A Smart-Contract–Based Validation Framework for Secure and Auditable Federated Learning in Dementia

Elif Calik
ADAPT Research Centre,
School of Computer Science,
University of Galway,
Galway, Ireland
email: elif.calik@universityofgalway.ie

Ayse Keles
ADAPT Research Centre,
School of Computer Science,
University of Galway,
Galway, Ireland
email: ayse.keles@universityofgalway.ie

Malika Bendechache
ADAPT Research Centre,
School of Computer Science,
University of Galway,
Galway, Ireland
email: malika.bendechache@universityofgalway.ie

***Abstract*— Typically arising from neurodegenerative diseases (most notably Alzheimer's disease, Lewy body dementia, frontotemporal dementia, and Parkinson's disease), dementia diagnosis and prognosis increasingly leverage Machine Learning (ML) across heterogeneous data modalities—including neuroimaging, structured Electronic Health Records (EHRs), and emerging digital biomarkers. However, strict privacy regulations and institutional barriers impede data pooling, while cross-site heterogeneity undermines model robustness. We present a permissioned-blockchain-enabled Federated Learning (FL) framework that addresses these challenges through a validation-first design grounded in a consensus Minimum Dataset (MDS) for dementia. Participating hospitals map local EHR/imaging fields to the shared MDS and perform local rule checks before training. A lightweight smart contract records hash-anchored validation receipts on-chain, creating an immutable audit trail without exposing raw data. Provenance links are maintained off-chain using content-addressed objects (e.g., InterPlanetary File System - IPFS) for EHR/imaging pointers and model artifacts; their identifiers are anchored on-chain to ensure verifiable version control. Local models are trained exclusively on validated records; contributions are aggregated (e.g., weighted Federated Averaging (FedAvg)) to produce a hash-verified global model distributed uniformly to all sites. The framework is technology-agnostic, minimizes operational overhead by storing only digests on chain, and directly targets key clinical adoption requirements: secure collaboration, reproducibility, and output compatibility under non-Independent Identically Distributed (non-IID) conditions. By standardizing input pre-training and enforcing transparent, tamper-evident exchanges, the approach aims to improve aggregation stability, cross-site generalizability, data quality, and information-governance compliance in dementia ML workflows.**

Keywords*— *Federated learning; blockchain; smart contract; data quality; neurodegenerative disease; dementia minimum-dataset.

I. INTRODUCTION

Dementia, which is marked by an irreversible decline in cognitive functions, such as memory, language, and decision-making, and is typically driven by progressive neurodegeneration, remains difficult to diagnose and forecast at scale. Early and precise diagnosis is crucial for timely and effective interventions, especially during the initial stages of the disease [1]. This need has fueled a growing interest in using Machine Learning (ML) models to assist clinicians with disease classification, progression prediction, and personalized care [2]. In ML applications, diverse data modalities are increasingly leveraged, including neuroimaging (e.g., Magnetic Resonance Imaging (MRI) to measure brain atrophy, a hallmark of neurodegenerative diseases) [3], with models, such as Support Vector Machines (SVMs) achieving high Alzheimer's Disease (AD) classification accuracy [4], structured clinical data from EHRs capturing demographics, labs, and diagnoses for progression prediction (e.g., Mild Cognitive Impairment (MCI) to AD) [5] , and emerging digital biomarkers, such as spontaneous speech [6] and wearable device data for real-time monitoring and early detection of cognitive decline [7].

However, a major obstacle is the scarcity and lack of diversity in datasets, which hinders the development of generalizable ML models. Pooling large volumes of patient data into a single repository for training is often impossible due to strict privacy regulations and institutional barriers, creating what are known as "data islands" or silos [8]. FL has emerged as a promising solution to these challenges [9][10]. Its core principle is to facilitate collaborative model training across decentralized institutions without the need for centralizing or explicitly exchanging raw patient data [11]. Instead, the ML algorithm is brought to the data, a process that inherently addresses patient privacy concerns and fosters trust [12]. In FL, a shared global model is trained on multiple local datasets, with the data samples remaining on their respective nodes [13]. Only model updates, such as the weights and biases of a deep neural network, are exchanged between local nodes and a central server or among the nodes themselves. This decentralized approach aligns with "Zero Trust", as it shares only model parameters without exposing raw data.

Two primary architectural paradigms exist in FL: client-server and peer-to-peer [14]. In the most common is the client-server or centralized model. A central server orchestrates the training process by distributing a global model to clients, collecting their locally trained model parameters, and aggregating them to create an improved global model. In contrast, the peer-to-peer architecture, nodes coordinate directly with their neighbours to obtain the global model by exchanging updates [14].

