Uncovering Stable and Domain-Specific Drivers of Cognitive Impairment via Time-Slice Causal Discovery

Siying Li (lisiying2024@stu.scu.edu.cn)[1,*] **Zhenda Xu (20241414900015@stu.scu.edu.cn)**[1,*]
JiaHui Xu (gawuei@stu.scu.edu.cn)[1] **Linna Wang (lenawang@stu.scu.edu.cn)**[1]
Li Lu (luli@scu.edu.cn)[1,†]

[1]College of Computer Science, Sichuan University, Chengdu, Sichuan, China
[*]These authors contributed equally to this work.
[†]Corresponding author.

Abstract

Cognitive impairment is a major public health challenge in aging populations. Identifying stable, actionable drivers is crucial for targeted intervention. Causal discovery has been leveraged for its interpretability to explore these factors. However, causal modeling approaches that integrate long-span, multi-wave data to specifically address the heterogeneity of cognitive impairment remain underdeveloped, thus inherently failing to reveal domain-specific causal pathways from long-term, sparse longitudinal data. To address this, we introduce a framework that decomposes cognition into distinct domains and then applies time-slice causal discovery to sparse, multi-wave data. Applied to 8-year CHARLS data, the framework decomposed cognition into Mental State and Contextual Memory, identifying both domain-specific and shared causal drivers. These drivers were validated through robust predictive modeling, external generalization (CFPS), and interpretability analyses, demonstrating their utility. Our findings provide new insights for heterogeneous prevention and targeted early intervention.

Keywords: Computer Science; Causal reasoning; Cognition of Time; Cognitive architectures; Computational Modeling

Introduction

With the acceleration of global population aging, cognitive impairment among middle-aged and elderly people is gradually becoming a major health problem (Roh et al., 2021). Cognitive impairment not only significantly impairs quality of life and functional independence (Song et al., 2023) but is also associated with increased risks of disability, hospitalization, and mortality (Fogg et al., 2019), posing a substantial challenge for healthcare systems (Chao et al., 2022). However, cognitive impairment is frequently underdiagnosed in routine clinical settings due to its insidious onset, heterogeneous progression patterns, and practical barriers including time constraints and insufficient clinician confidence in recognizing early cognitive changes (Mattke et al., 2023). Therefore, identifying potential drivers of cognitive impairment is crucial for improving early detection and risk assessment.

Yet, a key obstacle to identifying these drivers is the heterogeneous nature of cognitive decline itself. It manifests in distinct patterns across etiologies: for instance, Alzheimer's disease typically presents with early memory impairment (Didic et al., 2011), vascular cognitive impairment may primarily affect executive function (Black, 2011), and Lewy body dementia often involves visuospatial deficits (Johnson et al., 2005). Widely used composite assessment tools like the MMSE (Cockrell & Folstein, 2002) and MoCA (Nasreddine et al., 2005) summarize performance across multiple domains into a single score, which may risk masking etiology-specific deficits and obscuring distinct decline pathways. Thus, moving beyond composite scores by decomposing global assessments into constituent dimensions is necessary (Harvey, 2019).

To then uncover the directional mechanisms within these distinct dimensions, causal discovery offers a powerful framework, as it models relationships beyond mere association (Rojas-Saunero et al., 2023; Yu et al., 2025). Previous studies have applied causal discovery algorithms to investigate cognitive impairments, demonstrating the potential of these methods to uncover underlying mechanisms and identify meaningful drivers of cognitive decline (Fehr et al., 2023; Shen et al., 2020). However, causal modeling approaches that integrate long-span, multi-wave data to specifically address the heterogeneity of cognitive impairment remain underdeveloped, leaving distinct progression pathways largely unexplored.

To address this, we introduce a framework integrating cognitive domain decomposition with time-slice causal discovery, designed to extract causal features from sparse longitudinal data. We leverage the CHARLS dataset (Zhao et al., 2014) and adopt the framework proposed by Luo et al. (2021) to decompose composite cognitive assessment scores into two mechanistic dimensions: Mental State and Contextual Memory. These dimensions enable a more fine-grained characterization of cognitive function, facilitating cross-wave causal analysis and predictive modeling of cognitive impairment. Building on these aims, the contributions of this study are threefold:

- **Identify factors**: We identify potential factors associated with Mental Status and Contextual Memory using longitudinal causal discovery methods applied to multi-wave data.
- **Uncover causal pathways**: We uncover causal pathways among these factors, revealing how upstream variables may influence downstream cognitive outcomes.
- **Validate utility**: We validate the identified causal features through predictive modeling, external generalization (CFPS), and interpretability analysis, collectively supporting their relevance for cognitive impairment research.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

