around 80–100 documents each.

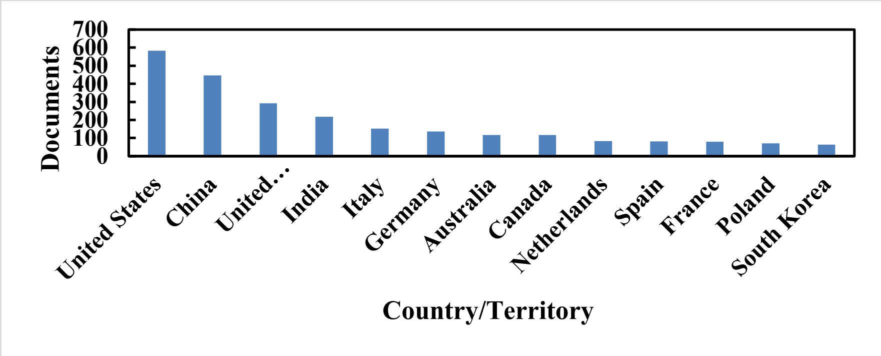

Fig. 4. Country-wise distribution of research publications on Digital Twin and stroke prediction.

Need for Software-Driven and Agent-Based Approaches

Healthcare systems need solutions that can scale effectively, affordably, and adaptively. A purely software-based implementation of a digital twin: (1) decreases reliance on the constant deployment of IoT hardware; (2) implies improved compatibility and integration with traditional EHRs; and (3) aids fast model update and policy simulation. Agent-based modeling increases the effectiveness of the framework by modeling patients as autonomous agents who change their behavior, lifestyle, and reactions to interventions over time—critical for stroke risk factors that depend on dynamic variables such as medication adherence, exercise, stress, and environment.

3. Digital Twin in Healthcare: Current Landscape

The concept of the Digital Twin originated in aerospace and manufacturing, where virtual replicas of physical assets were employed to optimize design and monitor performance for predicted failure [24]. This concept has evolved into dynamic virtual representations of patients, organs, or physiological systems that evolve with continuous real-world data inputs [25–28]. A Digital Twin in medicine is not a static model but a living simulation integrating biological, behavioral, and environmental factors. Its essential characteristics are personalization, adaptability, and predictive capability. By integrating heterogeneous data such as clinical records, imaging, wearable devices, and patient-reported outcomes, Digital Twins can mirror the real-time status of a patient and simulate potential health trajectories [30–35].

Applications of Digital Twins in healthcare are diverse and promising. In cardiology, patient-specific heart models have been developed to simulate electrophysiological behavior, optimize treatment strategies for arrhythmias, and predict outcomes of surgical interventions. In neurology, Digital Twins offer insights into disease progression, treatment response, and rehabilitation planning for neurodegenerative disorders like Parkinson's and Alzheimer's disease. In oncology, Digital Twins represent tumor growth and response to chemotherapy or radiotherapy, helping personalize treatment plans.

While Digital Twins are being increasingly applied in various medical domains, their application in stroke prediction remains limited. Stroke is a highly complex disorder caused by an interplay of genetic, physiologic, lifestyle, and environmental risk factors—making it a perfect candidate for the Digital Twin approach. However, most published research has focused on stand-alone prediction models rather than dynamic integrated simulations. There is an evident gap in using Digital Twins to integrate various risk factors, continuously update individual health profiles, and provide real-time personalized stroke risk assessments.

3.1 Hybrid Digital Twin Architecture

A digital twin framework should include two highly interconnected components: (1) Machine Learning Component — predicts basic stroke risk by employing supervised models based on clinical data; (2) Agent-Based Module — simulates individual patient behavior and environment interaction, dynamically adjusting risk levels.

3.2 Data Integration

An important innovation in this research consists in combining different data sources: clinical information (age, blood pressure, diabetes, cholesterol levels); behavioral variables (smoking behavior, exercise levels, patient compliance); and environmental and contextual data such as stress indicators and lifestyle trends. This multi-source data integration takes place within a software-based framework.