A major technical challenge in FL is data heterogeneity, which violates the independent and identically distributed (IID) assumption of traditional ML, especially prevalent in healthcare, where different hospitals may serve distinct patient populations or use varying equipment and diagnostic criteria [15]. Statistical heterogeneity in FL arises from the non-IID nature of healthcare data, where distributions vary significantly across institutions or devices due to differences in patient populations, clinical practices, and data acquisition protocols. Such variability leads to skewed data distributions that challenge model convergence, reduce generalizability, and limit the effectiveness of standard aggregation algorithms like FedAvg [16].

A central requirement for clinical adoption of FL is the ability to ensure both secure collaboration and consistent model performance across heterogeneous healthcare institutions. Blockchain-enhanced FL has already demonstrated significant promise by addressing privacy, secure parameter sharing, transparency, and accountability

through encrypted model updates, secure aggregation, peer-to-peer transfer protocols, and immutable audit trails [17]. However, a critical gap in blockchain-enabled FL for healthcare is the lack of robustness to heterogeneity, which is especially pronounced in dementia care, where data sources range from cognitive tests and neuroimaging to speech, gait, and wearables sensor data [7]. In this consortium setting, a permissioned blockchain can be used not as a compute substrate, but as a shared, tamper-evident coordination and provenance layer across independent hospitals. A classic centralized architecture (e.g., a single coordinator with a conventional database and logs) can orchestrate FL, but it implicitly concentrates trust: one operator becomes responsible for enrollment, event ordering, and the integrity of validation and model-version records, and disputes about "who submitted what and when" ultimately depend on that operator's logs. In contrast, anchoring validation receipts and model version identifiers on a permissioned ledger creates a jointly verifiable audit trail that no single site can unilaterally rewrite, while still keeping patient-level content off-chain. We adopt a permissioned design to align with hospital governance (authenticated membership, controlled read/write access) and to keep on-chain activity lightweight—only fixed-size digests and content identifiers (CIDs) are recorded—while model artifacts and pointers remain off-chain.

To address these gaps, our framework introduces a smart contract–based mechanism that standardizes data representation and enforces interoperability prior to local training. By harmonizing input features through a lightweight, verifiable smart contract framework, we reduce distributional discrepancies between participating institutions, thereby facilitating more robust aggregation and improving cross-site generalizability. This combined use of blockchain for secure collaboration and smart contracts for heterogeneity management establishes (i) robustness of aggregation to noise, bias, and malicious behavior, (ii) consistency and reproducibility of global outputs across heterogeneous sites, and (iii) a traceable, accountable process that meets information-governance expectations.

The rest of this paper is structured as follows. Section II surveys the relevant state of the art, Section III details the proposed research methodology, and Section IV presents the conclusions and discusses avenues for future work.

II. Related Work

Integrating blockchain with FL in AI-assisted brain imaging offers a promising direction for secure, collaborative, and regulation-compliant analytics in neurology [18]-[20]. Blockchain complements FL by adding a tamper-proof ledger, transparent coordination through smart contracts, and decentralized control over model sharing. Together, these features address long-standing issues, such as a lack of trust, weak data integrity, limited auditability, and vulnerability to adversarial manipulation that can undermine multi-institutional learning in healthcare.

The strength of Bhatia et al. [18] lies in its explicit operationalization of "data quality" as a measurable, pre-aggregation criterion. Their decentralized data-evaluation mechanism deploys miners to execute submitted local models against a concealed validation set and admit only those updates that exceed a predefined accuracy threshold. This "quality-gating" approach protects aggregation from data poisoning and label-noise attacks, stabilizes performance under heterogeneity, and preserves fidelity to the clinical signal. Its limitation, however, is the reliance on curator-held validation data and threshold design, which necessitate careful governance to avoid bias toward particular case mixes or scanner characteristics.

While the design of Imboccioli et al. [19] does not intrinsically judge the "quality" of an update by its predictive merit, it enhances the verifiability, provenance, and process fidelity that underpin trust in distributed clinical AI. Rather than evaluating accuracy explicitly at the gate, their framework orchestrates training phases via an immutable state machine, anchors cryptographic digests of model parameters on-chain, and coordinates storage of ciphered parameters (e.g., via IPFS) for reliable retrieval. This phased, auditable mechanism compels every participant—including the aggregator—to follow the same rules, thereby inhibiting out-of-protocol actions, covert parameter tampering, or selective sharing. A notable implication for output compatibility is that every collaborating hospital retrieves an identical, hash-verified global model, which curtails divergence due to inconsistent versions or update timing.