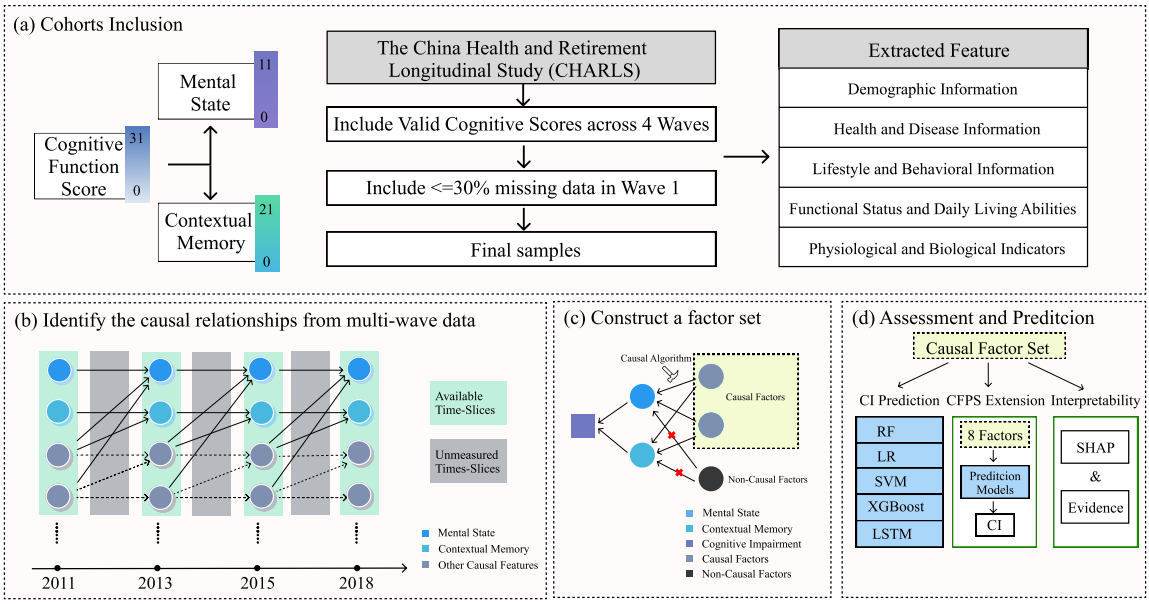

Figure 1: Our Domain-Specific Time-Slice Causal Framework are as follows: (a) describes the process of calculating cognitive function scores and selecting research samples; (b) and (c) show the acquisition of causal relationships with cognitive domain specificity from multi-wave temporal data; (d) illustrates three methods for validating recognition features, including predictive modeling, external generalization (CFPS), and interpretability analysis.

Notations and Preliminary

Notations

Let $\mathcal{D} = \{(X_i, Y_i)\}_{i=1}^{n}$ denote the longitudinal dataset with n individuals observed over t time points. For each individual i, we have $X_i \in \mathbb{R}^{t\times d}$ representing the feature matrix over t time points, where each column represents a measured attribute (demographic, physical health, and social factors), and $Y_i \in \mathbb{R}^{t\times 3}$ representing the outcome matrix, where each row consists of Mental State Y^{mental}, Contextual Memory Y^{memory}, and Cognitive Impairment status Y^{cog}. To identify temporal causal structures, we construct a three-time-slice dataset:

$$\mathcal{D}^{(3)} = \left\{\left(X_j^{(t_1)}, X_j^{(t_2)}, X_j^{(t_3)}, \mathbf{Y}_j^{(t_1)}, \mathbf{Y}_j^{(t_2)}, \mathbf{Y}_j^{(t_3)}\right)\right\}_{j=1}^{m}, \quad (1)$$

where $1 \le t_1 < t_2 < t_3 \le t$ denote time indices. Here, $X_j^{(t_i)} \in \mathbb{R}^d$ are feature vectors, and

$$\mathbf{Y}_j^{(t_i)} = \left(Y_j^{\text{mental},(t_i)}, Y_j^{\text{memory},(t_i)}\right) \in \mathbb{R}^2, \quad i \in \{1,2,3\} \quad (2)$$

denote intermediate cognitive outcomes. Y^{cog} is treated as a downstream clinical endpoint and is excluded from causal structure learning to avoid label leakage.

The identified causal relationships are represented by a directed acyclic graph (DAG) $G = (V, E)$, where V includes observed features and $(Y^{\text{mental}}, Y^{\text{memory}})$, and E denotes causal edges. For each node $v \in V$, we define:

$$\text{An}(v) = \{u \in V \mid u \to \cdots \to v \text{ in } G\}, \quad (3)$$

$$\text{De}(v) = \{w \in V \mid v \to \cdots \to w \text{ in } G\}. \quad (4)$$

Problem Definition

Task 1: Causal Feature Discovery for Cognitive Mechanisms Given the three-time-slice dataset $\mathcal{D}^{(3)}$, our objective is to identify a stable set of causal ancestors for the joint intermediate cognitive outcomes. Formally, let $\text{An}(Y^{\text{mental}})$ and $\text{An}(Y^{\text{memory}})$ denote the sets of causal ancestors for Mental State and Contextual Memory, respectively. We define the unified causal feature set for downstream tasks as:

$$\mathcal{A}^{\text{cog}} = \text{An}(Y^{\text{mental}}) \cup \text{An}(Y^{\text{memory}}). \quad (5)$$

This formulation ensures that the selected features are limited to those with a direct or indirect causal link to either of the core cognitive domains.

Task 2: Validation of Mechanistic Stability Given the identified causal feature set $\mathcal{A}^{\text{cog}}$, we aim to validate the stability and practical utility of the discovered mechanisms through predictive modeling and cross-dataset evaluation.

