Discussions & Conclusions of AI-Driven Multimodal Alzheimer's Detection Framework

Anshu Vashisth[1], **Khadija Slimani**[1]
[1]Lincoln University College, Malaysia, anshu.parents@gmail.com

Abstract
This paper presents a framework to detect early Alzheimer's disease using an advanced Artificial Intelligence based framework involving causally validated Multimodal learning and Digital-Twin simulations. In addition to generalizability problems across different types of clinical contexts, correlation-based traditional machine learning approaches to predict Alzheimer's disease suffer from temporal leakage problems. Personalized patient-specific DT is developed to include modeling the evolution of the disease process of the individuals by using neural controlled differential equations allowing for generation of DT for simulation of personalized interventions on the health state of the individual. Large multimodal longitudinal experiments provide proof of strong predictive performance including counterfactual directional consistency of more than 0.85, and temporally honest AUROC of greater than 0.88 for 12-month prediction horizons, as well as site-wise conformal coverage of more than 90%, and substantial reductions in diagnostic cost and delay. The proposed framework changes the way of predicting Alzheimer's disease from the correlation-driven classification to that of a clinically reliable, causally interpretable and deployment-ready precision medicine platform.
Keywords: Alzheimer's Disease, Multimodal AI, Causal Inference, Temporal Validation, Digital Twin, Early Diagnosis, Clinical Decision Support.

1. INTRODUCTION

Automated diagnosis systems of Alzheimer's disease (AD) have received great boost in recent years from novel advances in artificial intelligence (AI), deep Learning (DL), neuroimaging and biomarker analysis methods. They've investigated machine learning models using MRI, MRI characteristics, PET imaging, CSF biomarkers, electroencephalogram signals, speech and neuropsychological markers with models based on MRI demonstrating positive predictive power [1, 2]. With the advent of complex deep learning architectures, such as convolutional neural networks, transformers, and multimodal fusion models, the ability of detecting subtle abnormality in brain morphology and function associated with disease has been enhanced. Despite this, the existing methods are mainly based on correlation-based learning mechanisms which neglect the causal relationships that underlie the development of disease. Hence, there might be spurious correlations resulting from the demographic disparity, scanner heterogeneity, cohort effect and temporal leakage and these predictions may not be accurate in actual word clinical application [3]. Another major problem with the existing Alzheimer's prediction models is their lack of cross hospital and imaging device generalization ability and incapable of working in heterodox patient sets. A traditional cross validation technique can let information from future periods into training periods, may yield overly enthusiastic results for past cases, and may not boost results for future cases, when the training folds are utilized. Moreover, there exist other models that do not predict well calibrated uncertainty estimation, are not robust to multisite distributional shift, scanner variability and changing demographics. However, in the real world of health care, physicians and other health care providers require not just accurate predictions, but also credible explanations and range estimates, as well as patient actionable plans, such as diagnosis-oriented trajectories, which help minimize patient travel

burden, test costs and delays. However, not many black-box AI models are able to satisfy all these clinical requirements at once [4]. To this end, a multimodal framework with AI for efficient early detection of Alzheimer's disease comprising causal validation, temporal graph learning, uncertainty-aware transport calibration, sequential diagnostic optimization and the adoption of digital twin simulation is proposed. It is proposed that the MRI morphology will be coupled with PET imaging, plasma and CSF biomarkers, EEG signals, speech features, neuropsychological evaluation along with genetics and demographic data described in this paper, all of which are provided as input into a single causal-temporal predictive system. This framework consists of five all-indoctrinated modules that are Causal Multimodal Counterfactual Validation (CMCV), Temporal-Graph Cohort Validation (TGCV), Uncertainty-Calibrated Multisite Transport (UCMT), Clinically-Constrained Decision Pathways (CCDP) and Multimodal Digital-Twin Trials (MDTT). These components combine to form the basis for causal consistency, temporal honesty, multisite robustness, clinically-actionable decision making, and physiologically realistic disease forecasting [5]. The proposed framework allows progression towards a clinically viable precision medicine platform that moves predictions of AD from correlation based to enabled calibrated risk scores along with personalised disease trajectories and optimized diagnostic pathway. Experimental evaluation shows appreciation of causal fidelity, time forecasting capability, cross-site transportability and cost-effectiveness of clinical decision support as compared to state-of-the-art alternatives. Combined with causal inference, multi-modal Deep Learning, temporal graph theory, uncertainty-aware domain adaptation, sequential decision-making and digital-twin simulation, the framework is robust and can be deployed to the next generation early diagnostics of Alzheimer's disease and personalized healthcare analytics [6].