Third, Rajit et al. [20] approach safeguards against in-transit manipulation and supports traceability of each contribution, thereby reinforcing both data quality (through integrity guarantees) and output compatibility (through consistent application of an auditable, deterministic aggregator). They emphasize secure, auditable integration of client contributions with a focus on handling non-IID data distributions often observed across medical centers. Their architecture transforms each client's weight updates into SHA-256 hashes appended to the blockchain, enabling immutable provenance and tamper-evident transport. Aggregation proceeds via a weighted FedAvg scheme, producing a global model that reflects the relative information content contributed by each site. Importantly, the authors examine non-IID scenarios and introduce an accelerated aggregation variant to expedite convergence while maintaining accuracy—an attractive property for time-sensitive clinical deployments.

Taken together, these studies illustrate instructive trade-offs. Accuracy-based gating [18] directly controls predictive quality but hinges on governance of validation data and thresholds. Protocol-centric integrity [19] ensures observability, reproducibility, and non-repudiation but leaves performance vetting to external mechanisms. Hash-anchored, weighted aggregation [20] prioritizes provenance and distributional robustness, yet presumes acceptable upstream quality in local training and curation.

Building on these insights, this study advances a blockchain-enhanced FL framework for dementia imaging that treats data quality and output compatibility as first-class design objectives. Unlike prior work, our approach introduces smart contract–driven coordination to integrate quality control with verifiability, thereby aligning predictive performance, process integrity, and clinical trust.

To sum up, we propose a dementia-oriented blockchain-enabled FL workflow that couples validation-first interoperability (via an MDS) with tamper-evident provenance and versioning, aiming to reduce representation-driven heterogeneity prior to training while keeping patient data local.

III. Material and Methods

In multidisciplinary healthcare research, a persistent challenge is the limited opportunity for synchronous, in-person collaboration with clinicians, who typically operate on a 24/7 schedule. In domains dealing with complex and

sensitive data, establishing minimum datasets (MDS) for diagnosis is therefore a critical, consensus-driven step that enables coordinated progress. However, disease-specific minimum datasets are not always available. In this context, to enhance input quality for FL and to promote output consistency, we developed a dementia MDS. A smart-contract–based proposed framework was subsequently provided.

When constructing the MDS, we systematically reviewed the International Classification of Diseases 11th (ICD-11) [21] sub-diagnostic groups associated with dementia: (A) Alzheimer disease—(i) early-onset, (ii) late-onset, (iii) mixed with cerebrovascular disease, (iv) mixed with other non-vascular aetiologies, (v) unspecified onset; (B) Vascular dementia; (C) Dementia with Lewy bodies; (D) Frontotemporal dementia; (E) Substance/medication-induced dementia—(i) alcohol, (ii) sedative/hypnotic/anxiolytic, (iii) volatile inhalants, (iv) other specified; (F) Dementia due to other diseases—(i) Parkinson disease, (ii) Huntington disease, (iii) exposure to heavy metals/other toxins, (iv) human immunodeficiency virus (HIV), (v) multiple sclerosis, (vi) prion disease, (vii) other specified cause, (viii) unspecified cause. Note: Code (G)—behavioural/psychological symptoms in dementia—is a supplementary code and not a primary aetiology.

