RESEARCH PAPER

Review of Federated digital twins powered by AI for ethical and customized Alzheimer's clinical trials.

Madhumitha N [*], Devipriya .D[1] , Anbuselvi .S[2], Anbu Kumaresan[3], Nafisa Farheen. H [4]
*Assistant Professor, Department of Pharmaceutical Chemistry, PERI College of Pharmacy, Mannivakkam, 600048.
[1] 4th year student. B.Pharmacy , PERI College of Pharmacy, Mannivakkam, 600048
[2]Department of Nursing, PERI College of Nursing, Mannivakkam, 600048
[3]Department of Chemistry, PERI College of Arts and Science, Mannivakkam, 600048.
[4] Department of Pharmacology, VISTAS- Pallavaram, Tamilnadu 600117
***Corresponding author: madhumitha481@gmail.com,**

Abstract

Alzheimer's disease (AD) affects more than 50 million people worldwide and carries a projected global economic burden approaching US$14.5 trillion by 2050, yet roughly 98% of AD drug candidates fail during clinical development. Large mandatory placebo arms, marked patient heterogeneity, slow recruitment, and multi-year follow-up periods are principal contributors to this attrition, while fragmented, institution-locked patient data and cross-border privacy regulation (GDPR, HIPAA) constrain the multi-site collaboration that AI-driven solutions require. A technically based digital twin engine that couples amyloid aggregation, tau propagation, neuroinflammation, and blood-brain barrier permeability modules to individual multimodal patient data (MRI, PET, CSF and blood biomarkers, APOE genotype, cognitive assessments, wearables, and electronic health records) is integrated with federated learning (FedAvg aggregation with secure aggregation and differential privacy) across multi-hospital nodes. While SHAP- and attention-based explainability, subgroup fairness validation, and anomaly detection support regulatory credibility under emerging FDA and EMA AI guidance, each patient's digital twin creates a personalized synthetic control trajectory, reducing reliance on randomized placebo assignment. A methodical approach to internal and external validation against the ADNI, AIBL, OASIS-3, and NACC cohorts is described. The methodology is anticipated to improve the representation of APOE ε4 carriers and underrepresented demographics, lower placebo-arm exposure and recruiting barriers, and provide statistically effective, patient-matched synthetic control arms in line with recent regulatory precedent. By uniting federated learning, mechanistic digital twins, and explainable AI, the proposed framework offers a conceptually coherent, ethically motivated pathway toward faster, fairer, and more personalised Alzheimer's clinical trials, though prospective multi-site validation remains a necessary next step before clinical or regulatory deployment.

Keywords: Alzheimer, AI-driven, Neuro inflammation, Federated, Clinical trials
How to cite this article: Madhumitha N * Devipriya .D1 Anbuselvi .S2 Anbu Kumaresan3 Nafisa Farheen. H 4. Review of Federated digital twins powered by AI for ethical and customized Alzheimer's clinical trials.. Int J Drug Deliv Technol. 2026;16(90s):1469-1482. DOI: 10.25258/ijddt.16.90s.147.

Introduction

Alzheimer's disease (AD) has moved from a peripheral concern of geriatric medicine to one of the defining public health challenges of the twenty-first century. Global estimates place the number of people living with dementia at more than 50 million, a figure projected to nearly triple by mid-century as populations age [1,2] The Global Burden of Disease framework, which tracks incidence, mortality, and disability-adjusted life years across more than 200 countries, shows that dementia-related mortality and disability have risen steadily since 1990, with the steepest increases occurring in low- and middle-income countries where diagnostic infrastructure and specialist neurology services remain scarce.[3,4] The economic consequences are equally stark.

Alzheimer's disease and other dementias will cost the world a total of about US$14.5 trillion by 2050, according to a macroeconomic modeling study that covered 152 countries. This is due to formal long-term care costs, decreased GDP growth in the fastest-aging economies, and lost labor force participation among patients and unpaid caregivers.[1] The World Alzheimer Report 2024 also revealed that this burden is exacerbated for women, who make up the majority of both patients and unpaid caregivers, by stigma, delayed diagnosis, and unequal access to care.[5]

Despite the enormous clinical need, the Alzheimer's disease (AD) therapeutic pipeline has historically experienced an exceptionally high rate of failure. Evidence from previous analyses indicates that approximately 98–99% of AD drug candidates entering clinical trials do not achieve regulatory approval, a markedly higher failure rate than the nearly 90% observed across therapeutic areas overall.[6,7] This pattern has begun to change only in recent years with the accelerated and subsequent full approvals of the anti-amyloid monoclonal antibodies aducanumab and lecanemab, followed by donanemab, which collectively ended a prolonged period of nearly two decades without the approval of a disease-modifying therapy for AD.[8,9] Nevertheless, these approvals have not fully addressed the structural challenges of AD drug development. The 2024 drug-development pipeline showed a decline compared with 2023, with fewer active clinical trials, fewer investigational compounds, and a reduced number of

novel chemical entities entering development, highlighting the persistence of inefficiencies in trial design and execution, Furthermore, the average time required to develop a single AD therapy—from preclinical investigation through regulatory review—continues to be approximately 13 years.[9]

The high rate of attrition in Alzheimer's disease (AD) drug development is driven not only by biological complexity but also by fundamental limitations in clinical trial design. A major challenge is the continued reliance on conventional placebo-controlled, parallel-group trials, which require large numbers of participants to be allocated to control arms. According to an analysis of the 2022 AD clinical pipeline, the approximately 172 active trials would collectively require more than 50,000 participants if fully enrolled, with nearly half expected to receive placebo rather than an investigational therapy despite the irreversible progression of the disease during the study period These design constraints also contribute to substantial recruitment difficulties. Across clinical research, around 80% of trials fail to achieve enrolment within the planned timeframe, while nearly half of participating clinical sites recruit fewer participants than originally targeted, further delaying study completion [10].

In Alzheimer's disease (AD), recruitment is further complicated by high screening-failure rates. Evidence indicates that many otherwise eligible candidates either withdraw because of the possibility of being assigned to a placebo group or fail to meet biomarker-based inclusion criteria, particularly amyloid-positivity requirements, which exclude a considerable proportion of willing participants.[11,12]

Another major challenge is the pronounced clinical and biological heterogeneity of AD. Individuals classified within the same mild-to-moderate Alzheimer's disease category often differ substantially in amyloid and tau pathology, cerebrovascular comorbidities, APOE genotype, rates of cognitive decline, and neuropsychiatric manifestations.[13] Such variability increases the dispersion of clinical outcome measures, including the Alzheimer's Disease Assessment Scale–Cognitive Subscale (ADAS-Cog), thereby reducing statistical power and necessitating larger sample sizes to reliably detect treatment effects. This, in turn, further intensifies recruitment demands.