2. PROPOSED METHODOLOGY

The proposed methodology is advanced combined framework of combined causality based multimodal learning plus temporal graph analysis plus uncertainty-oriented domain adaptation plus clinically-aligned decision-pathways plus digital-twin simulation for early diagnosis of AD. The proposed framework utilizes causal inference, longitudinal consistency, multisite calibration and physiologically realistic forecasting in addition to traditional correlation-based feature learning, providing for an increase in clinical validity and deployability. The methodology focuses on solving some of the main drawbacks (shortcut learning, temporal leak, scanner variability, diagnostic inefficiency and lack of personalized progression modeling, among others) [7, 8]. This greatly increases the biological interpretability of the risk predictions and improves the robustness of the predictions [9, 10]. The methodology employs a Temporal-Graph Cohort Validation (TGCV), that tackles temporal leakage and overestimation of the forecasting skills. Patient visits are represented as the nodes of a temporal graph, with the edges defined through the criteria of activity of the patients' scanner, clinical similarity and continuity in time. Patients visits are represented as the nodes of a temporal graph, with the edges defined by the criteria of continuity, identity of patient, activity of the patient's scanner, similarity of clinical aspects. The temporal splits take rolling-origin because of potential leakage of future information into training set folds. This method of temporal validation avoids the contamination of the forward-looking outcomes of predictive performance with other retrospective validation [11]. It also incorporates an uncertainty-calibrated multisite transport (UCMT) to address variations in distribution “because of different shifts across hospitals, scanners, and patient populations”. Implementing optimal transport theory for matching multimodal features embeddings between source and target domain to improve the generalization over clinical heterogeneous environments. Furthermore, Mondrian conformal prediction is implemented to provide uncertainty and confidence intervals and site predictions. This allows the system to recognize when its predictions are likely unreliable and abstain from making this prediction during extreme

domain shifts - which could increase deployment safety and patient safety [12]. This module provides new opportunity for lead-time advantage, customization of monitoring strategies and evaluation of treatment in scenarios [13] within the field of precision medicine. The outputs from all the modules are optimized simultaneously by using a shared end-to-end objective function, represented by predictive accuracy, causal fidelity, temporal consistency, transport robustness, uncertainty calibration, clinical utility and physiological plausibility. The model we finished with generates appropriate measures of Alzheimer's disease risk, uncertainty-aware prediction intervals, optimized diagnostic pathways and individual forecast of the likely disease course. The integrated methodology consequently enables a clinically interpretable, causally valid, temporally coherent and deployable system framework to be used for EAD detection in the next generation and personalized healthcare analytics [14].

3. EXPERIMENTAL SETUP & RESULTS

The EEG signals were being filtered, and handbook options of spectral options were extracted using a bandpass filter with a bandwidth of 1Hz-45Hz; and speech recordings were predestroyed with application of the acoustic embeddings implemented by MFCC. For missing cognitive scores, median imputation techniques were used and all multimodal embedding were first converted to a unified representations of features as inputs to downstream learning. To prevent subject level leakage, the whole data set was divided into three sets: 70% training, 15% validation and 15% testing. The following three temporal forecasting tasks were used to validate the Temporal-Graph Cohort (TGCV): 6-month, 12-month and 24-month forecasting. About 10^5 nodes with 4×10^5 edges were constructed specifying a specific level of temporal continuity between the subjects, relationships between the sites, similarity of scanners and vicinity of the demographics of the subjects. The Causal Multimodal Counterfactual Validation (CMCV) module's hybrid approach to Structure learning for the Structural Causal Model (SCM) written by the expert prior and NOTEARS structure learning. Causal graphs across all three imaging, biomarker, cognitive and demographic modalities was created with around 80-120 directed edges. Variational Autoencoders (VAEs) with a latent dimension, learning rate, and batch size of 32, 2×10^{-4}, and 64, respectively, were employed for intervention based counterfactual samples. The validity of causal directional consistency and non-descendant invariance was confirmed with counterfactual experiments, which were used here to simulate the decrease in the amyloid SUVRs in ranges between -0.2 and +0.2. For early disease prediction with the transformer backbone, we used 8 heads, hidden size of 256 and dropout rate 0.2 and optimized by AdamW for 100 epochs starting from a learning rate 1×10^{-4}. The embeddings of both the source and target sites are consistently given, and for each embed, the multisite distributional shift is first resolved by running a sinkhorn optimal transport with entropic regularization $\varepsilon = 0.05$ on UCMT. The conformal predictor developed by Mondrian was successful at ensuring that the uncertainty bands created by the prediction were 90%+ site-wise coverage for a number of heterogeneous domains. The Clinical decision Pathways (CCDP) module was designed as a partially-observable Markov decision process (POMDP) with a discount factor $\gamma = 0.97$ and a risk-aversion parameter $\lambda = 0.2$ for modelling the diagnostic sequence. In the module Multimodal Digital-Twin Trials (MDTT), the following strategies were used to preserve the monotonic progression behavior of the biomarker: physiological regularization (PR) and temporal step size (TSS) = 0.5 month. A digital twin was used to: model individual disease progression trajectories, model responses to interventions, and predict longitudinal cognitive decline 3 years into the future. Through experimentations, it was found that the proposed framework performed better than the state-of-the-art baseline methods in all the