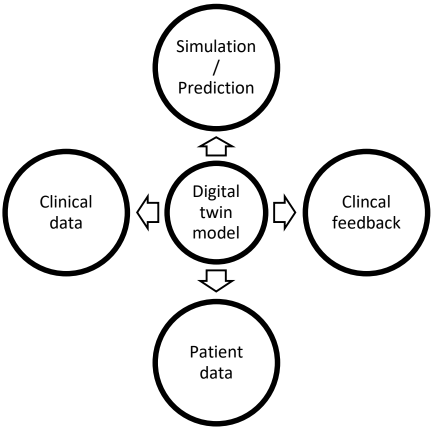

Fig. 5. Framework of a Digital Twin Model for Healthcare showing integration of patient data, clinical data, and wearables feeding into the digital twin model for simulation/prediction and clinical feedback.

4. Stroke Prediction Techniques: A Review of Existing Models

The challenge of stroke risk prediction has long been a concern in the medical research community owing to the complex nature of the phenomenon. A number of models have been developed over the years from traditional statistics to machine learning and deep learning algorithms, each with distinct advantages and constraints [63].

Innovative Components

The outstanding innovative features of the proposed model lie in: (1) Hybrid ML & ABM Integration — a new approach integrating ML for predictive capability and ABM for adaptability; (2) Software-Driven Digital Twin — a model that does not require hardware dependence unlike IoT-based DT models; (3) Adaptive Risk Score — risk scores for stroke are updated adaptively using simulated patient-agent scenarios; (4) Scenario Simulation — enables simulation of the effect of lifestyle modifications or treatment changes on a patient's stroke risk.

4.1 Traditional Statistical Models

The initial attempts to predict stroke were based on statistical models. The Cox proportional hazards regression model and multivariate regression techniques are popular in identifying and quantifying risk factors. Established risk scores such as FSRP and CHA_2DS_2-VASc utilize clinical variables like age, hypertension, diabetes, and smoking to estimate stroke risk. These models are relatively straightforward to interpret, cheap, and widely useful at the population level; however, their reliance on predefined variables limits their predictive power since they do not capture nonlinear interactions and dynamic interplay among diverse risk factors.

4.2 Machine Learning Approaches

To overcome the limitations of statistical approaches, machine learning has garnered growing attention in stroke risk prediction. Support Vector Machines (SVM), Random Forests (RF), and Artificial Neural Networks (ANN) have demonstrated capabilities in processing large amounts of data and identifying intricate variable patterns. SVM has been demonstrated in exploring high-dimensional datasets, RF in addressing nonlinearities, and ANNs in exploring underlying patterns within patient datasets. Machine learning has numerous advantages over conventional risk scores because of its accuracy in exploring electronic health records, medical imaging, and genetic factors . However, machine learning tends to be a "black box," resulting in a lack of understanding of the underlying prediction logic.

4.3 Deep Learning Hybrid Methods

Deep learning has been utilized for stroke prediction and analysis using imaging data like CT and MRI scans. CNNs have shown potential in predicting early symptoms of ischemia or hemorrhage. RNNs and LSTM networks have proven effective in modeling temporal data patterns from patients, such as continuous measurement of vital signs. Combinations of machine learning and deep learning algorithms have also been proposed, integrating clinical predictors with image and genetic data to design more holistic predictive models. Although highly promising, these algorithms remain computationally intensive and often less interpretable, raising concerns about clinical acceptance and trustworthiness .

4.4 Comparative Analysis and Baseline Models

Traditional methods remain important as reference models due to their interpretability and ease of application. Machine learning and deep learning models give better results related to accuracy and versatility but tend to lack interpretability and are highly data-dependent. In the evaluation of model effectiveness, universally accepted models such as the Framingham Stroke Risk Profile, CHA_2DS_2-VASc, and machine learning baselines including Random Forests, SVMs, and ANN classifiers serve as important references to enable researchers to identify whether improvements are possible with

Digital Twin concepts incorporating Agent-Based Modeling . Table 1 compares these models based on their tools, data sources, and applicability.

Table 1. Comparison of Stroke Prediction Model Types.

