Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Multi-Agent AI Framework for Predicting Alzheimer's Disease Trajectory Through Sleep-Dependent Memory Consolidation Biomarkers: A Conceptual Architecture

Bahaa Masry

mo.bahaa.masry@gmail.com

Damascus University

Anas Shahin

Syrian Virtual University

Research Article

Keywords: Alzheimer's disease, multi-agent system, sleep biomarkers, memory consolidation, cognitive decline prediction, EEG, wearable sensors, neuroimaging, trajectory modeling, neural ODE, conceptual framework

Posted Date: April 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9270054/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A Multi-Agent AI Framework for Predicting Alzheimer's Disease Trajectory Through Sleep-Dependent Memory Consolidation Biomarkers: A Conceptual Architecture

Bahaa Masry[1] and Anas Shahin[2]

[1]Faculty of Medicine, Damascus University, Damascus, Syria
[2]Faculty of Information Technology, Syrian Virtual University, Damascus, Syria

Corresponding author: mo.bahaa.masry@gmail.com

Abstract

Background: Alzheimer's disease (AD) affects over 55 million people worldwide and is projected to nearly triple by 2050, yet clinical diagnosis typically occurs after substantial irreversible neuronal damage. Emerging evidence demonstrates that disruptions in sleep-dependent memory consolidation—including alterations in sleep spindle–slow oscillation coupling, rapid eye movement (REM) architecture, and circadian integrity—serve as early biomarkers detectable years before clinical onset. Current AI approaches predominantly rely on single-modality analyses and binary classification schemes, failing to capture the temporal dynamics of cognitive decline.

Methods: Following design science research methodology, we developed a multi-agent AI framework integrating four data modalities through specialized agents: a Sleep Electrophysiology Agent (EEG/polysomnography), a Wearable Data Agent (actigraphy, heart rate variability), a Cognitive Assessment Agent (neuropsychological and digital biomarkers), and a Neuroimaging Agent (structural MRI volumetrics). A Central Orchestrator Agent performs attention-based multimodal fusion and generates individualized decline trajectory predictions using neural ordinary differential equations (Neural ODEs). A proof-of-concept (POC) implementation was evaluated on synthetic data ($N = 500$ subjects, 10 timepoints over 5 years) calibrated to published biomarker effect sizes.

Results: Comparative analysis against existing approaches demonstrates that no published framework combines all four modalities with sleep biomarkers as the central feature set within a trajectory prediction paradigm. The POC achieved 89.3% risk staging accuracy across four decline categories, with RMSE of 0.037 and trajectory-level Pearson correlation of 0.56. Misclassifications were predominantly adjacent-tier.

Cross-modal attention analysis revealed clinically interpretable modality weighting: the Cognitive Assessment Agent received highest attention (0.46–0.72), while the Neuroimaging Agent contribution increased progressively from cognitively normal to AD groups, consistent with advancing structural neurodegeneration.

Conclusions: The proposed framework establishes an architecturally validated foundation for sleep-based AD trajectory prediction. The POC demonstrates technical feasibility, clinically meaningful modality prioritization, and actionable risk staging. Validation on real clinical datasets following the proposed phased roadmap is essential to establish clinical utility.

Keywords: Alzheimer's disease, multi-agent system, sleep biomarkers, memory consolidation, cognitive decline prediction, EEG, wearable sensors, neuroimaging, trajectory modeling, neural ODE, conceptual framework

1. Background

1.1. The Alzheimer's disease burden

Alzheimer's disease (AD) represents one of the most pressing neurodegenerative challenges of the 21st century. The World Health Organization estimates that over 55 million people live with dementia globally, with approximately 10 million new cases annually and AD accounting for 60–70% of all cases [1]. Projections indicate this burden will reach 139 million by 2050, with annual global costs exceeding US$1.3 trillion [1, 2]. In the United States alone, 6.9 million individuals aged 65 and older live with AD dementia, with costs totaling $360 billion annually [2].

Despite decades of research investment, clinical diagnosis of AD typically occurs at a stage when substantial neurodegeneration has already taken place, severely limiting the efficacy of available therapeutic interventions. Current gold-standard diagnostic approaches—including amyloid positron emission tomography (PET) scanning ($3,000–$6,000 per scan), cerebrospinal fluid (CSF) biomarker analysis, and comprehensive neuropsychological assessment—are costly, invasive, and inaccessible to most populations globally [3]. The NIA-AA Research Framework acknowledges these limitations while emphasizing the need for accessible, non-invasive biomarker alternatives [3]. The emergence of plasma biomarkers such as phosphorylated tau-217 (p-tau217) has shown promise (AUC 0.92–0.97 for AD pathology detection [4, 5]), yet these remain laboratory-dependent and unsuitable for continuous monitoring.

1.2. Sleep disruption and Alzheimer's disease pathophysiology

A growing body of evidence has established a bidirectional relationship between sleep disruption and AD pathology (Fig. 2). During non-rapid eye movement (NREM) sleep, the glymphatic system facilitates clearance of metabolic waste—including amyloid-β (Aβ) and tau proteins—from the brain interstitium. The glymphatic system, discovered by Iliff et al. [6], operates via paravascular channels lined by astrocytic endfeet, facilitating CSF-interstitial fluid exchange. Xie et al. [7] demonstrated that this clearance is primarily active during sleep, with the interstitial space expanding by approximately 60% to enable convective flow. Chronic sleep disruption impairs this clearance, promoting Aβ accumulation and tau phosphorylation, which in turn damages sleep-regulating neural circuits, creating a pathological positive feedback loop [8, 9].

Sleep spindles—bursts of oscillatory activity at 11–16 Hz generated by thalamocortical circuits—play a critical role in memory consolidation by facilitating hippocampal-to-neocortical information transfer [10]. Helfrich et al. [11] demonstrated that precise temporal coupling between slow oscillations and sleep spindles deteriorates with age and medial prefrontal atrophy, predicting next-day memory failure. Mander et al. [10] established the mechanistic cascade linking prefrontal atrophy to reduced NREM slow waves and impaired hippocampus-dependent memory.

Lucey et al. [8] showed that decreased NREM slow-wave activity correlates with elevated CSF tau and tau-PET uptake in cognitively normal and early AD individuals, positioning specific sleep oscillations as markers of tau-related neurodegeneration. Winer et al. [9] provided evidence that sleep EEG features serve as non-invasive proxies for PET-measured Aβ and tau burden, with decreased prefrontal spindle density predicting higher medial temporal tau.

From an epidemiological perspective, Pase et al. [12] reported in the Framingham Heart Study (321 participants, mean age 67, 12-year follow-up) that each 1% reduction in REM sleep conferred a 9% increased dementia risk (HR 0.91, 95% CI 0.86–0.97). Lim et al. [13] found that higher sleep fragmentation was associated with a 1.5-fold increased risk of incident AD (HR 1.50, 95% CI 1.09–2.07) in a 6-year prospective study of 737 older adults. Musiek et al. [14] demonstrated that circadian rest-activity rhythm fragmentation occurs even in preclinical (amyloid-positive but cognitively normal) AD.

These sleep-derived electrophysiological signatures thus serve as non-invasive proxies for PET-measured AD pathology, forming the scientific foundation for the proposed framework.

1.3. Existing AI approaches to AD prediction

Deep learning approaches to AD prediction have primarily focused on structural MRI classification. Wen et al. [15] conducted a definitive reproducible evaluation of CNN architectures for AD classification, revealing that over 50% of published studies suffered from data leakage inflating accuracy. Under rigorous evaluation, performance ranges are typically 75–90% accuracy for AD vs. normal classification and 65–80% for MCI-to-AD conversion [16]. Qiu et al. [17] developed a multimodal deep learning framework combining demographics, neuropsychological testing, and neuroimaging, achieving performance comparable to practicing neurologists.

EEG-based approaches have been systematically reviewed by Tzimourta et al. [18] (49 studies, 2009–2020) and Pérez-Valero et al. [19], highlighting EEG as a low-cost, portable alternative for AD screening. However, sleep-specific EEG biomarkers remain underutilized in predictive models. Attention-based fusion mechanisms for multimodal medical data have been reviewed by Li et al. [20], identifying cross-modal transformers as the most promising direction.

Digital biomarkers measured via smartphones and wearables represent an emerging paradigm. Dagum [21] demonstrated that 7 days of passive smartphone data could predict neuropsychological test scores. Butler et al. [22] reported the largest consumer wearable cognition study to date (23,004 participants, 24 months, *Nature Medicine*), demonstrating feasibility of remote MCI classification with over 90% adherence at one year. Kourtis et al. [23] mapped AD manifestations to consumer sensor capabilities across gait, sleep, voice, and typing modalities.