The prolonged course of AD also presents significant practical and ethical challenges for clinical research. Because disease progression occurs over months to years, pivotal trials commonly require follow-up periods of at least 18 months, during which participants allocated to placebo continue to experience disease progression without active therapeutic intervention.[7,14] In addition, central nervous system drug development is constrained by the blood–brain barrier (BBB), which limits the delivery of therapeutic agents to the brain. Estimates suggest that more than 98% of small-molecule compounds and almost all large-molecule biologics fail to achieve meaningful brain penetration because of BBB transport restrictions, preventing many otherwise promising candidates from reaching pharmacologically effective concentrations at their intended targets.[7,15]Collectively, the need for large placebo cohorts, substantial patient heterogeneity, recruitment challenges, extended follow-up periods, high per-participant costs, and BBB-related limitations make AD clinical trials among the longest, costliest, and most resource-intensive programmes in contemporary drug development.

Artificial intelligence (AI) has increasingly emerged as a promising approach to addressing several longstanding challenges in Alzheimer's disease (AD) clinical research. Recent advances in machine learning and deep learning have enabled the identification of eligible participants from electronic health records, prediction of individual patterns of cognitive decline using baseline biomarker profiles, and detection of subtle neuroimaging and fluid-biomarker changes that may precede the onset of clinical symptoms by several years.[16,17] These capabilities have enhanced patient stratification and improved the precision of disease monitoring across different stages of AD progression.

Alongside these advances, the concept of the digital twin has gained considerable attention within healthcare. Initially developed in the manufacturing and aerospace industries as a dynamic computational representation of physical systems, the digital twin has been adapted for medical applications as a continuously updated virtual model of an individual patient.[18-20] By integrating clinical, biological, and imaging data, digital twins can simulate disease progression, estimate treatment responses, and assist clinical decision-making without requiring direct intervention on the patient. Their application has already been investigated in cardiovascular disease, oncology, and diabetes, and they are increasingly being considered as a strategy for generating synthetic or external control arms that may reduce reliance on placebo-assigned participants in clinical trials.[21-23]

A complementary technological advancement is federated learning (FL), a distributed machine-learning framework that enables multiple institutions to collaboratively develop a shared model while retaining patient data within their respective local servers.[24,25] This approach addresses one of the major obstacles in Alzheimer's research: the fragmentation of valuable datasets across hospitals, memory clinics, and national research registries. Because neuroimaging, genetic, and cognitive data are often stored in isolated institutional repositories, cross-border data sharing is frequently restricted by privacy regulations such as the General Data Protection Regulation (GDPR) in the European Union and the Health Insurance Portability and Accountability Act (HIPAA) in the United States.[26,27] By allowing collaborative model training without

centralized data transfer, federated learning offers a practical pathway for leveraging diverse multicenter datasets while maintaining compliance with these regulatory requirements.

Research Gap

Current research has largely explored digital twins, federated learning, and synthetic control arms as independent areas rather than as components of a unified framework. Existing digital twin studies in Alzheimer's disease (AD) predominantly rely on centrally pooled datasets from single research consortia, such as the Alzheimer's Disease Neuroimaging Initiative (ADNI), which limits the broader applicability of their findings and does not address the data-sharing challenges that federated approaches are designed to overcome.[28–30] In contrast, federated learning research in AD has primarily focused on diagnostic classification—distinguishing cognitively normal individuals, patients with mild cognitive impairment, and those with Alzheimer's disease—while giving comparatively little attention to modelling the longitudinal and mechanistic processes that underpin disease progression and treatment response, which are essential features of a true digital twin.[29,31,32]

To date, no published framework appears to integrate federated multi-institutional model training, mechanistic modelling of amyloid pathology, tau dynamics, and blood-brain barrier (BBB) interactions, together with synthetic control-arm generation, into a single architecture specifically designed to reduce placebo exposure while preserving patient privacy and maintaining regulatory credibility. This unmet need provides the rationale for the conceptual framework presented in Section 4, which proposes an AI-driven Federated Digital Twin (FDT) architecture tailored to the ethical, clinical, and statistical requirements of Alzheimer's disease clinical trials.

Artificial Intelligence in Alzheimer's disease:

Artificial intelligence (AI) in Alzheimer's disease (AD) employs different model families depending on the clinical task and data type. Traditional machine-learning algorithms, such as logistic regression, random forests, and support vector machines, are widely used for clinical and cognitive data because of their efficiency and interpretability.[16] Deep-learning models, particularly three-dimensional convolutional neural networks (CNNs), have become the standard for structural MRI analysis, where they identify hippocampal and medial temporal atrophy associated with disease progression.[33,34] Recurrent neural networks (RNNs), including LSTM and GRU architectures, are primarily used to model longitudinal changes in cognitive performance across repeated assessments.[35] More recently, graph neural networks (GNNs) and transformer-based models have improved the analysis of network-level brain degeneration and long-range spatial-temporal relationships, enhancing AD classification from neuroimaging data.[36–40] Generative models, including GANs and diffusion models, are also emerging for data augmentation, privacy-preserving synthetic data generation, and synthetic control-arm development in clinical trials.[22,29] **[Table -1]**

Table.1: AI methods used in Alzheimer's disease research

AI Method	Data Modality	Primary Application	Ref
Logistic regression / Random forest / SVM	Clinical, cognitive, tabular EHR	Risk stratification, MCI-to-AD conversion prediction	[16,41]
Convolutional Neural Networks (CNN)	Structural/functional MRI, PET	Atrophy detection, diagnostic classification	[33,34, 40]
Recurrent Neural Networks (RNN/LSTM)	Longitudinal cognitive scores, EHR sequences	Trajectory forecasting, progression modelling	[35]
Graph Neural Networks (GNN)	Brain connectivity graphs, multimodal fusion	Network-level degeneration, multimodal diagnosis	[36,37, 38]
Transformers / Vision Transformers	MRI, multi-site harmonised imaging	Global context modelling, multi-site domain adaptation	[31,39, 40]
Generative AI (GAN, diffusion, VAE)	Imaging, tabular, synthetic patient data	Data augmentation, synthetic control	[29,22, 17]

		arms	

Digital Twins in Healthcare

A digital twin is a continuously updated virtual representation of a physical entity, such as an individual patient, that integrates multimodal data to simulate disease progression and predict clinical outcomes.[19] Originally developed in industrial engineering for predictive maintenance, the concept has since been adapted to healthcare, where digital twins are classified as population- or patient-level, static or dynamic, and organ-specific or whole-body systems.[19,20,42] Their applications now extend to cardiovascular risk prediction, oncology treatment planning, diabetes management, and in-silico clinical trial design.[21,43,44] Unlike conventional prognostic models, digital twins evolve as new clinical, imaging, or wearable data become available, enabling real-time simulations that support personalised treatment decisions and outcome forecasting.[26,44]

Federated Learning and Privacy-Preserving AI

Federated learning (FL) is a distributed machine-learning approach that enables multiple institutions to collaboratively train a shared model without exchanging raw patient data.[24,25] In this framework, each institution trains the model locally on its own dataset and sends only model updates to a central server, which combines them using algorithms such as Federated Averaging (FedAvg) to improve the global model over successive training rounds.[24,45–47] The privacy of this approach is further strengthened through secure aggregation and differential privacy. Secure aggregation employs cryptographic methods that allow the server to aggregate client updates without accessing individual contributions, while differential privacy introduces carefully calibrated noise to model updates to reduce the risk of identifying individual patients.[15,26,48–50] Together, these techniques enable collaborative AI development across healthcare institutions while preserving patient confidentiality and supporting compliance with data privacy requirements.