Materials and Methods

Datasets and Preprocessing

Datasets CHARLS is a publicly available, nationally representative longitudinal cohort study of residents aged 45 and above, spanning 28 provinces in China (Gong et al., 2022). The study employed a multi-stage, stratified probability proportional sampling method to randomly select over 17,000 participants.

Study population This study used data from four waves of CHARLS collected in 2011, 2013, 2015, and 2018, including

individuals with valid cognitive function scores in all waves. Participants aged 45 years or older at baseline were included, and participants with data missing more than 30% in the first round were excluded to ensure data quality and reliability.

Data collection This study included 248 features across five categories: Demographic Information, Health and Disease Information, Lifestyle and Behavioral Information, Functional Status and Daily Living Abilities, and Physiological and Biological Indicators. Cognitive function was assessed using two dimensions, Mental State and Contextual Memory, corresponding to orientation and numerical/visual/spatial abilities, and word recall ability, respectively. The total cognitive score is the sum of these two dimensions, ranging from 0 to 31, and binary cognitive impairment labels were derived using standardized cutoffs (Hu et al., 2022; Luo et al., 2021).

Preprocessing To ensure data quality, missing values were imputed using within-subject forward/backward filling, and all continuous features were standardized to zero mean and unit variance to maintain temporal consistency and facilitate stable modeling.

Method

Causally-Informed Feature Discovery for Cognitive Mediators To systematically identify features with stable causal influence on cognitive outcomes, we adopt an overlapping window design based on the well-established theoretical framework of time-slice causal discovery. Within each three-time-slice (t_1, t_2, t_3) analysis window, we follow the causal discovery procedure below:

the objective is to control for historical information and identify features at time t_2 that potentially have causal influence on a future cognitive outcome y at time t_3. For a feature $x_i^{(t_2)}$, we define its causal ancestor set with respect to $y^{(t_3)}$ as:

$$\mathrm{An}^{t_2 \to t_3}(y) = \left\{ x_i^{(t_2)} \,\middle|\, x_i^{(t_2)} \not\perp\!\!\!\perp y^{(t_3)} \mid \mathrm{An}_i^{(t_2)} \right\}, \tag{6}$$

where the conditioning set $\mathrm{An}_i^{(t_2)}$ is constructed from observed variables at times t_1 and t_2, excluding $x_i^{(t_2)}$ itself and any variables marginally independent of it. The core principle is: if $x_i^{(t_2)}$ remains statistically associated with the future outcome $y^{(t_3)}$ after controlling for a sufficient conditioning set constructed from its past (non-descendant) information, we consider it to potentially have causal influence on cognitive outcomes. Based on this, our three-time-slice dataset is:

$$\mathcal{D}^{(3)} = \left\{ \left(X_j^{(t_1)}, X_j^{(t_2)}, X_j^{(t_3)}, \mathbf{Y}_j^{(t_1)}, \mathbf{Y}_j^{(t_2)}, \mathbf{Y}_j^{(t_3)} \right) \right\}_{j=1}^{m}. \tag{7}$$

Unlike standard constraint-based methods (e.g., PCMCI, tsFCI), for each feature $X_i^{(t_2)}$, we evaluate the conditional independence $X_i^{(t_2)} \perp \mathbf{Y}^{(t_3)} \mid \mathrm{An}_i^{(t_1,t_2)}$ using an adaptive GAM framework (Hastie, 2017) specifically tailored for mixed-type variables. To accommodate mixed-type data, we adopt a three-tier modeling procedure: (1) Variable Classification: Features with cardinality ≤ 4 are treated as categorical, while others are modeled via smooth splines to capture potential non-linearity. (2) Complexity Control: The basis dimension k is dynamically adjusted as $k = \min(k_{\text{default}}, \max(3, \lfloor n/3 \rfloor))$ to prevent overfitting while maintaining sufficient degrees of freedom. (3) Fault-Tolerant Fitting: Models are primarily fitted using Restricted Maximum Likelihood (REML). In cases of convergence failure, the system automatically reverts to a robust Linear Model (LM) baseline, supplemented by residual validity checks to ensure numerical stability.

This adaptive framework effectively balances flexibility for mixed-type data with robustness against model misspecification. Leveraging these properties, we applied the sliding window method to the CHARLS dataset containing 4 waves of data for causal discovery. Specifically, we construct two overlapping three-wave observation windows:

- Window 1: t_1 = 2011 (Wave 1), t_2 = 2013 (Wave 2), and t_3 = 2015 (Wave 3).
- Window 2: t_1 = 2013 (Wave 2), t_2 = 2015 (Wave 3), and t_3 = 2018 (Wave 4).

For each window, independent causal discovery is performed on two primary cognitive domains: Y^{mental} and Y^{memory}.

To identify stable causal features across temporal intervals, we define the unified causal feature set $\mathcal{A}^{\text{cog}}$ by aggregating the causal ancestors of the multi-dimensional cognitive outcome $\mathbf{Y}$ identified across all observation windows:

$$\mathcal{A}^{\text{cog}} = \bigcup_{w \in \mathcal{W}} \mathrm{An}^{w}(\mathbf{Y}) \tag{8}$$

where $\mathcal{W} = \{2011 \to 2015, 2013 \to 2018\}$ represents the set of sliding windows. This formulation naturally encompasses causal drivers for both Mental State and Contextual Memory.

