Joint independent component analysis for hypothesizing spatiotemporal relationships between longitudinal gray and white matter changes in preclinical Alzheimer's disease

Cai, Leon, Rheault, Francois, Kerley, Cailey, Aboud, Katherine, Beason-Held, Lori, et al.

Leon Y. Cai, Francois Rheault, Cailey I. Kerley, Katherine S. Aboud, Lori L. Beason-Held, Andrea T. Shafer, Susan M. Resnick, Lori C. Jordan, Adam W. Anderson, Kurt G. Schilling, Bennett A. Landman, "Joint independent component analysis for hypothesizing spatiotemporal relationships between longitudinal gray and white matter changes in preclinical Alzheimer's disease," Proc. SPIE 12032, Medical Imaging 2022: Image Processing, 120321H (4 April 2022); doi: 10.1117/12.2611562

SPIE.

Event: SPIE Medical Imaging, 2022, San Diego, California, United States

Joint independent component analysis for hypothesizing spatiotemporal relationships between longitudinal gray and white matter changes in preclinical Alzheimer's disease

Leon Y. Cai[a], Francois Rheault[b], Cailey I. Kerley[b], Katherine S. Aboud[c], Lori L. Beason-Held[d], Andrea T. Shafer[d], Susan M. Resnick[d], Lori C. Jordan[e], Adam W. Anderson[a,f,g], Kurt G. Schilling[f,g], and Bennett A. Landman[a,b,f,g]

[a]Department of Biomedical Engineering, Vanderbilt University, Nashville, TN, USA; [b]Department of Electrical and Computer Engineering, Vanderbilt University, Nashville, TN, USA; [c]Vanderbilt Brain Institute, Vanderbilt University, Nashville, TN, USA; [d]Laboratory of Behavioral Neuroscience, National Institute on Aging, National Institutes of Health, Baltimore, MD, USA; [e]Departments of Pediatrics and Neurology, Vanderbilt University Medical Center, Nashville, TN, USA; [f]Department of Radiology and Radiological Sciences, Vanderbilt University Medical Center, Nashville, TN, USA; [g]Vanderbilt University Institute of Imaging Science, Vanderbilt University, Nashville, TN, USA

ABSTRACT

Characterizing relationships between gray matter (GM) and white matter (WM) in early Alzheimer's disease (AD) would improve understanding of how and when AD impacts the brain. However, modeling these relationships across brain regions and longitudinally remains a challenge. Thus, we propose extending joint independent component analysis (jICA) into spatiotemporal modeling of regional cortical thickness and WM bundle volumes leveraging multimodal MRI. We jointly characterize these GM and WM features in a normal aging (n=316) and an age- and sex-matched preclinical AD cohort (n=81) at each of two imaging sessions spaced three years apart, training on the normal aging population in cross-validation and interrogating the preclinical AD cohort. We find this joint model identifies reproducible, longitudinal changes in GM and WM between the two imaging sessions and that these changes are associated with preclinical AD and are plausible considering the literature. We compare this joint model to two focused models: (1) GM features at the first session and WM at the second and (2) vice versa. The joint model identifies components that correlate poorly with those from the focused models, suggesting the different models resolve different patterns. We find the strength of association with preclinical AD is improved in the GM to WM model, which supports the hypothesis that medial temporal and frontal thinning precedes volume loss in the uncinate fasciculus and inferior anterior-posterior association fibers. These results suggest that jICA effectively generates spatiotemporal hypotheses about GM and WM in preclinical AD, especially when specific intermodality relationships are considered *a priori*.

Keywords: independent component analysis, Alzheimer's disease, spatiotemporal modeling, longitudinal, cortical morphometry, fiber tractography, white matter pathway shape analysis, multimodal MRI

1. INTRODUCTION

Machine learning has demonstrated marked success in capturing important features while accomplishing medical imaging tasks, but biophysical interpretation of those features remains an open problem [1]. This has limited the ability of machine learning to help investigators form hypotheses regarding underlying pathologic processes. To begin to address this gap in neuroimaging, we previously investigated joint independent component analysis (ICA) for capturing interpretable, simultaneously occurring, gray matter (GM) and white matter (WM) changes in the brain and found it generated plausible hypotheses regarding parallel GM and WM changes in symptomatic Alzheimer's disease (AD) cross-sectionally [2]. In doing so, we investigated joint ICA as a tool for hypothesizing multimodal spatial changes in GM and WM with structural and diffusion magnetic resonance imaging (MRI), respectively, in relation to symptomatic AD.

Medical Imaging 2022: Image Processing, edited by Olivier Colliot, Ivana Išgum, Proc. of SPIE Vol. 12032, 120321H · © 2022 SPIE
1605-7422 · doi: 10.1117/12.2611562

However, our prior study had two primary gaps. The first is that it did not offer specific insight into the temporal relationships of those changes, which in turn prevents the technique from elucidating the progression of GM and WM changes in AD across time. The second is that it focused solely on symptomatic AD. At present, the early diagnosis of AD forms the foundation for care, as it allows additional time for patients to seek treatment, plan care, and is associated with better outcomes [3], [4]. As such, it would be advantageous to understand GM and WM changes and how they progress in AD *before* symptom onset for early intervention. We propose to fill these gaps by extending our joint ICA framework for spatiotemporal modeling of GM and WM features in a retrospective study of preclinical AD.