The MDS includes key sociodemographic variables [22]: age (treated as a primary covariate for dementia risk and progression), sex (included because phenotypes and biomarker levels can differ by sex), and total years of formal education (as a proxy for cognitive reserve). To capture cognitive status, we incorporated Mini-Mental State Examination (MMSE) (global cognition), Montreal Cognitive Assessment (MoCA) (sensitive to mild cognitive impairment), Clinical Dementia Rating—Global Score (CDR_global) (ordinal staging of cognitive/functional impairment), Clinical Dementia Rating—Sum of Boxes (CDR_sb) (a composite sensitive to subtle clinical change), and Instrumental Activities of Daily Living (IADL_total) (functional autonomy in complex daily tasks). Neuropsychiatric symptoms are assessed using the Geriatric Depression Scale—Short Form (GDS_total) for depressive symptoms and the Neuropsychiatric Inventory Questionnaire (NPIQ_total) for the overall burden of behavioural/psychological symptoms. Comorbidities included are hypertension (HT), reflecting vascular risk associated with cognitive decline; diabetes mellitus (DM), reflecting metabolic risk; atrial fibrillation (AF), reflecting cardioembolic/low-perfusion risk; and history of stroke/transient ischaemic attack, capturing prior cerebrovascular events. Vital signs and laboratory measures comprise systolic blood pressure (SBP) (vascular risk profiling), diastolic blood pressure (DBP) (related to cerebral perfusion), body mass index (BMI) (nutritional/metabolic status relevant to prognosis), thyroid-stimulating hormone (TSH) (screens for reversible cognitive effects), vitamin B12 (B12) (deficiency can cause treatable cognitive impairment), haemoglobin A1c (HbA1c) (medium-term glycaemic control), and low-density lipoprotein cholesterol (LDL) (marker of atherosclerotic burden). Imaging variables: MRI Protocol / Modality Name (mri_protocol_name), Slice orientation and acquisition plane (acq_plane), Scanner Manufacturer (scanner_vendor), Series completeness (srs_comp), Field of View (FOV), and Matrix size (matrix_size). Blood-based biomarkers [23][24] include plasma phosphorylated tau-181 (plasma_ptau181) (associated with Alzheimer pathology), plasma neurofilament light chain (plasma_nfl) (marker of neuro-axonal injury and prognosis), and the amyloid-beta 42/40 ratio (abeta42_40) (a peripheral amyloid signature with diagnostic/prognostic value). The apolipoprotein E ε4 (apoe_e4) carrier status indicates whether at least one ε4 allele is present [22]-[25]. Finally, digital measures comprise gait speed (GS) (a functional digital marker linked to cognitive decline), mean daily step count (MDSC) (a proxy for physical activity and overall health), and speech pause rate (SPR) (a digital marker of language/fluency changes). Table 1 presents the structure of the constructed MDS, summarizing each variable by category, description, variable name, requirement level, and data type. Consistent with the requirement stratification of the MDS, variables are organized into a core, flexible, and extensible framework: variables designated as “Yes” constitute the minimum indispensable elements for consistent characterization; “Conditional” items provide methodological flexibility where equivalent instruments are acceptable; “Recommended” variables enhance clinical and analytic utility when available; and “Optional” measures allow scalable enrichment without compromising baseline interoperability across heterogeneous resource environments.

We built a practical, end-to-end workflow that lets hospitals collaborate on dementia models without moving patient data. The idea is simple: every hospital prepares its data in the same way, proves on a blockchain that it meets agreed-upon rules, trains locally, and then shares only model files (never raw data) for aggregation. In our proposed implementation, the coordination and provenance layer is realized using a permissioned Hyperledger Fabric network, where participating hospitals join as authenticated members under a shared governance model. The workflow logic is implemented as a Fabric chaincode written in TypeScript, exposing functions to (i) record MDS-validation receipts as fixed-size hashes and metadata, (ii) register model-version identifiers per training round, and (iii) enforce role-based access control and monotonic versioning to prevent unauthorized submissions or replay. Model artifacts themselves will remain off-chain (e.g., stored via content-addressed storage such as IPFS), while the ledger stores only cryptographic digests and CIDs to keep on-chain activity lightweight and auditable.

As shown in Figure 1, the proposed smart contract–based framework follows a validation-first pipeline: (i) each hospital validates its data against the agreed minimum dataset (MDS) and anchors validation hashes on the Hyperledger Fabric ledger; (ii) EHR and imaging artifacts remain off-chain in IPFS and are referenced via CIDs linked to those hashes; (iii) local training is performed strictly on records with on-chain validation receipts; (iv) local model updates are stored in IPFS and their versions are anchored on-chain; and (v) the aggregator builds the global model, stores it in IPFS, and anchors the released version on-chain so that every site retrieves the same, hash-verified artifact.