1.4. Multi-agent systems in healthcare

Multi-agent systems (MAS) employ autonomous specialized agents that collaborate to solve complex problems [24]. In healthcare, Fdez-Olivares et al. [24] proposed multi-agent planning for comorbidity management where each agent encapsulates a single-disease guideline. Bertoncelli et al. [25] developed PredictMed-epilepsy, a MAS combining deep learning with patient observation agents for real-time neurological monitoring (82% accuracy, 102 children). Ivaşcu and Negru [26] achieved 95.25% accuracy using multi-agent wearable sensor monitoring. Shaik et al. [27] proposed multi-agent deep reinforcement learning for physiological signal monitoring, outperforming single-agent baselines. Despite these precedents, no published MAS framework addresses sleep-based AD trajectory prediction.

1.5. Study aims and contributions

This paper presents a conceptual multi-agent AI framework addressing the limitations identified above through four primary contributions:

1. **Multimodal Integration Architecture:** A novel multi-agent system integrating four complementary data modalities—EEG/polysomnography (PSG), wearable sensors, cognitive assessments, and structural neuroimaging—each processed by a specialized AI agent with clearly defined inputs, processing pipelines, and output specifications.
2. **Sleep-Centric Biomarker Foundation:** Centering sleep-dependent memory consolidation biomarkers as the primary predictive feature set, grounded in established neurobiological mechanisms linking sleep disruption to AD pathogenesis.
3. **Trajectory Prediction Paradigm:** Generating individualized decline trajectory predictions using Neural ODEs rather than binary classifications, providing clinically actionable longitudinal risk profiles.
4. **Proof-of-Concept Validation:** A complete implementation demonstrating architectural feasibility on synthetic data calibrated to published biomarker distributions, achieving 89.3% risk staging accuracy with interpretable cross-modal attention weights.

The framework targets individuals with subjective cognitive decline (SCD) [28] and mild cognitive impairment (MCI) [29, 30] within a hybrid deployment model combining continuous home monitoring with periodic clinical assessment.

2. Methods

2.1. Framework development methodology

This conceptual framework was developed following design science research (DSR) methodology [31], which is particularly suited for developing and evaluating IT artifacts intended to solve identified organizational problems. The problem space was defined through systematic literature review across three domains: (1) sleep neuroscience and AD pathophysiology, (2) existing AI-based AD prediction approaches, and (3) multi-agent system architectures in healthcare. Key gaps identified include the absence of sleep-centric multimodal integration, predominance of binary classification

over trajectory modeling, and lack of hybrid deployment architectures balancing clinical accuracy with ecological validity.

Framework objectives were derived from clinical requirements analysis and technical capability assessment. Clinical objectives include early detection in preclinical/prodromal stages, actionable risk stratification, longitudinal monitoring capability, and interpretable outputs for clinical decision-making. Technical objectives include multimodal data integration, graceful degradation under missing data, scalable architecture, privacy-preserving design, and modular agent updating.

2.2. Sleep biomarker evidence base

Table 1 summarizes the key sleep biomarkers associated with AD pathology and their evidence base, forming the scientific foundation for the proposed framework's feature extraction.

Table 1: Summary of sleep and related biomarkers associated with Alzheimer's disease pathology.

Biomarker	**AD Association**	**Evidence**	**Agent**	**References**
Spindle–SO coupling	$\downarrow$ Coupling predicts $\uparrow$ tau burden	$r = -0.46$ to -0.54 with CSF tau	SEA	[8, 9, 11]
Sleep spindle density	$\downarrow$ Prefrontal spindles $\rightarrow \uparrow$ MTL tau	Significant in MCI and early AD	SEA	[9, 11]
NREM slow-wave activity	$\downarrow$ SWA correlates with $\uparrow$ tau-PET	Tracks tau $>$ amyloid	SEA	[8, 10]
REM sleep percentage	Each 1% $\downarrow \rightarrow$ 9% $\uparrow$ dementia risk	HR 0.91 (95% CI 0.86–0.97)	SEA	[12]
Sleep fragmentation	$\uparrow$ Fragmentation $\rightarrow$ 1.5$\times$ AD risk	HR 1.50 (95% CI 1.09–2.07)	SEA/WDA	[13]
Circadian amplitude	$\downarrow$ Amplitude in preclinical AD	Differentiates Aβ+ from Aβ–	WDA	[14]
Rest-activity stability	$\downarrow$ Stability in MCI vs. controls	Actigraphy-measurable	WDA	[32]
HRV (parasympathetic)	$\downarrow$ HRV across dementias	$g = -0.36$, meta-analysis	WDA	[33, 34]
Overnight memory retention	$\downarrow$ Consolidation in MCI	Linked to spindle-SO coupling	CAA	[10, 11]
Hippocampal volume	$\downarrow$ Volume predicts conversion	MCI$\rightarrow$AD rate: 16.5%/year	NIA	[3, 35]

2.3. Proposed framework architecture

The proposed framework consists of five core components: four modality-specific processing agents and one Central Orchestrator Agent (COA). Each processing agent ingests, preprocesses, and extracts predictive features from its designated data stream. The COA performs cross-modal fusion, resolves conflicting signals, and generates trajectory predictions. Figure 1 illustrates the overall system architecture.

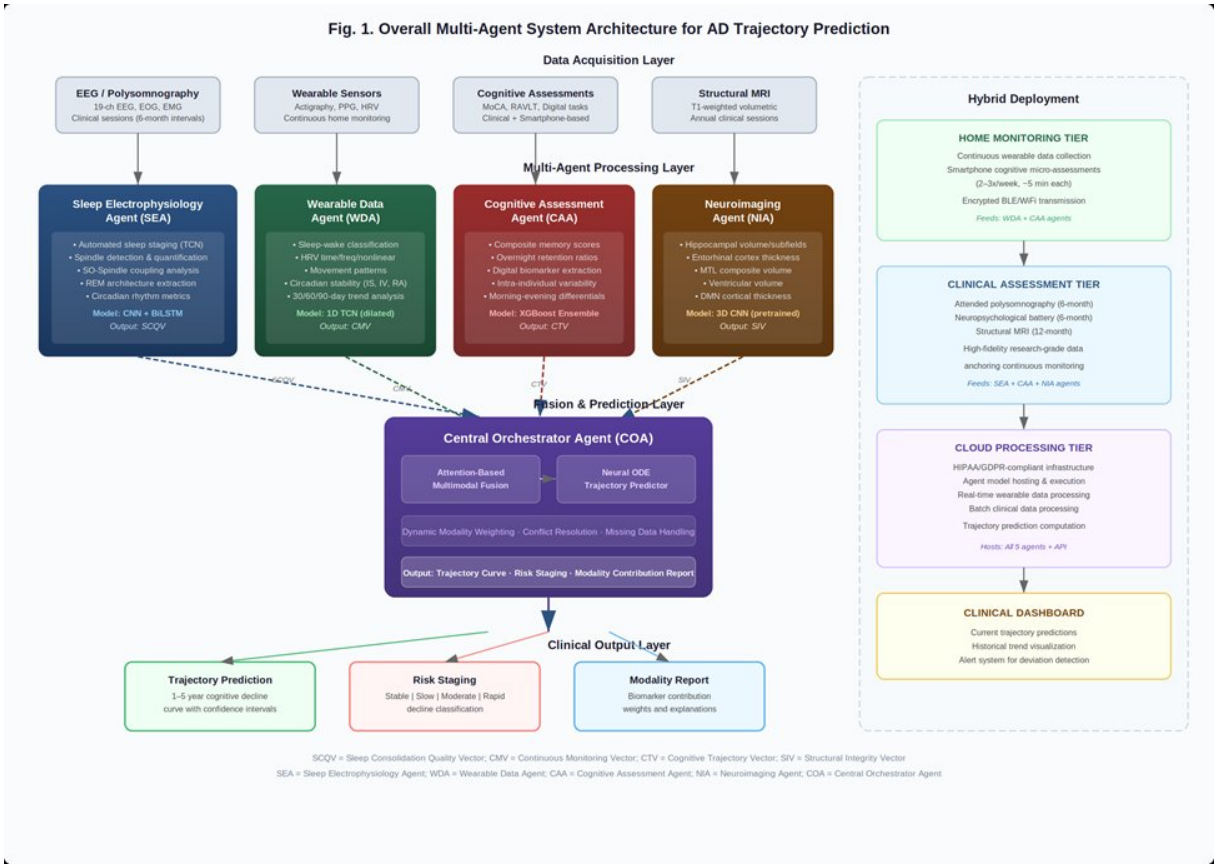


Figure 1: Overall multi-agent system architecture for AD trajectory prediction, showing the four modality-specific agents (SEA, WDA, CAA, NIA), the Central Orchestrator Agent with attention-based fusion and Neural ODE trajectory prediction, the hybrid three-tier deployment model, and the clinical output layer.