Specifically, using structural MRI to characterize the thickness of cortical GM regions and diffusion MRI to identify WM bundles and compute their volumes, we investigate how best to use joint ICA to generate hypotheses regarding the spatiotemporal intermodality relationships between GM and WM in preclinical AD. First, we gather longitudinal imaging data from two time points in a cohort of normal aging and an age- and sex-matched cohort of preclinical AD. Second, using these data, we characterize three different approaches to spatiotemporal joint ICA: (1) considering GM and WM at both time points, (2) considering only GM at the first time point and WM at the second, and (3) considering only WM at the first and GM at the second. We evaluate the reproducibility of these approaches in cross validation, assess their plausibility in identifying known GM and WM changes in AD, and characterize how these different approaches are suited for investigating different neuroanatomical phenomena. Last, we explore the hypotheses supported by these approaches regarding the neuroanatomical effects of AD before symptom onset.

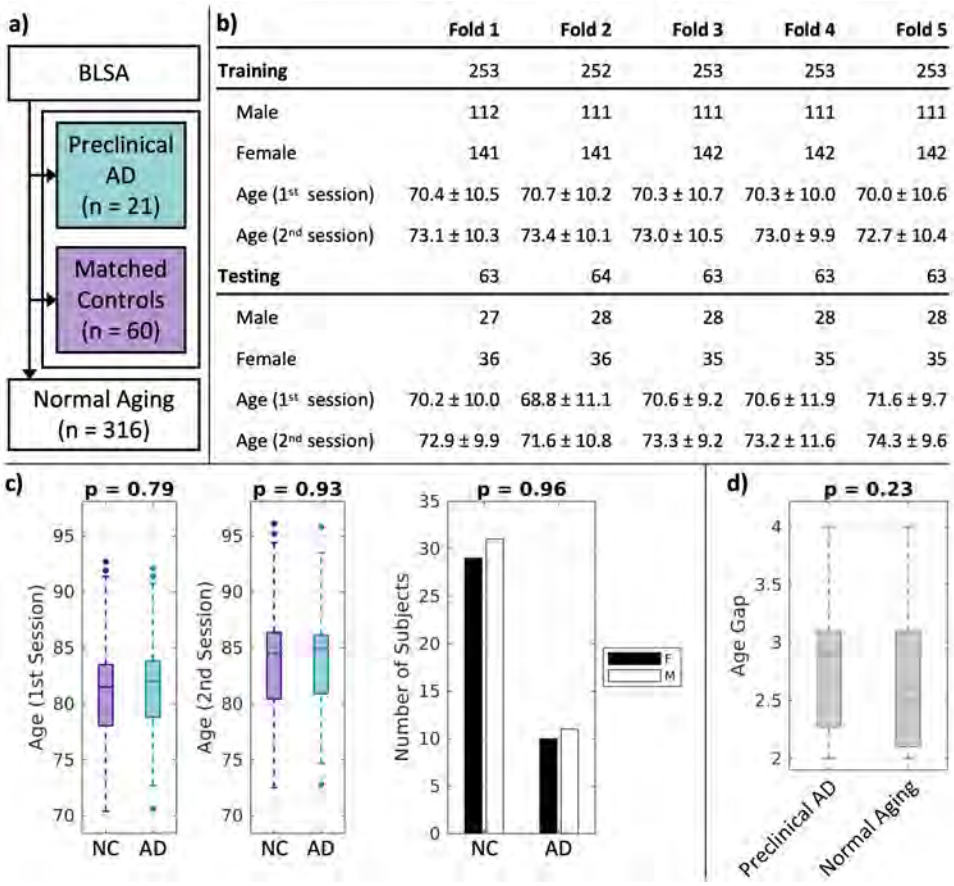

	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5
Training	253	252	253	253	253
Male	112	111	111	111	111
Female	141	141	142	142	142
Age (1st session)	70.4 ± 10.5	70.7 ± 10.2	70.3 ± 10.7	70.3 ± 10.0	70.0 ± 10.6
Age (2nd session)	73.1 ± 10.3	73.4 ± 10.1	73.0 ± 10.5	73.0 ± 9.9	72.7 ± 10.4
Testing	63	64	63	63	63
Male	27	28	28	28	28
Female	36	36	35	35	35
Age (1st session)	70.2 ± 10.0	68.8 ± 11.1	70.6 ± 9.2	70.6 ± 11.9	71.6 ± 9.7
Age (2nd session)	72.9 ± 9.9	71.6 ± 10.8	73.3 ± 9.2	73.2 ± 11.6	74.3 ± 9.6

Figure 1. Demographics of the normal aging and preclinical AD cohorts. (a) Both cohorts were extracted from the Baltimore Longitudinal Study of Aging (BLSA), using participants who had been scanned twice 3±1 years apart. (b) The normal aging cohort was evaluated in cross-validation. Ages are indicated as mean ± standard deviation. (c) The preclinical AD cohort consists of participants who were scanned twice before developing AD symptoms. Normal controls (NC) were matched 3:1 for sex, age at first session, and age at second session. (d) The age gaps between imaging sessions between the two cohorts were not significantly different. Statistical significance was evaluated with the Wilcoxon rank-sum test at 0.05 significance.

2. METHODS

2.1 Study design and image processing

We acquired de-identified MRI data for analysis from the Baltimore Longitudinal Study of Aging (BLSA) [5], [6], utilizing two cognitively unimpaired imaging sessions from each participant, spaced 3±1 years apart, all scanned on one 3T Philips scanner. We first excluded all participants who were not imaged at least twice while cognitively unimpaired with 3±1 years between two of their imaging sessions. We then identified those participants who were clinically diagnosed with Alzheimer's disease by their last visit (n = 21, Figure 1a), excluding participants who develop co-morbid neurologic disease, like Parkinsonism or vascular dementia, or other cognitive deficits. From the remaining cognitively unimpaired participants, we match those with preclinical AD at a ratio of 3:1 for sex, age at first imaging session, and age at second imaging session (Figure 1c) [7], thus forming an age- and sex-matched "preclinical AD cohort" (n = 81 total, Figure 1a). The remaining cognitively unimpaired participants (n = 316) form the "normal aging cohort" investigated presently (Figure 1a). This cohort was split in stratified 5-fold cross validation by sex and age (Figure 1b). Differences in median age-gap between sessions for the preclinical AD and normal aging cohorts were demonstrated to be not statistically significant (Figure 1d). Statistical significance was assessed with the Wilcoxon rank-sum test at 0.05 significance.

