Modeling Alzheimer's disease: Bayesian copula graphical model from demographic, cognitive, and neuroimaging data

Lucas Vogels[1] (iD), **Reza Mohammadi[1]** (iD), **Marit Schoonhoven[1]** (iD), **Ş Ilker Birbil[1]** (iD), **Martin Dyrba[2]** (iD) **and for the Alzheimer's Disease Neuroimaging Initiative**[*]

Abstract

Background: The early detection of Alzheimer's disease (AD) requires an understanding of the relationships between a wide range of features. Conditional independencies and partial correlations are suitable measures for these relationships, because they can identify the effects of confounding and mediating variables.
Objective: To estimate conditional dependencies and partial correlations between relevant features in AD using a Bayesian approach to Gaussian copula graphical models (GCGMs). This approach has two key advantages. First, it includes binary, discrete, and continuous variables. Second, it quantifies the uncertainty of the estimates. Despite these advantages, Bayesian GCGMs have not been applied to AD research yet.
Methods: We design a GCGM to find the conditional dependencies and partial correlations among brain-region specific gray matter volume and glucose uptake, amyloid-beta levels, demographic information, and cognitive test scores. We applied our model to 1 022 participants, including healthy and cognitively impaired, across different stages of AD.
Results: We found that aging reduces cognition through three indirect pathways: hippocampal volume loss, posterior cingulate cortex (PCC) volume loss, and amyloid-beta accumulation. We found a positive partial correlation between being woman and cognition, but also discovered four indirect pathways that dampen this association in women: lower hippocampal volume, lower PCC volume, more amyloid-beta accumulation, and less education. We found limited relations between brain-region specific glucose uptake and cognition, but discovered that the hippocampus and PCC volumes are related to cognition.
Conclusions: This study shows that the use of GCGMs offers valuable insights into AD pathogenesis.

Introduction

Alzheimer's disease (AD) is clinically characterized by amyloid-beta and tau accumulation. AD is the primary cause of dementia, a condition defined by a decline in cognitive and executive functions.[1] Despite extensive research, underlying causes of pathological changes in AD remain unknown.[2]

Traditional research in AD typically selects a set of (dependent) variables and assesses the statistical associations with other variables (predictors).[3,4] Although this approach can lead to valuable insights, it overlooks the complex network of relations underlying AD pathogenesis. For instance, predictor variables may influence each other (a

[1]Amsterdam Business School, University of Amsterdam, Amsterdam, The Netherlands
[2]German Center for Neurodegenerative Diseases (DZNE), Rostock, Germany

[*]Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (https://adni.loni.usc.edu). As such, the investigators within the ADNI provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: https://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf

Corresponding author:
Lucas Vogels, Amsterdam Business School, Plantage Muidergracht 12, 1018 TV Amsterdam, The Netherlands.
Email: l.f.o.vogels@uva.nl

phenomenon known as collinearity), or a confounding variable may obscure the true relationship between predictors and dependent variables, potentially leading to unreliable results.

Beyond this traditional approach, literature therefore increasingly treats relations between variables from a network perspective. A popular example of this perspective is the field of brain connectivity, which studies the network of relationships between different brain regions.[5] Brain connectivity has three subfields: structural, functional, and effective connectivity. Structural connectivity uncovers the physical relations among brain regions,[6] for instance using diffusion tensor imaging.[7] Functional connectivity involves the covariance, correlation, or other statistical dependency among brain regions.[8] Effective connectivity discovers the causal relations among brain regions[9]; a common statistical technique to study effective connectivity is Dynamic Causal Modeling.[10]

The present article studies the network of associations between regional gray matter volume and glucose metabolism, but goes beyond the standard brain connectivity approach by also including demographic factors, cognitive scores, and the global amyloid-beta accumulation. In total, we included 19 features that have been shown to play a role in the pathogenesis of AD. Network models between these variables are not new in AD literature and are predominantly causal of nature. A popular tool for causal inference is Structural Equation Modeling (SEM), a family of statistical techniques that, given the causal relations between variables, can estimate the effect size of each relationship.[11] Mediation analysis is part of the SEM family and studies the contribution of a third variable to the causal dependency between two other variables. It is a common technique in AD research.[12–14] Other applications of SEM in AD include path analysis,[15] factor analysis,[16] and structural regression.[17] For causal inference in AD research, there are various alternatives to SEM including Structural Causal models,[18] System Dynamic models,[19] and Dynamic Causal Modeling.[10]

This article considers undirected graphs. The resulting relations are therefore not causal of nature, as opposed to the directed and causal relations studied by SEMs. Undirected graphs have three main advantages over classical SEMs. First, undirected networks allow cycles. Often a variable A causes B and B causes A. We refer to such relation as a cycle. Cycles are common in AD pathology. SEMs, however, assume the network is acyclic and are therefore not suitable for inference on cyclic relations. Second, causal relations are hard to detect and generally need large sample sizes and longitudinal data. These data are costly to obtain. Inference on undirected networks is possible with smaller sample sizes and does not require longitudinal data. Lastly, undirected networks do not require prior knowledge about the structure of the graph. They can, therefore, be used without limitations on the structure. SEMs do need a causal relation graph as input. SEMs are able to refine the causal

structure of this input graph, but these refinements suffer from a poor performance.[18] SEMs are therefore often restricted to variables and relations of which the causal structure is available a priori. Undirected networks, like the graphical model estimated in this article, can facilitate to form hypotheses about the structure of a causal network. Thus, they have the potential to improve the performance of SEMs and other causal inference methods.

Formally, in an undirected network, an edge between two variables represents some measure of statistical dependency. A common choice for this dependency is the Pearson's correlation or the covariance. However, these metrics can lead to spurious associations by overlooking confounding or mediating factors. Ideally one wants a correlation that is corrected for such factors. Such a correlation exists and is called the partial correlation. Like the Pearson correlation, the partial correlation takes on a value in the range from -1 to 1. Under the assumption that all variables are multivariate normally distributed, a zero partial correlation implies conditional independence. Two variables A and B are conditionally independent when, given all other variables, there is no relation between A and B. Or, put differently, if we keep all other variables fixed, knowing the value of variable A, does not give any information about the value of variable B. One can depict conditional dependencies in an undirected graphical model, where nodes denote random variables and an edge between two nodes/variables is included if and only if the variables are conditionally dependent.[20,21] Such graphical models depict confounding and mediating pathways between large numbers of variables in one figure. They constitute, therefore, a powerful tool to uncovering the complex associations involved in AD, as previously demonstrated.[22]

Conditional dependencies, partial correlations, and graphical models are commonly estimated using a frequentist approach. This approach can estimate partial correlations and predict whether any two variables are conditionally dependent or not. In the field of brain connectivity in AD, there are two articles applying frequentist graphical models to uncover conditional dependencies between brain region-specific gray matter volume and glucose uptake.[23,24] Although frequentist approaches are relatively straightforward, they are not able to provide any uncertainty around the estimated conditional dependencies and partial correlations, a concept referred to as model uncertainty. To overcome this limitation, we employed a Bayesian approach, which offers a key advantage: it not only estimates conditional (in)dependence and partial correlations but also quantifies the uncertainty of these estimates.

Quantifying this uncertainty is important. Conditional independence, and statistical associations in general, are not black and white; based on the data, relations can be opaque. Consider, for example, the ambiguous role that sex plays in the pathology of AD.[25] When answering the question "Are two variables conditionally dependent?", a

simple yes or no might overlook the complexity of the relation. Instead, the answer "There is conditional dependence with 75%", gives more color. Uncertainty also allows researchers to state that, based on the data, it is not clear if some variables are conditionally dependent. Quantifying uncertainty becomes even more pivotal if conditional dependence graphs are used as a starting point to develop a causal model, in which mistakes in the starting graph might trickle down to spurious causal results.[18]

Despite the advantage that model uncertainty provides, the Bayesian approach is still uncommon in the field of AD research. This is largely due to the dated notion that the Bayesian approach is complex and inefficient. Although this claim is historically true, recent advancements have made Bayesian methods more computationally efficient,[26,27] making their application to AD research possible. For example, Bayesian methods have been applied to estimate conditional dependence between brain regions[22,28] and to determine the uncertainty of functional connectivity estimates.[29] All these applications, however, used so-called Gaussian graphical models (GGMs). In GGMs, all variables are assumed to follow a multivariate normal distribution. In the present study, we employed Gaussian copula graphical models (GCGMs). This methodology enabled the integration of non-normal variables, such as discrete (e.g., amyloid-beta levels), binary (e.g., sex) and continuous variables (e.g., MRI, FDG-PET).