Validation of Causal Features via Prediction and Interpretation To validate the utility of the identified causal features, we compare predictive performance across three feature sets: (1) $\mathcal{A}^{\text{cog}}$ (causally-identified features); (2) all available features; and (3) non-causal features (all features excluding $\mathcal{A}^{\text{cog}}$). This comparison assesses whether causal discovery yields features that retain substantial predictive power while achieving significant dimensionality reduction.

We employed a suite of standard machine learning classifiers to ensure robustness. Two evaluation settings are implemented: (1) Cross-sectional prediction: Models are trained and tested independently for each wave (2011, 2013, 2015, 2018); (2) Longitudinal prediction: LSTM networks leverage historical sequences to forecast future cognitive impairment. All target labels are binary cognitive impairment indicators.

Beyond predictive performance, we further assess interpretability and generalization robustness through SHAP-based analysis and cross-setting validation, respectively.

Experiments and Results

Experimental Setup

Dataset Division The dataset includes longitudinal records of 9,345 individuals with up to four waves of follow-up. For

Table 1: Baseline characteristics of CHARLS participants with and without cognitive impairment.

Variable	No CI (n=33,695)	CI (n=3,685)	P-value
Categorical Data, n(%)			
Gender:			
Female	17,140 (50.9)	2,692 (73.1)	<0.001
Male	16,555 (49.1)	993 (26.9)	<0.001
Education Level:			
≤Lower sec.	29,606 (87.9)	3,670 (99.6)	<0.001
Upper sec. & voc.	3,645 (10.8)	11 (0.3)	<0.001
Tertiary	444 (1.3)	4 (0.1)	<0.001
Exercise Regularly:			
Yes	30,782 (91.4)	3,117 (84.6)	<0.001
No	2,913 (8.6)	568 (15.4)	<0.001
Numerical Data, Median [Q1, Q3]			
Age, years	60 [54, 66]	66 [60, 72]	<0.001
Depression Score	7 [3, 11]	10 [5, 16]	<0.001
Eyesight	4 [4, 6]	4 [3, 5]	<0.001
Self-rated Health	3 [3, 3]	3 [2, 3]	<0.001
BMI, kg/m^2	23.6 [21.4, 26.2]	22.8 [20.3, 25.4]	<0.001
Income (k yuan)	12 [5, 24]	3 [1, 10]	<0.001

Note: CI = Cognitive Impairment; Lower sec. = Lower secondary; Upper sec. & voc. = Upper secondary and vocational training; Income = Individual Income. P-values for all three cognitive domains—Cognitive Impairment, Mental State, and Contextual Memory—were all less than 0.001.

cognitive impairment prediction, data splitting depends on the setting: under the cross-sectional setting, samples are randomly divided 8:2 for training and testing; under the temporal setting, all observations of an individual are assigned to either training or testing to avoid temporal leakage. The training set is used for model fitting, and the testing set for out-of-sample evaluation.

Baseline To evaluate the identified features, we compare predictive performance using all features, causal features, and non-causal features under both cross-sectional and temporal settings. We also assess feature correlations and cross-dataset generalizability.

Evaluation Metrics We report Accuracy, Area Under the Receiver Operating Characteristic Curve (AUC), Recall, and F1-score for prediction performance, and use SHapley Additive exPlanations (SHAP) values to interpret feature contributions.

Results

Population We pooled observations across all four survey waves (2011–2018), resulting in a total sample of 37,380 for analysis. As shown in Table 1, the cognitive impairment (CI) group differed significantly from the cognitively normal group across all examined demographic and health characteristics (all $p < 0.001$). Specifically, the CI group was notably older (median 66 vs. 60 years), predominantly female (73.1% vs. 50.9%), and had substantially lower education levels (99.6% with ≤lower secondary vs. 87.9%). Furthermore,

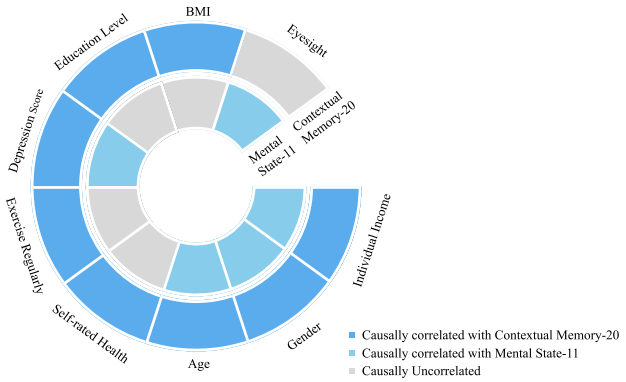

Figure 2: Nine potential causal features discovered using the time-slice algorithm. Blue denotes causal relationships; gray denotes no causal relationships.

the CI group reported higher depressive symptoms, lower income, lower BMI, and poorer self-rated health and eyesight. Statistical tests (Chi-square, Mann–Whitney U, and Pearson correlation) confirmed robust associations between these features and Cognitive Impairment, Mental State, and Contextual Memory in the studied population.