From each imaging session, we evaluated gray matter (GM) and white matter (WM) macrostructure using T1 weighted (T1w) MRI and diffusion tensor MRI (DTI), respectively. The T1w MRI were acquired at 1.0x1.0x1.2mm (sagittal, coronal, and axial) resolution. On each T1w MRI, we performed a cortical shape analysis to measure the thickness in 98 cortical regions defined by the BrainCOLOR atlas segmented with multi-atlas labeling [8]–[10]. The DTI were all 64-direction b = 700 s/mm^2 images with corresponding b = 0 s/mm^2 volumes acquired at 2.2mm isotropic resolution and were preprocessed with the PreQual pipeline prior to analysis [11]. On each DTI, we performed a bundle shape analysis using the DSIStudio software package to compute the volume of each of 43 WM bundles identified with fiber tractography [12]. In summary, these operations yielded 98 GM features and 43 WM features for each participant at each of two time points.

2.2 Approaches to spatiotemporal joint independent component analysis

We modeled these features using spatiotemporal joint ICA in three configurations. The first was to consider both GM and WM at both time points (the "joint GM and WM" approach, Figure 2a). The second was a focused approach, considering only GM at the first time point and WM at the second (Figure 2b), and the third was the opposite, considering only WM at the first time point and GM at the second (Figure 2c). All approaches were trained in cross-validation on the normal aging cohort and applied to the preclinical AD cohort.

For the joint GM and WM approach, we first concatenated the features from the first time point, z-scored each one, and performed principal component analysis (PCA), retaining principal components to 95% variance explained. The means, standard deviations, and principal components used to normalize the features from the first time point were then used to normalize the corresponding features in the second time point. The normalized features from the two time points were subsequently concatenated and reconstruction ICA was performed to determine independent components (ICs) and their associated loadings [13]. Similarly for the focused approaches, we concatenated the features from the two modalities across time points, z-scored each one, and performed PCA to 95% variance explained. Reconstruction ICA was then performed to identify ICs and their associated loadings. Selection of the number of ICs is detailed in section 2.4.

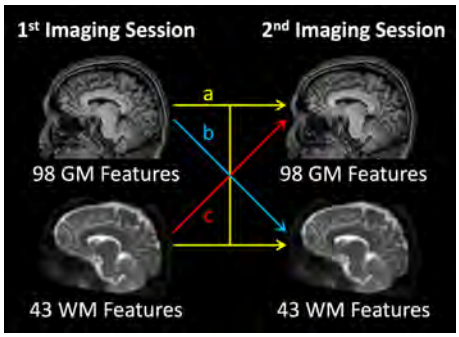

Figure 2. Three approaches to spatiotemporal joint independent component analysis. (a) In the "joint GM and WM" approach, both GM and WM features from both time points are considered in the model. (b) In the focused GM to WM approach, only GM features from the first time point and WM features from the second are considered. (c) In the focused WM to GM approach, only WM features from the first time point and GM features from the second are considered.

Since ICs identified with reconstruction ICA were not guaranteed to be in the same order or follow the same sign convention, ICs were matched within and between approaches prior to further analysis:

- To match ICs within approaches, for any given approach identifying k ICs, we designated fold 1 as the reference fold and sought to best match the ICs from folds 2-5 with those from fold 1. To achieve this for a given fold, we performed Pearson correlation for all the ICs with all the ICs from the reference fold in data space, yielding k^2 correlations. We matched a given fold's ICs 1-to-1 to the reference ICs by selecting matches that maximized the sum of the absolute correlations across matches by solving the corresponding linear assignment problem [14]. Briefly, the first match was designated with the largest absolute correlation. The matched IC and its reference were then taken out of consideration and the process was recursed until all k ICs had been matched. After matching, ICs and their associated loadings were negated if they were negatively correlated with their reference IC.
- To match ICs between approaches, we first matched ICs within each approach. We then applied Pearson correlation on the ICs for all folds of one approach with all folds of the other. For each fold combination, we performed the same matching scheme that we did for within-approach matching. We used majority vote across all fold combinations to determine the final matches. To correlate the ICs from the joint GM and WM approach to those from the focused approaches, only the corresponding joint GM and WM features were used.

Last, in order to understand what patterns in GM and WM each IC represented, after matching the reference ICs were backprojected from PCA space to data space, split into their constituent GM and WM parts, and visualized as detailed in section 2.3.

2.3 Visualizing independent components

We rescaled all ICs from 0 to 1 prior to visualization, considering the GM and WM ICs independently. First, we centered each IC by subtracting its mean value. Next, we divided each IC by the magnitude of its largest feature, positive or negative, producing ICs bounded between -1 and 1. Last, we rescaled each IC to between 0 and 1 by dividing by two and adding ½. For the joint GM and WM case, we rescaled GM features from both time points and the WM features from both time points together so that they could be compared directly. We visualized all ICs on a rendering of the same cortical multi-atlas segmentation [10] and WM bundle parcellation on a representative T1 image of an unassociated, de-identified, healthy subject on a Philips 3T scanner acquired at Vanderbilt University with Institutional Review Board approval (#160268).

