

Artificial Intelligence for Tracking Cognitive Deviation in Aging Populations: A Comprehensive Review of Techniques, Challenges, and Ethical Concerns

Dipali Prakash Patil^{1,*}, Chandrani Singh²

Abstract

Population aging is accelerating worldwide, and with it the burden of cognitive health conditions such as mild cognitive impairment (MCI), Alzheimer's disease (AD), and dementia. Detecting and monitoring cognitive change early is central to timely intervention, yet conventional diagnostic tools often miss the subtle signals that appear before overt symptoms. Artificial intelligence (AI) has emerged as a promising complement to clinical assessment because it can work through high-dimensional data and surface patterns that standard analysis tends to overlook. This review synthesizes recent research on AI methods – machine learning (ML), deep learning (DL), hybrid architectures, and reinforcement learning – applied to cognitive-deviation tracking in older adults across neuroimaging, electrophysiology, speech and language, wearable-sensor, and cognitive-assessment data. Following PRISMA 2020 guidance, we organize the evidence thematically around model families, data modalities, and application focus. Convolutional neural networks and transformer-based architectures show consistently strong performance, particularly in multimodal settings, though reported accuracy varies widely with dataset composition and staging definitions. Persistent limitations include a lack of demographic diversity in training data, weak sensitivity to preclinical change, opacity of many models, and unresolved questions of privacy, consent, and explainability. We conclude that translating these methods into practice will depend on inclusive datasets, improved early-detection methods, explainable AI, and privacy-preserving data-sharing infrastructure such as federated learning. The review is intended as a practical reference for researchers, clinicians, and policymakers working at the intersection of AI and healthy cognitive aging.

Keywords: Cognitive deviation, artificial intelligence, aging populations, early detection, explainable AI, deep learning, dementia

INTRODUCTION

By the World Health Organization's projection, one in six people worldwide will be aged 60 or older

*Author for Correspondence

Dipali Prakash Patil

E-mail: dipalipatil.siom@sinhgad.edu

¹Assistant Professor, Department of Master Computer Applications, STES Sinhgad Institute of Management, Vadgaon (Bk), Pune, Maharashtra, India

²Director, Department of Master Computer Applications, STES Sinhgad Institute of Management, Vadgaon (Bk), Pune, Maharashtra, India

Received Date: May; 13, 2026

Accepted Date: June 08, 2026

Published Date: August 31, 2026

Citation: Dipali Prakash Patil, Chandrani Singh. Artificial Intelligence for Tracking Cognitive Deviation in Aging Populations: A Comprehensive Review of Techniques, Challenges, and Ethical Concerns. Research & Reviews: A Journal of Neuroscience. 2026; 16(2): 27–37p.

by 2030, a share that continues to climb through mid-century [1]. That demographic shift brings a parallel rise in age-related neurocognitive conditions such as MCI, dementia, and AD. Because these conditions tend to progress quietly, with few obvious early signs, conventional diagnostic tools often fail to catch them until later stages. A late diagnosis is costly: it narrows the window for effective treatment, lowers quality of life, and adds strain to caregivers, families, and health systems already under pressure. Good care for older adults, therefore, depends heavily on catching cognitive decline early and accurately. Standard screening tools – neuropsychological tests, structured interviews, and brain imaging – were not built for this: they are sensitive mainly to

more advanced change, take time and trained clinicians to administer, and scale poorly to population-level screening [2].

This is where AI enters the picture, not as a replacement for clinical judgment but as a complementary set of tools. ML and DL methods have shown real promise in working through the kind of high-dimensional data in which early cognitive change tends to hide [2–4]. They can draw simultaneously on structural and functional neuroimaging, wearable-sensor streams, speech and language patterns, and digital cognitive tests, looking for signals a single clinical snapshot would miss. Rather than the periodic check-ins that define most current care, AI-based monitoring can in principle be more continuous, less invasive, and more scalable. Natural language processing (NLP) methods, for instance, can extract linguistic markers of cognitive impairment from spontaneous speech [5, 6]; sensor-based methods can do something similar for gait and movement [7]. Bringing AI into cognitive health also raises real questions, however – around data privacy and consent, algorithmic fairness, model interpretability, and how any of this is embedded into clinical workflows – that must be addressed before AI’s potential in this space can be realized.