2.3.1. Sleep Electrophysiology Agent (SEA)

The SEA processes multi-channel EEG and polysomnography data acquired during periodic clinical sleep studies, serving as the primary analytical component reflecting sleep-dependent memory consolidation's central predictive role.

Input: Multi-channel EEG (minimum 19 channels, 10–20 montage), electrooculography (EOG), electromyography (EMG), and respiratory signals from attended polysomnog-

raphy sessions conducted at 6-month intervals.

Feature extraction: The SEA implements hierarchical feature extraction. First-level automated sleep staging uses a temporal convolutional network (TCN) [36] trained on American Academy of Sleep Medicine (AASM) criteria. Second-level microstructural features include: sleep spindle density, frequency, amplitude, and duration; slow oscillation parameters; spindle–slow oscillation phase-amplitude coupling strength and timing precision [11]; REM sleep duration, latency, density, and muscle atonia integrity; sleep fragmentation index, arousal frequency, and cyclic alternating pattern parameters; and circadian rhythm metrics derived from sleep architecture across sessions [14].

Agent model: A hybrid architecture combining a convolutional neural network (CNN) for spatial feature extraction across electrode channels with a bidirectional long short-term memory (BiLSTM) network [37] for temporal pattern recognition. The agent outputs a Sleep Consolidation Quality Vector (SCQV).

2.3.2. Wearable Data Agent (WDA)

The WDA processes continuous data from consumer-grade wrist-worn wearable devices, providing ecological monitoring between clinical visits. Recent evidence demonstrates the feasibility of large-scale remote cognitive health monitoring via such devices [22, 23].

Input: Triaxial accelerometry, photoplethysmography (PPG)-derived heart rate and heart rate variability (HRV), skin temperature, and ambient light exposure, collected continuously.

Feature extraction: Actigraphy-derived sleep-wake classification and quality metrics; nightly movement patterns; HRV time-domain (RMSSD, SDNN), frequency-domain (LF/HF ratio), and nonlinear measures (sample entropy, detrended fluctuation analysis) [33, 34]; circadian rhythm stability via interdaily stability, intradaily variability, and relative amplitude from 14-day windows [14]; and longitudinal trends over rolling 30-, 60-, and 90-day windows.

Agent model: A 1D temporal convolutional network (TCN) [36] with dilated causal convolutions captures both short-term nightly patterns and long-term trends. Output: Continuous Monitoring Vector (CMV).

2.3.3. Cognitive Assessment Agent (CAA)

Input: Standardized neuropsychological battery (MoCA, RAVLT, Trail Making Test, digit span) administered at clinical visits, supplemented by smartphone-based cognitive micro-assessments (reaction time, paired associates learning, spatial navigation) performed 2–3 times weekly [21, 22].

Feature extraction: Age- and education-adjusted composite memory scores; sleep-dependent memory consolidation efficiency (pre/post-sleep performance comparison); overnight retention ratios; digital biomarkers including intra-individual variability, learning curves, and longitudinal trends; morning–evening performance differentials.

Agent model: XGBoost gradient-boosted ensemble [38], selected for interpretability with structured tabular data. Output: Cognitive Trajectory Vector (CTV).

2.3.4. Neuroimaging Agent (NIA)

Input: T1-weighted structural MRI scans acquired at 12-month intervals.

Feature extraction: Automated segmentation of bilateral hippocampal volume and subfields (CA1, CA3, dentate gyrus, subiculum); entorhinal cortex thickness; medial temporal lobe composite volume; ventricular volume; and cortical thickness of the default mode network. Volumes normalized to intracranial volume and z-scored against age-matched normative databases [35, 39].

Agent model: 3D CNN pre-trained on large-scale datasets (ADNI [39], UK Biobank) and fine-tuned for AD-relevant feature extraction [15]. Output: Structural Integrity Vector (SIV).

2.3.5. Central Orchestrator Agent (COA)

Fusion mechanism: The COA implements an attention-based fusion architecture [20] learning modality-specific weights dynamically based on data quality, recency, and predictive relevance. This allows adaptive prioritization—weighting sleep electrophysiology data more heavily when spindle-SO coupling degrades, or neuroimaging when structural changes become pronounced.

Trajectory prediction: The fused multimodal representation is processed by a Neural ODE architecture [40], modeling continuous-time dynamics of cognitive decline. Unlike discrete-time models, Neural ODEs generate predictions at arbitrary future time points and handle irregularly sampled longitudinal data [41]. The model outputs an

individualized trajectory curve projecting decline over a 1–5 year horizon with confidence intervals.

Output specification: Three outputs for clinical interpretation: (1) a predicted cognitive trajectory curve on a standardized composite scale; (2) a risk staging classification (stable, slow decline, moderate decline, rapid decline); and (3) a modality contribution report detailing which biomarkers drive the prediction. Figure 4 illustrates example trajectory outputs.

Table 2 provides complete technical specifications for all agents.

Table 2: Technical specifications of framework agents.

Agent	Input Source	Key Features	Model Architecture	Output Vector	Schedule
SEA	19-ch EEG/PSG	Spindle-SO coupling, REM, SWA, CAP	CNN + BiLSTM	SCQV	6-month
WDA	Wrist wearable	Actigraphy, HRV, circadian metrics	1D TCN (dilated)	CMV	Continuous
CAA	NP battery + app	Memory scores, retention, RT variability	XGBoost ensemble	CTV	Hybrid
NIA	T1-weighted MRI	Hippocampal vol., EC thickness, DMN	3D CNN (pretrained)	SIV	12-month
COA	SCQV + CMV + CTV+ SIV	Multimodal fusion, conflict resolution	Attention + Neural ODE	Trajectory	Per visit

2.4. Neurobiological foundation

The scientific rationale for centering sleep-dependent memory consolidation biomarkers in AD prediction rests on the bidirectional pathological cascade illustrated in Fig. 2. Sleep disruption impairs glymphatic Aβ clearance [6, 7], promoting amyloid and tau accumulation [8, 9], which damages thalamocortical circuits essential for sleep generation [10], creating a self-reinforcing cycle. The framework's agents are positioned to detect changes at each stage of this cascade.