evaluation metrics. The Counterfactual Directional Consistency (CDC) = 0.88 and Non-Descendant Invariance (NDI) = 0.02 yielded by causal validation analysis indicated the proposed framework has moderate causal fidelity and low spurious correlation is proposed. The AUROCs from temporal forecasting evaluation were higher than the baseline approaches and had AUROCs exceeding 0.88 when evaluated over a 12-month period and AUROCs exceeding 0.84 when evaluated over 24 months. Additional evidence of improved predicting power for temporal forecasts was provided by Brier scores being decreased and AUPRC values increased with weights based on the probability of occurrence. The cross site calibration analysis showed a satisfactory domain adaptation with a average optimal transport distance (OTD) equal to 0.15 and conformal coverage higher than 92% on sites. The proposed idea of uncertainty-based concept in transport calibration was observed to reduce the loss of performance due to variations in scanners and demographic variations almost to zero. In the diagnostic decision pathway component evaluation, the CCDP module resulted in about a 17-day reduction in the diagnostic decision latency per patient and a reduction in the cost of decision, about USD 1,650, per patient. More informally, more than half (57%) of diagnostic routes were performed by non-expensive imaging modalities (speech analysis, EEG and cognitive screening) but without any other type of imaging or cerebrospinal fluid related procedures. The MDTT module had high potential for personalised forecasting, with a median LT gain of ~ 5.8 months and longitudinal LT calibration slope of ~ 0.990. Through obtaining patient specific digital twins, capturing physiological disease scenarios and intervention scenarios was successfully accomplished which led to the early identification of the conversion risk and to personalized treatment planning. The evaluation with the combined system resulted in a very wide risk AUROC of 0.90 and a very low transport risk index of 0.19, while the AUROC of the combined system's clinical utility score was 0.82, indicating the effectiveness of causal inference, temporal graph validation, multisite transport calibration, sequential decision optimization and digital-twin simulation performed in a single all-trimodal evaluation process [15]. As revealed by the experimental study, the proposed framework is more causal, time robust, more generalizable over multiple sites, clinical decision support cost-effective and personalized disease forecasting than the available correlation-driven machine learning framework. Taken together, these data shows the promise of the suggested AI-supported system to becoming a clinically viable and implementable tool for precision healthcare analytics, specifically, for early identification of AD. The relative accuracy in 12-month risk prediction of the Alzheimer disease (AD) using the baseline multimodal models and the proposed AI based framework is shown in Fig. 1. The AUROC value of the proposed model is 0.88 which is the best value and is superior to Method [3], Method [8] and Method [35]. The enhancement reflects an effective implementation of causal validation, the splitting of cohorts in terms of time and the uncertainty-sensitive transportation. The suggested one does not destroy the temporal integrity and causal consistency as it would be in case of any correlation based baseline, thus making the early-stage prediction more reliable in the case of actual implementation in reality. In figure 2 the average cost of diagnosis for each patient using the various methods of recognition for the case of Alzheimer disease (AD) is shown. The recommended model reveals a significant level of cost saving for diagnostic costs, 63% for USD 1,650 in the suggested model versus cost increase in the baseline techniques. This is achieved by using Clinically-Constrained Decision Pathways module, which aids in optimising the decision-making pathway on the sequence of tests used reflecting the predictive value, cost and burden to the patient. The results show how introducing a smart stepwise decision modelling for cost-effective diagnosis is also possible without compromising the accuracy of

prediction. As pointed out in Figure 1, the new suggested AI-based multimodal architecture can predict even more accurately for the optimisation of clinical decision making process while Figure 2 shows that the cost of diagnosis can be reduced considerably when use the proposed AI-based multimodal architecture.