To visualize GM cortical thickness features, we created four 3-dimensional renderings, medial (MED) and lateral (LAT) views of both the right (R) and left (L) hemispheres, and colored the corresponding region in the rendering based on the IC's value in that region. We used a diverging colormap where blues indicated lower values below the mean and yellows indicated higher values above the mean. To visualize WM bundle volume features, we created 10 3-dimensional renderings, R and L views of all 43 bundles categorized into 5 sets: (1) "major" (MAJ) bundles, including the corpus callosum, cortical spinal tracts, and arcuate fasciculi; (2) anterior-posterior (AP) association bundles, including the inferior frontooccipital fasciculi, inferior longitudinal fasciculi, and superior longitudinal fasciculi; (3) superior-inferior (SI) association bundles, including the frontal and parietal aslant tracts, uncinate fasciculi, and vertical occipital fasciculi; (4) thalamic projections (TP), including the anterior, inferior, posterior, and superior thalamic radiations and optic radiations; and (5) the cingulum and fornix (CF).

2.4 Determining the number of independent components

In order to determine a stable number of ICs to investigate, we ran the joint GM and WM approach in cross-validation with a target number of ICs ranging from 2 to 50. Subsequently, we plotted the median interfold Pearson correlations for each identified IC after matching. We then visually evaluated for clusters of highly reproducible components defined as having interfold Pearson correlations greater than 0.75, and selected the number of components that, on average, achieved this threshold (Figure 3). This resulted in the selection of 8 ICs for further analysis.

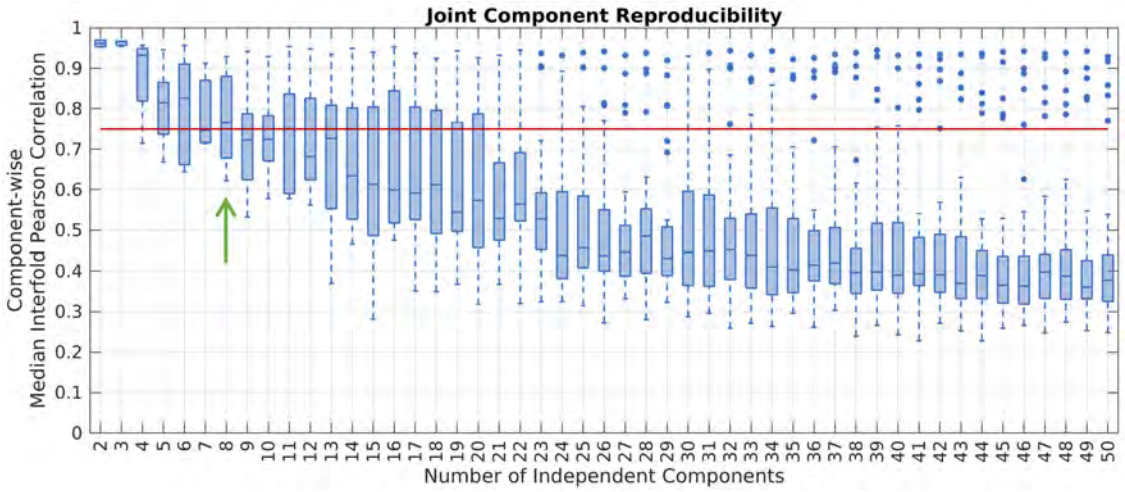

Figure 3. Reproducibility of joint GM and WM components. We ran the joint GM and WM configuration with 2 to 50 independent components in cross-validation on our normal aging cohort. We identified a cluster of reproducible components with component-wise interfold Pearson correlation > 0.75 (red line). We selected 8 components with a median Pearson correlation of > 0.75 for continued analysis (green arrow).

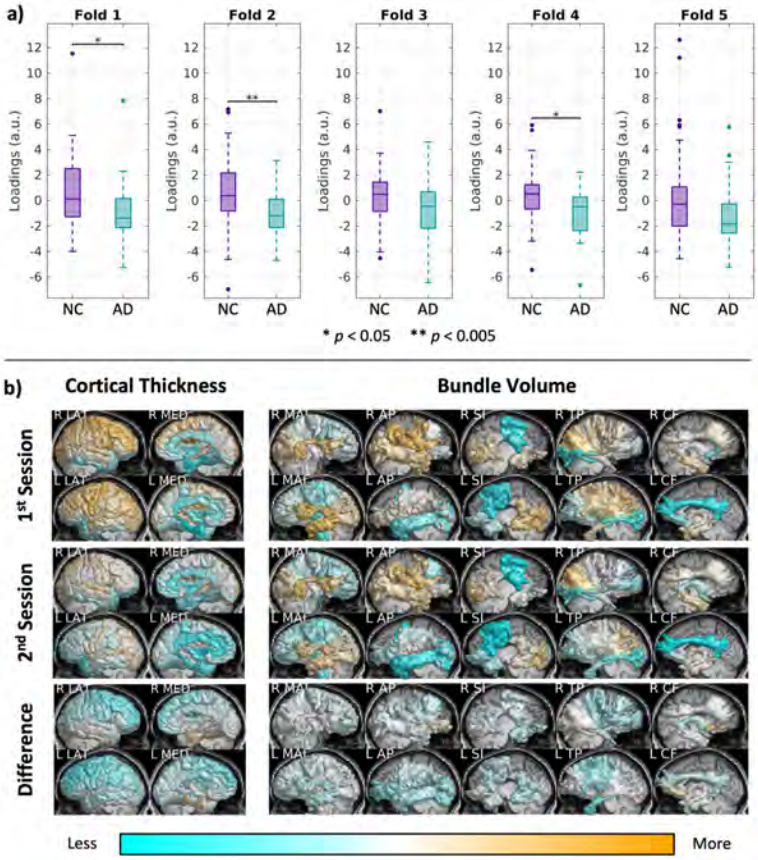

Figure 4. Plausibility of joint GM and WM components. (a) Of the 8 components identified with training on the normal aging cohort, we identified component number 7 as potentially inversely associated with preclinical AD. Statistical significance was evaluated with the Wilcoxon rank-sum test at 0.05 significance for each fold. (b) Visualization of this component demonstrates this association identifies decreased cortical thickness in the medial temporal lobe coupled with decreased volume in the anterior cingulum compared to the rest of the brain in preclinical AD.