This study presents a Bayesian approach to GCGMs. Our methodology has three capabilities. (i) the estimation of an undirected network depicting the conditional dependencies and partial correlations between variables, (ii) the quantification of the uncertainty of these estimates and, (iii) the inclusion of all types of variables whether normally distributed or not. Bayesian GCGMs have been successfully applied in other domains,[30–32] but remain, to the best of our knowledge, unexplored in the AD domain.

The application of GCGMs to AD, allowed us to study the conditional dependence structure among 19 relevant features. Specifically, our primary objective was to identify the conditional dependence pathways between demographic variables and cognition (i.e., memory and executive function scores), and between neuro-imaging variables (i.e., gray matter volume, glucose uptake along with amyloid-beta levels) and cognition. Furthermore, we explored the conditional dependency pathways through which demographic variables influence gray matter volume and glucose uptake, and investigated the conditional dependencies between brain region-specific volume and glucose metabolism.

Methods

Subjects

We obtained the data for this study from the Alzheimer's Disease Neuroimaging Initiative (ADNI), which provides a public database for AD research including clinical, neuropsychological, neuroimaging, and biomarker data. A complete description of the ADNI and up-to-date information is available at https://adni.loni.usc.edu. For this study, we selected the baseline examinations of the ADNI-GO and ADNI-2 phases. The final dataset included 1 022 participants selected based on the availability of concurrent T1-weighted structural MRI, FDG-PET, amyloid-sensitive AV45-PET data, neuropsychological assessments, and blood-based *APOE4* genotyping.

After initial quality control and data preparation, the final sample included 345 cognitively normal control subjects (CN), 297 patients with amnestic early mild cognitive impairment (EMCI), 205 patients with amnestic late mild cognitive impairment (LMCI), and 175 patients with AD. We included two cognitive composite scores assessing memory (ADNI-MEM) and executive function (ADNI-EF). ADNI-MEM is a weighted average of seven different memory tests.[33] Similarly, ADNI-EF combines eleven different executive functions tests in a single score.[34] These composite scores provide a higher robustness than individual test scores and a sound psychometric ability to differentiate between different cognitive profiles. Both ADNI-MEM and ADNI-EF range between −3 and 3, a higher score indicating a better performance. Demographics and cognitive test scores of the different diagnostic groups are summarized in Table 1.

Data preparation and feature extraction

In this study, we included 19 variables: three demographic variables, two composite cognitive test scores, the number of *APOE4* alleles, the amyloid stage, the gray matter volume of six brain regions, and the glucose metabolism of the same six regions. Supplemental Table 1 in the Supplemental Material lists all the variables along with their abbreviations and their data type (continuous, discrete, binary or categorical). In this subsection we discuss how we obtained and prepared these variables.

The six brain regions were selected based on our hypotheses and a priori literature findings of early involvement in neurodegenerative processes in AD: hippocampus, caudate, putamen, thalamus, posterior cingulate cortex (PCC), and precuneus.[35] All six regions are defined by the Harvard-Oxford atlas.[36] For each region, we took the average of the left and right brain regions. This avoided the number of variables becoming too large, which would hinder the interpretability of the results. In order to test the robustness of our model, we ran it a second time, this time including the left and right regions separately.

This study used three biomarkers, all associated with AD: amyloid-beta accumulation,[1] glucose metabolism,[37] and gray matter volume.[38] They were measured using AV45-PET, FDG-PET, and T1 weighted MRI scans,

Table 1. Subject characteristics per diagnosis group. Values denote the mean and standard deviation (in parentheses).

	CN	EMCI	LMCI	AD
Sample size	345	297	205	175
Female (% of total)	53%	44%	43%	42%
Age (y)	74.6 (6.5)	71.6 (7.4)	74.1 (8.1)	75.1 (8.0)
Education (y)	16.5 (2.7)	16.0 (2.7)	16.2 (2.9)	15.9 (2.7)
ADNI-MEM	1.1 (0.60)	0.58 (0.60)	0.04 (0.69)	−0.91 (0.59)
ADNI-EF	0.83 (0.76)	0.48 (0.77)	0.14 (0.87)	−0.84 (0.90)

respectively. The data resulting from these scans were processed as in earlier article.[22,35,39] The MRI scans were segmented into gray matter, white matter, and cerebrospinal fluid and spatially normalized to an aging/AD-specific reference template using SPM8 (Wellcome Centre for Human Neuroimaging, University College London) and VBM8 (Structural Brain Mapping Group, University of Jena) toolboxes, and the DARTEL algorithm.[40] FDG- and AV45-PET scans were co-registered to the T1 scan and spatially normalized by applying the deformation fields of the T1 scan. To reduce the effect of partial volume signal arising from the low spatial resolution of PET scans, we applied partial volume correction using a three-compartment model and the MRI-derived tissue segments.[41] We scaled the regional gray matter volumes proportionally by the total intracranial volume, the regional FDG-PET values by pons uptake, and the regional AV45-PET values by whole-cerebellum uptake. As amyloid-beta spreading has been shown to follow a specific sequence, we captured the amyloid-beta accumulation in a single five-level amyloid score reflecting the global severity of amyloid-beta deposition.[39] Glucose metabolism and gray matter volume were obtained for each of the six brain regions of interest.

We included the three demographic variables that are well-known to modulate the risk of AD: age, sex, and education.[1] Sex is represented as a binary variable, equal to one for women and zero for men. Education is expressed in the number of years of formal education received. The number of *APOE4* alleles is strongly associated with amyloid-beta production and AD[1] and, therefore, included as a variable in our analysis. We coded this variable as binary, i.e., zero for patients with no *APOE4* allele and one for patients with at least one *APOE4* allele. Lastly, we included memory and executive function composite scores named ADNI-MEM and ADNI-EF, respectively.

Modeling: Bayesian inference in GCGMs

We used a Bayesian framework within GCGMs. This framework combines three essential components: graphical models for estimating conditional dependencies and partial correlations, Bayesian methods for determining the uncertainty of those estimates, and the Gaussian copula for

accommodating diverse data types (continuous, discrete, and binary). In this section, we provide a detailed explanation of each component.

Graphical models[20] represent conditional dependencies between variables in the form of a graph G, in which each node represents a variable, and edges connect pairs of variables that are conditionally dependent. GGMs, also known as Markov random fields, assume that all variables $Z_1, .., Z_p$ come from a multivariate normal distribution with mean 0 and unknown covariance matrix Σ. A parameter of interest is the unknown precision matrix $\mathbf{K} = \Sigma^{-1}$ with entries k_{ij}, because a simple transformation of the precision matrix $\mathbf{K}$ gives the partial correlations. Moreover, the sparsity pattern of $\mathbf{K}$ directly encodes the conditional dependence structure. That is, $k_{ij} = 0 \Leftrightarrow Z_i \text{and} Z_j \text{areconditionallyindependent}$. GGMs leverage this relationship to recover the graphical model G.