Causally-informed Feature Analysis Figure 2 shows the $\mathrm{An}(Y^{\mathrm{mental}})$ and $\mathrm{An}(Y^{\mathrm{memory}})$ derived from four-wave longitudinal data using three-time-slice causal discovery strategy. The inferred causal ancestors reveal distinct pathways associated with the two intermediate cognitive outcomes: Mental State is directly linked to Eyesight, whereas Contextual Memory is directly linked to BMI, Education Level, Exercise Regularly, and Self-rated Health. Beyond outcome-specific pathways, they also reveals a set of shared upstream factors—including Depression Score, Age, Gender, and Individual Income—that influence both cognitive domains. These patterns collectively indicate the presence of both outcome-specific and shared mechanisms, providing insight into potential intervention targets and pathways underlying cognitive impairment.

To assess the stability and directness of these relationships, we conducted pairwise conditional independence tests across consecutive waves and visualized the results as cross-wave causal graphs (Figure 3). The graphs reveal two patterns: direct influence, where features have edges directly to cognitive outcomes at $t+1$, and indirect influence, where features affect cognition via mediators along ancestral paths. Features showing consistent direct influence across transitions (e.g., $\mathrm{Age}_t \rightarrow \mathrm{Contextual\ Memory}_{t+1}$, $\mathrm{Exercise\ Regularly}_t \rightarrow \mathrm{Mental\ State}_{t+1}$) represent stable and critical drivers, highlighting potential intervention targets. Persistent indirect influencers, such as $\mathrm{Individual\ Income}_t$ impacting $\mathrm{Mental\ State}_{t+1}$ via Depression Score, indicate upstream determinants shaping cognitive trajectories. These patterns provide insights for designing early, targeted intervention strategies.

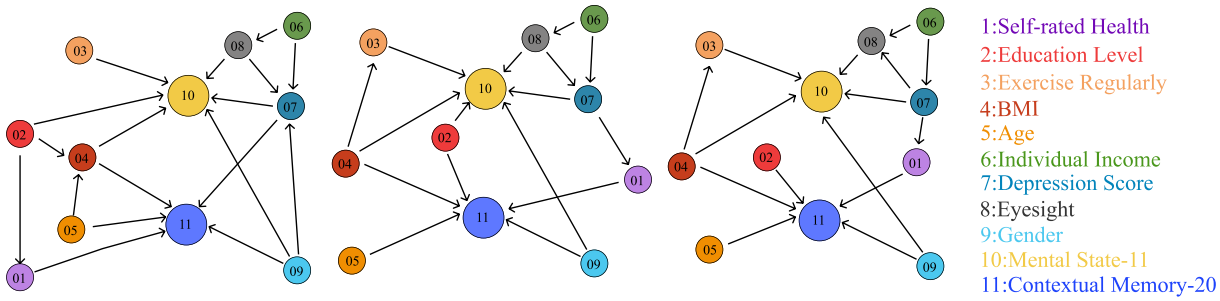

Figure 3: Causal discovery of nine potential features, Mental State, and Contextual Memory across two data waves. The panels represent causal graphs identified for three periods: 2011–2013, 2013–2015, and 2015–2018, with arrows indicating cause-effect relationships.

Table 2: Performance of LSTM models under different variable settings. Values are reported as mean (95% confidence interval).

Metric	All	Causal	Non-C
Accuracy	0.7884	**0.7945**	0.7929
AUC	0.6662	**0.6859**	0.6714
Recall	0.9337	**0.9740**	0.9435
F1-score	0.8743	**0.8819**	0.8777

Note: Non-C = Non-Causal. All values represent mean performance across bootstrap samples. 95% confidence intervals are omitted for brevity.

Prediction Performance on CHARLS Dataset We evaluate the predictive performance of causal features under both temporal and cross-sectional settings. For temporal prediction across four waves, LSTM models trained using only causal features consistently outperform models using all or non-causal features, particularly in Accuracy, AUC, Recall, and F1-score (Table 2). This demonstrates the stability and informativeness of the identified causal features for longitudinal cognitive prediction. For cross-sectional prediction using single-wave data, Table 3 summarizes the performance of SVM, XGBoost, Random Forest, and Logistic Regression across different feature sets. Predictions based on causal features are generally closer to those obtained using all features, compared with predictions based on non-causal features. This suggests that the selected causal features capture the majority of predictive information while excluding less relevant variables, highlighting their potential for more parsimonious modeling.

Cross-Dataset Validation on CFPS To assess the generalizability of our identified causal features, we conducted external validation using the China Family Panel Studies (CFPS) dataset (Xie & Hu, 2014). Cognitive impairment in CFPS was measured using wave-specific cognitive tests (word recognition and math; memory and number series). We standardized and summed test scores, defining cognitive impairment as scores below the median (Pu et al., 2023).