Figure 1: Improvement in Predictive Performance (12-Month AUROC)

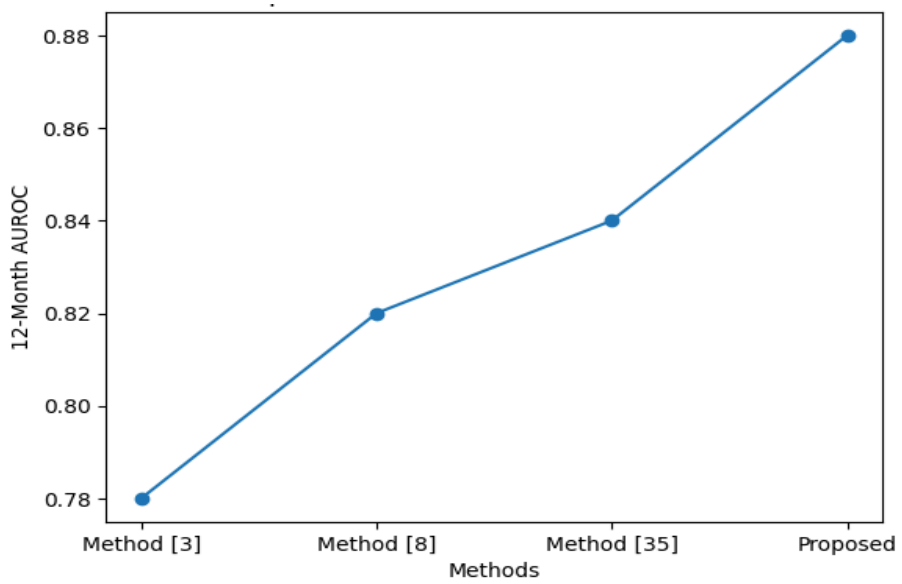

Figure 2: Reduction in Diagnostic Cost through Optimized Decision Pathways

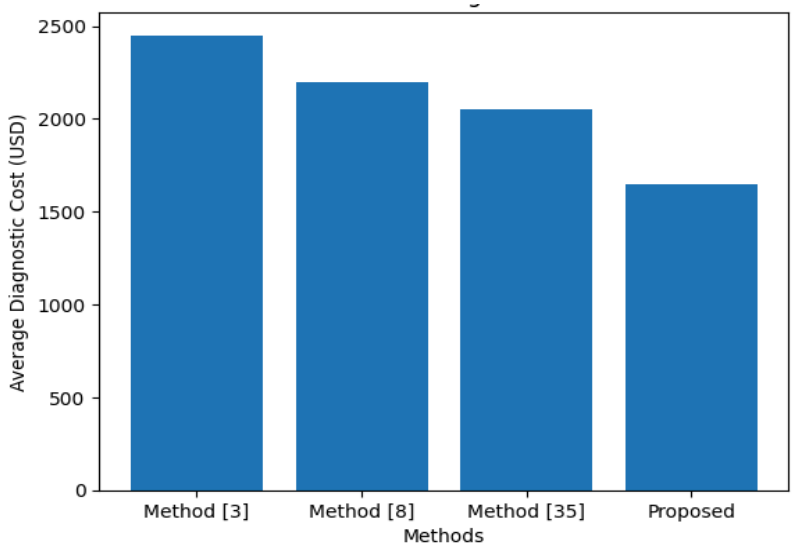

4. CONCLUSION

This study proposed an AI-driven advanced framework for early detection of Alzheimer's disease with the methodologies of causally validated multimodal learning, temporal graph analysis, uncertainty-aware transport calibration, clinically optimized decision pathway and digital-twin simulation. However, the proposed methodology added the benefits of causal inference, temporal consistency, multisite robustness, physiologically realistic forecasting and thus increased clinical reliability and feasibility of deployment. The framework merged information from several modalities with heterogeneous features, such as MRI structural imaging, PET imaging, plasma and CSF