The contribution of this review is threefold. First, it consolidates recent AI methods for cognitive-deviation tracking across the full range of data modalities, rather than focusing on a single imaging type as many prior surveys do. Second, it pairs a technical synthesis with an explicit treatment of fairness, generalizability, and ethics, which remain under-examined in the modeling literature. Third, it distills the findings into a conceptual pipeline and a set of concrete research priorities intended to be usable by both technical and clinical readers.

ROLE OF AI IN COGNITIVE HEALTH

What AI brings to cognitive health assessment, fundamentally, is the ability to pick out complex, non-linear patterns in high-dimensional data that standard analysis – or an unaided clinical eye – tends to miss. Convolutional neural networks (CNNs), recurrent neural networks (RNNs), and, more recently, transformer-based architectures have all been used for classification, progression prediction, and subtyping in AD and MCI [2, 4, 8]. Paired with imaging modalities, such as MRI, PET, or functional MRI, these models support both diagnosis and prognosis, sometimes picking up structural or functional signatures of disease well before they become clinically obvious [2, 9]. AI has also found its way into digital phenotyping, using wearables, voice recordings, and mobile cognitive tests to track behavioral change over time [7].

None of this is without friction. A central problem is generalizability: many AI models are trained on datasets more homogeneous than the populations they are eventually meant to serve, under-representing variation in age, ethnicity, socioeconomic status, and language. External validation studies confirm this concern directly – Bron et al. [10] showed that MRI-based classifiers can lose accuracy when moved from a research cohort to an independent clinical population, and that the drop is larger for the harder task of predicting MCI-to-AD conversion. The “black-box” character of many AI systems compounds the issue, raising legitimate questions about transparency, accountability, and interpretability that matter more, not less, when the population involved is older adults facing a vulnerable health condition [11].

Privacy and consent deserve equal weight. Collecting and analyzing sensitive cognitive and behavioral data calls for governance that holds up against frameworks such as the General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA). If AI is to deliver on its promise here, the next generation of systems will need inclusivity, fairness, and interpretability designed in from the start, in the spirit of the broadly shared AI-governance principles – transparency, justice, non-maleficence, responsibility, and privacy – identified across the global landscape of AI ethics guidelines [12, 13]. Getting there will require sustained collaboration among clinicians, computer scientists, ethicists, and policymakers [11, 14].

Objectives and Scope

This review pulls together recent research on applying AI to monitor cognitive deviation in older adults, organized around four objectives:

- Review recent AI-enabled research applying ML, DL, and NLP techniques to cognitive health assessment in older adults.
- Compare how different AI approaches perform across data modalities – neuroimaging, speech and behavioral data, and sensor-based inputs – for detecting cognitive decline.
- Assess the limitations of current AI applications with respect to data diversity, methodological constraints, and ethical considerations.
- Discuss future research directions centered on explainability, inclusivity, early detection, and responsible AI deployment in real clinical settings.

METHODOLOGY

Review Design and Protocol

This review follows the structure recommended by the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) 2020 guidelines [15], a widely adopted framework for study identification, screening, eligibility assessment, and reporting. Following PRISMA is intended to keep the process transparent and to reduce bias in how studies were selected and interpreted. The focus throughout is on recent AI applications for identifying and monitoring cognitive deviation in older adults, with particular attention to literature reflecting current model architectures and data modalities, given how quickly the field moves. Figure 1 summarizes the selection process.




Figure 1. PRISMA 2020 flow diagram of the study identification, screening, eligibility, and inclusion process. Record counts are indicative of the search and screening workflow described in Section 2.