The aim of GGMs is to use observations of the variables to estimate the precision matrix $\mathbf{K}$ and conditional dependence graph G. These observations are denoted by the $n \times p$ matrix $\mathbf{Z}$, which contains n observations of each of the variables $Z_1, .., Z_p$. With the data $\mathbf{Z}$, one can estimate the precision matrix $\mathbf{K}$ and conditional dependence graph G. This is commonly done using a frequentist approach, such as the graphical lasso.[42] This approach renders a single estimate of G and $\mathbf{K}$, also called a point estimate. In contrast, the Bayesian approach estimates an entire distribution called the posterior. It is given by $P(G, \mathbf{K}|\mathbf{Z})$ and denotes the probability, that, given the data $\mathbf{Z}$, the true conditional dependence graph equals G and the true precision matrix equals $\mathbf{K}$. Due to the posterior, we can go beyond frequentist point estimates and make claims such as: "variable A and B are conditionally dependent with a probability of 60%", or "the partial correlation between A and B is between 0.1 and 0.2 with a probability of 90%". In other words, the posterior provides model uncertainty.

Before obtaining any data, a researcher can already have a belief about G and $\mathbf{K}$. For example, based on literature one might expect the presence of an *APOE4* allel to be conditionally dependent with the amyloid stage. In Bayesian statistics, such beliefs are captured in a distribution called the prior.[43] In GGMs, this prior is denoted by $P(\mathbf{K}, G)$. When no prior information is available, one can choose an uninformative prior that deems every conditional dependence equally likely. In this study, we selected a prior probability

of conditional dependence of 20% for all variable pairs. Now, the posterior reflects how the data $\mathbf{Z}$ update this prior belief. This is captured in Bayes' formula given by

$$P(G, \mathbf{K}|\mathbf{Z}) \propto P(\mathbf{Z}|\mathbf{K})P(\mathbf{K}, G),$$

in which $\propto$ denotes equality up to a constant. Notice in Bayes' formula how the posterior is a combination of the data, given by the likelihood $P(\mathbf{Z}|\mathbf{K})$, and the prior belief, given by $P(\mathbf{K}, G)$. Estimating the posterior is commonly done with Markov Chain Monte Carlo (MCMC) algorithms. Such algorithms iteratively obtain samples $(G, \mathbf{K})$ of this posterior. These samples can then be used to make a variety of uncertainty claims about G and $\mathbf{K}$. For example, if 30% of the sampled graphs contain an edge between variable A and B, then the probability of conditional dependence between A and B is 30%.

So far, we have assumed that the data $\mathbf{Z}$ come from a multivariate normal distribution. This assumption, however, limits the applicability of GGMs in practical settings, as real-world data often include non-Gaussian variables. For instance, in our study, the dataset comprises non-Gaussian continuous variables, binary variables (*APOE4* and sex), ordinal variables (amyloid stage), and discrete variables (age and education). To overcome this limitation, we employed Gaussian copula graphical models (GCGMs). These models can handle mixed variable types while preserving the theoretical advantages of GGMs. They provide a robust representation of conditional dependency structures and have been successfully applied in neuroscience and brain connectivity studies.[32,44] In GCGMs, the observed variables $Y_1, , Y_p$ are transformed into Gaussian variables $Z_1, , Z_p$. The partial correlations and conditional dependence structure is then calculated for these transformed variables $Z_1, , Z_p$. Continuous variables (brain-region-specific glucose uptake, gray matter volume, ADNI-MEM and ADNI-EF), were transformed once, before the start of the MCMC chain, with a semiparametric transformation.[45] This transformation is such that the conditional dependencies and partial correlations between the transformed variables reflect those of the observed variables.[45] We treated age, albeit a discrete variable, as a continuous variable too. Discrete or categorical variables (sex, education, *APOE4* presence and amyloid stage) were transformed at every iteration in the MCMC algorithm according to the copula framework.[46] The estimated partial correlations and conditional dependencies resulting from this discrete variable transformation reliably represent the underlying structure of observed data, but do not come with a theoretical guarantee.[46] In this study we used 120,000 MCMC iterations, discarding the first 20,000 iterations as burn-in. Convergence diagnostics confirmed that the remaining 100,000 iterations were sufficient to provide reliable and stable estimates.

We refer interested readers to a more detailed explanation of Bayesian Gaussian copula graphical models.[31] For guidance on implementing this approach, we refer readers to the R package BDgraph.[47] To ensure reproducibility, the R scripts used to produce our results can be found on the GitHub page https://github.com/lucasvogels33/Modeling-AD-Bayesian-GCGM-from-Demographic-Cognitive-and-Neuroimaging-Data.

Results

This section presents the results. It contains the estimated conditional dependency networks and the corresponding estimated partial correlations, but also showcases the uncertainty of these estimates. We also briefly discuss how the results change as the disease progresses and what the impact is of considering both the left and right side of each brain region.

Figure 1(a) displays all conditional dependencies in a network. In such a network, an edge between a pair of conditionally dependent variables (e.g., age and executive function) is also called a direct pathway. Some pairs of variables are connected via two or more edges. We refer to such connections as indirect conditional dependence pathways. We observe conditional dependency pathways (both direct and indirect), between age and cognition (Figure 1(b)), as well as between sex and cognition (Figure 1(c)). We observe in Figure 1(d) that the amyloid stage and brain-region specific gray matter volume are conditionally dependent with cognition, but report limited conditional dependency between brain-region specific glucose uptake and cognition. Figure 1(e) shows that both old age and being a woman are predominantly negatively partially correlated with brain-region specific volume and glucose metabolism. Lastly, Figure 1(f) depicts the ten conditional dependencies between brain-region specific volume and glucose uptake.

The conditional dependence networks in Figure 1 reveal what variables were likely to be conditionally dependent. They do not reveal, however, the strength of this dependence. These are given by the partial correlations and are shown, alongside the Pearson correlations in Figure 2. Pearson correlations were set to zero when the *p*-value exceeds 0.05, while partial correlations were set to zero when the corresponding edge inclusion probability is below 50%. The Pearson correlation heatmap (Figure 2(a)) is denser compared to the partial correlation heatmap (Figure 2(b)). The average absolute Pearson correlation was 0.17, while the average absolute partial correlation was only 0.07. Moreover, just 27% of Pearson correlations were set to zero versus 65% of partial correlations. Strong Pearson correlations (with an absolute value greater than 0.25) were fairly common (26%), while strong partial correlations were less frequent (7%).