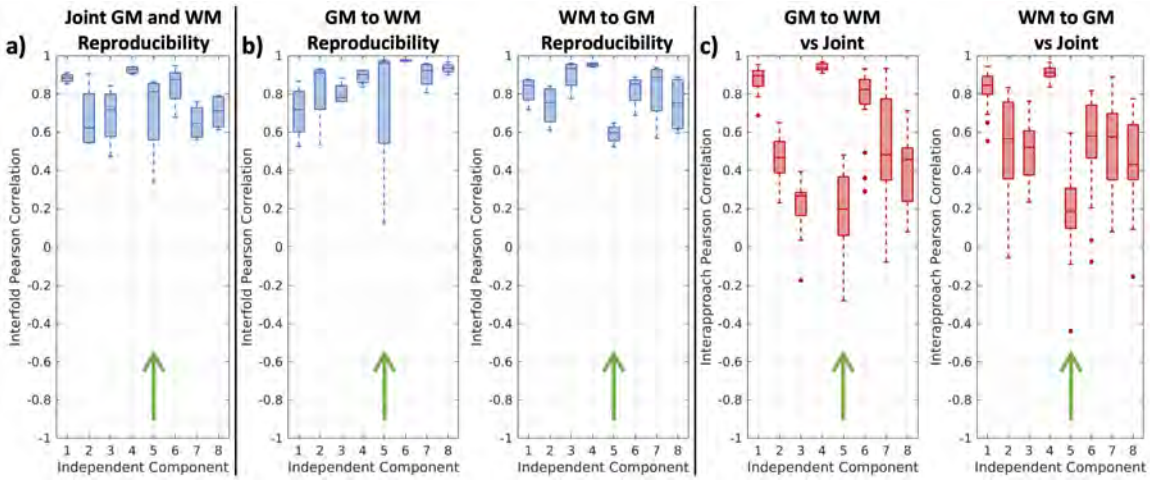

Figure 5. Pearson correlations between joint GM and WM components with those identified in focused approaches: GM at the first and WM at the second (GM to WM) and vice versa (WM to GM). (a) Interfold Pearson correlations of the 8 joint GM and WM components identified at both time points. This is equivalent to an expanded version of the distribution denoted with the green arrow in Figure 3. (b) The focused GM to WM and WM to GM approaches exhibit similar interfold reproducibility to the joint GM and WM approach in cross validation. (c) Many of the 8 components identified with the focused approaches demonstrate poor Pearson correlations with those identified in the joint GM and WM approach, suggesting the focused approaches captured different spatiotemporal patterns in GM and WM. The most extreme example of this is component 5 (green arrows).

3. RESULTS

3.1 Plausibility of spatiotemporal joint independent component analysis

First, we investigated the plausibility of a spatiotemporal joint ICA by evaluating whether any of the 8 joint GM and WM components trained on the normal aging cohort were statistically associated with preclinical AD. After projecting the preclinical AD cohort onto the ICs identified for each fold, we determined that one of the ICs, number 7, was consistently inversely associated with preclinical AD across the five folds (Figure 4a). We evaluated for statistically significant differences in each fold with the Wilcoxon rank-sum test at 0.05 significance. To understand the spatiotemporal GM and WM phenomena captured by this IC, we visualized both the GM thicknesses and WM bundle volumes (Figure 4b). We found that this IC demonstrates decreased cortical thickness in the medial temporal lobe coupled with decreased volume in the anterior cingulum compared to the rest of the brain in preclinical AD. This result suggests that these regions are preferentially affected early in AD with the extension that these GM and WM phenomena occur in parallel. We find this to be consistent with the current literature [15]–[17].

3.2 Differences between approaches

After determining the plausibility of a spatiotemporal joint ICA, we evaluated whether the ICs captured by the joint GM and WM approach differed from those captured by the focused approaches, hypothesizing that ICs identified when both GM and WM features are included at both time points will be dominated by parallel GM and WM phenomena as opposed to GM changes that precede WM and vice versa. We plotted the within approach and between approach Pearson correlations after matching (Figure 5a and 5b). We found that the three different approaches all exhibit comparable reproducibility at 8 ICs, with the median interfold Pearson correlations across ICs hovering between 0.7 and 0.8. On the other hand, we find that many of the 8 ICs identified by the focused approaches capture different information than those captured by the joint GM to WM approach (Figure 5c). Specifically, we identify IC number 5. We find that the IC captured by the focused approaches correlate poorly with those from the joint GM and WM approach, suggesting the focused approaches capture different neuroanatomical phenomena.

To investigate further, we evaluated how the fifth IC identified by each approach associated with preclinical AD and found that the strength of association varies between the three approaches (Figure 6). The joint GM and WM approach identified a component that trends with preclinical AD across 4/5 folds with one statistically significant different fold. The GM to WM approach identified 3/5 folds that are statistically associated with preclinical AD and one trending fold. The WM to

GM approach identified 1 or 2 folds may trend with preclinical AD and that the remainder visually exhibit no association. Statistical significance was determined for each fold with the Wilcoxon rank-sum test at 0.05 significance. We interpret these results to show (1) that the GM to WM approach identified the component that associated the most with preclinical AD, followed by the joint GM and WM approach, and finally the WM to GM; and (2) that the joint GM to WM approach likely consists of a mixture of the GM to WM and WM to GM phenomena, suggesting the specificity of a focused approach may be higher.