Model Type	Examples (Tools/Algorithms)	Data Used	Strengths	Weakness	Applicability
Traditional Statistical	Cox Regression, Framingham Stroke Risk Profile, CHA_2DS_2-VASc	Clinical variables (age, BP, diabetes, smoking, history)	Easy to interpret; low cost; validated in populations	Limited accuracy; ignores nonlinear interactions; poor adaptability	Baseline screening in clinics and population studies
Machine Learning (ML)	Support Vector Machine (SVM), Random Forests, ANN	Clinical + EHR data; lifestyle factors	Higher accuracy than traditional scores; captures complex patterns	Data-dependent; limited interpretability ("black box")	Hospital-level risk prediction; EHR-based screening
Deep Learning (DL)	CNN, RNN, LSTM	Imaging (CT, MRI), sequential health data	Excellent in image recognition; captures temporal patterns	Requires large datasets; computationally intensive	Early stroke detection from imaging; monitoring systems
Hybrid Approaches	ML + DL, multimodal models	Clinical + Imaging + Genetic + Behavioral	Comprehensive modeling; integrates heterogeneous data	Complex; high cost; interpretability challenges	Precision medicine; advanced research models

5. Stroke Agent-Based Modeling in Healthcare

Agent-Based Modeling (ABM) encompasses a computational methodology that simulates the actions and interactions of independent entities, here called agents, within a specified environment. Each agent is an individual unit—a patient, a cell, or a healthcare provider—possessing unique characteristics, behaviors, and decision-making rules. In contrast to classical mathematical models that rely on population aggregated data, ABM captures heterogeneity by modeling individual variability and local interactions. This bottom-up modeling paradigm enables researchers to observe how complex system-level patterns emerge from the collective behaviors of agents .

5.1 Applications of ABM in Simulating Physiological and Behavioral Processes

ABM has been used extensively in various fields within the healthcare industry, from the biological scale up to the population scale. Physiological ABM models include the immune system response to diseases and infections. At the population scale, ABM has demonstrated its effectiveness in modeling patient behavior, delivery of healthcare, and spreading of diseases among communities. In epidemiology, for example, agent-based modeling has played an integral role in grasping the outbreak dynamics of diseases like influenza and COVID-19, using agents with properties such as age, health condition, mobility, and social structures to track contact patterns and demonstrate how diseases propagate or are affected by vaccination or quarantine

5.2 Potential of ABM in Modeling Stroke-Related Risk Interactions

Stroke is a complex clinical entity with influences from genetic background, lifestyle factors, comorbidities, and environmental factors. ABM is uniquely suited to represent these dynamic interactions and heterogeneities. Agents can be assigned risk attributes based on familial history or genetic predispositions. Lifestyle variables such as diet, exercise, alcohol consumption, and smoking can be incorporated, and agents modify their behaviors based on changing environmental conditions, social influences, or healthcare interventions. Comorbidities including diabetes, obesity, and cardiovascular disease further define an agent's risk profile. Environmental factors such as air quality, socioeconomic status, and access to healthcare can be expressed as external conditions that affect agent behavior and health outcomes. By simulating these interactions, ABM can reproduce the nonlinear pathways leading to stroke onset, progression, and recurrence.

5.3 How ABM Complements Digital Twin Development

Digital Twins in healthcare aim to produce a personal and dynamic model for each individual. ABM enhances this by allowing the model to analyze dynamics between various factors on both the individual and system levels. Implementing ABM in the digital twin model makes it possible to simulate the trajectory of every individual's health while considering the variability within their genes, lifestyle, and environment. The adaptability offered by ABM ensures that the digital twin remains a living representation of the patient as new data become available through patient reports, medical files, or imaging studies. ABM also enables investigation of 'what-if' scenarios, allowing health practitioners to predict how changes in diet, medication, or lifestyle can affect a patient's likelihood of having a stroke .

ABM further promotes the scalability of Digital Twins. While regular Digital Twins might deal with specific patients, ABM makes it possible to create population-level digital models with unique properties, providing insights for both personalized treatment and population health policies. ABM-enhanced Digital Twins can allow government bodies to analyze the effect of community interventions on stroke incidence.

The majority of current stroke prediction models are based on traditional statistical techniques (Cox regression), standalone ML models (SVM, Random Forest, deep learning), or hardware-dependent digital twins using wearable sensors. The approach

Table 2. Comparison of Stroke-Prediction Studies and Author Contributions (parameters, methods, outcomes).