TABLE I. MINIMUM DATASET FOR DEMENTIA DIAGNOSIS

Category	Variable description	Variable name	Requirement level	Type
Clinical Diagnosis (ICD-11)	Primary ICD-11 code	icd11_dementia_code	Recommended	string
Socio Demographics	Age in years at baseline	age_years	Yes	integer
	Biological sex	sex	Yes	categorical
	Total years of formal education	education_years	Yes	integer
Cognitive Status	Mini-Mental State Examination	mmse_total	Conditional (MMSE or MoCA)	integer
	Montreal Cognitive Assessment	moca_total	Conditional (MMSE or MoCA)	integer
Clinical Staging and Function	Clinical Dementia Rating - Global Score	cdr_global	Conditional (CDR or CDR-SB)	ordinal
	Clinical Dementia Rating - Sum of Boxes	cdr_sb	Conditional (CDR or CDR-SB)	float
	Instrumental Activities of Daily Living	iadl_total	Yes	integer
Neuropsychiatric Symptoms	Geriatric Depression Scale - Short Form	gds_total	Conditional (GDS or NPI-Q)	integer
	Neuropsychiatric Inventory Questionnaire	npiq_total	Conditional (GDS or NPI-Q)	integer
Comorbidities	Hypertension	ht	Yes	binary
	Diabetes Mellitus	dm	Yes	binary
	Atrial Fibrillation	af	Yes	binary
	History of Stroke / Transient Ischaemic Attack	stroke_history	Yes	binary
Vitals and Labs	Systolic Blood Pressure	sbp_mmhg	Recommended	integer
	Diastolic Blood Pressure	dbp_mmhg	Recommended	integer
	Body Mass Index	bmi_kg_m2	Recommended	float
	Thyroid-Stimulating Hormone	tsh_mIU_L	Recommended	float
	Vitamin B12	b12_pg_mL	Recommended	float
	Hemoglobin A1c	hba1c_percent	Recommended	float
	Low-Density Lipoprotein Cholesterol	ldl_mmol_L	Recommended	float
Imaging	MRI Protocol / Modality Name	mri_protocol_name	Yes	string
	Slice orientation and acquisition plane	acq_plane	Yes	string
	Scanner Manufacturer	scanner_vendor	Yes	string
	Series completeness	srs_comp	Yes	boolean
	Field of View	fov	Yes	float
	Matrix_size	matrix_size	Optional	list
Blood Biomarkers	Plasma Phosphorylated Tau-181	plasma_ptau181_pg_mL	Optional	float
	Plasma Neurofilament Light Chain	plasma_nfl_pg_mL	Optional	float
	Amyloid-beta 42/40 Ratio	abeta42_40_ratio	Optional	float
Genetics	Apolipoprotein E e4 Carrier Status	apoe_e4_carrier	Optional	binary
Digital Measures	Gait Speed	gait_speed_m_s	Optional	float
	Mean Daily Step Count	daily_step_count_mean	Optional	integer
	Speech Pause Rate	speech_pause_rate	Optional	float

We next describe each step in detail in the following section (Figure 1).

Step 1 — Agree on the minimum dataset (MDS): Before anything else, participating hospitals align on a compact list of variables that are most informative for dementia (e.g., age, sex, education, cognitive and functional scales, key comorbidities, selected labs, MRI descriptors, and optional blood/digital biomarkers). We encode this list as a machine-readable JSON schema that defines types (integer, float, categorical), valid ranges, and simple "either/or" rules (for example, MMSE or MoCA; CDR-Global or CDR-Sum of Boxes). This schema becomes the sole source of truth for what "valid input" means across all sites.

Step 2 — Validate inputs with smart contracts (on-chain): Each hospital checks its EHR entries and related imaging metadata against the MDS. Rather than trusting a central service, we use a permissioned blockchain with a lightweight smart contract that exposes a validate() function. To keep the contract lightweight, schema conformance checks are executed locally using the agreed JSON schema and validator version, and the chaincode records a signed, time-ordered receipt that binds (a) the schema/version identifier, (b) a digest of the validation outcome, and (c) the submitting member identity to a training round. This creates an immutable receipt that the input met the agreed-upon rules, without revealing sensitive content.

Step 3 — Keep pointers off-chain, link them on-chain: To preserve privacy, the actual EHR locators (where the record lives in the hospital system) and imaging identifiers are stored off-chain in the IPFS. IPFS gives each object a CID. We then store the hash of that CID on the blockchain alongside the validation event. Later, anyone can prove that the same object was used, simply by recomputing its hash—yet no clinical details appear on the chain.

Step 4 — Train locally on validated caches: Only records that have a valid on-chain receipt are admitted to the hospital's training cache. Each institution trains its own model on-premises; no patient-level data leaves the site. This step supports both client–server FL and peer-to-peer FL.

Step 5 — Share models via IPFS, anchor versions on chain: After local training, the site packages its model weights (and, if desired, optimizer state) and uploads them to IPFS. The resulting CID is hashed and committed to the blockchain (e.g., via commitModel()), creating a permanent, tamper-evident version tag. The aggregator (or peers) read these on-chain events to discover the latest models and fetch them from IPFS.