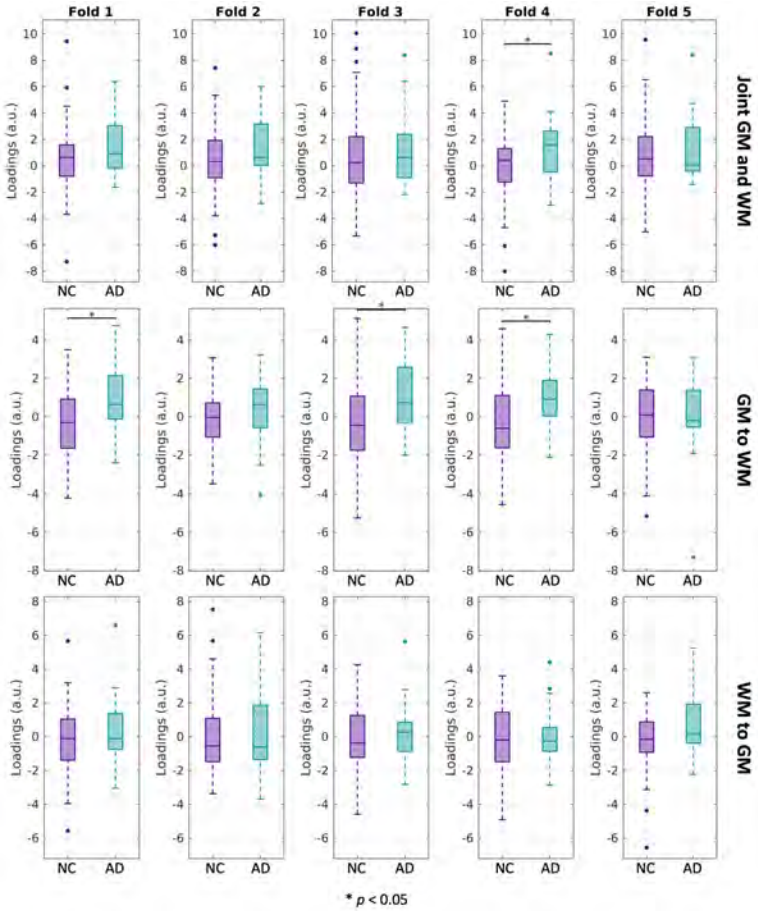

Figure 6. Differences in association with preclinical AD for component 5 identified with the joint GM and WM approach versus the focused approaches. The associations of this component with preclinical AD from strongest to weakest are in the GM to WM focused approach, the joint GM and WM approach, and the WM to GM approach. Statistical significance was evaluated with the Wilcoxon rank-sum test for each fold at 0.05 significance.

3.3 Spatiotemporal hypotheses about preclinical Alzheimer's disease

Based on the differences between the joint and focused approaches, we visualize the ICs captured by the joint GM and WM approach and the GM to WM approach to generate spatiotemporal hypotheses regarding GM thickness and WM volumes in relation to preclinical AD (Figure 7).

First, we confirm visually that the information captured in these ICs indeed paint different neuroanatomical pictures. The joint GM and WM approach generates the hypothesis that temporal cortical thickness decreases *in parallel* with WM bundle volume most notably in the anterior and posterior corpus callosum, cingulum, and AP and SI association fibers (Figure 7a). Notably, the intermodality temporal relationships remain unclear in this visualization. Meanwhile, the GM to WM approach generates the hypothesis that medial temporal and frontal thinning *precedes* relative loss of volume of the uncinate fasciculus and inferior AP association fibers in preclinical AD (Figure 7b). Last, considering the different strengths of association with preclinical AD (Figure 6), we interpret that the GM to WM hypothesis is likely stronger than that of the joint GM and WM approach.

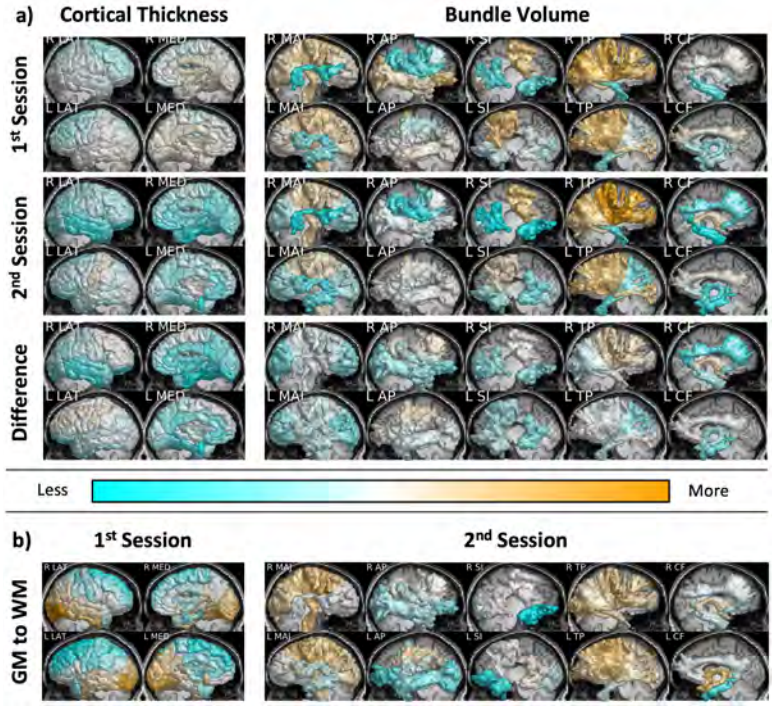

Figure 7. Formulating spatiotemporal hypotheses about GM and WM relationships in preclinical AD. (a) Decreasing temporal thickness in parallel with decreases in bundle volume most notably in the anterior and posterior corpus callosum, cingulum, and anterior-posterior association fibers is identified in the joint approach and weakly associated with preclinical AD. The temporal relationships between the GM and WM changes remains unclear. (b) The focused GM to WM approach is more strongly associated with preclinical AD and suggests medial temporal and frontal thinning precedes relative loss of volume of the uncinate fasciculus and inferior anterior-posterior association fibers.