Data Sources and Search Strategy

The literature search drew on three databases widely indexed for biomedical and technological research: PubMed (biomedical and clinical studies on cognitive disorders and AI in healthcare), IEEE Xplore (engineering and computer-science research on AI algorithms and wearable technologies), and Scopus (multidisciplinary peer-reviewed work spanning healthcare, neuroscience, and computational science). Search terms combined AI-related concepts (“artificial intelligence,” “machine learning,” “deep learning,” “neural networks”) with cognitive-condition terms (“cognitive impairment,” “mild cognitive impairment,” “dementia,” “Alzheimer’s”), task descriptors (“detection,” “monitoring,” “prediction”), and population terms (“aging,” “elderly”), joined with Boolean operators. Searching across disciplinarily distinct databases in this way is a recognized way of reducing publication bias and improving how completely a systematic review captures the relevant literature [16].

Inclusion and Exclusion Criteria

A clear set of inclusion and exclusion criteria, consistent with recommended systematic-review practice [15], guided study selection throughout, as summarized in Table 1.

Table 1. Inclusion and exclusion criteria applied during study selection.

Criterion type	Applied criteria
Inclusion – AI application	Uses AI, ML, DL, or a related computational technique to detect, monitor, or predict cognitive health deviation (dementia, MCI, or AD).
Inclusion – population & evidence	Older adults (generally 60+) or age-related cognitive impairment; reports empirical results or validation (accuracy, AUC, sensitivity, specificity); peer-reviewed.
Exclusion – study type	Animal-only or unvalidated-synthetic-data studies; editorials and purely conceptual papers (used only for background).
Exclusion – evidence gaps	No empirical validation or performance metrics; duplicates; non-peer-reviewed preprints.

Study Selection and Quality Appraisal

Screening followed the staged approach PRISMA recommends [15]: an initial title-and-abstract pass to remove clearly irrelevant records, a full-text review against the criteria in Table 1, and a quality-appraisal step built around established criteria – clarity of aims, methodological transparency, sample characteristics, and reporting of performance metrics. The intent of this staged process is a reasonably rigorous and reproducible synthesis, rather than an unfiltered aggregation of search hits. The synthesis that follows is organized around AI model family, data modality, and application focus (early diagnosis, progression monitoring, or risk prediction), which makes patterns – and gaps – that cut across individual studies easier to see.

COMPARATIVE REVIEW OF AI TECHNIQUES

AI has meaningfully expanded what is possible in detecting, tracking, and predicting cognitive deviation in older adults. ML and DL methods, in particular, are gaining ground because they can handle high-dimensional, multimodal data and pick up on non-linear patterns that classical statistical methods tend to miss [2, 3]. Figure 2 organizes the main model families discussed in this section.

Machine Learning Models

Conventional algorithms – support vector machines (SVMs), random forests, and gradient-boosted trees – such as XGBoost [17] – remain the workhorses for structured data: clinical assessments, demographic variables, and cognitive test scores. They tend to be computationally efficient, reasonably interpretable, and effective even on moderate-sized datasets, which makes them a practical fit for feature-based clinical decision support. Early MRI-biomarker work using SVMs demonstrated that structural neuroimaging features could separate AD and MCI patients from controls with useful accuracy [9], and interpretable gradient-boosting frameworks that pair XGBoost with SHAP have since been used to handle the class imbalance that arises as cognition progresses from normal to MCI to AD, while keeping the model’s reasoning legible to clinicians [18]. Their main weakness is that they

generally struggle with unstructured, high-dimensional input – raw neuroimaging, EEG, or speech – where spatial or temporal structure carries much of the diagnostic signal, and they usually depend on substantial manual feature engineering.




Figure 2. Taxonomy of AI model families applied to cognitive-deviation tracking, with representative algorithms under each family.