Author	Region	Data Sources	Key Parameters	Model/Method	Main Contribution	Key Limitations
Boehme et al. (2017)	USA	Longitudinal clinical	Age, SBP, TC/HDL, DM, smoking	Cox risk score	First population risk score for stroke	Limited to aggregate linear effects
Chun et al. (2021)	China	EHR + labs	BMI, SBP, LDL, HbA1c, smoking	Random Forest	ML exceeds classical scores on EHR	Single-center; class imbalance
Qiu et al. (2023)	EU	EHR + lifestyle	Diet, activity, alcohol, BP	SVM (RBF)	Added lifestyle improves risk	Limited external validation
Chen et al. (2024)	USA	EHR + sociodemographic	SES index, insurance, comorbids	Gradient Boosting	Social determinants quantified	Missingness handling ad-hoc
Jo et al. (2023)	Korea	MRI + CT + EHR	Lesion vol., NIHSS, BP	CNN + LR	Imaging + clinical fusion	Small imaging cohort
Mondal et al. (2023)	India	EHR	BP, DM, lipids, family hx	ANN (3 layers)	LMIC-specific calibration	Limited feature breadth
Chun et al. (2021)	UK	Biobank + genetics	PRS, BP, BMI, AF	XGBoost	Genetics improves discrimination	Potential ancestry bias
Lip et al. (2022)	EU	EHR + meds	Med adherence, BP trend	LSTM (temporal)	Captures temporal trends	Requires dense time series
Grout et al. (2024)	USA	EHR + environment	PM2.5, heat index, ZIP-SES	RF + SHAP	Explainability of drivers	Spatial confounding
Jiang et al. (2023)	EU	Imaging + labs	Plaque metrics, LDL, CRP	Hybrid CNN+GBM	Multimodal fusion pipeline	Compute intensive

Author	Region	Data Sources	Key Parameters	Model/Method	Main Contribution	Key Limitations
Heo et al. (2019)	USA	EHR	BP control trajectories	Bayesian survival	Survival-time risk estimates	Fewer behavioral vars
Bhute et al. (2025)	India	EHR + SES	BP, DM, tobacco, income	CatBoost	Robust to missing data	Limited imaging
Singhai et al. (2024)	India	EHR + lifestyle	Tobacco, diet, BMI, BP	TabNet	Interpretable tabular DL	Needs larger validation
Zimmerman et al. (2025)	USA	EHR + SDOH + env.	Housing, food access	GBM + SHAP	Equity-aware explanations	Data sparsity in SDOH

proposed herein combines ML and ABM to capture both prediction and behavior, supports dynamic patient-specific simulations, and offers greater scalability and flexibility due to its software-centric design. Table 2 compares stroke prediction studies across different regions, highlighting data sources, parameters, models, and contributions.

6. Conceptual Architecture for a Software-Only Digital Twin

The concept of a Digital Twin in healthcare has generally been associated with advanced hardware systems where data streams from IoT devices, wearable sensors, and clinical monitoring tools feed into virtual models. While effective in producing high-fidelity simulations, this architecture's reliance on costly physical devices restricts scalability and limits accessibility, especially in resource-constrained environments. To address these challenges, a software-only Digital Twin architecture is proposed, designed to function without dependency on IoT-based infrastructure

The framework consists of four primary modules: data integration, agent-based simulation, predictive analysis, and feedback. At the base level, the framework retrieves information from multiple digital sources including electronic health records, lab results, medical imaging studies, prescription information, and lifestyle reports. This information is standardized in a unified framework for interoperability. The simulation module models stroke risk factors as autonomous agents. The machine learning engine analyzes outcomes from these simulations in real-time, producing personalized stroke risk values. Feedback channels enable the interpretation of predictions and simulations to adjust strategies accordingly

6.1 Multi-source Data Integration

Risks for stroke are multifaceted and depend on clinical history, behavior patterns, and environmental factors. The key advantage of the proposed Digital Twin is its ability to harmonize diverse variables for a holistic patient profile

- Clinical data: demographics, blood pressure, cholesterol levels, glucose levels, comorbid conditions (hypertension and diabetes), and imaging records (CT and MRI scans).
- Behavioral data: lifestyle factors such as smoking, alcohol consumption, eating habits, physical activity, and medication compliance, measured using surveys or apps.
- Environmental data: external conditions including air quality, socioeconomic status, geographic location, and access to healthcare facilities.