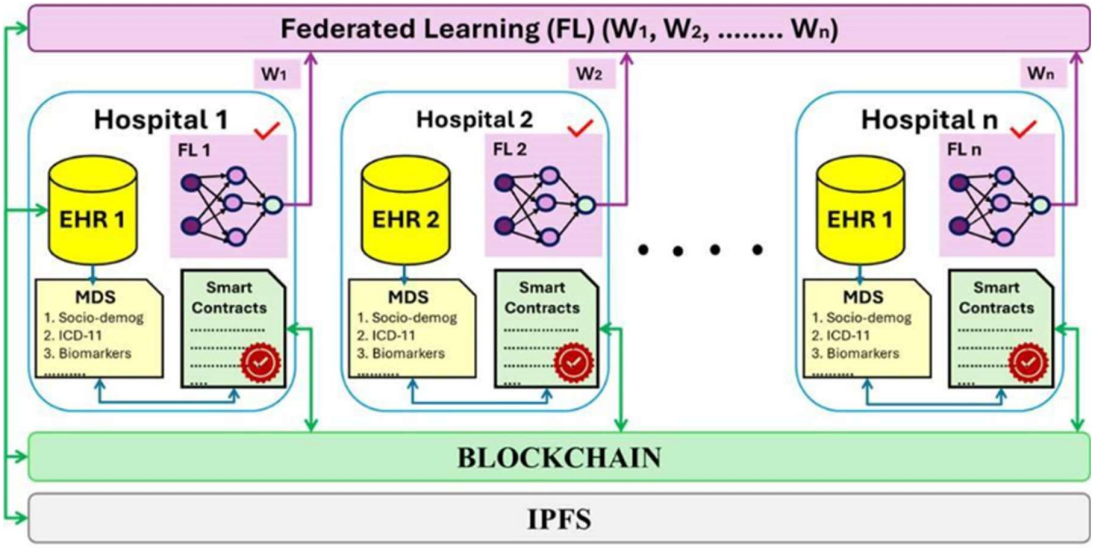

Figure 1. Proposed Smart Contract-based framework

Step 6 — Aggregate and redistribute a single, verified global model: The default aggregator is weighted FedAvg, but robust alternatives can be used without changing the attestation layer. The aggregated global model is also stored in IPFS, and its version hash is written on the chain. This guarantees every site downloads the same, hash-verified global model, avoiding silent version drift.

On the other hand, threat model, attack surface, and evaluation plan can be articulated as follows: We consider a consortium setting in which participating hospitals are authenticated members of a permissioned network, yet a subset of clients may behave maliciously. Assets include patient-level data (kept on premises), the correctness of MDS validation outcomes, model updates, and the integrity of global model releases and version lineage. Adversaries may attempt (i) poisoning or backdoor updates, (ii) replay of stale submissions, (iii) unauthorized contract invocation, and (iv) denial-of-service against validation or model-commit endpoints; additionally, the network and off-chain storage are not assumed confidential or highly available. Under this threat model, the ledger layer primarily targets integrity, accountability, and reproducibility by anchoring tamper-evident validation receipts and model-version identifiers, while privacy against inference attacks and robustness against poisoning require complementary mechanisms (e.g., secure aggregation, differential privacy, and robust aggregation). From an implementation perspective (i.e., Hyperledger Fabric, TypeScript chaincode), we explicitly bound the smart-contract attack surface by enforcing role-based access control, monotonic round/version checks to prevent replay or rollback, fixed-size inputs to limit on-chain payloads, and rate limits to reduce DoS risk. Because IPFS provides integrity via content addressing but not confidentiality or guaranteed availability, off-chain artifacts are protected through encryption and consortium pinning/replication policies, supported by operational controls for membership management, key rotation, and revocation. Finally, we will evaluate the framework through (a) learning outcomes (AUC/F1, sensitivity/specificity, calibration, and convergence), (b) systems overhead (end-to-end round time, Fabric endorsement/commit latency and throughput, communication volume per round, ledger growth, and IPFS retrieval/pinning costs), and (c) adversarial robustness by injecting realistic attacks (label-flipping, backdoor triggers, sign-flip/scale, and replay) at varying fractions of malicious clients; we will report attack success rate, performance degradation, and "when defenses are enabled" detection accuracy and robustness of the final global model.

The key contributions of the proposed Smart Contract-based frameworks are as follows.

i. Security and privacy: Raw data never moves; only fixed-size digests and CIDs are recorded on the permissioned ledger.
ii. Auditability and reproducibility: Every key step (input validation, model commits, global releases) leaves a signed, time-ordered trail.
iii. Consistency across sites: Because the same schema gates inputs and the same versioned model is redistributed, results are easier to reproduce and compare.
iv. Heterogeneity control: Standardizing inputs before training reduces avoidable site-to-site differences, helping aggregation work better.