Deep Learning Models

Deep learning has become central to cognitive health research precisely because it handles the kind of high-dimensional, unstructured data – neuroimaging, EEG, speech, longitudinal behavioral measures – that classical models struggle with. CNNs, RNNs, and transformer-based architectures have all been applied to diagnosing and tracking early cognitive impairment [2, 4]. Comprehensive reviews report that CNN-based models dominate the recent MRI-based detection literature and that a large share of studies achieve accuracies above 90%, though this figure should be read cautiously: it depends heavily on dataset characteristics, staging definitions, and validation methodology, and comparing accuracy figures directly across studies is not always meaningful [2, 19]. EEG-based deep learning has followed a similar trajectory, with CNN and recurrent architectures increasingly applied to the automated detection of AD and MCI from electrophysiological signals [20, 21].

Transformer architectures, originally built for NLP [4], are increasingly repurposed for neuroimaging and time-series data. Their self-attention mechanism weighs relevant features across spatial or temporal context all at once, rather than step by step, which can improve sensitivity to subtle, early-stage change relative to purely convolutional or recurrent approaches [2]. The main trade-offs are well known: DL models are data-hungry, computationally demanding, and often difficult to interpret – a real barrier in healthcare settings where clinicians need to understand why a prediction was made [11]. These limitations motivate the interpretability tools discussed in Section 3.6.

Hybrid and Reinforcement Learning Models

As the patterns behind cognitive deviation grow more complex, hybrid architectures and reinforcement learning (RL) approaches are drawing more attention. Hybrid models combine two or

more paradigms – a tree-based model handling structured clinical variables, say, while a deep network processes unstructured imaging in parallel – to play to each method’s strengths; this can sharpen sensitivity and specificity, usually at the cost of more computation and less explainability. RL, where an agent learns from reward and penalty signals, is increasingly explored for adaptive cognitive training and assessment, such as digital exercises, that adjust difficulty in real time to a user’s performance. Both directions remain comparatively early in their translation to clinical settings, with validation and reproducibility the main open problems.

Data Modalities for Cognitive Tracking

How well AI performs depends heavily on the nature, quality, and combination of the underlying data. Because cognitive impairment manifests across biological, behavioral, and neurological domains, researchers increasingly draw on multiple modalities. Neuroimaging (MRI, PET, fMRI) remains the most heavily used, carrying rich spatial and volumetric information that CNN-based models can mine for subtle atrophy or functional change [2, 9, 19]. Electrophysiological signals (EEG, MEG) capture temporal brain dynamics and are commonly analyzed with recurrent or convolutional models [20, 21]. Structured cognitive scores – such as the MMSE and MoCA – are typically handled with classical ML [17]. Speech and language data are increasingly mined with NLP for linguistic markers – reduced lexical diversity, slower speech rate, weakened semantic coherence [5] – and a recent systematic review of 51 NLP studies found that combining linguistic and acoustic features gave the best diagnostic performance (87% average accuracy, AUC 0.89), against 83% for linguistic features alone [6]. Wearable and behavioral data offer a non-invasive, longitudinal window into cognitive status: recent work combining waist-worn accelerometer data with ML classifiers showed that gait-derived features can meaningfully discriminate cognitive status in older adults, albeit in modest samples [7]. Figure 3 shows the approximate distribution of these modalities across the reviewed studies.




Figure 3. Approximate distribution of the primary data modalities across the reviewed studies. Percentages are indicative and sum to slightly over 100% because some studies employed more than one modality.

Combining modalities – MRI with EEG, say, or cognitive scores with speech – generally improves accuracy relative to single-modality approaches, and multimodal deep learning that fuses imaging, genetic, and clinical data has been shown to outperform single-modality baselines for AD-stage classification [8]. Fusion is not free, however: harmonizing heterogeneous data types, handling missing data, and managing added computational cost all remain challenges, and transformer-based and ensemble architectures are increasingly used to address them [2, 4].

Representative Findings

Table 2 summarizes representative, independently verifiable studies and systematic reviews discussed above, illustrating the range of AI techniques and data modalities currently applied to cognitive health monitoring in older adults. It is intended as an illustrative cross-section, not an exhaustive inventory.