We applied GCGMs to AD primarily to discover the conditional dependencies related to cognition (i.e., memory and executive function). Figure 3 presents the ten variables that were most likely conditionally dependent with memory. For each variable, the figure displays the
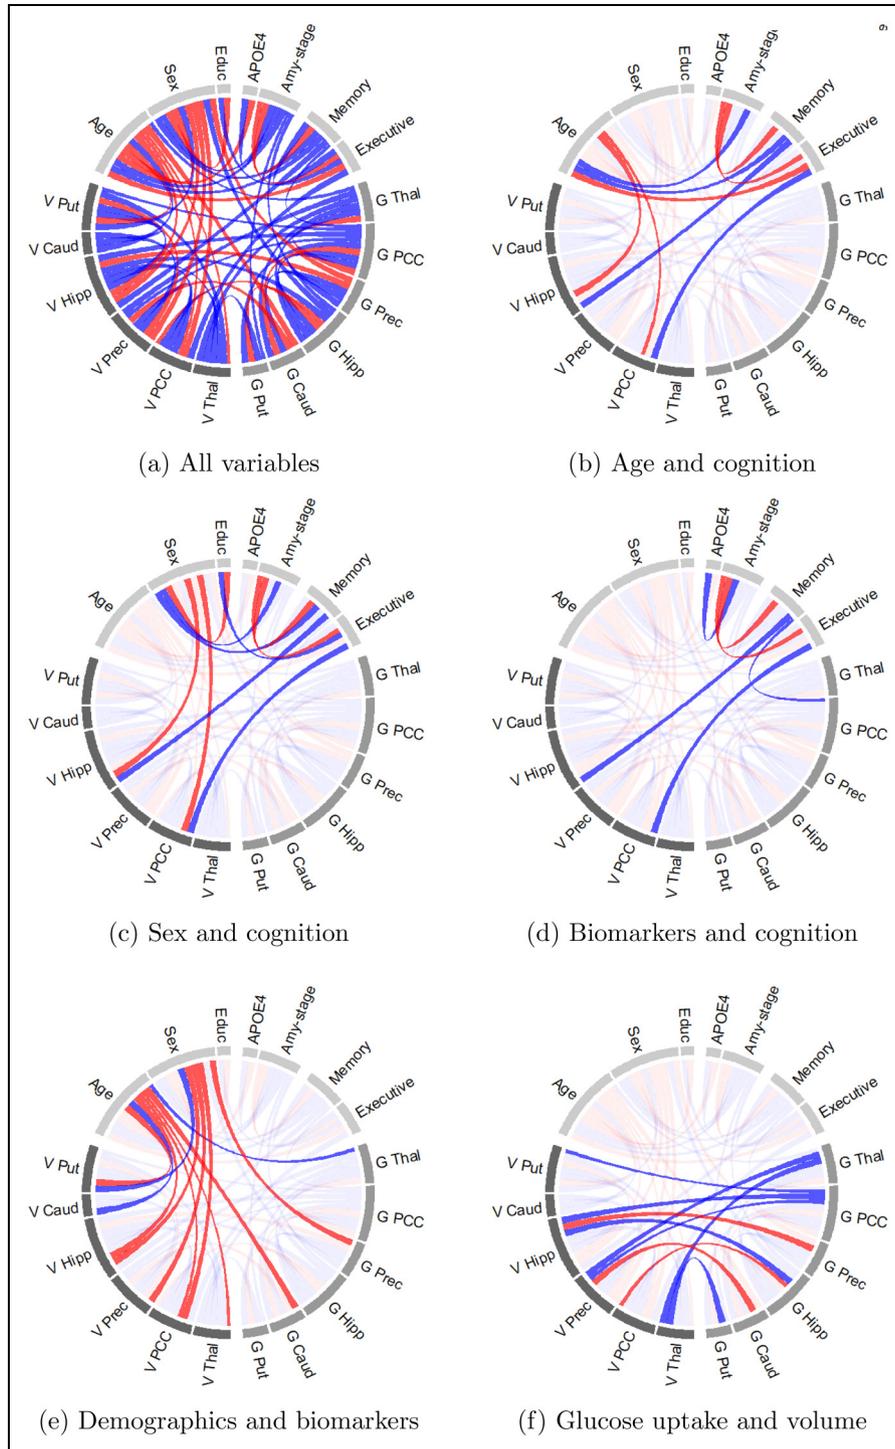
Figure 1. Visualization of conditional dependence among brain-region specific glucose uptake (G), brain-region specific gray matter volume (V), and demographic variables. An edge between variables indicates a conditional dependence with a probability of at least 50%. The width of the edges denotes the size of this probability ranging from 50% to 100%. A blue (red) edge denotes a positive (negative) partial correlation. (a) All variables; (b) Age and cognition; (c) Sex and cognition; (d) Biomarkers and cognition; (e) Demographics and biomarkers; (f) Glucose uptake and volume.

probability of conditional dependence (left) and the mean partial correlation (right). Memory was partially correlated with hippocampus volume (0.25), being a woman (0.3),

amyloid stage (−0.25), and age (0.1). Figure 4 presents the same information for executive function. Executive function was partially correlated with PCC volume (0.1)
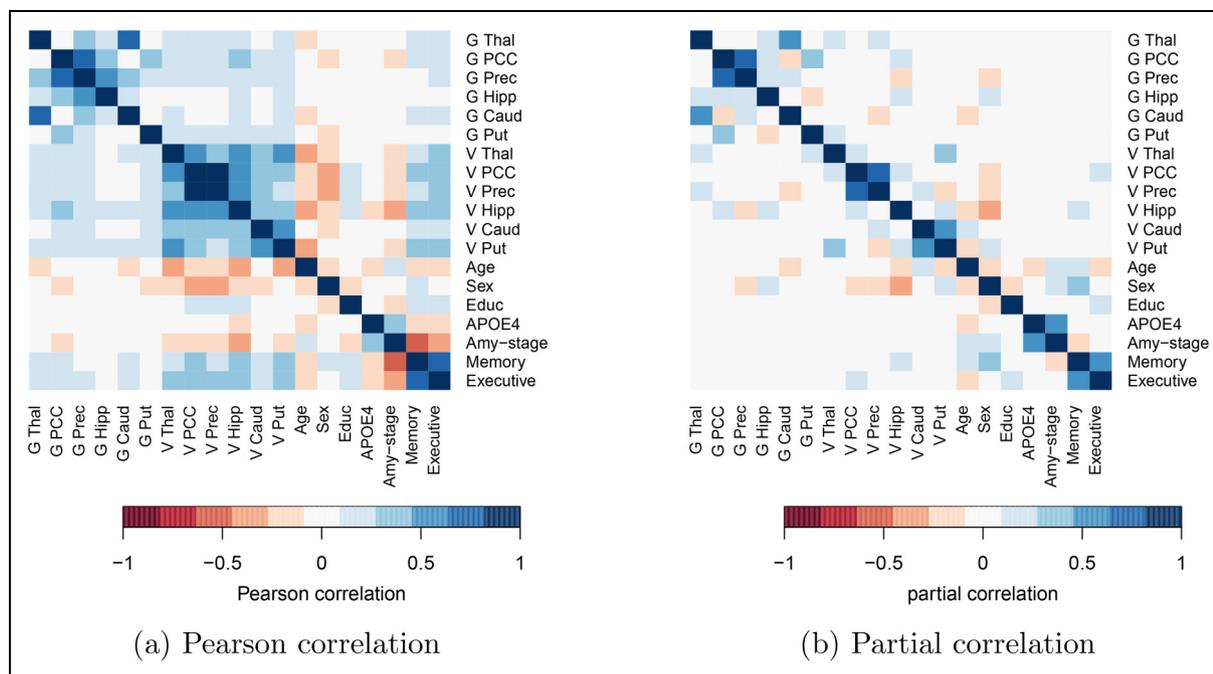
Figure 2. Pearson correlations (left) and partial correlations (right) among brain-region specific glucose uptake (G), brain-region specific gray matter volume (V), and demographic variables. Pearson correlations are set to zero when p-value >0.05. Partial correlations are set to zero when the corresponding edge inclusion probability is smaller than 50%. (a) Pearson correlation; (b) Partial correlation.

and age (−0.1). For both education and amyloid stage, we reported a 75% probability of conditional dependence with executive function.

The main advantage of the Bayesian approach is its ability to compute the uncertainty of its estimates. We showcase this in Figure 5 which presents the probability density plots of four selected partial correlations. The spike at zero in some of these plots denotes the probability that the partial correlation is zero, i.e., the probability that the variables were conditionally independent. The partial correlations depicted in the plots have low standard deviations (≤ 0.06). In fact, across all pairs of variables, the average standard deviation was just 0.02 with a maximum of 0.06. The estimated means of the partial correlations, as reported in Figure 2(b), were therefore likely to be close to the true partial correlations.

So far, we have considered the results of 1 022 patients combined over all four disease stages: healthy, EMCI, LMCI, and AD. Next, we examined how pathways vary across different disease stages by analyzing each stage separately. The resulting networks are shown in Supplemental Material 2. We observe that the disease-stage specific networks are sparser than the combined model, with approximately 80% of variable pairs being conditionally independent in each of the four stage specific models, compared to 65% in the combined model. The average partial correlation value also decreased from 0.07 in the combined model to around 0.05 in each of the disease-stage specific

models. This increased sparsity can be attributed to reduced within-group variance. This reduced variance made it harder for the model to deduce conditional dependencies. The reduced variance was observed in all variables, but was particularly present in the memory and executive function variables. The composite memory score, for example, ranged between −3 and 3 in the combined model. Among AD patients, however, it ranged between −3 and 0.5. This is why, for this disease stage, memory and executive function had almost no conditional dependencies left, see Supplemental Figure 1(e) in the Supplemental Material. In later disease stages, reduced variance also caused age and sex to lose conditional dependencies with *APOE4* and the amyloid score. The hippocampus volume was an exception to this trend and had more conditional dependencies as the disease progresses. Of the pathways that we found for the combined model, some continued to exist in all four disease stages. Examples are the conditional dependencies between age and brain volume, as well as between sex and brain volume. Other pathways of the combined model only appeared in one disease stage, predominantly in the EMCI stage. These included the pathway linking age to memory via the amyloid stage, or the pathway linking sex and memory through the hippocampus volume.