IV. Conclusion and Future Work

We introduced a pragmatic, end-to-end workflow for multi-institutional dementia modeling that combines FL with a permissioned blockchain and an MDS-driven validation layer. The design secures collaboration without centralizing patient-level data, provides deterministic provenance via on-chain hashes and off-chain content addresses, and reduces avoidable distributional drift by standardizing inputs prior to local training. Anchoring both input validation and model versioning on the chain yields a reproducible lineage for all contributions and ensures every site receives the same, hash-verified global model. Collectively, these properties address persistent blockers to clinical deployment—trust, auditability, and cross-site consistency—while remaining lightweight and compatible with existing hospital governance.

In our future work, we will conduct prospective, multi-site pilots to quantify accuracy, calibration, convergence, and operational burden under realistic non-IID conditions. We will also validate the proposed architectural framework through simulation studies and/or prototype implementations, enabling a clearer assessment of feasibility and any measurable performance gains over baseline FL. To make

these comparisons interpretable, we will isolate the effect of the MDS validation layer by evaluating (i) standard FL without schema gating, (ii) MDS-gated FL using a classic architecture (conventional logging rather than on-chain receipts), and (iii) the full permissioned-ledger design with on-chain validation receipts and model version anchoring. We will harden privacy and robustness with secure aggregation, differential privacy, and robust aggregators, guided by a formal threat model and adversarial testing. We will broaden modality support (in the dementia use-case; speech, gait, wearables), provide interoperability tooling (EHR/Fast Healthcare Interoperability Resources (FHIR) mappings, validators), and mature lifecycle operations (drift monitoring, version pinning/rollback) to enable sustained clinical deployment.

COMPETING INTERESTS:

The authors declare no conflicts of interest.

ACKNOWLEDGMENT

This research was funded by Digi+ Marie Skłodowska-Curie Action (MSCA) CoFund under grant number 101081609, and supported by Taighde Éireann - Research Ireland under Grant number 13/RC/2106/P_2 (ADAPT Centre), and the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Action (MSCA) grant agreement No. 101152004-101109961. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

REFERENCES

[1] B. Xin, D. Zhang, H. Fu, and W. Jiang, "Association between multimorbidity and the risk of dementia: a systematic review and meta-analysis," Arch. Gerontol. Geriatr., vol. 131, Art. no. 105760, Apr. 2025, doi: 10.1016/j.archger.2025.105760.

[2] S. Dattola, A. Ielo, G. Varone, A. Cacciola, A. Quartarone, and L. Bonanno, "Frontotemporal dementia: a systematic review of artificial intelligence approaches in differential diagnosis," Front. Aging Neurosci., vol. 17, Art. no. 1547727, 2025, doi: 10.3389/fnagi.2025.1547727.

[3] D.-H. Shih, Y.-H. Wu, T.-W. Wu, Y.-K. Wang, and M.-H. Shih, "Classifying dementia severity using MRI radiomics analysis of the hippocampus and machine learning," IEEE Access, vol. 12, pp. 160030–160051, 2024, doi: 10.1109/ACCESS.2024.3483833.

[4] S. Klöppel et al., "Automatic classification of MR scans in Alzheimer's disease," Brain, vol. 131, no. 3, pp. 681–689, 2008, doi: 10.1093/brain/awm319.

[5] J. Pan, Z. Fan, G. E. Smith, Y. Guo, J. Bian, and J. Xu, "Federated learning with multi-cohort real-world data for predicting the progression from mild cognitive impairment to Alzheimer's disease," Alzheimers Dement., vol. 21, no. 4, Art. no. e70128, Apr. 2025, doi: 10.1002/alz.70128.

[6] X. Ouyang, "Design and deployment of multi-modal federated learning systems for Alzheimer's disease monitoring," in Proc. 21st Annu. Int. Conf. Mobile Syst., Appl. Serv. (MobiSys), Helsinki, Finland, 2023, pp. 612–614, doi: 10.1145/3581791.3597505.

[7] G. Cornelius, W. Hodgson, R. Maguire, and K. Egan, "Wearable technology, smart home systems, and mobile apps for the self-management of patient outcomes in dementia care: systematic review," J. Med. Internet Res., vol. 27, Art. no. e65385, 2025, doi: 10.2196/65385.

[8] A. Mitrovska, P. Safari, K. Ritter, B. Shariati, and J. K. Fischer, "Secure federated learning for Alzheimer's disease detection," Front. Aging Neurosci., vol. 16, Art. no. 1324032, Mar. 2024, doi: 10.3389/fnagi.2024.1324032.

[9] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," Knowl.-Based Syst., vol. 216, Art. no. 106775, Mar. 2021, doi: 10.1016/j.knosys.2021.106775.