Table 2. Representative studies across AI approaches and data modalities for cognitive-deviation tracking.

Study	AI approach	Data modality	Reported finding	Key limitation
Battineni et al. [19]	Multiple ML/DL (meta-analysis)	Structural MRI	Large heterogeneity in AD-stage accuracy across studies	Cross-study comparability limited by reporting standards.
Malik et al. [2]	CNN, RNN, hybrid DL (116-study review)	MRI, PET, EEG, MEG	DL architectures dominate recent AD/MCI detection	Reliance on small, imbalanced public datasets (ADNI, OASIS).
Aviles et al. [20]	ML/DL systematic review	EEG	Growing use of CNN/recurrent models for EEG-based AD detection	Electrode setup and cohort heterogeneity.
Shankar et al. [6]	NLP (linguistic + acoustic)	Speech / language	87% avg. accuracy (AUC 0.89) across 51 studies	Datasets concentrated in monolingual clinical speech.
Venugopalan et al. [8]	3D-CNN + autoencoder fusion	MRI + genetic + clinical	Multimodal fusion outperforms single-modality baselines	ADNI-only; limited external validation.
Shen et al. [7]	Logistic regression, LightGBM	Accelerometer (gait)	Gait features classify cognitive status in older adults	Small sample; single-site validation.
Bron et al. [10]	SVM vs. CNN (external validation)	Structural MRI	Performance generalizes to external cohort with modest drop	Larger drop for MCI-conversion prediction.

Emerging Techniques

Several methodological trends are shaping where the field is headed. Transformers [4] are increasingly applied to cognitive tracking because they model long-range dependencies well, which suits sparse, irregular, or longitudinal data – such as fragmented electronic health records or EEG. Multimodal fusion models – that integrate neuroimaging, biomarkers (amyloid- β , tau), genetic risk factors (APOE4), and neuropsychological testing are becoming state of the art for personalized assessment [8]. Explainable AI (XAI) – most notably SHAP [22] and LIME [23] – is increasingly built into diagnostic pipelines to visualize which features drove a prediction, supporting clinician trust and, in some jurisdictions, regulatory compliance [11, 18]. Finally, time-series and longitudinal modeling is shifting the field from single-timepoint classification toward trajectory modeling: recurrent and multimodal architectures have been used to predict MCI-to-AD conversion from sequences of clinical and imaging data, capturing within-person change over time [24].

THEMATIC SYNTHESIS AND RESEARCH GAPS

Dataset Diversity and Generalizability

One of the most persistent problems in AI-based cognitive-deviation tracking is the lack of demographic, geographic, and socioeconomic diversity in training and testing data. Much of the published literature draws on data from high-income countries and populations with comparatively good access to healthcare and digital technology, which limits how well these models generalize to under-represented groups. That lack of representativeness can translate into meaningful algorithmic bias, since models may pick up patterns that do not transfer across populations. The point generalizes beyond cognitive health: Obermeyer et al. [25] showed that a widely deployed healthcare risk-prediction algorithm carried substantial racial bias – not because it used race directly, but because it relied on healthcare cost as a proxy for need, and cost itself reflects unequal access to care. It is a clear illustration of how bias can enter a clinical AI system indirectly, through a proxy variable, without any demographic variable ever appearing as an explicit input. Closing this gap will likely take concerted dataset diversification paired with privacy-preserving methods, such as federated learning [26] and transfer learning, that let models generalize without compromising privacy.

Early-Detection Limitations

Despite genuine progress, catching cognitive abnormality at the preclinical or very mild stage remains hard. Many models classify moderate-to-severe impairment well but lose sensitivity for subtler, preclinical change. Part of this comes down to the data itself: public datasets – such as ADNI and OASIS – skew toward later-stage cases, so models see fewer early-stage examples during training [2, 19], and the early-stage signal in MRI, EEG, and cognitive scores tends to be faint. External-validation work reinforces the point, showing that predictive performance for MCI-to-AD conversion – an inherently early-detection task – drops more than diagnostic performance when models move to new cohorts [10]. Incremental learning, transfer learning, and longitudinal modeling are the most promising routes to improvement here.