We also conducted a separate analysis in which we included the left and right brain regions separately. The resulting network with 31 variables is shown in the
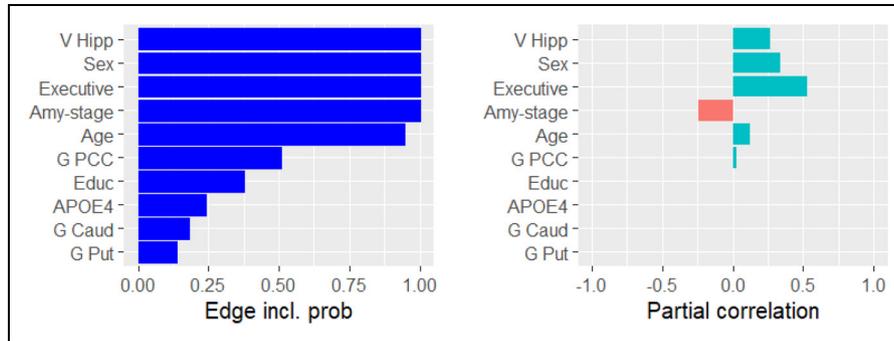
Figure 3. The probability of being conditionally dependent with memory (left) and the partial correlation with memory (right).
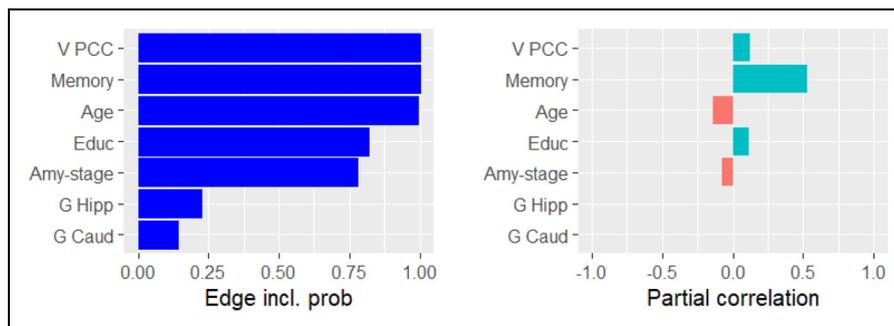
Figure 4. The probability of being conditionally dependent with executive function (left) and the partial correlation with executive function (right).

Supplemental Figures 2(a) and 3 and was sparser compared to the original, merged model. Specifically, 77% of conditional dependencies had a lower probability than 0.5, compared to 65% in the combined model. Additionally, the average partial correlation decreased to 0.047 from 0.07 in the original model. Our previous conclusions remain largely unchanged. Notably, sex continued to influence memory through education, amyloid-beta, hippocampal volume (V Hipp), and posterior cingulate cortex volume (V PCC), while age still impacted memory through amyloid-beta. However, some associations no longer held: age was no longer conditionally dependent with hippocampal volume, and amyloid-beta no longer associated with executive function. Separating the left and right brain region analysis, also provided extra insights: the volumes of the left and right brain regions were highly partially correlated, as is their glucose metabolism. Furthermore, the left hippocampal volume was specifically linked to memory, while the left PCC showed a gender-related association. Finally, the association between volume and metabolism was notably stronger within the same brain regions.

Discussion

This section interprets the results of the previous section and compares them with the existing results in the literature.

We also discuss the limitations of our work and end with a conclusion.

Interpretation of results

The interpretation of all conditional (in)dependencies and partial correlations goes beyond the scope of this paper. Instead, we interpret here those associations that we deem relevant.

We first elucidate the conditional dependencies (direct and indirect) between the demographic variables (age, sex, and education) and cognition (memory and executive function), starting with age. It is well-established in the literature that age is the biggest risk factor for AD.[1] In line with this, we found negative Pearson correlations between age and cognition (Figure 2(a)) and a negative partial correlation between age and executive function (Figure 4). However, we found a surprising positive partial correlation between age and memory. Specifically, we estimated that age and memory are conditionally dependent with 92% and that their partial correlation is between 0.05 and 0.2 with a 90% probability (Figure 5(a)). At first glance, this result suggested a possible limitation in the model. However, we found a similar result using a linear regression using the same 19 variables of our model with memory as a dependent variable: the resulting regression coefficient of
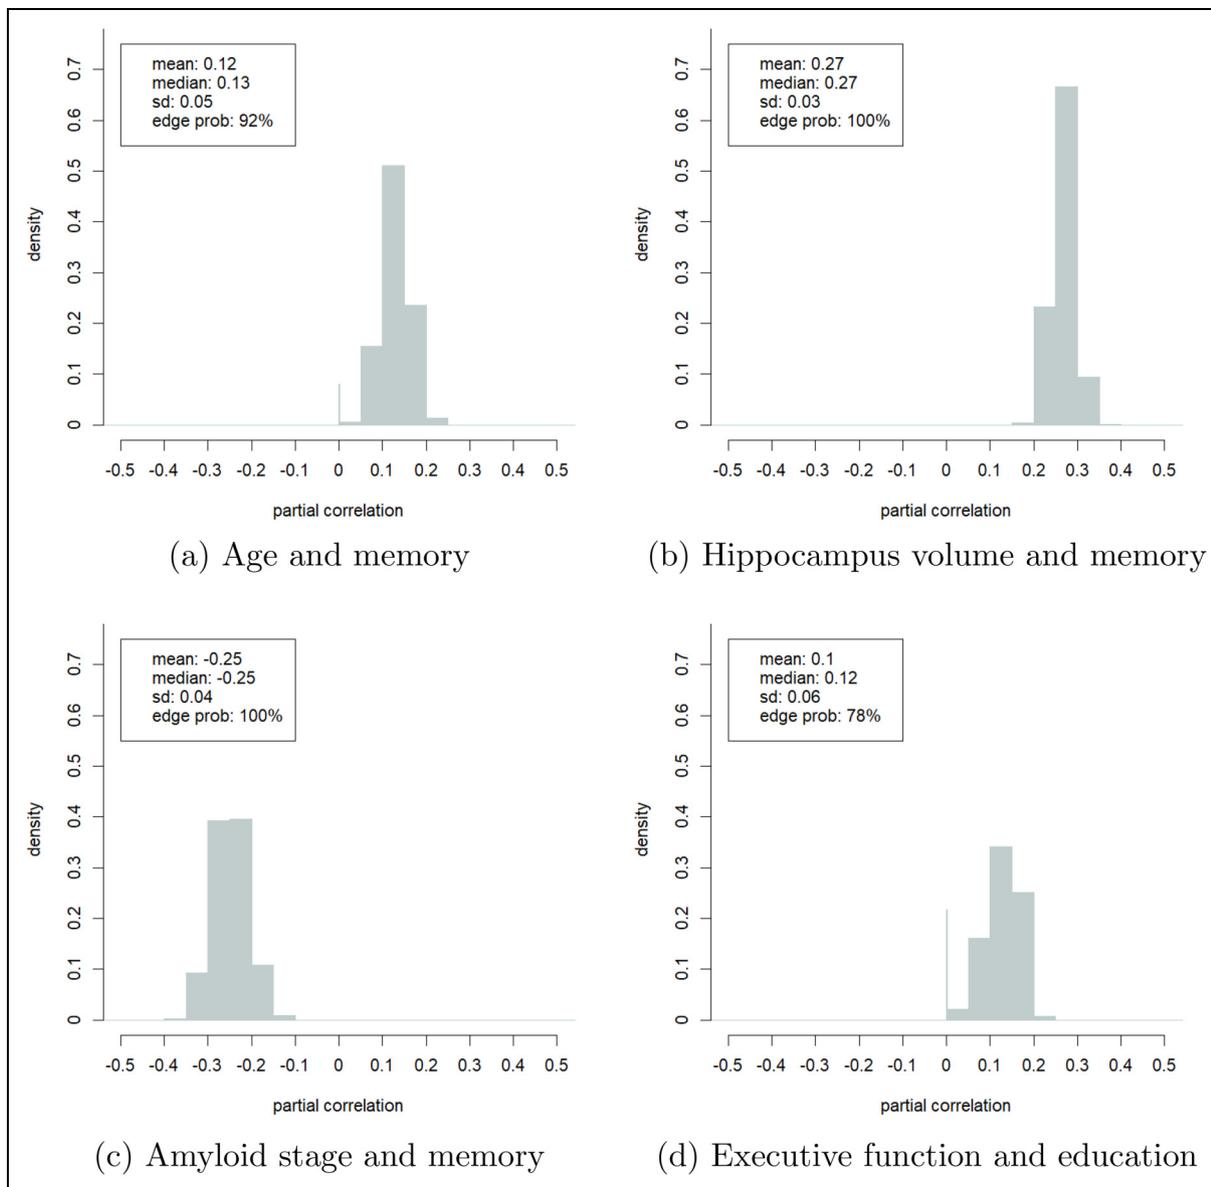
Figure 5. Estimated probability density plots of the partial correlations of four selected pairs of variables. (a) Age and memory; (b) Hippocampus volume and memory; (c) Amyloid stage and memory; (d) Executive function and education.