Why Federated Learning Fits Alzheimer's Research

Alzheimer's disease (AD) research is particularly well positioned to benefit from federated learning (FL) for three key reasons. First, large collaborative initiatives such as the Alzheimer's Disease Neuroimaging Initiative (ADNI) have demonstrated the value of harmonised multi-site datasets spanning MRI, PET, cerebrospinal fluid biomarkers, genetics, and cognitive assessments.[51,52] However, the reliance on centralised data sharing has also exposed practical limitations, as many valuable datasets remain inaccessible because of legal, institutional, and logistical barriers.[53] Second, FL is designed for multi-institutional collaboration and has already shown promise in AD through federated neuroimaging classification, domain-adaptation transformers, and privacy-preserving models using homomorphic encryption.[29–31] Finally, evolving data protection regulations, including the General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA), make cross-border sharing of identifiable health data increasingly challenging.[26,27] By allowing institutions to retain local control of patient data while contributing to a shared global model, FL provides a practical and privacy-preserving foundation for collaborative digital-twin frameworks in Alzheimer's clinical research.

Ethical Challenges in Alzheimer's Clinical Trials

The Placebo Dilemma: Randomised placebo-controlled trials remain the gold standard for establishing treatment efficacy, but their use in Alzheimer's disease (AD) raises important ethical concerns because the disease is progressive and irreversible. Participants assigned to placebo continue to experience measurable and permanent cognitive and functional decline throughout trials that often extend for 18 months or longer.[7,14] This concern is well recognised by patients and caregivers, with studies consistently identifying fear of placebo assignment, together with the burden of invasive procedures such as lumbar puncture and the need for a study partner, as major barriers to trial participation.[54,55] Consequently, recruitment may become less representative, as individuals who perceive a greater chance of benefiting from active treatment are often less willing to accept the uncertainty associated with randomisation.

Privacy Challenges

The data most valuable for Alzheimer's disease (AD) research are also among the most sensitive in healthcare. Neuroimaging modalities such as structural and functional MRI and amyloid/tau PET scans can be inherently identifiable, as facial features may sometimes be reconstructed from raw MRI images, creating significant storage and governance challenges.[53] Genetic information, particularly APOE genotyping and broader sequencing data, presents additional privacy concerns because inherited risk extends beyond the individual to biological relatives.[56,57] Electronic health records further contain sensitive clinical, medication, and psychiatric information that could lead to discrimination if disclosed. Given the persistent stigma associated with Alzheimer's disease, breaches of such data may reveal not only a patient's clinical history but also an early or pre-symptomatic risk status that they may not yet have chosen to share with others.[5]

Bias and Representation

Alzheimer's disease (AD) clinical trials have historically underrepresented racially and ethnically diverse populations, with participants disproportionately drawn from non-Hispanic White, highly educated, and higher-income groups.[46,58] This imbalance is further amplified during amyloid-based eligibility screening, as analyses of the A4 and

AHEAD trials found that individuals from underrepresented racial and ethnic groups were more frequently excluded because of differences in measured amyloid-positivity rates.[46,58] Genetic diversity also presents a significant challenge. Although the APOE ε4 allele is the strongest common genetic risk factor for late-onset AD, most evidence has been generated from European-ancestry populations, and studies in more diverse cohorts indicate that its association with AD varies across ancestries.[59,60] Consequently, models trained on limited demographic populations may not generalise accurately to broader patient groups. The scarcity of rare genetic variants, including familial and Down syndrome-associated AD mutations, further highlights the need for multi-institutional and multinational collaboration, making federated learning a practical approach for improving model diversity while keeping patient data within its institution of origin.[3]

Regulatory Considerations

Regulatory agencies in the United States and Europe have increasingly established structured frameworks for the responsible use of artificial intelligence (AI) in drug development. In 2025, the U.S. Food and Drug Administration (FDA) introduced draft guidance outlining a risk-based credibility assessment framework that requires developers to clearly define an AI model's regulatory purpose and context of use before determining the level of evidence needed to support its application.[61–63] Complementary FDA guidance on AI in pharmaceutical manufacturing further reflects a broader effort to ensure consistent oversight across the product lifecycle.[62]

Similarly, the European Medicines Agency (EMA), in collaboration with the Heads of Medicines Agencies, launched an AI Workplan extending to 2028 that addresses AI applications from early drug discovery through regulatory evaluation.[64] The publication of the FDA–EMA Guiding Principles of Good AI Practice in Drug Development in 2026 further indicates growing international alignment in AI governance.[63] Across both regulatory frameworks, explainability has emerged as a key expectation, with regulators emphasising that increasingly complex AI models require stronger evidence demonstrating that their predictions are interpretable, transparent, and clinically justified.[31,65]

Proposed AI-Driven Federated Digital Twin Framework

The framework presented in this section is a conceptual synthesis of the technical foundations discussed in Sections 2 and 3. It is intended as an architectural blueprint for future empirical validation rather than a description of an implemented system. **[Table-2]**

Overall System Architecture

The proposed Federated Digital Twin (FDT) framework adopts a six-layer architecture that progresses from institution-level data management to clinical trial decision support.

Layer 1 – Data Acquisition: Participating hospitals, memory clinics, and research consortia retain multimodal patient data—including neuroimaging, fluid biomarkers, genetic information, cognitive assessments, wearable-device data, and electronic health records—within their own institutions, ensuring that raw patient data remain local.

Layer 2 – Data Harmonisation: Each site standardises its data to a common format to account for differences in imaging protocols, PET tracers, cognitive assessments, and EHR coding. Drawing on harmonisation strategies similar to the Uniform Data Set (UDS3) used in ADNI and NACC, this layer also identifies incomplete or inconsistent records for exclusion or imputation before model training begins.[17,66]

Layer 3 – Local AI Models: Harmonised data are used to train a local version of a shared hybrid CNN–GNN–Transformer model for several training epochs before only model parameters are prepared for sharing.