Interpretability and Clinical Trust

Even with strong predictive performance, many modern models – deep networks especially – remain largely opaque about their internal reasoning. That opacity sits uneasily with the demands of clinical trust, regulatory approval, and adoption; clinicians are reasonably wary of leaning on predictions they cannot explain, all the more so in a domain as sensitive as cognitive health [11]. XAI techniques – such as SHAP [22] and LIME [23] – are increasingly folded into diagnostic pipelines to provide feature-level interpretability, and interpretable gradient-boosting frameworks have been proposed specifically for AD staging [18]. A related body of work frames algorithmic output as something that augments, rather than replaces, clinician judgment.

Ethical, Privacy, and Data-Sharing Concerns

Applying AI to cognitive and neurological healthcare raises serious ethical, privacy, and governance questions. Neuroimaging, behavioral data, and speech recordings are all highly sensitive and demand strong safeguards; current applications do not always adequately handle anonymity, informed consent, or data ownership. Reviews of ethical design in dementia-focused assistive technology have repeatedly found that explicit ethical consideration is often missing from the design process itself, not just from deployment [14], and analyses of the global AI ethics landscape show convergence on high-level principles but persistent divergence on how they are actually implemented [13]. This points to a need for globally coordinated ethical norms and secure, privacy-preserving data-sharing infrastructure that supports collaborative research without compromising individual privacy.

Fragmented Modalities and Integration Challenges

Much existing research is confined to a single modality, which limits the depth and generalizability of the resulting models. Multimodal AI models – combining neuroimaging, behavioral, speech, and clinical data – offer a route to more comprehensive assessment, but practical obstacles remain: synchronizing heterogeneous formats, the scarcity of large annotated multimodal datasets, and the computational cost of fusion. Getting more out of multimodal AI will likely require standardized frameworks for data harmonization, interoperability, and fusion architectures that scale without a substantial hit to performance [8]. Table 3 consolidates these gaps and the corresponding opportunities.

Table 3. Persistent research gaps and corresponding opportunities for AI-based cognitive-deviation tracking.

Research gap	Description	Future opportunity
Dataset bias	Limited diversity in training data	Global, inclusive data-sharing consortiums; federated learning.
Early detection	Weak sensitivity to preclinical markers	Incremental, self-supervised, and longitudinal learning.
Black-box models	Limited explainability of DL systems	Embedded XAI (SHAP, LIME, attention maps).
Ethical standards	Inconsistent privacy and consent frameworks	Coordinated, AI-specific health-ethics guidelines.
Multimodal integration	Fusing heterogeneous data remains difficult	Interoperable, scalable multimodal fusion platforms.

FUTURE RESEARCH DIRECTIONS

Bringing AI into cognitive-deviation tracking for older adults holds real promise for earlier detection, better diagnosis, and more personalized intervention, but progress depends on sustained work across several fronts. First, more inclusive, demographically representative datasets are essential; cross-institutional efforts supported by privacy-preserving approaches – such as federated learning [26, 27] – could help build open, well-stratified data resources. Second, models built specifically for early-stage change – using transfer learning, self-supervised and incremental learning, and better temporal modeling [24] – could widen the window for effective intervention. Third, transparent and explainable AI is close to a precondition for clinical integration; building interpretability tools – such as SHAP [22] and LIME [23] into model design from the outset – keep output aligned with clinical reasoning. Fourth, rigorous ethical and data-security standards that center informed consent – particularly given how vulnerable many older adults are – will be needed, alongside closer engagement with bioethics committees and regulators [13, 14]. Finally, multimodal integration and real-world clinical validation remain underdeveloped; testing systems in actual clinical workflows, not just experimental settings, will be essential for impact at scale, including in resource-limited and rural settings. Figure 4 situates these priorities within an end-to-end pipeline.