6.2 Agent-Based Simulation of Stroke Risk Dynamics

The agent-based modeling layer is the most critical layer of the proposed architecture. Each risk factor—such as high blood pressure, smoking, and obesity—is represented by an autonomous agent with set behaviors and interaction rules. Agents interact with each other to simulate the accumulation, change, and degradation of risks over time [90]. This dynamic modeling enables simulation of what-if situations. For example, modifying agents related to smoking cessation and blood pressure management enables analysis of how stroke risk affects that particular individual.

6.3 Machine Learning Prediction Engine within the Twin

The predictive power of the Digital Twin rests on the machine learning engine embedded within the framework. This layer uses algorithms like Random Forests, Gradient Boosting, and Deep Neural Networks to analyze results produced from agent-based simulations. These models improve their predictions over time by learning from larger datasets [91–93]. The predictive engine offers personalized and dynamic risk scores that change in real-time, departing from current models that make static predictions. If a patient improves medication adherence and physical activity levels, the risk scores are recalculated to reflect reduced risks.

6.4 Advantages of Eliminating IoT Devices

The proposed software-only approach offers several key advantages:

- Scalability: By removing reliance on physical devices, the model can be implemented across diverse healthcare systems regardless of technological infrastructure.
- Cost Reduction: IoT-based systems involve huge financial investments in devices, maintenance, and transmission infrastructure. The software approach minimizes these costs, making advanced predictive tools more accessible.
- Accessibility: A model relying on routinely collected health data is inclusive, allowing consideration across socioeconomic groups where wearable technologies are impractical or unaffordable.
- Ease of Integration: Since most healthcare systems already maintain digital records, little disruption is involved in incorporating the software-only Digital Twin into existing EHR platforms.

- Sustainability: Reduced hardware dependence decreases the ecological footprint due to lower electronic waste and energy requirements for producing and maintaining IoT devices.

Through the promotion of cost-effectiveness, inclusiveness, and environmental sustainability, the proposed model is consistent with global healthcare goals including equity and efficiency

7. Evaluation and Benchmarking

The development process of the software-based Digital Twin stroke prediction system should be accompanied by a rigorous evaluation process to ensure acceptable accuracy and predictable performance across diverse healthcare contexts

7.1 Framework for Validating the Digital Twin Model

Validation of the Digital Twin in stroke prediction should be framed based on both internal and external validation procedures. Internal validation tests the model on data used during development phases (cross-validation, bootstrapping) to identify overfitting. External validation entails applying the model on data from other centers or geographic locations to determine generalization ability

The evaluation framework should incorporate three distinct phases:

7.1.1 Baseline Performance Assessment: Compare the Digital Twin against existing stroke risk models such as the Framingham Stroke Risk Profile, CHA_2DS_2-VASc score, or recent machine learning benchmarks to establish whether the proposed system offers tangible improvements in predictive performance.

7.1.2 Dynamic Validation: Since the Digital Twin adjusts to new data inputs as time progresses, it must be validated under conditions that reflect real-world updates.

7.1.3 Clinical Simulation Testing: Clinically simulate the model to verify whether it performs well when healthcare practitioners utilize it as an aid in decision-making processes.

7.2 Performance Metrics: Accuracy, Sensitivity, Specificity, AUC

To assess the performance of the Digital Twin, the following predictive modeling performance measures must be considered:

7.2.1 Accuracy: The proportion of correct predictions for the total number of instances. This gives a general idea about the model but lacks relevance in medical applications with imbalanced databases where non-stroke cases predominate.

7.2.2 Sensitivity (Recall): Indicates how well a model predicts people who suffer from stroke risks—a high priority in stroke prediction since missing high-risk individuals means lost prevention opportunities.

7.2.3 Specificity: Reflects the model’s ability to correctly identify patients who are not at risk of stroke (true negatives), important to avoid overburdening healthcare systems with unnecessary testing.

7.4 Scalability and Adaptability in Diverse Healthcare Settings

The software-based Digital Twin model is adaptable and scalable. Unlike IoT-based systems, the software-based system relies on existing digital health records and self-reported data, enabling implementation across all sizes of healthcare environments. The adaptability test should affirm the proposal's effectiveness across different population demographics (young vs. old, urban vs. rural). The model should also be tested for its potential to integrate with existing healthcare workflows and decision-support tools. Table 3 outlines key performance parameters.

Table 3. Key Evaluation Metrics for Stroke Prediction Models.