Layer 4 – Federated Aggregation: Model updates are securely transmitted to a coordinating server, where Federated Averaging (FedAvg), supported by secure aggregation and differential privacy, combines local contributions into an updated global model without exposing patient-level information.[24,45–50]

Layer 5 – Digital Twin Engine: The federated model is then combined with each patient's baseline data to generate an individual digital twin capable of simulating future amyloid, tau, cognitive, and functional trajectories under different treatment scenarios.

Layer 6 – Clinical Trial Decision Layer: Predictions from individual digital twins support higher-level trial functions, including synthetic control-arm generation, adaptive participant selection, subgroup-specific treatment-effect estimation, and safety monitoring through an explainable AI interface.

This layered architecture is designed to preserve patient privacy while enabling collaborative model development. Data governance and harmonisation remain institution-specific, model training occurs without centralising sensitive information, and clinical decision-support outputs are generated only after the federated model has been established, reducing the influence of site-specific variation on overall performance.

Table: 2- Multimodal Data Sources, Key Variables, and Their Roles in Alzheimer's disease Digital Twin Development

Data Source	Modality	Key Variables	Role in the Digital Twin
MRI	Imaging	Hippocam	Atrophy

(structural/functional)		pal/medial temporal volume, cortical thickness, connectivity	trajectory, network-level degeneration modelling
PET (amyloid, tau)	Imaging	SUVR, regional amyloid/tau burden	Biological staging, target-engagement simulation
CSF biomarkers	Fluid	Aβ42/40 ratio, p-tau, total tau	Early biological confirmation, amyloid-cascade staging
Blood-based biomarkers	Fluid	Plasma p-tau217, GFAP, NfL	Non-invasive screening, longitudinal monitoring
APOE genotype	Genetic	ε2/ε3/ε4 status	Risk stratification, subgroup simulation
Cognitive assessments	Clinical	MMSE, ADAS-Cog, CDR-SB	Functional outcome and endpoint
			prediction
Wearables	Digital biomarker	Gait, speech, sleep, activity, keystroke dynamics	Continuous remote monitoring, early functional-decline detection
Electronic Health Records	Structured/unstructured	Comorbidities, medications, visit history	Contextualisation, eligibility screening, safety monitoring

Biological Modelling Within the Digital Twin

A defining feature of the proposed framework is that the digital twin is designed as a mechanistically informed model rather than a purely statistical predictor. This approach addresses concerns that data-driven models, despite achieving high predictive accuracy, may underrepresent biologically important processes; for example, reinforcement-learning models of AD progression have been shown to undervalue amyloid accumulation despite its established role in disease pathology.[67]

The framework integrates four interconnected biological modules. The amyloid module models Aβ plaque formation resulting from β- and γ-secretase cleavage of amyloid precursor protein, with emphasis on the aggregation-prone Aβ1–42 species and current Alzheimer's disease staging criteria.[13,68] The tau module represents the prion-like spread of hyperphosphorylated tau across anatomically connected brain regions, making use of the same connectivity graph employed by the GNN architecture.[68] The neuroinflammation module captures the reciprocal relationship between inflammatory responses and amyloid-tau pathology, reflecting its emerging role as a major pathological axis in AD.[68] Finally, the synaptic dysfunction module links these molecular processes to clinical outcomes by modelling the functional decline measured through MMSE, ADAS-Cog, and CDR-SB scores. Together, these coupled modules aim to produce predictions that are both biologically interpretable and clinically meaningful.

Blood-Brain Barrier Module

The framework incorporates a dedicated blood-brain barrier (BBB) module to estimate whether therapeutic agents can achieve effective

concentrations within the brain. Drug permeability is predicted from physicochemical properties, including molecular weight, lipophilicity, and charge, acknowledging the well-established limitation of passive BBB diffusion for molecules exceeding approximately 400–500 Da.[17] The model further accounts for active transport mechanisms, including efflux proteins such as P-glycoprotein and influx transporters such as GLUT1, which influence brain drug delivery.[69] Receptor-mediated transcytosis through transferrin and insulin receptors is modelled as an alternative pathway for delivering larger biologics, including anti-amyloid and anti-tau antibodies.[45,69] To capture inter-patient variability, BBB integrity is adjusted using imaging and biomarker-derived indicators of ageing, cerebrovascular disease, and AD-related barrier dysfunction.[70,71]

Federated Training Workflow

The federated learning component operates through a recurring five-stage workflow. First, each institution trains the shared model on harmonised local data. Second, the difference between the updated local parameters and the previous global model is calculated as a weight update, while patient data remain local. Third, these updates are transmitted through secure aggregation, preventing access to individual institutional contributions.[26,48] Fourth, the coordinating server combines the updates using Federated Averaging (FedAvg), with optional server-side momentum to improve convergence under differential privacy.[16,24] Finally, participating institutions receive the updated global model, which is combined with each patient's baseline information to create or refine personalised digital twins while preserving local data ownership.

Personalised Clinical Trial Simulation

Once generated, each digital twin can simulate clinically relevant outcomes across multiple dimensions of trial design. The framework first projects disease progression under standard care by modelling future amyloid, tau, and neurodegenerative changes. These biological predictions are then translated into expected cognitive trajectories, including MMSE, ADAS-Cog, and CDR-SB outcomes. A drug-response module, informed by BBB permeability and pharmacokinetic-pharmacodynamic modelling, estimates how different treatment regimens may alter an individual's projected course. The framework also includes adverse-event prediction, incorporating comorbidities and genetic factors such as APOE ε4 status to identify patients at greater risk of amyloid-related imaging abnormalities (ARIA) during anti-amyloid therapy.[72] Finally, each twin can generate a synthetic control trajectory by simulating the patient's expected outcome under standard care alone, providing an individualised counterfactual that forms the basis for synthetic control-arm designs in future Alzheimer's disease clinical trials.