age was positive (0.011) and significant (p-value <0.001). The positive partial correlation and regression coefficient suggested that there was a confounding variable not in our model through which age positively associated with memory. A potential candidate for this confounder is lifestyle. As noted in literature, 45% of AD cases can be attributed to 14 lifestyle related risk factors,[1] of which only education was included in our model. The older patients in our dataset might be the ones with a better lifestyle, and therefore a better memory.

The positive partial correlation between age and memory also suggested that the main indirect pathways through which age contributes to memory decline were captured

in our model. Figure 1(b) shows two such pathways. Aging increased amyloid-beta deposition and reduced gray matter volume of the hippocampus and PCC, impairing memory and executive function. Both associations were confirmed in the literature.[1,48,49] The GCGM model allowed us to estimate the uncertainty of these estimates. For example, with more than 95% certainty we know that the partial correlation between the amyloid stage and memory was between -0.15 and -0.35 (Figure 5(c)). Similarly, we observed that with almost 100% certainty the partial correlation between memory and the gray matter volume of the hippocampus ranged between 0.2 and 0.35 (Figure 5(b)).

Next, we look at the relation between sex and cognition. We found a positive partial correlation between being a woman and memory (Figure 2(b)). The Pearson correlation between sex and memory was weaker (Figure 2(a)). This suggested that there may exist indirect conditional dependency pathways between sex and cognition, each of which dampening the positive effect of the observed direct partial correlation. Figure 1(c) depicts four such pathways. First, women in our dataset have spent, on average, one year less in education. This effect was also observed in literature.[1,50,51] The density plot in Figure 5(d) suggested that education and executive function were conditionally dependent with a probability of 78% with a partial correlation ranging between 0 and 0.2. Second, we observed a positive partial correlation between being a woman and amyloid accumulation. An association also found by others.[52] This association may provide evidence that amyloid-beta deposition is associated with our brain's immune system,[53] which was reported to be stronger in women.[54] Third, we found that women have smaller hippocampal volumes. A conclusion also found in a large study with 18 600 individuals.[55] Lastly, we found a negative partial correlation between sex and PCC gray matter volume. We found a 40% chance that memory was conditionally dependent on education (Figure 3) and a 78% chance that executive function was conditionally dependent on education (Figure 4). This is in line with the prevailing notion that education improves cognition and reduces the risk of developing AD.[1]

Next, we identified the conditional dependency pathways between neuro-imaging variables (gray matter volume, glucose uptake and amyloid-beta accumulation) and cognition. Figure 1(d) depicts all such dependencies. We reported a conditional dependency between the hippocampal volume and memory with a probability of one, see Figure 3. This supported the established notion that the hippocampus is associated with memory.[56] Figure 4 shows an almost certain conditional dependency between the PCC volume and executive function, confirming a theory that PCC is involved in attention.[57] Our findings suggested that the high Pearson correlations between brain-region specific volume and cognition (Figure 2(a)) were solely due to each region's conditional dependence with the hippocampus and PCC volume. Corrected for other variables, we found no relation between glucose metabolism and cognition. The link between brain-region specific glucose uptake and cognition has not been widely studied, although there is evidence for an association between PCC glucose uptake and cognition.[58] We found a conditional dependence between these two variables with a 50% probability, but the corresponding partial correlation was close to zero (Figure 3). We observed that the presence of at least one allele of the *APOE4* gene was partially correlated with amyloid-beta accumulation, which, in turn, was conditionally dependent with both executive function and memory. This came as no surprise, as this supported the widely accepted hypothesis that the presence of *APOE4* alleles is linked to amyloid-beta accumulation and cognitive decline.[48]

Third, we investigated the conditional dependencies between demographic variables (sex and age) and brain-region specific glucose uptake and volume. Figure 1(e) depicts such dependencies. We found that age and gender showed more conditional dependencies with brain-region specific volume than with brain-region specific glucose uptake. Age was negatively partially correlated with the volume of the putamen, hippocampus, PCC and thalamus. Age induced atrophy in these regions was also described in literature.[59] We reported negative partial correlations between being female and the volume of three brain regions: the hippocampus, precuneus and PCC. We also found a positive partial correlation between being a woman and the volume of the putamen, an association also reported by others.[60] Concerning glucose uptake, age was only negatively partially correlated with the caudate. This aligns with earlier findings that the precuneus cortex, hippocampus, thalamus, and putamen are among the regions whose metabolism is least affected by aging.[61]

Lastly, we looked at the conditional dependencies between brain-region specific volume and glucose metabolism. They are depicted in Figure 1(f). We found ten such conditional dependencies. Supplemental Figure 4 in the Supplemental Material shows that these conditional dependencies increased in number as the disease progresses. Among EMCI patients, we found almost no conditional dependencies between brain volume and glucose uptake, whereas in AD patients several such conditional dependencies can be observed. Among AD patients, the volume and glucose uptake of the hippocampus were positively partially correlated. The same was true for the thalamus and the putamen. This was in line with the hypothesis that in AD reduced glucose metabolism precedes neuronal loss and brain volume reduction.[37]

Limitations

Our Bayesian approach to GCGMs has three main limitations. First, it becomes slow for practical applications when the number of variables p and the number of observations n increases. The GCGM discussed in this paper had $p = 19$ variables and $n = 1022$ observations and ran within 10 minutes. A larger model, however, with more than 100 variables and/or more than 5000 observations would be infeasible for a Bayesian GCGM. Bayesian uncertainty evaluation relies on an MCMC algorithm that iteratively samples new graphs and precision matrices, both of a dimension $p \times p$. Increasing p therefore exponentially increases the running time and memory requirements of the model. Moreover, at every MCMC iteration, the algorithm needs to resample every non-continuous variable n

times. This resampling allows for the inclusion of binary, discrete and ordinal variables, but leads to a lack of scalability in the number of observations.

The second limitation of the GCGM is its linearity assumption. More specifically, partial correlations represent the strength of a linear relationship with the effect of other variables removed. Partial correlations can be over-, or underestimated, when the underlying relationship is not linear.