Metric	Definition	Clinical Significance
Accuracy	Proportion of correct predictions (true positives + true negatives) over all cases	Provides an overall measure of correctness, but may be misleading with imbalanced datasets
Sensitivity (Recall)	Proportion of actual stroke cases correctly identified (true positives / all positives)	Ensures high-risk patients are detected early, minimizing missed opportunities for prevention
Specificity	Proportion of non-stroke cases correctly identified (true negatives / all negatives)	Reduces false alarms, preventing unnecessary interventions and healthcare burden
Precision	Proportion of predicted high-risk cases that are truly at risk (true positives / predicted positives)	Important for minimizing false positives and building clinician trust
F1 Score	Harmonic mean of precision and recall	Balances sensitivity and precision, especially useful in imbalanced clinical datasets
AUC-ROC	Area under the ROC curve; measures ability to discriminate between stroke and non-stroke cases across thresholds	High AUC indicates reliable discrimination, crucial for clinical decision-making
Calibration	Agreement between predicted risk probabilities and observed outcomes	Ensures predictions are realistic and interpretable for guiding patient-level interventions

8. Research Gaps and Future Directions

The use of Digital Twin models and Agent-Based Modeling approaches in stroke prediction represents a promising shift. Despite the vast opportunities, limitations still persist in adoption and implementation. Addressing these limitations is crucial to ensure that Digital Twin technology is not only technically feasible but also ethically acceptable within the health sciences

8.1 Current Limitations in Digital Twin and ABM Adoption

Though Digital Twins have become popular in healthcare, their development is largely in the experimental phase for most use cases, including stroke prediction. The most significant challenge is posed by the heterogeneity and fragmentation of data. Healthcare data is normally scattered across different systems, formats, and organizations, hindering standardization and creating challenges for the interoperability of clinical, behavioral, and environmental data

Another constraint is the lack of adequate longitudinal big data, especially regarding the complicated interactions of risk factors. Stroke, impacted by the dynamic interactions of genetics, lifestyle, and comorbid conditions, demands longitudinal data that changes according to patient conditions. Current approaches to modeling are limited in that they are static representations and therefore inadequate for dynamic predictions. Finally, there lacks integration between Digital Twins and healthcare workflows, with many medical practitioners finding it challenging to interpret and apply AI-driven outcomes effectively without user-friendly interfaces and EHR integration.

8.2 Ethical and Data Privacy Considerations

Ethical challenges arising with Digital Twin adoption are equally significant. The technology requires highly sensitive patient data, ranging from genomics to behavioral and environmental indicators. Data privacy, data security, and patient consent are underlying considerations. Cyber-attacks, data breaches, and secondary use of health data are sources of serious concern. Transparency and interpretability are also key ethical aspects; clinicians and patients must be able to understand how predictions are generated. Black-box models, especially in deep learning, pose challenges in accountability when predictions influence critical medical decisions.

8.3 Integration with Cloud Computing, Federated Learning, and Blockchain

Further development of Digital Twins will need to build on advanced technologies to overcome current limitations. Cloud computing offers the scalable infrastructure required to support intensive ABM and ML computations, enabling continuous model updates accessible to healthcare providers worldwide without local high-performance computing resources.

Federated learning addresses the challenge of fragmented datasets while preserving privacy. Instead of centralizing patient data, federated learning allows models to be trained locally at different hospitals or institutions, with only model updates shared across networks. This would enable Digital Twins to learn from global populations, improving generalizability and reducing bias in stroke prediction.

Blockchain technology can further consolidate trust and security in Digital Twin ecosystems. With blockchain, decentralized, tamper-proof ledgers can assure data integrity, secure sharing, and patient-controlled access rights over data. In a Digital Twin context, blockchain could provide significant regulation on data use for transparency and accountability

8.4 Future Scope: Precision Medicine, Continuous Updates, Adaptive Learning

The future envisions Digital Twins in stroke treatment that provide precision medicine, where predictions and treatment actions will be personalized. Using the results of genomic, proteomic, and metabolomic analysis together with traditional clinical information, the vision foresees development of highly individualized risk profiles that not only predict stroke occurrence but also pinpoint specific preventive strategies.

Continuous updates represent another interesting area. Reinforcement learning and adaptive AI could be combined in digital twins to enhance predictions and interventions by learning from patient outcomes. If a proposed lifestyle change leads to decreased blood pressure and subsequently reduced stroke risk, this success could be adapted for a digital twin to predict similar outcomes for similar patients. If interventions do not work, a corresponding adjustment could be made by the digital twin. There is also potential in applying Digital Twins not only at the personal level but also in population-level health management, helping determine risk areas, allocate resources, and strategize prevention programs.