Synthetic Control Arms

Traditional Placebo Arms

Randomised placebo-controlled trials remain the regulatory standard for evaluating treatment efficacy because randomisation balances both measured and unmeasured confounding factors across study groups. However, this design poses significant challenges in Alzheimer's disease (AD), where participants assigned to placebo continue to experience progressive cognitive and functional decline without access to disease-modifying treatment.[7,14] Placebo arms also increase recruitment and retention demands, effectively requiring larger study populations while discouraging participation among patients and caregivers concerned about receiving inactive treatment.[10,54] These challenges become even more pronounced in rare forms of AD, such as early-onset and autosomal-dominant familial Alzheimer's disease, where the limited number of eligible participants can make conventional two-arm trials difficult to conduct.[3]

Digital Twin Control Arms

Digital twin-based synthetic control arms provide an alternative by replacing some or all placebo participants with personalised computational counterfactuals. Instead of assigning individuals to inactive treatment, each participant's expected outcome under standard care is estimated using a digital twin trained on historical and federated longitudinal datasets.[19,28,30] Unlike traditional external control arms that compare outcomes with a separate historical cohort, digital twins generate patient-specific predictions based on individual biomarker, genetic, and clinical characteristics.[23,73] This approach aligns with the FDA's 2025 draft guidance on AI-supported externally controlled studies, provided that comparator populations are appropriately matched and the underlying models demonstrate sufficient credibility.[31,61] Similar regulatory experience with synthetic and external controls in oncology and rare diseases further supports the feasibility of extending these methods to neurological disorders.[31,32]

By reducing the proportion of participants exposed to placebo, digital twin control arms can improve the ethical profile of AD trials while facilitating recruitment, as observed in other therapeutic areas where comparable approaches have already been explored.[20,22]

Statistical Benefits

In addition to ethical advantages, well-validated synthetic control arms can improve statistical efficiency. Because each participant is compared with an individualised counterfactual rather than a population-average control, prognostic adjustment methods can reduce residual variance in treatment-effect estimates and potentially lower sample-size requirements.[28,30] Similar approaches in neurological and cardiovascular research have reported sample-size reductions of approximately 10–

30% while maintaining appropriate Type I error control.[20] Smaller study populations can also accelerate interim analyses and shorten recruitment timelines, potentially reducing trial costs.[20,32] Nevertheless, these benefits depend on rigorous validation of digital twin performance, as poorly calibrated synthetic controls may introduce bias, underscoring regulators' emphasis on robust evidence for increasingly complex AI models.[31]

Explainable and Trustworthy AI

Explainability

AI-generated predictions to be acceptable in regulated Alzheimer's disease (AD) trials, they must be transparent to clinicians, statisticians, and regulators. SHAP (SHapley Additive exPlanations) has become one of the most widely used post-hoc explainability methods, quantifying the contribution of each input feature to an individual prediction using principles from cooperative game theory.[12] In AD research, SHAP has been applied to MRI-based CNN models, multimodal XGBoost models, and disease-progression models.[33,54,67,74] The transformer components of the proposed framework further enhance interpretability through attention mechanisms, which reveal the brain regions or time points most influential in model predictions.[39,40] Reviews of explainable AI methods, including SHAP, LIME, Grad-CAM, and layer-wise relevance propagation, consistently highlight their value in improving clinical trust and meeting emerging regulatory expectations for interpretable AI in drug development.[75,76]

Fairness

The framework treats subgroup validation as an essential part of model evaluation by assessing performance across APOE ε4 status, race and ethnicity, and sex.[46,58,60] To reduce demographic bias, federated learning can rebalance contributions from underrepresented sites during model aggregation and incorporate fairness-aware optimisation that penalises performance disparities between population groups.[48]

Privacy Protection

Privacy protection is achieved by combining differential privacy with secure aggregation during federated learning.[26,48–50] Differential privacy adds calibrated noise to model updates, limiting the risk of identifying individual patients, while secure aggregation ensures that the coordinating server receives only aggregated model updates rather than contributions from individual institutions.[26,48]

Anomaly Detection

Anomaly detection provides an additional safeguard for both data quality and model reliability. Autoencoders identify abnormal multimodal records by detecting unusually high reconstruction errors, helping to exclude corrupted scans, incomplete records, or out-of-distribution cases before training begins. For lower-dimensional clinical data, Isolation Forest and One-Class SVM offer efficient alternatives that can be deployed at individual federated sites without extensive computational resources. Beyond improving data quality, these methods can also identify patients whose observed clinical trajectory deviates markedly from their digital twin's prediction, potentially signalling adverse events, data-entry errors, or atypical disease progression requiring further clinical assessment.

Validation Strategy

Establishing the credibility of the proposed framework, consistent with the FDA's risk-based credibility assessment framework, requires a staged validation strategy combining both internal and external evaluation.[61,63] Internal validation would begin with k-fold cross-validation within each institution's harmonised dataset, followed by federated cross-validation, where entire institutions are excluded from individual training rounds to assess whether the global model can generalise to previously unseen sites.[28,30] [**Table-3]**

Table: 3 - Performance Evaluation Metrics for Digital Twin Model Validation and Prediction

Metric	**Purpose**	**Applicable Content**
Area Under the Curve (AUC)	Discriminative accuracy of diagnostic/risk classification	Local and global classification models
Root Mean Square Error (RMSE)	Magnitude of error in continuous trajectory prediction	Cognitive score and biomarker trajectory forecasts
Mean Absolute Error (MAE)	Interpretable average deviation of predicted vs. observed outcomes	Digital twin trajectory simulation
Calibration (slope/intercept)	Agreement between predicted probabilities and observed frequencies	Risk stratification, adverse-event prediction
Sensitivity	Proportion of true positives	Diagnostic classification,

	correctly identified	safety-signal detection

External validation would utilise four major Alzheimer's disease (AD) research cohorts with complementary strengths. ADNI provides extensive multimodal longitudinal data, including MRI, PET, cerebrospinal fluid biomarkers, genetics, and cognitive assessments across multiple centres.[51,52] AIBL contributes a large Australian cohort with a particular focus on early-stage AD and biomarker research.[77,78] OASIS-3 offers longitudinal neuroimaging and cognitive data spanning normal ageing to AD, while NACC provides harmonised clinical data from numerous U.S. Alzheimer's Disease Research Centers.[66,79–81] Evaluating the global model independently on each of these cohorts—without including them in the corresponding training phase—would provide a robust assessment of generalisability across diverse populations and research settings.[78,81]

Clinical and Pharmaceutical Applications

The proposed Federated Digital Twin (FDT) framework has applications across both pharmaceutical development and clinical practice.In drug development, it can support target prioritisation by simulating the effects of modulating key pathological pathways, including amyloid clearance, tau inhibition, and neuroinflammatory regulation, before candidate molecules enter experimental development. It can also facilitate dose optimisation through patient-specific blood-brain barrier (BBB) modelling and enable computational screening of BBB permeability before resource-intensive in vivo pharmacokinetic studies.[45,70]In precision medicine, the framework's multimodal and APOE-aware architecture enables more personalised treatment strategies. For example, APOE ε4 carriers receiving anti-amyloid antibodies could undergo enhanced monitoring for amyloid-related imaging abnormalities (ARIA), reflecting their established risk of treatment-related adverse events.[72] By integrating genetic, biomarker, and clinical information, the framework also supports prediction of individual treatment responses, advancing AD research toward biomarker-guided, subgroup-specific therapeutic approaches similar to those adopted in oncology.[17]