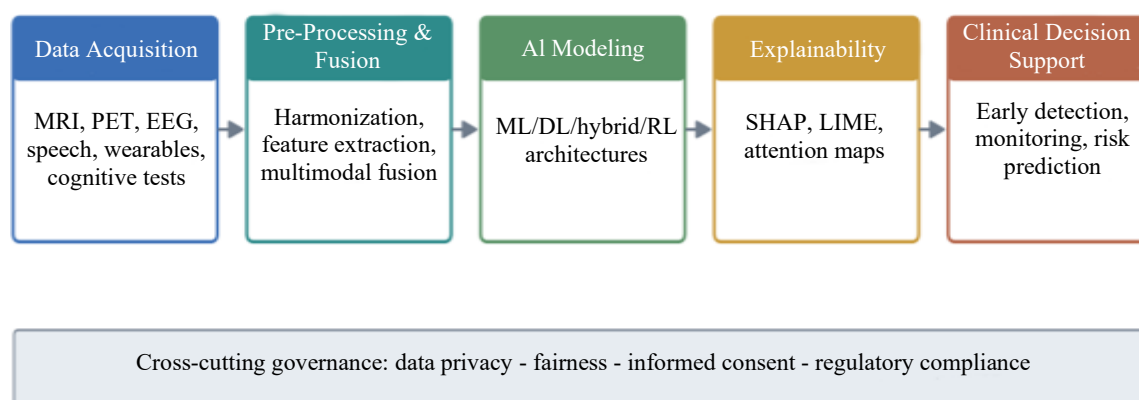


Figure 4. Conceptual pipeline for AI-based cognitive-deviation tracking, from data acquisition through clinical decision support, with a cross-cutting governance layer spanning privacy, fairness, consent, and regulatory compliance.

CONCLUSION

Given how quickly the global population is aging, the need for scalable, ethically sound approaches to monitoring cognitive health keeps growing. This review has synthesized recent work applying ML, DL, hybrid, and reinforcement-learning approaches to detecting and monitoring cognitive deviation in older adults across neuroimaging, electrophysiology, cognitive assessments, and speech and language data. The evidence suggests that AI – and CNN- and transformer-based models specifically – holds real potential for sharpening diagnostic accuracy and catching early markers of cognitive impairment. Meaningful limitations remain, however: a shortage of diverse, representative datasets; weak sensitivity to preclinical impairment; the continued opacity of many models; and unresolved ethical questions around privacy, consent, and regulation. Turning these findings into clinical impact will require building systems that are not only accurate but also explainable, fair, and workable across diverse healthcare settings – supported by ethically governed, diverse, multimodal data and validated in real clinical workflows. Pursued carefully, AI-based cognitive health monitoring has genuine potential to support timely, personalized care and to improve quality of life and cognitive outcomes for aging populations worldwide.

Author Contributions

Conceptualization, D.P.P. and C.S.; methodology, D.P.P.; investigation, D.P.P.; writing – original draft preparation, D.P.P.; writing – review and editing, C.S.; supervision, C.S. Both authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Conflicts of Interest

The authors declare no conflicts of interest.