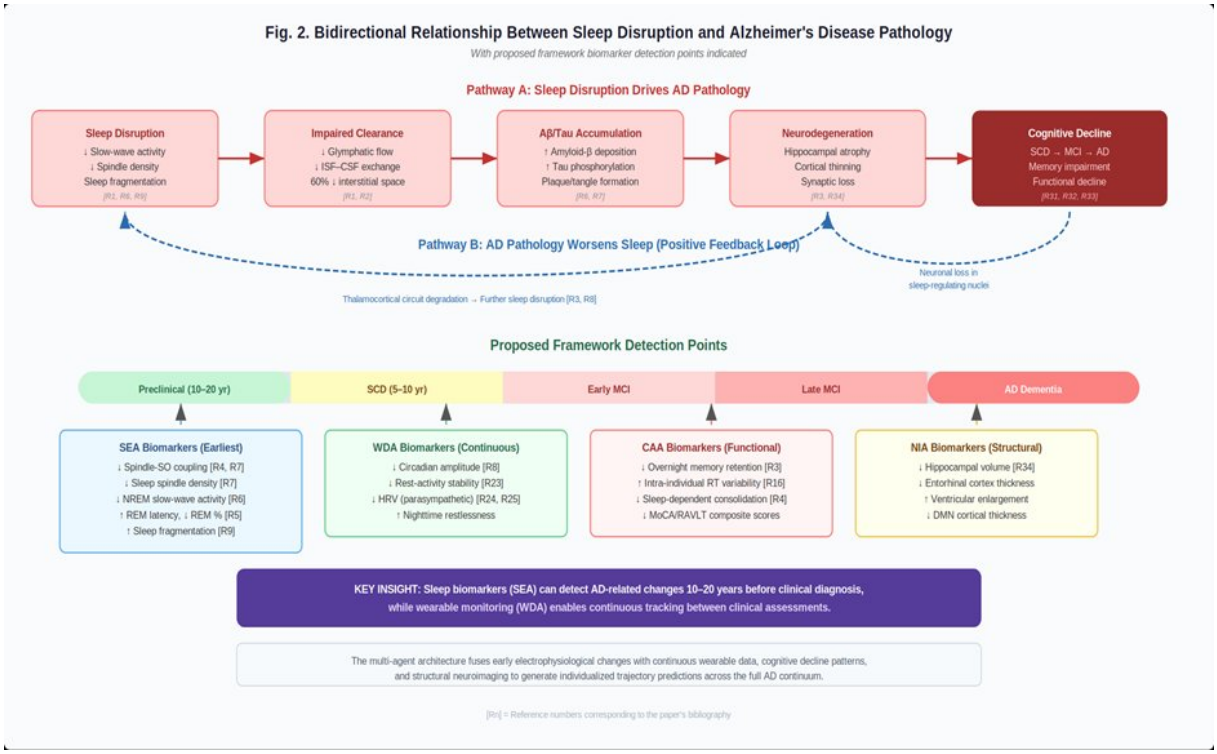


Figure 2: Bidirectional relationship between sleep disruption and Alzheimer's disease pathology, with proposed framework biomarker detection points mapped to the disease continuum. Pathway A shows how sleep disruption drives AD pathology through impaired glymphatic clearance. Pathway B shows how neurodegeneration worsens sleep through thalamocortical circuit degradation.

2.5. Hybrid deployment architecture

The framework operates within a three-tier hybrid deployment model (Fig. 3) balancing data quality with practical feasibility.

Home Monitoring Tier: Participants wear a consumer wrist-based device continuously, providing uninterrupted actigraphy, PPG, and temperature data. Smartphone cognitive micro-assessments are administered 2–3 times weekly (∼5 minutes each). Data is transmitted via encrypted Bluetooth/Wi-Fi. This tier prioritizes minimal participant burden and ecological validity [22, 23].

Clinical Assessment Tier: Participants attend a clinical facility at 6-month intervals for attended PSG and neuropsychological assessment, and at 12-month intervals for structural MRI. These sessions provide research-grade data anchoring the continuous home monitoring.

Cloud Processing Tier: All computation occurs on HIPAA/GDPR-compliant cloud infrastructure. The cloud hosts trained agent models, performs real-time wearable pro-

cessing, batch clinical data processing, and generates updated trajectory predictions. A clinical dashboard API provides longitudinal prediction visualization.

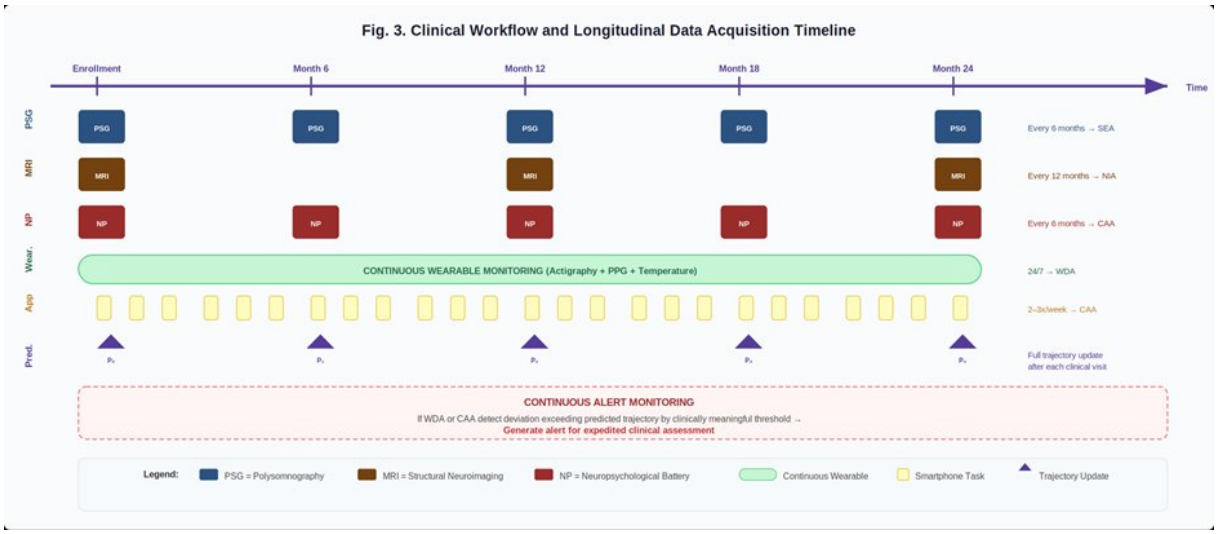


Figure 3: Clinical workflow and longitudinal data acquisition timeline showing acquisition schedules across all modalities and trajectory prediction updates.

2.6. Clinical workflow integration

Upon enrollment, baseline clinical data (PSG, neuropsychological battery, MRI) establishes an initial trajectory estimate. Between visits, the WDA and CAA continuously update feature representations. An alert system monitors for significant deviations from predicted trajectory; if detected decline exceeds the predicted rate by a clinically meaningful threshold, an alert recommends expedited clinical assessment.

Following each clinical visit, the COA recomputes the full multimodal trajectory incorporating updated clinical-grade data. Clinicians receive longitudinal reports showing trajectory evolution, deviation from prior predictions, and contributing biomarkers. Figure 4 illustrates example trajectory outputs across four risk staging categories.

3. Results

3.1. Comparative analysis

Table 3 compares the proposed framework against existing approaches to AD prediction, highlighting key differentiators in modality integration, prediction type, sleep biomarker inclusion, and deployment model. No existing approach combines all four

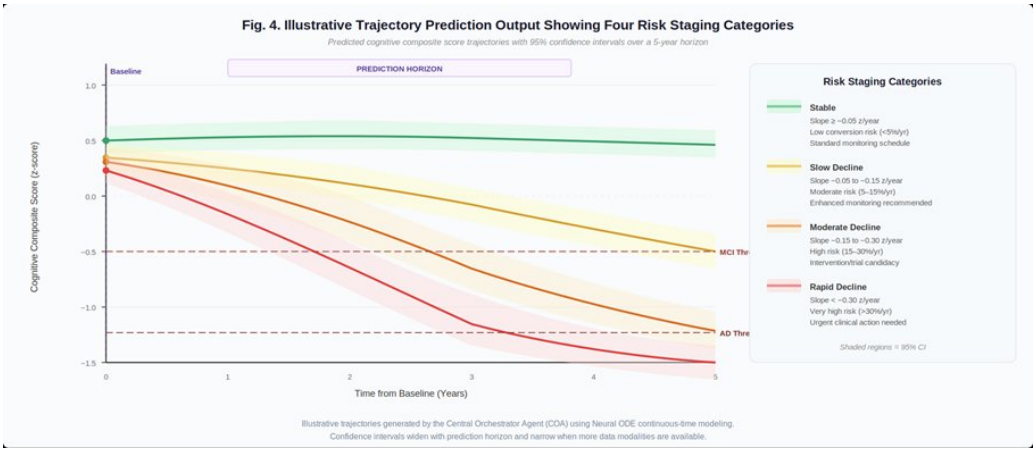


Figure 4: Illustrative trajectory prediction output showing four risk staging categories (stable, slow decline, moderate decline, rapid decline) with 95% confidence intervals over a 5-year prediction horizon.

modalities with sleep biomarkers as the central feature set within a trajectory prediction framework. The closest comparator—multimodal deep learning by Qiu et al. [17]—does not incorporate sleep data and outputs categorical labels rather than temporal trajectories.

Table 3: Comparative analysis of the proposed framework against existing approaches.