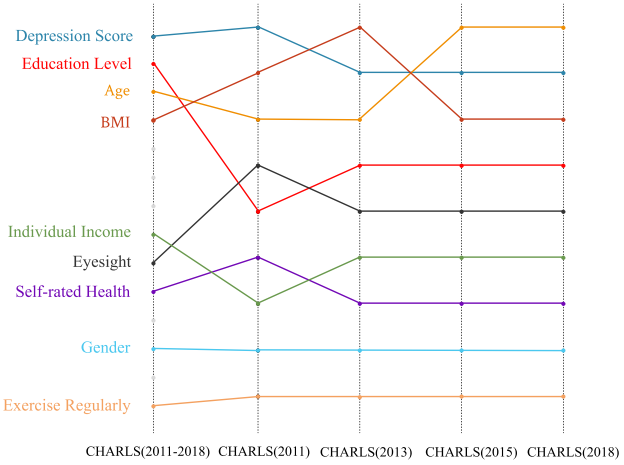

Figure 4: SHAP values for predicting cognitive impairment using an XGBoost model trained on four-wave longitudinal and cross-sectional data, ranked in descending order of importance.

We utilized three waves of CFPS data (2016, 2018, 2020) comprising 1,053 individuals aged 45 and above to ensure comparability with CHARLS. Eight of the nine causal features identified from CHARLS were successfully mapped to CFPS variables; vision was excluded due to measurement differences between the two datasets. As shown in Table 4, the identified causal features demonstrate consistent predictive performance across all three waves, with SVM and XGBoost achieving particularly strong results in 2018 (AUC > 0.87). The LSTM model, trained on all three waves, further validates the temporal stability of these features. This cross-dataset consistency supports the robustness and generalizability of our causal discovery approach.

Feature Interpretation Figure 4 shows feature importance rankings derived from SHAP values for XGBoost models trained on all features to predict cognitive impairment. Labeled points indicate the nine causally identified features, which consistently rank among the top contributors. Depression Score, BMI, and Age are persistently the highest, while

Table 3: Performance comparison of SVM, XGBoost, Random Forest (RF), and Logistic Regression (LR) models across different survey years.

Year	Metric	SVM			XGBoost			RF			LR		
		All	Causal	Non-C	All	Causal	Non-C	All	Causal	Non-C	All	Causal	Non-C
2011	Accuracy	0.8068	**0.7246**	0.6895	0.8549	**0.8438**	0.7130	0.8733	**0.8742**	0.6922	0.7246	**0.7025**	0.6329
	AUC	0.7263	**0.7130**	0.6571	0.7082	**0.7057**	0.6809	0.7693	**0.7361**	0.7029	0.7716	**0.7588**	0.7110
	Recall	0.8235	**0.7319**	0.6958	0.8856	**0.8655**	0.7210	0.8952	**0.8989**	0.6970	0.7270	**0.7037**	0.6312
	F1-score	0.8909	**0.8357**	0.8112	0.9203	**0.9139**	0.8281	0.9312	**0.9318**	0.8127	0.8348	**0.8191**	0.7673
2013	Accuracy	0.7645	0.7067	**0.7175**	0.8549	**0.8281**	0.7373	0.8388	**0.8672**	0.7294	0.6911	**0.6871**	0.6780
	AUC	0.6969	0.6068	**0.7304**	0.7082	0.6137	**0.7403**	0.7534	0.6640	**0.7601**	0.7435	0.6940	**0.7614**
	Recall	0.7842	**0.7218**	0.7208	0.8856	**0.8589**	0.7439	0.8622	**0.9013**	0.7336	0.6931	**0.6928**	0.6755
	F1-score	0.8631	0.8232	**0.8295**	0.9203	**0.9045**	0.8437	0.9101	**0.9277**	0.8377	0.8094	**0.8074**	0.7999
2015	Accuracy	0.8588	**0.7923**	0.7719	0.8924	**0.8733**	0.7373	0.8812	**0.9040**	0.7625	0.7575	**0.7561**	0.6802
	AUC	0.7872	**0.7778**	0.7232	0.7947	**0.7667**	0.7403	0.8203	**0.7879**	0.7687	0.8205	**0.8019**	0.7633
	Recall	0.8719	**0.7980**	0.7799	0.9079	**0.8885**	0.7439	0.8946	**0.9224**	0.7668	0.7580	**0.7574**	0.6778
	F1-score	0.9227	**0.8812**	0.8683	0.9422	**0.9313**	0.8437	0.9357	**0.9489**	0.8616	0.8579	**0.8571**	0.8035
2018	Accuracy	0.7186	**0.6968**	0.6442	0.8136	**0.7941**	0.6960	0.7774	**0.8127**	0.6889	0.6659	**0.6660**	0.6156
	AUC	0.7291	**0.7105**	0.6886	0.7396	0.6893	**0.7072**	0.7592	0.7159	**0.7229**	0.7589	**0.7229**	0.7213
	Recall	0.7284	**0.7035**	0.6457	0.8419	**0.8246**	0.7047	0.7958	**0.8428**	0.6943	0.6632	**0.6672**	0.6077
	F1-score	0.8274	**0.8111**	0.7706	0.8933	**0.8813**	0.8111	0.8687	**0.8929**	0.8051	0.7862	**0.7872**	0.7454

Note: Note: Non-C = Non-Causal. All values represent mean performance. Confidence intervals omitted for brevity.

Table 4: Validate the predictive performance of cognitive impairment labels on the CFPS dataset using identified causal features.