4. DISCUSSION AND CONCLUSIONS

In this work, we present an extension of our previous exploration of joint ICA into spatiotemporal modeling of cortical GM thickness and WM bundle volumes to better understand the neuroanatomical phenomena underlying preclinical AD. To our knowledge, this represents the first time ICA has been used for modeling multimodal structural and diffusion MRI in both space and time in pursuit of a better understanding of the biophysical changes associated with AD prior to symptom onset. Investigating one joint and two focused approaches, we identified that all three exhibit similar reproducibility at 8 components, and that the joint approach yields plausible neuroanatomical findings in preclinical AD in the setting of the literature. Further diving into the utility of the method, we found that the three different approaches capture different neuroanatomical phenomena. In the joint method, since both GM and WM were considered simultaneously in the model, the temporal relationships between the GM and WM changes remained unclear, suggesting that the joint approach is limited to describing parallel spatiotemporal phenomena. In the focused methods, however, the temporal relationships are explicitly designated, thus increasing the specificity of the hypotheses generated, especially in context of temporal intermodality relationships. Utilizing this characterization of spatiotemporal joint ICA, we generate the hypothesis that medial temporal and frontal thinning precedes relative loss of volume of the uncinate fasciculus and inferior AP association fibers in preclinical AD. Interestingly, we find that these WM bundles have previously been hypothesized to be affected early in AD due to their proximity to the temporal lobe, and thus find that spatiotemporal joint ICA reveals a novel temporal relationship between GM thickness and WM volumes [18]. We conclude that joint ICA is effective for generating hypotheses regarding spatiotemporal relationships in aging and preclinical AD, but that the technique is most effective when the relevant intermodality relationships are considered *a priori*, as the joint approach would necessarily need to also model parallel changes.

Understanding how AD affects the brain, especially early on in the disease, is critical to improving understanding of how to treat the disease. By leveraging multimodal neuroimaging and the interpretability offered by spatiotemporal joint ICA, this work seeks to formulate hypotheses regarding both local and global GM and WM changes and how they develop and

interact in preclinical AD. Our hope is that the methods presented here will allow the field to move toward validation and extension of the hypotheses uncovered as potential non-invasive biomarkers for AD or form the basis for mechanistic studies that could improve interventions for AD.

One limitation to this work is that, by definition, any one component captured by reconstruction ICA only tells part of the story. In order to reconstruct the entire input, multiple components are needed. This is further accentuated by the training and evaluation of the method presently. By training joint ICA on the normal aging cohort and interrogating the preclinical AD cohort, we seek to find signatures of, or a lack thereof, preclinical AD in a normal aging cohort. This indicates that ICs identified to associate with preclinical AD likely do not tell the entire story of preclinical AD, but only one statistically independent portion of it. This indicates that the hypotheses generated by the ICs will likely be weaker regarding preclinical AD. However, since the method is designed as a hypothesis-generating approach, we find it reassuring that even if spatiotemporal joint ICA cannot explain the entire story, it may elucidate part of it for further investigation.

This selection of this normal vs preclinical AD experimental design was in large part due to the limited number of preclinical AD participants with at least two cognitively unimpaired imaging sessions in the BLSA. In order to increase the sample size for model training, the normal aging cohort was used. This indicates that another limitation to this work is decreased power when evaluating for associations with preclinical AD. That being said, we find it reassuring that despite lower power, the focused GM to WM approach was still able to resolve statistically significant differences across a majority of folds when generating hypotheses.

Another limitation is that the ICA process is limited in its ability to compute biophysical values as it returns all components with zero mean and unit variance. This means that all visualization of ICs must be done relatively (i.e., more or less), as opposed with absolute thicknesses or volumes. For instance, we visualized ICs by scaling around the mean value included in each constituent part of each IC across time points to make temporal comparisons more obvious. However, there is no guidance as to whether this is the appropriate scaling, thus limiting the inference of the components to relative relationships as opposed to absolute ones. Additionally, this limitation means that any changes in time are relative. For instance, we evaluated imaging sessions spaced 3±1 years apart. Thus, a decrease across this time frame revealed by our analysis does not specify the rate at which these changes occur, for instance.

Future directions for this work include external validation of the hypotheses with entirely separate datasets or retraining of the model directly on a larger number of preclinical AD samples to better capture a more holistic view of the disease. Another area of exploration includes the incorporation of DTI microstructural features which have been hypothesized to be affected early in AD with GM changes [19]–[21]. Another opportunity is the exploration of other ICA implementations such as the ICASSO clustering-based method [22] as potential alternatives to reconstruction ICA. Last, the extension of this work into deep learning could move hypothesis generation from linear to nonlinear spatiotemporal relationships. For instance, instead of modeling each imaging session as a string of extracted features, using convolutional features would allow for image-level, nonlinear features to be characterized, potentially shortening the gap between investigation and the development of radiologically identifiable biomarkers for preclinical AD.

ACKNOWLEDGEMENTS

This work was conducted in part using the resources of the Advanced Computing Center for Research and Education at Vanderbilt University, Nashville, TN. This work was supported by the National Institutes of Health (NIH) under award numbers R01EB017230, T32EB001628, 1U34DK123895-01, 5UL1TR002243-04, 1R01MH121620-01, and T32GM007347; by ViSE/VICTR VR3029; and by the National Center for Research Resources, Grant UL1RR024975-01, and is now at the National Center for Advancing Translational Sciences, Grant 2UL1TR000445-06. This project was also supported by the National Science Foundation under award numbers 1452485 and 2040462. This research was conducted with the support from the Intramural Research Program of the National Institute on Aging of the NIH. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

REFERENCES

[1] A. Adadi and M. Berrada, “Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI),” *IEEE Access*, vol. 6, pp. 52138–52160, Sep. 2018, doi: 10.1109/ACCESS.2018.2870052.