Lastly, GCGMs are undirected. This comes with advantages. GCGMs can deal with cyclic networks, require relatively few observations and do not need prior knowledge of the structure. However, the undirected nature of GCGMs make them not suited for identifying causal effects, such as the impact of a treatment on amyloid-beta accumulation, or the effect of a lifestyle change on cognition. The undirected networks provided by GCGMs can, however, be used as a starting point for causal analysis.

The particular application of GCGMs described in this article also comes with three caveats, each opening an avenue for further research. First, apart from education, we did not include any lifestyle variables in our model. These variables account for an estimated 45% of AD cases[1] and are a probable cause for the surprising positive partial correlation between old age and memory. Second, Bayesian models allow for the inclusion of prior knowledge. For simplicity, we choose in this article for an uninformative prior in which every pair of variables is conditionally dependent with 20%. Future work could use AD literature to construct an informed prior, certainly for variable pairs whose association is known, such as *APOE4* and the amyloid stage. The use of an informed prior could lead to a more realistic conditional dependence structure. Lastly, this work did not consider temporal data, and therefore, did not cover the progression of AD over time. Future research could add the disease stage over time as a set of variables. This would allow the model to discover the most important factors contributing to the conversion of CN to EMCI to LMCI, and ultimately, to AD.

Conclusions

In this article we present a Bayesian approach to GCGMs. This approach has three capabilities: (i) the estimation of an undirected network depicting the conditional dependencies and partial correlations among variables, (ii) the quantification of the uncertainty of these estimates and, (iii) the inclusion of all types of variables whether normally distributed or not. Bayesian GCGMs have been successfully applied in other domains,[30–32] but remain, to the best of our knowledge, unexplored in the AD domain.

Our GCGM uncovered the partial correlations and conditional dependencies among demographic data, the global amyloid stage, cognitive test scores, brain-region specific gray matter volume, and brain-region specific glucose uptake. Our Bayesian approach enabled us to estimate the uncertainties of these estimates too.

As expected, we observed that the partial correlations and their corresponding conditional dependency networks were sparser than the commonly used Pearson correlation or covariance measures. Our study confirmed existing knowledge, but also opened up new hypotheses. We found three indirect pathways through which old age reduces cognition: hippocampal volume loss, PCC volume loss, and amyloid-beta accumulation. Moreover, we found that women performed better on cognitive tests, but also discovered four indirect pathways that dampen this effect: lower hippocampal volume, lower PCC volume, more amyloid-beta accumulation, and less education. We also found that the hippocampus and PCC volumes are conditionally dependent on cognition, but found limited conditional dependence between brain-region specific glucose uptake and cognition. We found that age and sex were more conditionally dependent with brain-region specific volume than with brain-region specific glucose uptake. Lastly, we discovered that the conditional dependence between brain-region specific volume and glucose metabolism increased as the disease progresses.

particular his suggestions on the visualization of the conditional dependencies improved the article.

ORCID iDs

Lucas Vogels https://orcid.org/0009-0005-3715-3937
Reza Mohammadi https://orcid.org/0000-0001-9538-0648
Marit Schoonhoven https://orcid.org/0000-0002-2268-7486
Ş Ilker Birbil https://orcid.org/0000-0001-7472-7032
Martin Dyrba https://orcid.org/0000-0002-3353-3167

Ethical considerations

Not applicable.

Consent to participate

Not applicable.

Consent for publication

Not applicable.

Author contribution(s)

Lucas Floris Oscar Vogels: Conceptualization; Formal analysis; Investigation; Methodology; Software; Visualization; Writing – original draft; Writing – review & editing.
Reza Mohammadi: Supervision; Writing – review & editing.
Marit Schoonhoven: Supervision; Writing – review & editing.
Ilker Birbil: Supervision; Writing – review & editing.
Martin Dyrba: Conceptualization; Supervision; Writing – review & editing.

Declaration of conflicting interest

The authors declared the following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: For this research we collaborated with Stefan Teipel. He is a Guest Editor of the supplemental issue, "Early Detection of Alzheimer's Disease: From Multiplex Assays and Imaging to AI-Based Functional Monitoring", to which this manuscript is submitted. This forms a potential conflict of interest. However, Dr. Teipel was not involved in handling or review of this submission. The authors declared no other potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Data availability

The data supporting the findings of this study are available upon request from the corresponding author. The data are not publicly available due to privacy and ethical restrictions. The R script used to produce our results can be found on the GitHub page https://github.com/lucasvogels33/Modeling-AD-Bayesian-GCGM-from-Demographic-Cognitive-and-Neuroimaging-Data.

Supplemental material

Supplemental material for this article is available online.

References

1. Livingston G, Huntley J, Liu K, et al. The lancet commissions dementia prevention, intervention, and care: 2024 report of the lancet standing commission. *Lancet* 2024; 404: 572–628.
2. Breijyeh Z and Karaman R. Comprehensive review on Alzheimer's disease: causes and treatment. *Molecules* 2020; 25: 5789.
3. Nemy M, Cedres N, Grothe M, et al. Cholinergic white matter pathways make a stronger contribution to attention and memory in normal aging than cerebrovascular health and nucleus basalis of Meynert. *Neuroimage* 2020; 211: 116607.
4. Teipel S, Dyrba M, Ballarini T, et al. Association of cholinergic basal forebrain volume and functional connectivity with markers of inflammatory response in the Alzheimer's disease spectrum. *J Alzheimers Dis* 2022; 85: 1267–1282.
5. Sporns O. Structure and function of complex brain networks. *Dialogues Clin Neurosci* 2013; 15: 247–262.
6. Sporns O, Tononi G and Kötter R. The human connectome: a structural description of the human brain. *PLoS Comput Biol* 2005; 1: e42.
7. O'Donnell L and Westin C. An introduction to diffusion tensor image analysis. *Neurosurg Clin N Am* 2011; 22: 185–196.
8. Warnick R, Guindani M, Erhardt E, et al. A Bayesian approach for estimating dynamic functional network connectivity in fMRI data. *J Am Stat Assoc* 2018; 113: 134–151.
9. Cao J, Zhao Y, Shan X, et al. Brain functional and effective connectivity based on electroencephalography recordings: a review. *Hum Brain Mapp* 2021; 43: 1–20.
10. Daunizeau J, David O and Stephan K. Dynamic causal modelling: a critical review of the biophysical and statistical foundations. *NeuroImage* 2009; 58: 312–322.
11. Kline R. *Principles and practice of structural equation modeling*. New York: The Guilford Press, 2023.
12. Caffo B, Chen S, Stewart W, et al. Are brain volumes based on magnetic resonance imaging mediators of the associations of cumulative lead dose with cognitive function? *Am J Epidemiol* 2008; 167: 429–447.
13. Yingxuan E, Yao X, Liu K, et al. Polygenic mediation analysis of Alzheimer's disease implicated intermediate amyloid imaging phenotypes. *AMIA Annu* 2021; 2020: 422–431.
14. Fan P, Li H, Xu H, et al. A chain mediation model reveals the association between depression and cognitive function in the elderly. *Sci Rep* 2024; 14: 31375.
15. van Oudenhoven F, Swinkels S, Hartmann T, et al. Modeling the underlying biological processes in Alzheimer's disease using a multivariate competing risk joint model. *Stat Med* 2022; 41: 3421–3433.

16. Neumann A, Kucukali F, Bos I, et al. Rare variants in IFFO1, DTNB, NLRC3 and SLC22A10 associate with Alzheimer's disease CSF profile of neuronal injury and inflammation. *Mol Psychiatry* 2022; 27: 1–10.

17. Neumann A, Ohlei O, Kucukali F, et al. Multivariate GWAS of Alzheimer's disease CSF biomarker profiles implies GRIN2D in synaptic functioning. *Genome Med* 2022b; 15: 79.

18. Shen X, Ma S, Vemuri P, et al. Challenges and opportunities with causal discovery algorithms: application to Alzheimer's pathophysiology. *Sci Rep* 2020; 10: 2975.