9. Conclusion

This review has examined the achievements of stroke treatment using advanced predictive tools, particularly Digital Twins and agent modeling. The literature indicates that despite substantial advances using statistical methods or artificial intelligence algorithms for stroke risk stratification, these methods often remain limited in fully grasping genetic, behavioral, and environmentally induced variables. While accuracy enhancements using deep learning or hybrid methods seem possible, they obviously lack general flexibility or individualization.

The software-based Digital Twin model introduced in this paper fills the mentioned gaps by overcoming reliance on IoT sensors and using a wide range of data—from patient data and life habits to environmental influences—through a single effective and economical solution. The combination of agent-based simulation adds a dynamic component for the interactions among risks. The novelty of this solution is that it can work purely with software components, making it more accessible, less costly to implement, and more broadly applicable, especially in developing or middle-income countries that cannot afford sophisticated devices.

Future prospects regarding the use of Digital Twins in stroke management suggest immense transformative power. Through effective proactive and personalized healthcare delivery, this approach has the potential to bring a paradigm shift from reactive to proactive healthcare. Its adaptability provides ample scope for application in precision medicine. Through the software-based and agent-based Digital Twin approach, there is immense scope to reduce the burden of strokes globally.

This paper proposes and illustrates the use of a newly developed software-based digital twin for predicting the risks associated with strokes by employing artificial intelligence combined with agent-based modeling. This hybrid model is designed to overcome the limitations of previous stroke prediction models and digital twins by being flexible and requiring minimal hardware support.

References

[1] Upadrista, D., Ghosh, R., Gupta, R.: Blockchain-enabled digital twin system for brain stroke prediction. Brain Informatics 12(1), 11 (2025). https://doi.org/10.1186/s40708-024-00203-7

[2] Allen, M., Mao, W., Das, A.: A digital twins machine learning model for forecasting disease progression in stroke patients. Applied Sciences 11(15), 7066 (2021). https://doi.org/10.3390/app11157066

[3] Hassan, R., Alzubaidi, L., Abdulkareem, M.: Predictive modelling and identification of key risk factors for stroke using machine learning. Scientific Reports 14, 12294 (2024).

[5] Gu, X., Tan, M., Kong, L., et al.: Identification of digital twins to guide interpretable AI for diagnosis and prognosis in heart failure. NPJ Digital Medicine 8, 11 (2025).

[6] Archbold, S., Ferguson, N.M., Hunter, E.: An agent-based model of the spread of behavioural risk factors for cardiovascular disease in city-scale populations. PLOS ONE 19(2), e0284564 (2024).

[7] Zhou, Y., Wang, L., Zhang, H., et al.: Machine learning-based prediction of 90-day prognosis and in-hospital mortality in hemorrhagic stroke patients. Scientific Reports 15, 3421 (2025).

[8] Gaikar, A., Gogate, U., Panchal, A.: Review on evaluation of stroke prediction using machine learning methods. International Journal for Research in Applied Science & Engineering Technology 11(4) (2023).

[9] Guo, Y.: A new paradigm of real-time stroke risk prediction and integrated care management in the digital health era. Thrombosis and Haemostasis 122(1), 5-7 (2022).

[10] Hussain, I., Hossain, M.A., Park, S.-J.: A healthcare digital twin for diagnosis of stroke. In: Proc. IEEE BECITHCON 2021, pp. 1-6 (2021).

[11] Dritsas, E., Trigka, M.: Stroke risk prediction with machine learning techniques. Sensors 22(13), 4670 (2022).

[12] Fernandes, J.N.D., Cardoso, V.E.M., Comesana-Campos, A., Pinheira, A.: Machine and deep learning in brain stroke diagnosis: A comprehensive review. Sensors 24(13), 4355 (2024).

[13] Mihai, S., Yaqoob, M., Hung, D.V., et al.: Digital twins: Enabling technologies, challenges, and future prospects. IEEE Communications Surveys & Tutorials 24(4), 123-156 (2022).

[14] Barricelli, B., Casiraghi, E., Fogli, D.: A survey on digital twin: Definitions, characteristics, and applications. IEEE Access 7, 167653-167671 (2019).