REFERENCES

1. World Health Organization. Ageing and health: Key facts. Geneva: World Health Organization; 2023 [cited 2026 Jul 7]. Available from: <https://www.who.int/news-room/fact-sheets/detail/ageing-and-health>
2. Malik I, Iqbal A, Gu YH, Al-antari MA. Deep learning for Alzheimer's disease prediction: A comprehensive review. *Diagnostics*. 2024;14:1281.
3. Esteva A, Robicquet A, Ramsundar B, Kuleshov V, DePristo M, Chou K, et al. A guide to deep learning in healthcare. *Nat Med*. 2019;25:24–29.
4. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, editors. *Advances in Neural Information Processing Systems*. Vol. 30. Red Hook (NY): Curran Associates; 2017. p. 5998–6008.
5. Fraser KC, Meltzer JA, Rudzicz F. Linguistic features identify Alzheimer's disease in narrative speech. *J Alzheimers Dis*. 2016;49:407–422.
6. Shankar R, Bunde A, Mukhopadhyay A. A systematic review of natural language processing techniques for early detection of cognitive impairment. *Mayo Clin Proc Digit Health*. 2025;3:100205.
7. Shen J, Nagata Y, Shimamoto T, Matsubara S, Nakamura M, Sato F, et al. Accelerometer-based gait analysis as a predictive tool for mild cognitive impairment in older adults. *Sensors*. 2025;25:7390.
8. Venugopalan J, Tong L, Hassanzadeh HR, Wang MD. Multimodal deep learning models for early detection of Alzheimer's disease stage. *Sci Rep*. 2021;11:3254.
9. Salvatore C, Cerasa A, Battista P, Gilardi MC, Quattrone A, Castiglioni I. Magnetic resonance imaging biomarkers for the early diagnosis of Alzheimer's disease: A machine learning approach. *Front Neurosci*. 2015;9:307.
10. Bron EE, Klein S, Papma JM, Jiskoot LC, Venkatraghavan V, Linders J, et al. Cross-cohort generalizability of deep and conventional machine learning for MRI-based diagnosis and prediction of Alzheimer's disease. *Neuroimage Clin*. 2021;31:102712.
11. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell*. 2019;1:206–215.
12. Floridi L, Cows J. A unified framework of five principles for AI in society. *Harv Data Sci Rev*. 2019;1:1.
13. Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. *Nat Mach Intell*. 2019;1:389–399.
14. Ienca M, Wangmo T, Jotterand F, Kressig RW, Elger B. Ethical design of intelligent assistive technologies for dementia: A descriptive review. *Sci Eng Ethics*. 2018;24:1035–1055.
15. Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*. 2021;372:n71.
16. Bramer WM, Rethlefsen ML, Kleijnen J, Franco OH. Optimal database combinations for literature searches in systematic reviews: A prospective exploratory study. *Syst Rev*. 2017;6:245.
17. Chen T, Guestrin X. XGBoost: A scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York (NY): ACM; 2016. p. 785–794.
18. Yi F, Yang H, Chen D, Qin Y, Han H, Cui J, et al. XGBoost-SHAP-based interpretable diagnostic framework for Alzheimer's disease. *BMC Med Inform Decis Mak*. 2023;23:137.
19. Battineni G, Chintalapudi N, Amenta F. Machine learning driven by magnetic resonance imaging for the classification of Alzheimer disease progression: Systematic review and meta-analysis. *JMIR Aging*. 2024;7:e59370.

20. Aviles M, Sánchez-Reyes LM, Álvarez-Alvarado JM, Rodríguez-Reséndiz J. Machine and deep learning trends in EEG-based detection and diagnosis of Alzheimer's disease: A systematic review. *Eng.* 2024;5:1464–1484.
21. Miltiadous A, Tzimourta KD, Giannakeas N, Tsipouras MG, Afrantou T, Ioannidis P, et al. Alzheimer's disease and frontotemporal dementia: A robust classification method of EEG signals and a comparison of validation methods. *Diagnostics.* 2021;11:1437.
22. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, editors. *Advances in Neural Information Processing Systems*. Vol. 30. Red Hook (NY): Curran Associates; 2017. p. 4765–4774.
23. Ribeiro MT, Singh S, Guestrin C. “Why should I trust you?”: Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York (NY): ACM; 2016. p. 1135–1144.
24. Lee G, Nho K, Kang B, Sohn KA, Kim D. Predicting Alzheimer's disease progression using multi-modal deep learning approach. *Sci Rep.* 2019;9:1952.
25. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science.* 2019;366:447–453.
26. Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. *NPJ Digit Med.* 2020;3:119.
27. Shastry KA, Vijayakumar V, Manoj Kumar MV, Manjunatha BA, Chandrashekhar BN. Deep learning techniques for the effective prediction of Alzheimer's disease: A comprehensive review. *Healthcare.* 2022;10:1842.