19. Uleman J, Melis R, Ntanasi E, et al. Simulating the multi-causality of Alzheimer's disease with system dynamics. *Alzheimers Dement* 2023; 19: 2633–2654.

20. Lauritzen S. *Graphical models*. vol. 17. Oxford: Oxford University Press, 1996.

21. Koller D and Friedman N. *Probabilistic graphical models: principles and techniques*. Cambridge, MA: MIT Press, 2009.

22. Dyrba M, Mohammadi A, Grothe MJ, et al. Gaussian graphical models reveal inter-modal and inter-regional conditional dependencies of brain alterations in Alzheimer's disease. *Front Aging Neurosci* 2020; 12: 99.

23. Ortiz A, Munilla J, Illan I, et al. Exploratory graphical models of functional and structural connectivity patterns for Alzheimer's disease diagnosis. *Front Comput Neurosci* 2015; 9: 132.

24. Titov D, Diehl-Schmid J, Shi K, et al. Metabolic connectivity for differential diagnosis of dementing disorders. *J Cereb Blood Flow Metab* 2017; 37: 252–262.

25. Mielke M. Sex and gender differences in Alzheimer's disease dementia. *Psychiatr Times* 2018; 35: 14–17.

26. Mohammadi A, Schoonhoven M, Vogels L, et al. Large-scale Bayesian structure learning for Gaussian graphical models using marginal pseudo-likelihood. 2024, Unpublished manuscript, Arxiv 2307.00127.

27. Vogels L, Mohammadi A, Schoonhoven M, et al. Bayesian structure learning in undirected Gaussian graphical models: literature review with empirical comparison. *J Am Stat Assoc* 2024; 119: 3164–3182.

28. Dyrba M, Grothe MJ, Mohammadi A, et al. Comparison of different hypotheses regarding the spread of Alzheimer's disease using Markov random fields and multimodal imaging. *J Alzheimers Dis* 2018; 65: 731–746.

29. Hinne M, Janssen R, Heskes T, et al. Bayesian estimation of conditional independence graphs improves functional connectivity estimates. *PLoS Comput Biol* 2015; 11: e1004534.

30. Dobra A and Lenkoski A. Copula Gaussian graphical models and their application to modeling functional disability data. *Ann Appl Stat* 2011; 5: 969–993.

31. Mohammadi A, Abegaz F, Heuvel E, et al. Bayesian modelling of Dupuytren disease by using Gaussian copula graphical models. *J R Stat Soc Ser C Appl Stat* 2016; 66: 629–645.

32. Chandra NK, Sitek KR, Chandrasekaran B, et al. Functional connectivity across the human subcortical auditory system using an autoregressive matrix-Gaussian copula graphical model approach with partial correlations. *Imaging Neurosci (Camb)* 2024; 2: 1–23.

33. Crane PK, Carle A, Gibbons LE, et al. Development and assessment of a composite score for memory in the Alzheimer's disease neuroimaging initiative (ADNI). *Brain Imaging Behav* 2012; 6: 502–516.

34. Gibbons LE, Carle A, Mackin RS, et al. A composite score for executive functioning, validated in Alzheimer's disease neuroimaging initiative (ADNI) participants with baseline mild cognitive impairment. *Brain Imaging Behav* 2012; 6: 517–527.

35. Grothe MJ and Teipel SJ. Spatial patterns of atrophy, hypometabolism, and amyloid deposition in Alzheimer's disease correspond to dissociable functional brain networks. *Hum Brain Mapp* 2016; 37: 35–53.

36. Desikan RS, Ségonne F, Fischl B, et al. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* 2006; 31: 968–980.

37. Butterfield D and Halliwell B. Oxidative stress, dysfunctional glucose metabolism and Alzheimer disease. *Nat Rev Neurosci* 2019; 20: 148–160.

38. Bolton C, Khan OA, Moore EE, et al. Baseline grey matter volumes and white matter hyperintensities predict decline in functional activities in older adults over a 5-year follow-up period. *Neuroimage Clin* 2023; 38: 103393.

39. Grothe MJ, Barthel H, Sepulcre J, et al. In vivo staging of regional amyloid deposition. *Neurology* 2017; 89: 2031–2038.

40. Ashburner J. A fast diffeomorphic image registration algorithm. *NeuroImage* 2007; 38: 95–113.

41. Müller-Gärtner H, Links JM, Prince JL, et al. Measurement of radiotracer concentration in brain gray matter using positron emission tomography: MRI-based correction for partial volume effects. *J Cereb Blood Flow Metab* 1992; 12: 571–583.

42. Friedman J, Hastie T and Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 2007; 9: 432–441.

43. Hoff P. *A first course in Bayesian statistical methods*. New York: Springer, 2009.

44. Solea E and Li B. Copula Gaussian graphical models for functional data. *J Am Stat Assoc* 2020; 117: 781–793.

45. Liu H, Lafferty J and Wasserman L. The nonparanormal: semiparametric estimation of high dimensional undirected graphs. *J Mach Learn Res* 2009; 10: 2295–2328.

46. Hoff PD. Extending the rank likelihood for semiparametric copula estimation. *Ann Appl Stat* 2007; 1: 265–283.

47. Mohammadi A and Wit EC. BDgraph: an R package for Bayesian structure learning in graphical models. *J Stat Softw* 2019; 89: 1–30.

48. Ma C, Hong F and Yang S. Amyloidosis in Alzheimer's disease: pathogeny, etiology, and related therapeutic directions. *Molecules* 2022; 27: 1210.

49. Pini L, Pievani M, Bocchetta M, et al. Brain atrophy in Alzheimer's disease and aging. *Ageing Res Rev* 2016; 30: 25–48.

50. Rocca W. Time, sex, gender, history, and dementia. *Alzheimer Dis Assoc* 2017; 31: 76–79.

51. Castro-Aldrete L, Moser M, Putignano G, et al. Sex and gender considerations in Alzheimer's disease: the women's brain project contribution. *Front Aging Neurosci* 2023; 15: 1105620.

52. Nemes S, Logan P, Manchella M, et al. Sex and APOE4 carrier effects on atrophy, amyloid PET, and tau PET burden in early-onset Alzheimer's disease. *Alzheimers Dement* 2023; 19: S49–S63.

53. Eimer W, Kumar D, Navalpur S, et al. Alzheimer's disease-associated beta-amyloid is rapidly seeded by herpesviridae to protect against brain infection. *Neuron* 2018; 99: 56–63.e3.

54. Klein S and Flanagan K. Sex differences in immune responses. *Nat Rev Immunol* 2016; 16: 626–638.

55. Dima D, Amirhossein M, Efstathios P, et al. Subcortical volumes across the lifespan: data from 18,605 healthy individuals aged 3–90 years. *Hum Brain Mapp* 2022; 43: 452–469.

56. Ezzati A, Katz M, Zammit A, et al. Differential association of left and right hippocampal volumes with verbal episodic and spatial memory in older adults. *Neuropsychologia* 2016; 93: 380–385.

57. Hahn B, Ross TJ and Stein EA. Cingulate activation increases dynamically with response speed under stimulus unpredictability. *Cereb Cortex* 2006; 17: 1664–1671.

58. Landau S, Harvey D, Madison C, et al. Associations between cognitive, functional, and FDG-PET measures of decline in AD and MCI. *Neurobiol Aging* 2011; 32: 1207–1218.

59. Raji C, Lopez O, Kuller L, et al. Age, Alzheimer disease, and brain structure. *Neurology* 2009; 73: 1899–1905.

60. Nunnemann S, Wohlschläger A, Ilg R, et al. Accelerated aging of the putamen in men but not in women. *Neurobiol Aging* 2009; 30: 147–151.

61. Berti V, Mosconi L and Pupi A. Brain: normal variations and benign findings in FDG PET/CT imaging. *PET Clin* 2014; 9: 129–140.