The framework also has the potential to improve clinical trial efficiency. AI-assisted pre-screening using harmonised electronic health records and biomarker data may reduce the high screen-failure rates currently seen in AD trials.[11,12] Digital twin-based synthetic control arms can minimise placebo exposure, addressing a major barrier to participant recruitment.[54] In addition, continuous monitoring through wearable technologies—including gait, speech, sleep, and activity tracking—can reduce the need for frequent in-person visits while providing clinically relevant measures of functional decline during long-term trials.[82,83]

Challenges and Limitations

No framework of this scope is free of practical and technical obstacles, several of which remain genuinely unresolved in the current literature. Several limitations of the proposed framework warrant consideration. First, although federated learning substantially strengthens privacy protection, it does not eliminate security risks. Model inversion and membership inference attacks remain recognised challenges, meaning that differential privacy and secure aggregation should be viewed as mechanisms for reducing—rather than completely removing—privacy risk.[49,50] Second, the biological modules described in Sections 4.3 and 4.4 rely on current knowledge of amyloid pathology, tau propagation, neuroinflammation, and blood-brain barrier (BBB) dynamics. As these mechanisms continue to be refined through on-going research, the biological validity of the digital twin will depend on the accuracy of the underlying assumptions. Finally, the proposed framework remains conceptual. None of its components has yet been prospectively evaluated in a real-world, multi-site Alzheimer's disease (AD) clinical trial, making the validation strategy outlined in Section 7 an essential next step. **[Table -4]**

Table: 4 – Key Challenges and Limitations in the Development and Implementation of Digital Twin Models

Challenge	Description	Proposed Situation
Data heterogeneity	Non-identical distributions of imaging protocols, cognitive tests, and demographics across sites (non-IID data)	Site-level harmonisation (Layer 2); personalised federated learning (e.g., FedPer-style local adaptation) rather than one fully shared model
Missing biomarkers	Not all sites collect CSF, PET, or genetic data; modality availability differs sharply across cohorts	Modality-agnostic architecture with learned imputation; graceful degradation when a modality is

		absent
Rare mutations	Autosomal-dominant and early-onset AD variants are too sparse for robust learning at any single site	Federated aggregation explicitly pools rare-variant carriers across sites without violating data-locality constraints
Model drift	Population characteristics, imaging technology, and standard-of-care therapy shift over time, degrading accuracy	Scheduled federated re-training cycles; continuous anomaly-based monitoring of prediction–observation divergence (Section 6.4)
Computational cost	Multi-modal CNN–GNN–transformer hybrids are resource-intensive; smaller sites may lack GPU infrastructure	Lightweight local model variants; cloud-based federated compute partnerships for resource-constrained sites
Regulatory uncertainty	AI-in-drug-development guidance remains in draft form in both major jurisdictions; digital-twin synthetic-control standards are still evolving	Early and iterative regulatory engagement (e.g., FDA CDER AI/ML consultation programme) at each stage of development

Future Directions

Several emerging developments could further strengthen the proposed framework. Advances in wearable technologies now enable continuous monitoring of gait, speech, sleep, and other behavioural signals that capture subtle functional changes between clinic visits. The emergence of wearable foundation models also suggests that broadly pretrained sensor models may soon support multiple biomarkers without requiring separate models for each application.[84–86] Integrating these continuous data streams with digital twins could enable real-time updates throughout a patient's disease course.

Adaptive clinical trial designs represent another promising direction. By combining digital twin-generated counterfactuals with interim trial data, future studies could modify eligibility criteria, randomisation ratios, or dosing strategies more efficiently while maintaining statistical rigour.[20,32] Incorporating multi-omics data—including genomic, transcriptomic, proteomic, and metabolomic information—could further improve the biological realism of digital twins alongside existing imaging and clinical biomarkers.[87]

Looking ahead, foundation models trained on large-scale multimodal neurological datasets may provide a shared backbone for disease-specific digital twins, reducing the need to train models from scratch during each federated learning cycle.[85] As digital twin technology matures, broader brain-wide models that capture overlapping mechanisms across Alzheimer's disease, vascular dementia, and related proteinopathies may support a unified platform for multiple neurodegenerative disorders.[88]

Conclusion

The limited success of many Alzheimer's disease clinical trials reflects not only therapeutic challenges but also shortcomings in trial design, participant recruitment, and data sharing. Large placebo groups expose participants to continued disease progression, recruitment often fails to reflect population diversity, and valuable clinical data remain fragmented across institutions. The proposed Federated Digital Twin (FDT) framework brings together federated learning, mechanistic digital twins, and explainable, privacy-preserving AI to address these interconnected challenges within a single conceptual architecture.

By enabling collaborative model training without transferring raw patient data, federated learning offers a practical solution to data fragmentation. Patient-specific digital twins provide an ethical alternative to extensive placebo allocation by generating personalised counterfactual predictions, while explainable AI, fairness auditing, and privacy-preserving techniques help align the framework with emerging regulatory expectations.

The framework also supports a more personalised approach to Alzheimer's research by integrating

amyloid, tau, genetic, and functional information to inform subgroup-specific treatment strategies. At the same time, its emphasis on secure aggregation, differential privacy, and local data governance demonstrates how sensitive neurological data can contribute to collaborative research without compromising institutional control.

Despite its promise, substantial work remains before this framework can be considered suitable for regulated clinical use. Prospective validation across independent cohorts such as ADNI, AIBL, OASIS, and NACC, comprehensive fairness assessments across demographic and genetic subgroups, and continued alignment with evolving FDA and EMA guidance will be essential. Nevertheless, as federated learning, digital twins, and AI continue to advance, their integration offers a compelling direction for developing Alzheimer's disease clinical trials that are more efficient, equitable, and patient-centred.[61,63,88]

Author contributions

Author of corresponding Assiatant Professor, Madhumitha. N (madhumitha481@gmail.com): writing—original draft, Devipriya. D-IV year, (devipriyadhanasekaran2710@gmail.com): review & editing, Assistant Professor, Anbuselvi .S,(cathy@gmail.com): Formal analysis, investigation. Assiatant Professor, Anbukumaresan, (anbukumaresan1992@gmail.com) :writing and review, Assistant Professor, Nafisa Farheen H, (nafisa.talwar@gmail.com) :writing and review.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declarations