Model	Accuracy	AUC	Recall	F1-score
		2016		
RF	0.6649	0.8035	0.6193	0.7980
SVM	0.7480	0.8077	0.7605	0.7528
XGBoost	0.6953	0.7638	0.6935	0.6974
		2018		
RF	0.7069	0.7765	0.7053	0.7100
SVM	0.8183	0.8786	0.8375	0.8223
XGBoost	0.8196	0.8840	0.8300	0.8222
		2020		
RF	0.7271	0.7932	0.7306	0.7292
SVM	0.7087	0.7677	0.6877	0.7032
XGBoost	0.7204	0.7887	0.7179	0.7203
		2016-2020 (3-wave)		
LSTM	0.7464	0.8134	0.7509	0.7489

Note: All values represent mean performance across bootstrap samples. 95% confidence intervals are omitted for brevity.

Education Level, Individual Income, and Eyesight remain in mid-to-high positions. This correspondence between causal identification and predictive importance supports the validity of our approach.

Extensive literature corroborates the causal pathways identified in our study. Sensory and psychological factors, such as visual impairment and depression, are well-documented drivers of functional decline and mental health deterioration (Kim & Park, 2023; Kupfer & Frank, 2003). The vulnerability to these conditions is further modulated by demographic and socioeconomic determinants; for example, gender-based disparities and lower income levels are consistently associated with adverse mental health trajectories (Gresenz et al., 2001; Kiely et al., 2019; O'Neal, 2024). Regarding cognitive performance, educational attainment and physical exercise are established protective factors that increase cognitive reserve and spatial memory (Cassilhas et al., 2012; Christensen et al., 2011). Conversely, advanced age, high BMI, and depressive symptoms are linked to structural neural alterations and accelerated memory decline (Cheke et al., 2017; Dillon, 2015; Small, 2002). Furthermore, socioeconomic status and self-rated health remain pivotal predictors of long-term memory maintenance (Bendayan et al., 2017; Herrmann & Guadagno, 1997). Collectively, this convergence of empirical evidence substantiates the theoretical validity of our causally-identified feature set.

Limitation Constrained by observational data, this study leaves potential unmeasured confounders unaccounted for. Therefore, the identified features require further validation through controlled clinical studies to confirm their underlying mechanisms and clinical significance.

Conclusion

Our framework bridges causal discovery with cognitive theory to extract domain-specific drivers of cognitive aging from sparse longitudinal data. We demonstrate that distinct subdomains possess divergent pathways: sensory deficits primarily cascade into Mental State decline, while lifestyle factors modulate Contextual Memory by influencing cognitive reserve. Consistent across multi-wave cohorts, our findings highlight the potential benefits of moving beyond generalized cognitive assessments. Future interventions targeting these specific pathways—such as sensory correction or lifestyle enrichment—can inform tailored prevention strategies.

References

Bendayan, R., Piccinin, A. M., Hofer, S. M., & Muniz, G. (2017). Are changes in self-rated health associated with memory decline in older adults? *Journal of Aging and Health*, *29*(8), 1410–1423.

Black, S. (2011). Vascular cognitive impairment: Epidemiology, subtypes, diagnosis and management. *Journal of the Royal College of Physicians of Edinburgh*, *41*(1), 49–56.

Cassilhas, R. C., Lee, K. S., Fernandes, J., Oliveira, M. G. M. d., Tufik, S., Meeusen, R., & De Mello, M. (2012). Spatial memory is improved by aerobic and resistance exercise through divergent molecular mechanisms. *Neuroscience*, *202*, 309–317.

Chao, Y.-T., Kuo, F.-H., Lee, Y.-S., Huang, Y.-H., Weng, S.-C., Chou, Y.-Y., Lin, C.-S., & Lin, S.-Y. (2022). Characteristics and outcome determinants of hospitalized older patients with cognitive dysfunction. *International Journal of Environmental Research and Public Health*, *19*(1), 584.

Cheke, L. G., Bonnici, H. M., Clayton, N. S., & Simons, J. S. (2017). Obesity and insulin resistance are associated with reduced activity in core memory regions of the brain. *Neuropsychologia*, *96*, 137–149.

Christensen, H., Anstey, K. J., Leach, L. S., & Mackinnon, A. J. (2011). Intelligence, education, and the brain reserve hypothesis. In *The handbook of aging and cognition* (pp. 133–188). Psychology Press.

Cockrell, J. R., & Folstein, M. F. (2002). Mini-mental state examination. *Principles and practice of geriatric psychiatry*, 140–141.

Didic, M., Barbeau, E. J., Felician, O., Tramoni, E., Guedj, E., Poncet, M., & Ceccaldi, M. (2011). Which memory system is impaired first in alzheimer's disease? *Journal of Alzheimer's Disease*, *27*(1), 11–22.

Dillon, D. G. (2015). The neuroscience of positive memory deficits in depression. *Frontiers in psychology*, *6*, 1295.

Fehr, J., Piccininni, M., Kurth, T., & Konigorski, S. (2023). Assessing the transportability of clinical prediction models for cognitive impairment using causal models. *BMC Medical Research Methodology*, *23*(1), 187.