[10] S. Zhan, L. Huang, G. Luo, S. Zheng, Z. Gao, and H.-C. Chao, "A review on federated learning architectures for privacy-preserving AI: lightweight and secure cloud–edge–end collaboration," Electronics (Basel), vol. 14, no. 13, Art. no. 2512, 2025, doi: 10.3390/electronics14132512.

[11] M. I. Sharif, M. Mehmood, M. P. Uddin, K. Siddique, Z. Akhtar, and S. Waheed, "Federated learning for analysis of medical images: a survey," J. Comput. Sci., vol. 20, no. 12, pp. 1610–1621, Oct. 2024, doi: 10.3844/jcssp.2024.1610.1621.

[12] R. Seyghaly, J. Garcia, and X. Masip-Bruin, "A comprehensive architecture for federated learning-based smart advertising," Sensors (Basel), vol. 24, no. 12, Art. no. 3765, 2024, doi: 10.3390/s24123765.

[13] H. Guan, P.-T. Yap, A. Bozoki, and M. Liu, "Federated learning for medical image analysis: a survey," Pattern Recognit., vol. 151, Art. no. 110424, Jul. 2024, doi: 10.1016/j.patcog.2024.110424.

[14] L. Shanmugam, R. Tillu, and M. Tomar, "Federated learning architecture: design, implementation, and challenges in distributed AI systems," J. Knowl. Learn. Sci. Technol., vol. 2, no. 2, pp. 371–384, 2023, doi: 10.60087/jklst.vol2.n2.p384.

[15] M. Nasajpour et al., "Federated learning in smart healthcare: a survey of applications, challenges, and future directions," Electronics (Basel), vol. 14, no. 9, Art. no. 1750, 2025, doi: 10.3390/electronics14091750.

[16] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artif. Intell. And Statist. (AISTATS), 2017.

[17] N. Nezhadsistani, N. S. Moayedian, and B. Stiller, "Blockchain-enabled federated learning in healthcare: survey and state-of-the-art," IEEE Access, vol. 13, pp. 119922–119945, 2025, doi: 10.1109/ACCESS.2025.3587345.

[18] L. Bhatia and S. Samet, "A decentralized data evaluation framework in federated learning," Blockchain: Res. Appl., vol. 4, no. 4, Art. no. 100152, 2023, doi: 10.1016/j.bcra.2023.100152.

[19] F. Imboccioli, G. Cialone, and S. Ferretti, "Decentralization of learning and trust in healthcare: blockchain-driven federated learning for Alzheimer's MRI image classification," in 2024 IEEE Int. Conf. Pervasive Comput. Commun. Workshops (PerCom Workshops), 2024, pp. 739–744, doi: 10.1109/PerComWorkshops59983.2024.10502820.

[20] S. Rajit, Z. F. Ananna, M. M. Ehsan, N. N. Punom, and S. Siddique, "Multi-class brain tumor classification of MRI image using federated learning with blockchain," in 2024 IEEE Region 10 Symp. (TENSYMP), 2024, pp. 1–8, doi: 10.1109/TENSYMP61132.2024.10752160.

[21] World Health Organization, "International classification of diseases (ICD): ICD-11 for mortality and morbidity statistics," [Online]. Available: WHO ICD-11. Browser: https://icd.who.int/browse11. Accessed: Mar. 8, 2026.

[22] W. Qi et al., "Mapping knowledge landscapes and emerging trends in AI for dementia biomarkers: bibliometric and visualization analysis," J. Med. Internet Res., vol. 26, Art. no. e57830, Aug. 2024, doi: 10.2196/57830.

[23] T. Sekimori, K. Fukunaga, D. I. Finkelstein, and I. Kawahata, "Advances in blood biomarkers and diagnosis approaches for neurodegenerative dementias and related diseases," J. Integr. Neurosci., vol. 23, no. 10, Art. no. 188, Oct. 2024, doi: 10.31083/j.jin2310188.

[24] F. Santos, V. Cabreira, S. Rocha, and J. Massano, "Blood biomarkers for the diagnosis of neurodegenerative dementia: a systematic review," J. Geriatr. Psychiatry Neurol., vol. 36, no. 4, pp. 267–281, Jul. 2023, doi: 10.1177/08919887221141651.

[25] E. Solje, A. Benussi, E. Buratti, A. M. Remes, A. Haapasalo, and B. Borroni, "State-of-the-art methods and emerging fluid biomarkers in the diagnostics of dementia—a short review and diagnostic algorithm," Diagnostics, vol. 11, no. 5, Art. no. 788, 2021, doi: 10.3390/diagnostics11050788.