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper

References

1. Chen S, Cao Z, Nandi A, Counts N, Jiao H, Prettner K, et al. The global macroeconomic burden of Alzheimer's disease and other dementias: estimates and projections for 152 countries or territories. Lancet Glob Health. 2024. doi:10.1016/S2214-109X(24)00264-X.
2. Zhang N, Chai S, Wang J. Assessing and projecting the global impacts of Alzheimer's disease. Front Public Health. 2025;12:1453489.
3. Safiullah M, Bhatti MI, Bashir MA, Ali M, Jha A, Fatima I, et al. Global burden of Alzheimer's disease and dementia: trends in mortality, risk factors and incidence across Sustainable Development Index regions from 1990 to 2021. Alzheimers Dement. 2025.
4. Global burden of Alzheimer's disease and other dementias in adults aged 65 years and older, 1991–2021: population-based study. 2025. Available from PMC: PMC12261920.
5. Alzheimer's disease International. World Alzheimer Report 2024: Global changes in attitudes to dementia. London: ADI; 2024.
6. Tesi RJ. Why do so many Alzheimer's clinical trials fail? Clinical Leader. 2023 Jul 19.
7. Clinical trials in Alzheimer's disease: a hurdle in the path of remedy. 2020. Available from PMC: PMC7152954.
8. Clinical trials of new drugs for Alzheimer disease: a 2020–2023 update. J Biomed Sci. 2023. Available from PMC: PMC10544555.
9. Cummings J, Zhou Y, Lee G, Zhong K, Fonseca J, Cheng F. Alzheimer's disease drug development pipeline: 2024. Alzheimers Dement (N Y). 2024;10:e12465.
10. The future of Alzheimer's disease research needs solutions that benefit trial patients. STAT [Internet]. 2022 Nov 22.
11. Understanding reasons for screen failures in Alzheimer's disease and related dementias (ADRD) trials: challenges and implications for recruitment. Alzheimers Dement. 2025. Available from PMC: PMC12725729.
12. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. Adv Neural Inf Process Syst. 2017;30:4765–74.
13. Genetically determined Alzheimer's disease research advances: the Down Syndrome & Autosomal Dominant Alzheimer's Disease 2024 Conference. Alzheimers Dement. 2024. Available from PMC: PMC12221810.
14. Challenges and proposed solutions to conducting Alzheimer's disease psychosis trials. Alzheimers Dement. 2024. Available from PMC: PMC11135469.
15. Drug delivery strategies to cross the blood-brain barrier in Alzheimer's disease: a comprehensive review on three promising strategies. 2025. PubMed PMID: 40393907.
16. Applications of deep learning in Alzheimer's disease: a systematic literature review of current trends, methodologies, challenges, innovations, and future directions. Artif Intell Rev. 2024.
17. AI-driven fusion of multimodal data for Alzheimer's disease biomarker assessment. Nat Commun. 2025.
18. Coorey G, Figtree GA, Fletcher DF, Redfern J. The health digital twin to tackle cardiovascular disease—a review of an emerging interdisciplinary field. NPJ Digit Med. 2022;5:126.
19. Health digital twins as tools for precision medicine: considerations for computation, implementation, and regulation. NPJ Digit Med. 2022.
20. Sun T, He X, Li Z. Digital twin in healthcare: recent updates and challenges. Digit Health. 2023;9:20552076221149651.

21. Advancing precision oncology with digital and virtual twins: a scoping review. 2024. Available from PMC: PMC11593079.
22. Bertolini D, Fuller F, Smith AM, Walsh JR, Zhuang R. Digital twins as synthetic controls in single-arm trials. arXiv:2605.12832. 2026.
23. Digital twins, synthetic patient data, and in-silico trials: can they empower paediatric clinical trials? ScienceDirect. 2025.
24. McMahan HB, Moore E, Ramage D, Hampson S, y Arcas BA. Communication-efficient learning of deep networks from decentralized data. Proc 20th Int Conf Artif Intell Stat (AISTATS). 2017;54:1273–82.
25. Rieke N, Hancox J, Li W, Milletarì F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119.
26. Bonawitz K, Ivanov V, Kreuter B, Marcedone A, McMahan HB, Patel S, et al. Practical secure aggregation for privacy-preserving machine learning. Proc ACM SIGSAC Conf Comput Commun Secur (CCS). 2017:1175–91.
27. Kairouz P, McMahan HB, Avent B, Bellet A, Bennis M, Bhagoji AN, et al. Advances and open problems in federated learning. Found Trends Mach Learn. 2021;14(1–2):1–210.
28. Federated learning and differential privacy for medical image analysis. Sci Rep. 2022;12:1953.
29. Mitrovska A, Safari P, Ritter K, Shariati B, Fischer JK. Secure federated learning for Alzheimer's disease detection. Front Aging Neurosci. 2024;16:1324032.
30. Stripelis D, Gupta U, Steeg GV, Ambite JL. A federated learning architecture for secure and private neuroimaging analysis. Patterns. 2024;5(8).
31. Zhu B, Lei B, Zhu Y, Liang E, et al. Federated domain adaptation via transformer for multi-site Alzheimer's disease diagnosis. IEEE Trans Med Imaging. 2023;42(12):3651–64.
32. Su C, Hou Y, Xu J, Zhang H, Zhang H, Bian J, et al. Federated learning with multi-cohort real-world data for predicting the progression from mild cognitive impairment to Alzheimer's disease. Alzheimers Dement. 2025.
33. Alzheimer's disease prediction via an explainable CNN using genetic algorithm and SHAP values. PLOS One. 2026.
34. Alzheimer's disease prediction via an explainable CNN using genetic algorithm and SHAP values [PMC version]. 2026. Available from PMC: PMC12810829.
35. Applications of deep learning in Alzheimer's disease: a systematic literature review of current trends, methodologies, challenges, innovations, and future directions. Artif Intell Rev. 2024. (RNN/LSTM trajectory modelling).
36. Graph neural networks in Alzheimer's disease diagnosis: a review of unimodal and multimodal advances. Front Neurosci. 2025.
37. Graph neural networks in Alzheimer's disease diagnosis: a review of unimodal and multimodal advances. Front Neurosci. 2025. Available from PMC: PMC12511118.
38. Discovering robust biomarkers of psychiatric disorders from resting-state functional MRI via graph neural networks: a systematic review. arXiv:2405.00577. 2024.
39. Topology-aware hybrid graph-transformer network for Alzheimer's disease diagnosis from structural magnetic resonance imaging. Front Comput Neurosci. 2026.
40. SFNet: a spatial-frequency domain deep learning network for efficient Alzheimer's disease diagnosis. arXiv:2507.16267. 2025.
41. Applying a statistical model-based AI method to identify prognostic factors for long-term cognitive decline in Alzheimer's disease: evidence from pooled placebo data of four phase III trials. ScienceDirect. 2026.
42. Digital twins for health: a scoping review. NPJ Digit Med. 2024;7:77.
43. Experiences of and recommendations on clinical trial design in Alzheimer's disease from the participant's point of view: a mixed-methods study in two clinical trial centers in the Netherlands. 2023. Available from PMC: PMC10071606.
44. Sheller MJ, Edwards B, Reina GA, Martin J, Pati S, Kotrotsou A, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Sci Rep. 2020;10:12598.
45. Wasielewska JM, Chaves JCS, Johnston RL, Milton LA, Hernández D, Chen L, et al. A sporadic Alzheimer's blood-brain barrier model for developing ultrasound-mediated delivery of Aducanumab and anti-Tau antibodies. Theranostics. 2022;12(15):6826-6847.
46. Racial and ethnic differences in plasma biomarker eligibility for a preclinical Alzheimer's disease trial. Alzheimers Dement. 2024.
47. APOE-ε4 allele[s]-associated adverse events reported from placebo arm in clinical trials for Alzheimer's disease: implications for anti-amyloid beta therapy. 2024. Available from PMC: PMC11285649.
48. Applying a statistical model-based AI method to identify prognostic factors for long-term cognitive decline in Alzheimer's disease: evidence from pooled placebo data