Fogg, C., Meredith, P., Culliford, D., Bridges, J., Spice, C., & Griffiths, P. (2019). Cognitive impairment is independently associated with mortality, extended hospital stays and early readmission of older people with emergency hospital admissions: A retrospective cohort study. *International journal of nursing studies*, *96*, 1–8.

Gong, J., Wang, G., Wang, Y., Chen, X., Chen, Y., Meng, Q., Yang, P., Yao, Y., & Zhao, Y. (2022). Nowcasting and forecasting the care needs of the older population in china: Analysis of data from the china health and retirement longitudinal study (charls). *The Lancet Public Health*, *7*(12), e1005–e1013.

Gresenz, C. R., Sturm, R., Tang, L., et al. (2001). Income and mental health: Unraveling community and individual level relationships. *Journal of Mental Health Policy and Economics*, *4*(4), 197–204.

Harvey, P. D. (2019). Domains of cognition and their assessment. *Dialogues in clinical neuroscience*, *21*(3), 227–237.

Hastie, T. J. (2017). Generalized additive models. *Statistical models in S*, 249–307.

Herrmann, D., & Guadagno, M. A. (1997). Memory performance and socio-economic status. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, *11*(2), 113–120.

Hu, X., Nie, Z., Ou, Y., Qian, Z., McMillin, S. E., Aaron, H. E., Zhou, Y., Dong, G., & Dong, H. (2022). Air quality improvement and cognitive function benefit: Insight from clean air action in china. *Environmental Research*, *214*, 114200.

Johnson, D. K., Morris, J. C., & Galvin, J. E. (2005). Verbal and visuospatial deficits in dementia with lewy bodies. *Neurology*, *65*(8), 1232–1238.

Kiely, K. M., Brady, B., & Byles, J. (2019). Gender, mental health and ageing. *Maturitas*, *129*, 76–84.

Kim, A. M., & Park, J.-H. (2023). Mental health and depressive mood in people with visual impairments. *Journal of Visual Impairment & Blindness*, *117*(4), 314–325.

Kupfer, D., & Frank, E. (2003). Comorbidity in depression. *Acta Psychiatrica Scandinavica*, *108*, 57–60.

Luo, Y., Zhong, Y., Pang, L., Zhao, Y., Liang, R., & Zheng, X. (2021). The effects of indoor air pollution from solid fuel use on cognitive function among middle-aged and older population in china. *Science of the Total Environment*, *754*, 142460.

Mattke, S., Batie, D., Chodosh, J., Felten, K., Flaherty, E., Fowler, N. R., Kobylarz, F. A., O'Brien, K., Paulsen, R., Pohnert, A., et al. (2023). Expanding the use of brief cognitive assessments to detect suspected early-stage cognitive impairment in primary care. *Alzheimer's & Dementia*, *19*(9), 4252–4259.

Nasreddine, Z. S., Phillips, N. A., Bédirian, V., Charbonneau, S., Whitehead, V., Collin, I., Cummings, J. L., & Chertkow, H. (2005). The montreal cognitive assessment, moca: A brief screening tool for mild cognitive impairment. *Journal of the American Geriatrics Society*, *53*(4), 695–699.

O'Neal, M. A. (2024). Women and the risk of alzheimer's disease. *Frontiers in Global Women's Health*, *4*, 1324522.

Pu, L., Pan, D., Wang, H., He, X., Zhang, X., Yu, Z., Hu, N., Du, Y., He, S., Liu, X., et al. (2023). A predictive model for the risk of cognitive impairment in community middle-aged and older adults. *Asian Journal of Psychiatry*, *79*, 103380.

Roh, M., Dan, H., & Kim, O. (2021). Influencing factors of subjective cognitive impairment in middle-aged and older adults. *International Journal of Environmental Research and Public Health*, *18*(21), 11488.

Rojas-Saunero, L. P., Young, J. G., Didelez, V., Ikram, M. A., & Swanson, S. A. (2023). Considering questions before methods in dementia research with competing events and causal goals. *American Journal of Epidemiology*, *192*(8), 1415–1423. https://doi.org/10.1093/aje/kwad090

Shen, X., Ma, S., Vemuri, P., & Simon, G. (2020). Challenges and opportunities with causal discovery algorithms: Application to alzheimer's pathophysiology. *Scientific reports*, *10*(1), 2975.

Small, G. W. (2002). What we need to know about age related memory loss. *Bmj*, *324*(7352), 1502–1505.

Song, R., Fan, X., & Seo, J. (2023). Physical and cognitive function to explain the quality of life among older adults with cognitive impairment: Exploring cognitive function as a mediator. *BMC psychology*, *11*(1), 51.

Xie, Y., & Hu, J. (2014). An introduction to the china family panel studies (cfps). *Chinese sociological review*, *47*(1), 3–29.

Yu, X., Lophatananon, A., Holmes, V., Muir, K., & Guo, H. (2025). Investigating causal networks of dementia using causal discovery and natural language processing models. *npj Dementia*, *1*, 4. https://doi.org/10.1038/s44400-025-00006-2

Zhao, Y., Hu, Y., Smith, J. P., Strauss, J., & Yang, G. (2014). Cohort profile: The china health and retirement longitudinal study (charls). *International journal of epidemiology*, *43*(1), 61–68.