[15] Croatti, A., Gabellini, M., Montagna, S., Ricci, A.: Integration of agents and digital twins in healthcare. Journal of Medical Systems 44(1), 161 (2020).

[16] Suhail, S., Hussain, R., Jurdak, R., et al.: Blockchain-based digital twins: Trends and challenges. ACM Computing Surveys 54(11s), 240 (2022).

[17] Lenzi, G.: Problems with stroke genetics studies. Neurological Sciences 24(1), 1 (2003).

[18] Ahmadvand, R., Sharif, S., Banad, Y.: Neuromorphic digital twin for multi-agent system control. In: Proc. IEEE AIxSET 2024, pp. 55-62 (2024).

[19] Massel, L., Massel, A., Tsybikov, A.R.: Agent-service approach to building digital twins. In: Proc. IEEE SmartIndustryCon 2024, pp. 102-108 (2024).

[20] Lip, G.Y.H., Genaidy, A., Tran, G., et al.: Improving dynamic stroke risk prediction in atrial fibrillation. European Heart Journal Quality of Care & Clinical Outcomes 8(5), 548-558 (2022).

[21] Lolak, B., Gebru, K., Maldonado, R., et al.: Explainable machine learning for stroke risk prediction. JMIR Cardio 7, e51074 (2023).

[22] Park, D., Jeong, E., Kim, H., et al.: Three-month outcome prediction in acute ischemic stroke using machine learning. Diagnostics 11(10), 1909 (2021).

[23] Alaka, S.A., Menon, B.K., Brobbey, A., et al.: Functional outcome prediction in ischemic stroke. Frontiers in Neurology 11, 889 (2020).

[24] Jo, H., Kim, C., Gwon, D., et al.: Combining clinical and imaging data for outcome prediction after stroke. Scientific Reports 13, 16926 (2023).

[25] Qiu, Y., Liu, C., Zhou, X., et al.: Rapid stroke risk prediction models using machine learning. BMJ Open 13(3), e068045 (2023).

[26] Yang, Y., Chen, W., Liu, H., Zhang, Y.: Stroke prediction for hypertensive patients using EMRs. JMIR Medical Informatics 9(11), e30277 (2021).

[27] Wang, X., Zhu, S., Sun, H., et al.: Machine learning for infarct outcome prediction: A meta-analysis. Frontiers in Neurology 13, 913673 (2022).

[28] Bojsen, J.A., Udesen, J., Skaarup, S.H., Thomsen, C.: AI for MRI stroke detection: A systematic review. Insights into Imaging 15(1), 71 (2024).

[29] Heseltine-Carp, W., Helman, A., Truran, D., Lievesley, M.: Machine learning to predict stroke risk from hospital data. International Journal of Medical Informatics 190, 105028 (2025).

[30] Jiang, B., Tong, E., Zaharchuk, G., Wintermark, M.: Deep learning in acute ischemic stroke imaging. Radiology 306(2), e240775 (2025).

[31] Sung, S.-M., Ryoo, S., Kim, J.-M., et al.: Early neurological deterioration prediction using machine learning. Journal of the Neurological Sciences 418, 117144 (2020).

[32] Yuan, N., Ouyang, D., He, B., et al.: Deep learning of ECGs to predict atrial fibrillation. JAMA Cardiology 8(10), 1065-1074 (2023).

[33] Gramegna, M., Maffia, P., D'Ascenzo, F.: AI-assisted risk prediction in atrial fibrillation. Heart Failure Reviews 29(1), 221-238 (2024).

[34] Mutke, M.A., Madai, V.I., Hilbert, A., et al.: Outcome prediction after mechanical thrombectomy. Frontiers in Neurology 13, 737667 (2022).

[35] Heo, J.H., Lee, Y., Nam, H.S.: Machine learning-based outcome prediction in acute stroke. Stroke 50(5), 1346-1351 (2019).

[36] Wang, M., Li, H., Chen, X., et al.: Radiomics-based prognosis modeling in ischemic stroke. Computer Methods and Programs in Biomedicine 235, 107510 (2023).

[37] Klug, J., Schmidbauer, M., Zollei, L., et al.: Dynamic outcome prediction after ischemic stroke. Communications Medicine 4, 00666 (2024).