of four phase III trials. ScienceDirect. 2026. (APOE ε4 status analysis).
49. Federated learning and differential privacy for medical image analysis. Sci Rep. 2022;12:1953. (Differential privacy mechanisms).
50. Apolipoprotein E and Alzheimer's disease pathology in a diverse autopsy study. Alzheimers Dement. 2025. Available from PMC: PMC12635860.
51. Machine learning workflow to explain black-box models for early Alzheimer's disease classification evaluated for multiple datasets. arXiv:2205.05907. 2022.
52. Using space-filling curves and fractals to reveal spatial and temporal patterns in neuroimaging data. arXiv:2501.12111. 2025.
53. Evaluating the Alzheimer's disease data landscape. 2020. Available from PMC: PMC7744022.
54. Interpreting artificial intelligence models: a systematic review on the application of LIME and SHAP in Alzheimer's disease detection. Brain Inform. 2024;11:9.
55. Explainable artificial intelligence in neuroimaging of Alzheimer's disease. 2025. Available from PMC: PMC11899653.
56. A Critical Review of the FDA's draft guidance on artificial intelligence in drug and biological product regulation. J Chem. 2026. doi:10.1155/joch/5202999.
57. Key takeaways from FDA's draft guidance on use of AI in drug and biological life cycle. DLA Piper Insights. 2025 Jan.
58. Cognitive screening biases in a secondary prevention Alzheimer's disease clinical trial. medRxiv. 2025.
59. Apolipoprotein E and Alzheimer's disease pathology in a diverse autopsy study. Alzheimers Dement. 2025. (Ancestry-specific APOE associations).
60. Apolipoprotein E and Alzheimer's disease pathology in a diverse autopsy study. Alzheimers Dement. 2025. (Diverse Brazilian autopsy cohort methodology).
61. US Food and Drug Administration. Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products [draft guidance]. Silver Spring (MD): FDA/CDER; 2025 Jan.
62. US Food and Drug Administration. Artificial intelligence and medical products: how CBER, CDER, CDRH, and OCP are working together. Silver Spring (MD): FDA; 2024 Mar (rev. 2025 Feb).
63. FDA's AI guidance: 7-step credibility framework explained. IntuitionLabs [Internet]. 2026 Apr (rev.).
64. European Medicines Agency, Heads of Medicines Agencies. AI workplan 2023–2028. Amsterdam: EMA/HMA; 2023 Dec.
65. A Critical Review of the FDA's draft guidance on artificial intelligence in drug and biological product regulation. J Chem. 2026. (Niazi SK, critique and EMA comparison).
66. AI-driven fusion of multimodal data for Alzheimer's disease biomarker assessment. Nat Commun. 2025. (Cohort harmonisation to UDS3 format).
67. Unifying interpretability and explainability for Alzheimer's disease progression prediction. arXiv:2406.07777. 2024.
68. Editorial: digital remote patient monitoring in neurodegenerative diseases. Front Digit Health. 2025;7:1725819.
69. Overcoming the blood–brain barrier in Alzheimer's disease: translational perspectives on advanced drug delivery platforms. Front Neurosci. 2026.
70. Wasielewska JM, Chaves JCS, White AR, Oikari LE. Modeling the blood–brain barrier to understand drug delivery in Alzheimer's disease. In: Huang X, editor. Alzheimer's Disease: Drug Discovery. Brisbane: Exon Publications; 2020. Chap. 7.
71. Deciphering Alzheimer's disease: molecular mechanisms, preclinical models and strategies to overcome blood-brain barrier. ScienceDirect. 2025.
72. Digital twins in clinical trials: virtual controls & FDA. IntuitionLabs [Internet]. 2026.
73. Interpreting artificial intelligence models: a systematic review on the application of LIME and SHAP in Alzheimer's disease detection. Brain Inform. 2024;11:9. (LIME/SHAP comparative review).
74. Explainable artificial intelligence in neuroimaging of Alzheimer's disease. 2025. Available from PMC: PMC11899653. (SHAP, LIME, Grad-CAM, LRP methods).
75. Convolutional Neural Networks for Classification of Alzheimer's Disease: Overview and Reproducible Evaluation. arXiv:1904.07773.
76. Alzheimer's disease progression and risk factors: a standardized comparison between six large data sets. 2020. Available from PMC: PMC6804515.
77. OASIS-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer disease. 2019 (with later applications through 2026).
78. GMN4AD: graph matching network for Alzheimer's disease diagnosis with test-time domain adaptation using multi-centered structure magnetic resonance imaging. arXiv:2606.13919. 2026.

79. Evaluating the Alzheimer's disease data landscape. 2020. Available from PMC: PMC7744022. (Cross-cohort modality availability comparison).
80. Development of neurodegenerative disease diagnosis and monitoring from traditional to digital biomarkers. Biosensors (Basel). 2025;15(2):102.
81. Digital remote patient monitoring in neurodegenerative diseases [editorial]. Front Digit Health. 2025;7:1725819.
82. From everyday to existential — the ethics of shifting the boundaries of health and data with multimodal digital biomarkers. arXiv:2511.09238. 2025.
83. Physical activities enable scalable foundation modelling for broad-spectrum health prediction. arXiv:2607.06954. 2026.
84. AnyPPG: an ECG-guided PPG foundation model trained on over 100,000 hours of recordings for holistic health profiling. arXiv:2511.01747. 2025.
85. From fragmented neurotechnologies to closed-loop brain health systems: integrating artificial intelligence, digital health, digital twins, and advanced materials for brain disorders. Front Bioeng Biotechnol. 2026.
86. Nerrise F, Schütz N, Zhao Q, Gould C, Milstein A, Schulman K, et al. A framework of digital biomarkers for neurodegenerative diseases. Nat Rev Bioeng. 2026. doi:10.1038/s44222-026-00433-7.