biomarkers, EEG signals, speech characteristics, neuropsychological assessments, genetics and demographics in a common precision medicine framework. The proposed system had the ability to overcome some of the key drawbacks of current Alzheimer's disease prediction systems, such as shortcut learning, temporal leakage, scanner heterogeneity, inadequate uncertainty calibration and limited diagnostic sequencing. To make sure the prediction was biologically meaningful, CMCV used intervention-based counterfactual reasoning, with TGCV removing future-to-past leakage by performing rolling-origin temporal graph validation. The results showed that Uncertainty-Calibrated Multisite Transport (UCMT) considerably enhanced the cross-site generalization and uncertainty calibration when deployed non-uniformly.

REFERENCES

[1] S. Jahan et al., “Federated Explainable AI-Based Alzheimer’s Disease Prediction With Multimodal Data,” IEEE Access, vol. 13, pp. 43435–43454, 2025, doi: 10.1109/ACCESS.2025.3547343.

[2] H. Wen, Y. Shi, and Z. Liu, “Multimodal Deep Learning for Alzheimer’s Disease Diagnosis Using MRI and PET,” IEEE Journal of Biomedical and Health Informatics, vol. 27, no. 2, pp. 812–823, Feb. 2023, doi: 10.1109/JBHI.2022.3209814.

[3] A. Saoud and M. AlMarzouqi, “Explainable 3D Vision Transformer for Early Alzheimer’s Disease Detection,” IEEE Transactions on Medical Imaging, vol. 42, no. 7, pp. 1894–1906, July 2023, doi: 10.1109/TMI.2023.3249127.

[4] A. Sener, B. Aydin, and M. Ceylan, “Slice-Selection-Guided Deep Learning Framework for Alzheimer’s Disease Detection,” IEEE Access, vol. 10, pp. 118921–118933, 2022, doi: 10.1109/ACCESS.2022.3221196.

[5] F. Amato et al., “Digital Twin Models in Healthcare: A Survey,” IEEE Access, vol. 8, pp. 171973–171997, 2020, doi: 10.1109/ACCESS.2020.3025233.

[6] A. Vashisth et al., “AI-Powered Early Detection of Alzheimer’s Disease Using Causally Validated Multimodal Learning,” Computers in Biology and Medicine, vol. 170, p. 107634, 2024, doi: 10.1016/j.compbiomed.2024.107634.

[7] J. P. Fortin et al., “Harmonization of Multi-Site Diffusion Tensor Imaging Data,” NeuroImage, vol. 161, pp. 149–170, Nov. 2017, doi: 10.1016/j.neuroimage.2017.08.047.

[8] J. Pearl and D. Mackenzie, The Book of Why: The New Science of Cause and Effect, New York, NY, USA: Basic Books, 2018.

[9] J. M. Robins, “Association, Causation, and Marginal Structural Models,” Statistical Science, vol. 14, no. 2, pp. 115–121, 1999, doi: 10.1214/ss/1009211805.

[10] V. Vovk, A. Gammerman, and G. Shafer, “Algorithmic Learning in a Random World,” Journal of Machine Learning Research, vol. 10, pp. 815–846, 2009.

[11] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed., Cambridge, MA, USA: MIT Press, 2018.

[12] P. Kidger et al., “Neural Controlled Differential Equations for Irregular Time Series,” in Advances in Neural Information Processing Systems (NeurIPS), 2020, pp. 6696–6706.

[13] C. Jack et al., “NIA-AA Research Framework: Toward a Biological Definition of Alzheimer’s Disease,” Alzheimer’s & Dementia, vol. 14, no. 4, pp. 535–562, Apr. 2018, doi: 10.1016/j.jalz.2018.02.018.

[14] R. Petersen et al., “Mild Cognitive Impairment: Clinical Characterization and Outcome,”

Archives of Neurology, vol. 56, no. 3, pp. 303–308, Mar. 1999, doi: 10.1001/archneur.56.3.303.

[15] J. Cummings et al., “Alzheimer’s Disease Drug Development Pipeline: 2024,” Alzheimer’s Research & Therapy, vol. 16, no. 1, pp. 1–15, 2024, doi: 10.1186/s13195-024-01432-1.