Approach	**Modalities**	**Prediction**	**Sleep**	**Deploy.**	**Output**	**Population**
CNN-MRI [15, 16]	MRI only	Binary/ternary	No	Clinical	Class label	AD vs. NC
EEG-ML [18, 19]	EEG only	Binary	Partial	Clinical	Class label	AD vs. NC
Multimodal DL [17]	MRI+NP+Demo	Multi-class	No	Clinical	Class label	AD spectrum
Digital biomarkers [22]	Wearable+App	MCI detection	Indirect	Home	Risk score	Community
MAS-Epilepsy [25]	EEG+Clinical	Detection	No	Hybrid	Alert	Epilepsy
Proposed framework	**EEG/PSG+ Wear.+NP+MRI**	**Trajectory**	**Central**	**Hybrid**	**Curve+ Stage**	**SCD/MCI**

3.2. Feasibility and implementation analysis

3.2.1. Data availability

The proposed framework can leverage several established datasets for initial development and validation: **ADNI** (Alzheimer's Disease Neuroimaging Initiative) [35, 39] provides over 2,000 participants with longitudinal MRI, PET, CSF biomarkers, and neuropsychological assessments; **Sleep Heart Health Study (SHHS)** provides over 6,000 participants with full PSG recordings and longitudinal health outcomes; **UK Biobank** provides 500,000 participants with wearable accelerometer data, structural MRI, and cognitive assessments; and **National Sleep Research Resource (NSRR)** provides aggregated PSG data from multiple studies, standardized for research use.

3.2.2. Privacy and regulatory considerations

The framework design incorporates privacy-preserving mechanisms: all data transmission uses TLS 1.3 encryption; data at rest uses AES-256; cloud infrastructure is designed for HIPAA/GDPR compliance. Federated learning approaches [42, 43] may enable multi-site model training without centralizing raw patient data—Sheller et al. [43] demonstrated 99% of centralized performance across 10 institutions.

3.3. Proof-of-concept evaluation

To demonstrate architectural feasibility and validate the multi-agent fusion design, we implemented a complete proof-of-concept (POC) of the proposed framework using synthetic data calibrated to published biomarker distributions.

3.3.1. Synthetic data generation

We generated longitudinal multimodal data for $N = 500$ subjects across $T = 10$ timepoints (6-month intervals, 5-year horizon), distributed across four diagnostic groups: cognitively normal (CN; $n = 150$), subjective cognitive decline (SCD; $n = 125$), mild cognitive impairment (MCI; $n = 150$), and early AD ($n = 75$). For each group, biomarker trajectories were calibrated to published effect sizes: spindle–SO coupling strength ($r = -0.46$ to -0.54 with CSF tau [8, 11]); REM sleep reduction (HR 0.91 per 1% decrease [12]); sleep fragmentation (HR 1.50 [13]); parasympathetic HRV reduction ($g = -0.36$ [33]); and hippocampal atrophy rates (3–5% annual in MCI, 16.5%/year for MCI$\rightarrow$AD conversion [35]). Group-specific decline rates, baseline distributions, and

inter-feature correlations were derived from these published parameters. All features were z-scored within modality.

3.3.2. Implementation details

All five agents were implemented in PyTorch with architectures matching the paper's specifications: the SEA uses Conv1D followed by bidirectional LSTM (8,176 parameters); the WDA uses a three-layer dilated TCN (4,792 parameters); the CAA uses a two-hidden-layer MLP (936 parameters; MLP substituted for XGBoost to enable end-to-end differentiability); the NIA uses a single-hidden-layer MLP (476 parameters; MLP substituted for 3D CNN as synthetic data consists of extracted volumetric features). The COA implements four-head cross-modal attention with learnable temporal encodings, followed by a Neural ODE trajectory predictor using the `torchdiffeq` library's `dopri5` adaptive solver (34,850 parameters). Total framework: 49,230 trainable parameters.

Training used AdamW optimization (lr $= 10^{-3}$, weight decay $= 10^{-5}$) with cosine annealing over 150 epochs and early stopping (patience = 20). The loss function combined MSE for trajectory fitting, heteroscedastic negative log-likelihood for uncertainty estimation, monotonicity regularization (penalizing biologically implausible cognitive score increases), and second-order smoothness regularization. Data was split 70/15/15 for train/validation/test. Training converged at epoch 133.

3.3.3. Trajectory prediction performance

Table 4 summarizes the test set performance. The framework achieved overall RMSE of 0.037 (on a normalized 0–1 scale) and MAE of 0.029, with trajectory-level Pearson correlation of 0.56. Per-group analysis reveals highest accuracy for CN (RMSE = 0.030) and MCI (0.033), with larger error for AD (0.056), consistent with greater phenotypic heterogeneity in the AD group.

Table 4: Proof-of-concept test set results ($N_{\text{test}} = 75$).

Metric	**Overall**	**Per Group (RMSE)**
RMSE	0.037	CN: 0.030
MAE	0.029	SCD: 0.039
Trajectory Correlation	0.563	MCI: 0.033
Risk Staging Accuracy	89.3%	AD: 0.056

3.3.4. *Risk staging classification*

The clinically critical risk staging task achieved 89.3% accuracy across four categories (stable, slow decline, moderate decline, rapid decline). Figure 5 shows the confusion matrix: misclassifications were predominantly adjacent-tier (e.g., moderate↔slow decline), with zero catastrophic errors (e.g., no rapid decline cases classified as stable). Crisis-sensitive staging—correct classification of rapid decline cases—achieved 87.5% sensitivity (14/16 correct), with the 2 misclassified cases assigned to the adjacent moderate decline tier.

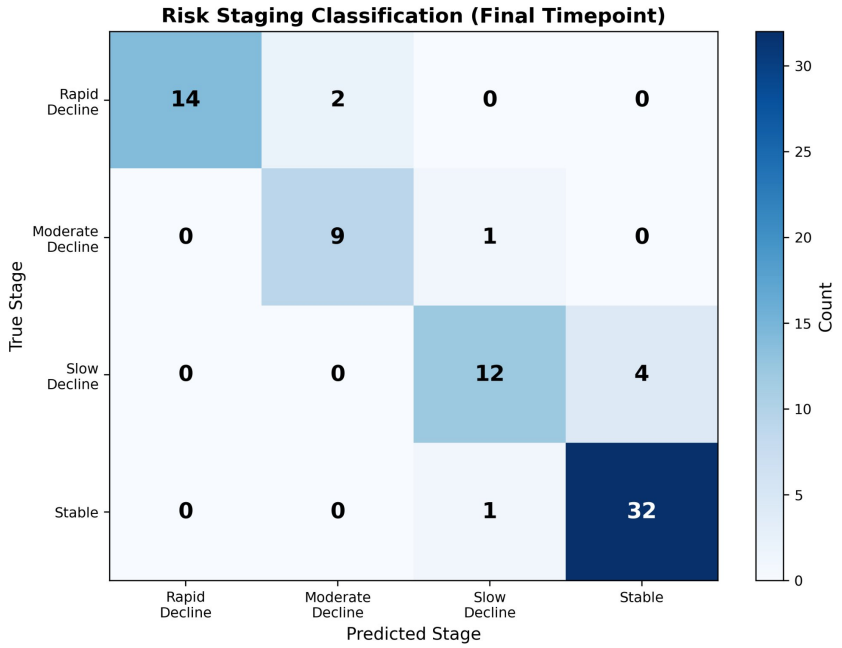


Figure 5: Risk staging confusion matrix at the final timepoint (year 4.5). Misclassifications cluster along the diagonal, indicating adjacent-tier errors rather than catastrophic misclassification.

3.3.5. *Group-level trajectory prediction*

Figure 6 displays predicted versus actual cognitive trajectories stratified by diagnostic group. The framework successfully captures group-level decline dynamics: near-flat trajectories for CN (mean decline ∼2 points over 5 years), gradual decline for SCD (∼5 points), clinically significant decline for MCI (∼12 points), and rapid decline for AD (∼20 points). Individual trajectory predictions (dashed lines) track actual trajectories (solid lines) with tightening confidence intervals for lower-variance groups (CN, SCD) and appropriately wider intervals for higher-variance groups (AD).