[2] L. Y. Cai *et al.*, “Joint cortical surface and structural connectivity analysis of Alzheimer’s Disease,” 2021.
[3] B. P. Leifer, “Early Diagnosis of Alzheimer’s Disease: Clinical and Economic Benefits,” *J. Am. Geriatr. Soc.*, vol. 51, pp. 281–288, 2003.
[4] A. Atri, “The Alzheimer’s Disease Clinical Spectrum: Diagnosis and Management,” *Med. Clin. North Am.*, vol. 103, no. 2, pp. 263–293, 2019, doi: 10.1016/j.mcna.2018.10.009.
[5] N. W. Shock, *Normal human aging: the Baltimore longitudinal study of aging*. Baltimore, Md: U.S. Dept. of Health and Human Services, Public Health Service, National Institutes of Health, National Institute on Aging, Gerontology Research Center, 1984.
[6] S. M. Resnick *et al.*, “One-year age changes in MRI brain volumes in older adults,” *Cereb. Cortex*, vol. 10, no. 5, pp. 464–472, 2000, doi: 10.1093/cercor/10.5.464.
[7] C. I. Kerley, “Data Preparation,” *pyPheWAS 4.0.3.*, 2021. https://pyphewas.readthedocs.io/en/latest/dataprep.html#maximizecontrols.
[8] I. Lyu, H. Kang, N. D. Woodward, M. A. Styner, and B. A. Landman, “Hierarchical spherical deformation for cortical surface registration,” *Med. Image Anal.*, vol. 57, pp. 72–88, 2019, doi: 10.1016/j.media.2019.06.013.
[9] P. Parvathaneni *et al.*, “Cortical Surface Parcellation Using Spherical Convolutional Neural Networks,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2019, vol. 11766 LNCS, pp. 501–509, doi: 10.1007/978-3-030-32248-9_56.
[10] A. J. Asman, A. S. Dagley, and B. A. Landman, “Statistical label fusion with hierarchical performance models,” in *Medical Imaging 2014: Image Processing*, 2014, vol. 9034, p. 90341E, doi: 10.1117/12.2043182.
[11] L. Y. Cai *et al.*, “PreQual: An automated pipeline for integrated preprocessing and quality assurance of diffusion weighted MRI images,” *Magn. Reson. Med.*, no. December 2020, pp. 1–15, 2021, doi: 10.1002/mrm.28678.
[12] F. C. Yeh, “Shape analysis of the human association pathways,” *Neuroimage*, vol. 223, no. September, 2020, doi: 10.1016/j.neuroimage.2020.117329.
[13] Q. V. Le, A. Karpenko, J. Ngiam, and A. Y. Ng, “ICA with reconstruction cost for efficient overcomplete feature learning,” 2011.
[14] I. S. Duff and J. Koster, “On algorithms for permuting large entries to the diagonal of a sparse matrix,” *SIAM J. Matrix Anal. Appl.*, vol. 22, no. 4, pp. 973–996, 2001, doi: 10.1137/S0895479899358443.
[15] R. Duara *et al.*, “Medial temporal lobe atrophy on MRI scans and the diagnosis of Alzheimer disease,” *Neurology*, vol. 71, no. 24, pp. 1986–1992, 2008, doi: 10.1212/01.wnl.0000336925.79704.9f.
[16] P. J. Visser, F. R. J. Verhey, P. A. M. Hofman, P. Scheltens, and J. Jolles, “Medial temporal lobe atrophy predicts Alzheimer’s disease in patients with minor cognitive impairment,” *J. Neurol. Neurosurg. Psychiatry*, vol. 72, no. 4, pp. 491–497, 2002, doi: 10.1136/jnnp.72.4.491.
[17] Y. Zhang *et al.*, “Diffusion tensor imaging of cingulum fibers in mild cognitive impairment and Alzheimer disease,” *Neurology*, vol. 68, no. 1, pp. 13–19, 2007, doi: 10.1212/01.wnl.0000250326.77323.01.
[18] S. Kitamura *et al.*, “Longitudinal white matter changes in Alzheimer’s disease: A tractography-based analysis study,” *Brain Res.*, vol. 1515, pp. 12–18, 2013, doi: 10.1016/j.brainres.2013.03.052.
[19] L. Clerx, P. J. Visser, F. Verhey, and P. Aalten, “New MRI Markers for Alzheimer’s Disease: A Meta-Analysis of Diffusion Tensor Imaging and a Comparison with Medial Temporal Lobe Measurements,” *J. Alzheimer’s Dis.*, vol. 29, pp. 405–429, 2012, doi: 10.3233/JAD-2011-110797.
[20] P. Selnes *et al.*, “Diffusion Tensor Imaging Surpasses Cerebrospinal Fluid as Predictor of Cognitive Decline and Medial Temporal Lobe Atrophy in Subjective Cognitive Impairment and Mild Cognitive Impairment,” *J. Alzheimer’s Dis.*, vol. 33, pp. 723–736, 2013, doi: 10.3233/JAD-2012-121603.
[21] L. Serra *et al.*, “Grey and White Matter Changes at Different Stages of Alzheimer’s Disease,” *J. Alzheimer’s Dis.*, vol. 19, pp. 147–159, 2010, doi: 10.3233/JAD-2010-1223.
[22] J. Himberg, A. Hyvärinen, and F. Esposito, “Validating the independent components of neuroimaging time series via clustering and visualization,” doi: 10.1016/j.neuroimage.2004.03.027.