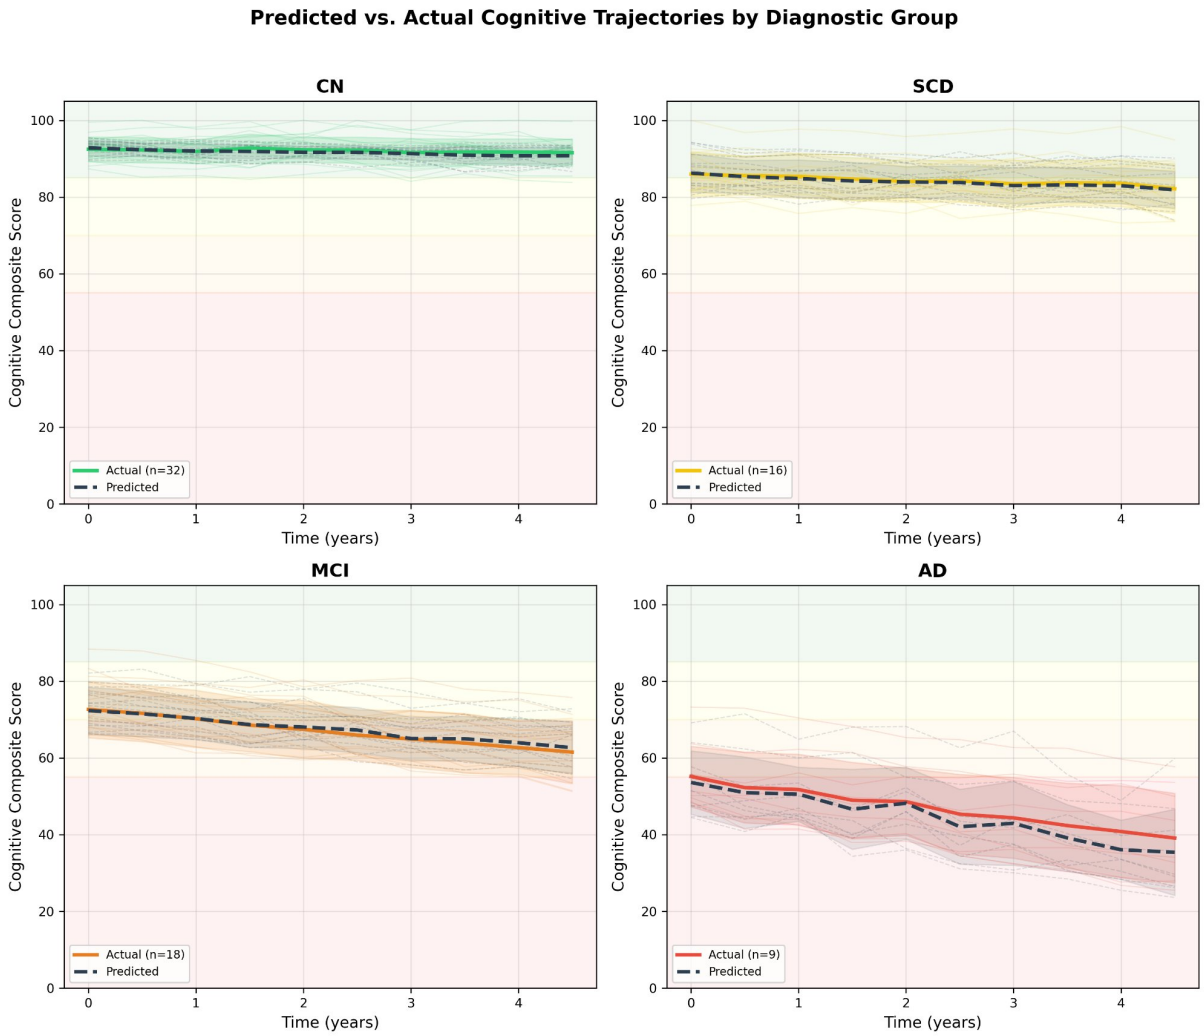


Figure 6: Predicted (dashed) versus actual (solid) cognitive composite score trajectories by diagnostic group, showing group means ± SD with individual traces. Color-coded staging bands: green (stable), yellow (slow decline), orange (moderate), red (rapid).

3.3.6. Cross-modal attention analysis

Figure 7 presents the cross-modal attention weights learned by the COA, averaged across subjects within each diagnostic group. The Cognitive Assessment Agent (CAA) consistently receives the highest attention weight across all groups (0.46–0.72), reflecting its direct correspondence with the prediction target. Notably, the Neuroimaging Agent (NIA) contribution increases progressively from CN to AD groups, consistent with the biological expectation that structural changes become more salient as neurodegeneration advances. The Sleep Electrophysiology Agent (SEA) maintains moderate attention across all groups, supporting the framework's premise that sleep biomarkers provide predictive information throughout the disease continuum.

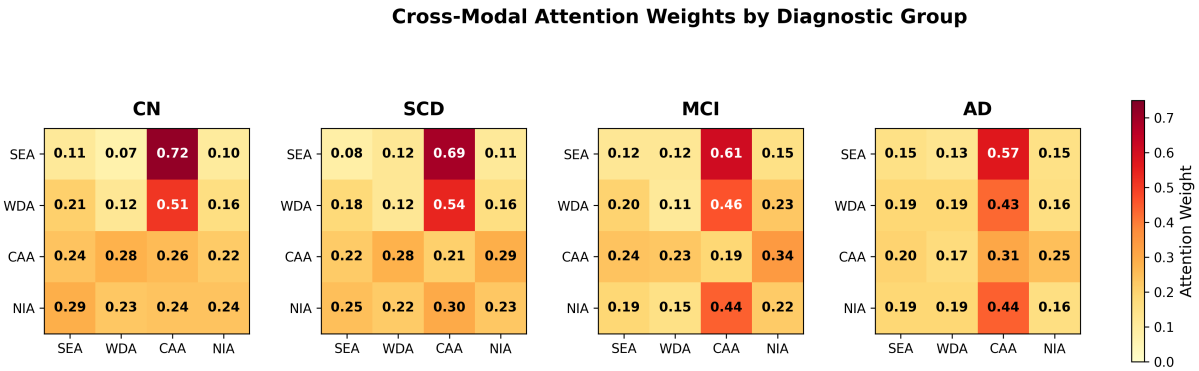


Figure 7: Cross-modal attention weight matrices by diagnostic group. Color intensity indicates learned attention between modality pairs. The CAA column shows consistently high weights, while NIA contribution increases from CN to AD.

3.3.7. Modality contribution dynamics

Figure 8 illustrates how modality contributions evolve over the 5-year prediction horizon. In the AD group, the NIA contribution grows from ∼20% at baseline to ∼35% at year 4.5, paralleling progressive structural atrophy. The CAA contribution shows an inverse pattern—slightly decreasing over time—consistent with floor effects as cognitive scores saturate at lower bounds. These dynamic attention shifts suggest the framework learns clinically meaningful modality weighting rather than static feature importance.

3.4. Validation roadmap

We propose a phased validation strategy:

Phase 1: Retrospective Feasibility (Months 1–12). Train individual agents on available public datasets (ADNI, SHHS, UK Biobank); evaluate single-agent predictive performance with cross-validation; develop fusion architecture with synthetic multimodal data.

Phase 2: Prospective Pilot (Months 12–24). Recruit 100 SCD/MCI participants from memory clinics; deploy hybrid monitoring protocol for 12 months; validate trajectory predictions against observed outcomes.

Phase 3: Multi-Site Validation (Months 24–48). Expand to 3–5 clinical sites with diverse demographics; implement federated learning for privacy-preserving training; conduct head-to-head comparison with standard clinical assessment.

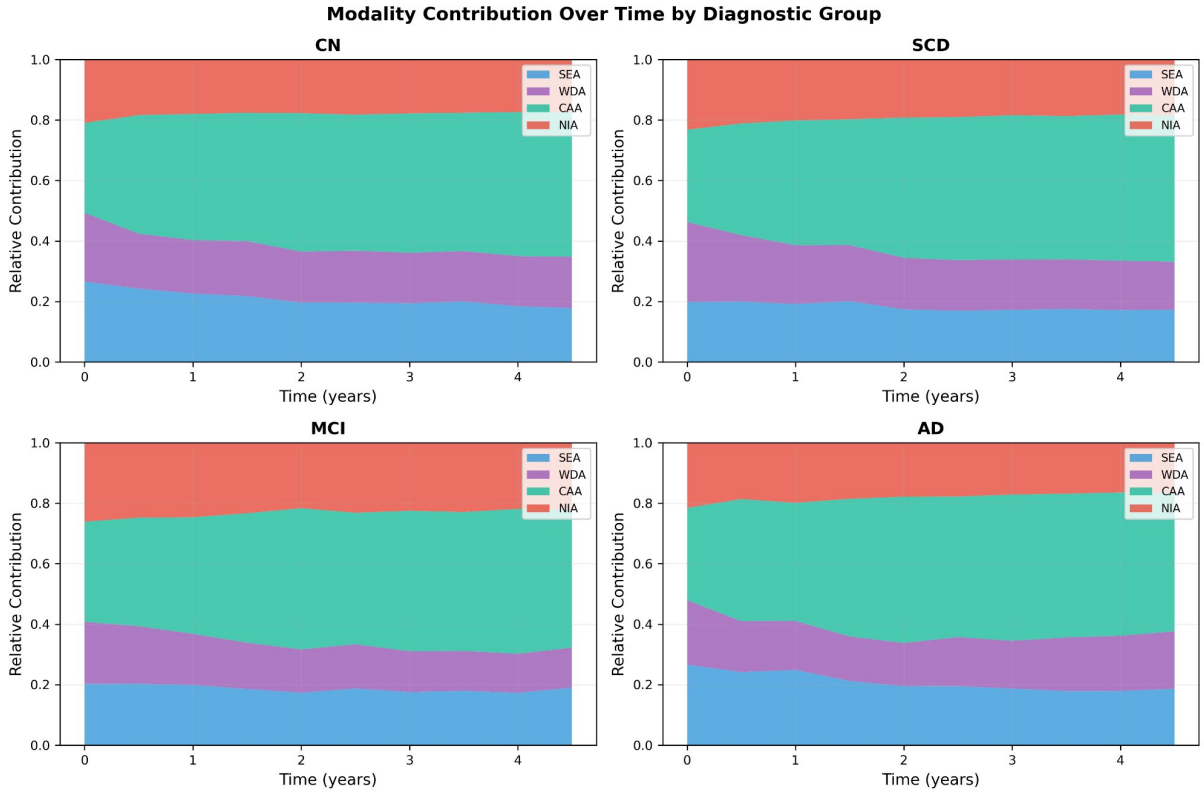


Figure 8: Modality contribution over the 5-year prediction horizon by diagnostic group, showing relative importance of each agent as disease progresses.

Phase 4: Clinical Integration (Months 48–60). Develop clinical decision support interface; conduct clinician usability studies; pursue regulatory pathway assessment (FDA breakthrough device).

4. Discussion

The proposed framework represents a conceptual and architectural advance in several dimensions. First, by centering sleep-dependent memory consolidation as the primary predictive mechanism, it is grounded in the known neurobiology of AD—the bidirectional relationship between sleep disruption and amyloid/tau pathology [6–9] provides a mechanistic rationale for why sleep biomarkers may detect AD-related changes earlier than conventional clinical markers.

Second, the multi-agent architecture offers practical advantages over monolithic models. Modular design allows individual agents to be updated independently as algorithms or sensor technologies advance. The system can operate with varying data availability, from minimal wearable-only monitoring to full four-modality assessment,

with the attention-based fusion mechanism adapting dynamically. The proof-of-concept evaluation validates this design: the learned attention weights demonstrate clinically meaningful modality prioritization that shifts with disease stage.

Third, the trajectory prediction approach provides more clinically useful output than binary classification. The Neural ODE architecture [40] models continuous-time dynamics from irregularly sampled observations [41], enabling precise patient stratification and optimal intervention timing. The POC staging accuracy of 89.3% with adjacent-tier error patterns suggests this approach can generate actionable clinical outputs.

The POC demonstrates architectural feasibility with clinically interpretable outputs. However, several limitations of the POC must be acknowledged. Synthetic data calibrated to published effect sizes validates architectural design but cannot establish clinical predictive validity—the model partially recovers patterns embedded in the data generation process. The moderate trajectory-level correlation (0.56) reflects both stochastic noise in the synthetic data and the inherent difficulty of continuous trajectory prediction versus discrete staging. Agent architectures are simplified relative to full specifications (MLP replaces XGBoost and 3D CNN for end-to-end differentiability). The synthetic data lacks complex confounding structure of real clinical populations (comorbidities, medication effects, dropout patterns). Validation on real datasets following the proposed roadmap is essential to establish clinical utility.

4.1. Limitations

Five broader limitations require acknowledgment: (1) while the POC demonstrates architectural feasibility, agent performance and trajectory accuracy on real clinical data must be established through prospective studies following the proposed validation roadmap; (2) the reliance on periodic clinical PSG represents a practical barrier that advances in home-based EEG devices may eventually mitigate; (3) five concurrent agents with Neural ODE trajectory modeling require careful evaluation for real-world scalability, though the POC implementation (49,230 parameters, sub-second inference) suggests feasibility; (4) sleep architecture varies with age, sex, ethnicity, and comorbid conditions, requiring diverse training data to avoid algorithmic bias; (5) generalizability across healthcare systems with varying resources warrants investigation.

4.2. Future directions

Building upon the proposed framework, several extensions merit investigation: (1) a fifth agent incorporating plasma biomarkers (p-tau217 [4, 5]) as point-of-care assays become available; (2) closed-loop interventions (e.g., auditory stimulation during slow-wave sleep) based on detected consolidation deficits; (3) disease-specific adaptation for Parkinson's disease, Lewy body dementia, and other neurodegenerative conditions with sleep-related prodromal markers; and (4) integration of emerging consumer EEG headbands for continuous sleep electrophysiology monitoring.

5. Conclusions

This paper presents a multi-agent AI framework for predicting Alzheimer's disease trajectory through sleep-dependent memory consolidation biomarkers. By integrating EEG/PSG, wearable sensor data, cognitive assessments, and structural neuroimaging through specialized AI agents coordinated by an attention-based central orchestrator with Neural ODE trajectory modeling, the framework addresses critical limitations of existing single-modality, binary-classification approaches. A proof-of-concept implementation on calibrated synthetic data demonstrates 89.3% risk staging accuracy, clinically meaningful attention-based modality weighting, and interpretable trajectory predictions across diagnostic groups. The hybrid deployment model balances clinical-grade data quality with continuous ecological monitoring, and the trajectory prediction output provides clinicians with individualized longitudinal risk profiles stratified into actionable staging categories. While validation on real clinical data is essential, this framework establishes a theoretically grounded, architecturally validated foundation for the next generation of sleep-based AD prediction systems.

Declarations

Ethics approval and consent to participate

Not applicable. This paper presents a conceptual framework and proof-of-concept evaluation on synthetic data; no human participants were involved.

Consent for publication

Not applicable.

Availability of data and materials

The proof-of-concept source code and synthetic data generation scripts are available from the corresponding author upon reasonable request.

Competing interests

The authors declare that they have no competing interests.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Authors' contributions

Bahaa Masry: Conceptualization, Methodology, Investigation, Writing – Original Draft, Writing - Review & Editing.

Anas Shahin: Conceptualization, Methodology, Software, Formal Analysis, Visualization, Writing – Review & Editing.

Acknowledgments

Not applicable.

References

[1] World Health Organization. *Global Status Report on the Public Health Response to Dementia*. WHO, Geneva, 2021.

[2] Alzheimer's Association. 2024 Alzheimer's disease facts and figures. *Alzheimers Dement*, 20(5):3708–3821, 2024.

[3] C. R. Jack Jr, D. A. Bennett, K. Blennow, M. C. Carrillo, B. Dunn, S. B. Haeberlein, et al. NIA-AA research framework: toward a biological definition of Alzheimer's disease. *Alzheimers Dement*, 14(4):535–562, 2018.

[4] S. Palmqvist, S. Janelidze, Y. T. Quiroz, H. Zetterberg, F. Lopera, E. Stomrud, et al. Discriminative accuracy of plasma phospho-tau217 for Alzheimer disease vs other neurodegenerative disorders. *JAMA*, 324(8):772–781, 2020.

[5] N. J. Ashton, W. S. Brum, G. Di Molfetta, A. L. Benedet, B. Arslan, E. Jonaitis, et al. Diagnostic accuracy of a plasma phosphorylated tau 217 immunoassay for Alzheimer disease pathology. *JAMA Neurol*, 81(3):255–263, 2024.

[6] J. J. Iliff, M. Wang, Y. Liao, B. A. Plogg, W. Peng, G. A. Gundersen, et al. A paravascular pathway facilitates CSF flow through the brain parenchyma and the clearance of interstitial solutes, including amyloid β. *Sci Transl Med*, 4(147):147ra111, 2012.

[7] L. Xie, H. Kang, Q. Xu, M. J. Chen, Y. Liao, M. Thiyagarajan, et al. Sleep drives metabolite clearance from the adult brain. *Science*, 342(6156):373–377, 2013.

[8] B. P. Lucey, A. McCullough, E. C. Landsness, C. D. Toedebusch, J. S. McLeland, A. M. Zaza, et al. Reduced non-rapid eye movement sleep is associated with tau pathology in early Alzheimer's disease. *Sci Transl Med*, 11(474):eaau6550, 2019.

[9] J. R. Winer, B. A. Mander, R. F. Helfrich, A. Maass, T. M. Harrison, S. L. Baker, et al. Sleep as a potential biomarker of tau and β-amyloid burden in the human brain. *J Neurosci*, 39(32):6315–6324, 2019.

[10] B. A. Mander, V. Rao, B. Lu, J. M. Saletin, J. R. Lindquist, S. Ancoli-Israel, et al. Prefrontal atrophy, disrupted NREM slow waves and impaired hippocampal-dependent memory in aging. *Nat Neurosci*, 16(3):357–364, 2013.

[11] R. F. Helfrich, B. A. Mander, W. J. Jagust, R. T. Knight, and M. P. Walker. Old brains come uncoupled in sleep: slow wave–spindle synchrony, brain atrophy, and forgetting. *Neuron*, 97(1):221–230.e4, 2018.

[12] M. P. Pase, J. J. Himali, N. A. Grima, A. S. Beiser, C. L. Satizabal, H. J. Aparicio, et al. Sleep architecture and the risk of incident dementia in the community. *Neurology*, 89(12):1244–1250, 2017.

[13] A. S. P. Lim, M. Kowgier, L. Yu, A. S. Buchman, and D. A. Bennett. Sleep fragmentation and the risk of incident Alzheimer's disease and cognitive decline in older persons. *Sleep*, 36(7):1027–1032, 2013.

[14] E. S. Musiek, M. Bhimasani, M. A. Zangrilli, J. C. Morris, D. M. Holtzman, and Y.-E. S. Ju. Circadian rest-activity pattern changes in aging and preclinical Alzheimer disease. *JAMA Neurol*, 75(5):582–590, 2018.

[15] J. Wen, E. Thibeau-Sutre, M. Diaz-Melo, J. Samper-González, A. Routier, S. Bottani, et al. Convolutional neural networks for classification of Alzheimer's disease: overview and reproducible evaluation. *Med Image Anal*, 63:101694, 2020.

[16] T. O. Frizzell, M. Glashutter, C. C. Liu, A. Zeng, D. Pan, S. G. Hajra, et al. Artificial intelligence in brain MRI analysis of Alzheimer's disease over the past 12 years: a systematic review. *Ageing Res Rev*, 77:101614, 2022.

[17] S. Qiu, M. I. Miller, P. S. Joshi, J. C. Lee, C. Xue, Y. Ni, et al. Multimodal deep learning for Alzheimer's disease dementia assessment. *Nat Commun*, 13(1):3404, 2022.

[18] K. D. Tzimourta, V. Christou, A. T. Tzallas, N. Giannakeas, L. G. Astrakas, P. Angelidis, et al. Machine learning algorithms and statistical approaches for Alzheimer's disease analysis based on resting-state EEG recordings: a systematic review. *Int J Neural Syst*, 31(5):2130002, 2021.

[19] E. Pérez-Valero, M. A. López-Gordo, C. Morillas, F. Pelayo, and M. A. Vaquero-Blasco. A review of automated techniques for assisting the early detection of Alzheimer's disease with a focus on EEG. *J Alzheimers Dis*, 80(4):1363–1376, 2021.

[20] Y. Li, M. E. H. Daho, P.-H. Conze, R. Zeghlache, H. Le Boîté, R. Tadayoni, et al. A review of deep learning-based information fusion techniques for multimodal medical image classification. *Comput Biol Med*, 177:108635, 2024.

[21] P. Dagum. Digital biomarkers of cognitive function. *npj Digit Med*, 1:10, 2018.

[22] P. M. Butler, J. Yang, R. Brown, M. Hobbs, A. Becker, et al. Smartwatch- and smartphone-based remote assessment of brain health and detection of mild cognitive impairment. *Nat Med*, 31(3):829–839, 2025.

[23] L. C. Kourtis, O. B. Regele, J. M. Wright, and G. B. Jones. Digital biomarkers for Alzheimer's disease: the mobile/wearable devices opportunity. *npj Digit Med*, 2: 9, 2019.

[24] J. Fdez-Olivares, E. Onaindia, L. Castillo, J. Jordán, and J. Cózar. Personalized conciliation of clinical guidelines for comorbid patients through multi-agent planning. *Artif Intell Med*, 96:167–186, 2019.

[25] C. M. Bertoncelli, S. Costantini, F. Persia, D. Bertoncelli, and D. D'Auria. PredictMed-epilepsy: a multi-agent based system for epilepsy detection and prediction in neuropediatrics. *Comput Methods Programs Biomed*, 236:107548, 2023.

[26] T. Ivaşcu and V. Negru. Activity-aware vital sign monitoring based on a multi-agent architecture. *Sensors*, 21(12):4181, 2021.

[27] Thanveer Shaik, Xiaohui Tao, Lin Li, Haoran Xie, Hui Dai, Fei Zhao, Zilu Liang, Guangyan Huang, Meng Han, Wei Xiang, Zenghui Wang, and Li Liu. Ai-driven multi-agent reinforcement learning framework for real-time monitoring of physiological signals in stress and depression contexts. *Brain Informatics*, 12 (1):14, 2025. doi: 10.1186/s40708-025-00262-1. URL `https://doi.org/10.1186/s40708-025-00262-1`.

[28] F. Jessen, R. E. Amariglio, M. van Boxtel, M. Breteler, M. Ceccaldi, G. Chételat, et al. A conceptual framework for research on subjective cognitive decline in preclinical Alzheimer's disease. *Alzheimers Dement*, 10(6):844–852, 2014.

[29] R. C. Petersen. Mild cognitive impairment as a diagnostic entity. *J Intern Med*, 256 (3):183–194, 2004.

[30] R. C. Petersen, B. Caracciolo, C. Brayne, S. Gauthier, V. Jelic, and L. Fratiglioni. Mild cognitive impairment: a concept in evolution. *J Intern Med*, 275(3):214–228, 2014.

[31] A. R. Hevner, S. T. March, J. Park, and S. Ram. Design science in information systems research. *MIS Q*, 28(1):75–105, 2004.

[32] E. J. Wams, G. K. Wilcock, R. G. Foster, and K. Wulff. Sleep-wake patterns and cognition of older adults with amnestic mild cognitive impairment (aMCI). *Curr Alzheimer Res*, 14(10):1030–1041, 2017.

[33] K. Y. Liu, T. Elliott, M. Knowles, and R. Howard. Heart rate variability in relation to cognition and behavior in neurodegenerative diseases: a systematic review and meta-analysis. *Ageing Res Rev*, 73:101539, 2022.

[34] Y.-C. Cheng, Y.-C. Huang, and W.-L. Huang. Heart rate variability in patients with dementia or neurocognitive disorders: a systematic review and meta-analysis. *Aust N Z J Psychiatry*, 56(1):16–27, 2022.

[35] R. C. Petersen, P. S. Aisen, L. A. Beckett, M. C. Donohue, A. C. Gamst, D. J. Harvey, et al. Alzheimer's Disease Neuroimaging Initiative (ADNI): clinical characterization. *Neurology*, 74(3):201–209, 2010.

[36] S. Bai, J. Z. Kolter, and V. Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv:1803.01271, 2018.

[37] X. Hu, S. Yuan, F. Xu, Y. Leng, K. Yuan, and Q. Yuan. Scalp EEG classification using deep Bi-LSTM network for seizure detection. *Comput Biol Med*, 124:103919, 2020.

[38] T. Chen and C. Guestrin. XGBoost: a scalable tree boosting system. In *Proc 22nd ACM SIGKDD*, pages 785–794, 2016.

[39] S. G. Mueller, M. W. Weiner, L. J. Thal, R. C. Petersen, C. R. Jack, W. Jagust, et al. Ways toward an early diagnosis in Alzheimer's disease: the ADNI. *Alzheimers Dement*, 1(1):55–66, 2005.

[40] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In *Proc NeurIPS*, pages 6571–6583, 2018.

[41] P. Wendland, C. Birkenbihl, M. Gomez-Freixa, M. Sood, M. Kschischo, and H. Fröhlich. Generation of realistic synthetic data using multimodal neural ordinary differential equations. *npj Digit Med*, 5(1):122, 2022.

[42] N. Rieke, J. Hancox, W. Li, F. Milletarì, H. R. Roth, S. Albarqouni, et al. The future of digital health with federated learning. *npj Digit Med*, 3:119, 2020.

[43] M. J. Sheller, B. Edwards, G. A. Reina, J. Martin, S. Pati, A. Kotrotsou, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci Rep*, 10:12598, 2020.