

Robust Measurement of Individual Localized Changes to The Aging Hippocampus

Jing Xie

Department of Computer Science, UC, Davis

xie@ucdavis.edu

Evan Fletcher

Department of Neurology, UC, Davis

evanfletcher@gmail.com

Baljeet Singh

Department of Neurology, UC, Davis

bjaysingh@gmail.com

Owen Carmichael

Department of Neurology, UC, Davis

ocarmichael@ucdavis.edu

November 30, 2011

Abstract

Alzheimer’s disease (AD) is characterized by a stereotypical spatial pattern of hippocampus (HP) atrophy over time, but reliable and precise measurement of localized longitudinal change to individual HP in AD have been elusive. We present a method for quantifying subject-specific spatial patterns of longitudinal HP change that aligns serial HP surface pairs together along the central HP axis, cuts slices off the ends of the HP that were not shared across baseline and follow-up surfaces, estimates weighted correspondences between baseline and follow-up HP, and finds a concise set of localized spatial change patterns that explains HP changes while down-weighting HP surface points whose estimated change is biologically implausible. Experiments on 320 elderly HP measured at one-year intervals suggest that the alignment, weighted correspondence estimation and down-weighting steps reduce the amount of implausible HP changes indicated among individual HP, increase the strength of association between HP change and cognitive function related to AD, and enhance the estimation of spatially-localized HP change patterns.

1 Introduction

Alzheimer’s Disease (AD) is the most common form of dementia in the elderly. Its pathological process preferentially damages the hippocampus (HP), a brain region crucial to a variety of cognitive functions. Because this HP damage is visible on structural magnetic resonance images (MRI) of the brain, longitudinal MRI-based measurements of HP change have been proposed as markers of the effects of AD pathology, with possible uses in clinical diagnosis of AD and quantification of treatment effects in clinical trials [2] [12][5]. Furthermore, HP atrophy associated with AD occurs with a stereotypical spatial progression: among HP sub-regions, AD-related damage begins in one region (called CA1) and spreads to another (subiculum) and yet another (CA2-4 and dentate gyrus). Therefore, numerous approaches for quantifying spatial patterns of localized HP atrophy have been proposed as measures that may be more specific to AD-related damage than total HP volume (for example, [6] [17]). These efforts have suggested that localized HP damage to CA1 and subiculum may be preferentially associated with the pathology and cognitive consequences of AD [4].

However, quantitative measurements of how individual hippocampi evolve over time in response to AD pathology, based on serial MRI, have been extremely rare. Most reports of localized HP changes related to AD are cross-sectional in nature, showing spatial patterns of differences between groups of HP that are each measured at a single time point [6][8]. The few reports of localized longitudinal HP change related to AD show population mean spatial maps of localized HP thickness at baseline and follow-up, without providing information about localized changes within any individual [15]. Such longitudinal localized HP change measurements at an individual level will be important for eventually applying localized HP change measures to AD clinical diagnosis and to clinical trials. In contrast to modeling longitudinal changes to individual closed surfaces such as the HP, several techniques measure individual longitudinal change trajectories at a voxel level

[13][16]; for each individual, dense deformation fields among series of images are estimated, with various methods used to encourage the deformation fields to be spatially and temporally smooth. These exploratory quantifications of change are unbiased, highly automated, and general, but the need to detect subtle, localized shape changes to complex-shaped regions require deformation quantification at a very fine scale, making the methods computationally complex.

The relative lack of scientific findings related to individual trajectories of longitudinal localized patterns of HP change is due to a series of unique technical challenges. The hippocampus is a small structure relative to typical voxel sizes in conventional structural MRI, especially when it is atrophied, and changes over short periods of time can be subtle, even in AD; this makes precise delineation of the HP in each image, and highly precise quantification of change, imperative. However, the HP is notoriously difficult to delineate in a reliable fashion; there are scores of manual tracing protocols [7] and even highly trained experts agree about where to trace the HP boundary to varying degrees on varying HP sub-regions [10] [3]. Therefore, beyond precisely aligning baseline HP to corresponding follow-up HP, any method for measuring localized HP changes related to AD must be robust to discrepancies between baseline and follow-up in what tissue was or was not included in the delineation.

In this paper, we present a robust method for measuring longitudinal, localized patterns of spatial change to individual aging HP. We align follow-up to baseline HP by deforming the follow-up HP along its central cylindrical axis, and cut off ends of the structure that are not shared between the two delineations. Then we estimate weighted correspondences between baseline and follow-up HP. After baseline alignment to a template, rate of change in localized HP radius is quantified at a grid of surface points on the template, and points whose change rates are biologically implausible with respect to the aging and AD processes (*i.e.*, regions whose local radius appears to increase with time) are down-weighted in subsequent quantification of localized spatial patterns of radius change. Those spatial patterns are estimated through a modification of Localized Components Analysis (LoCA) which searches for a concise set of localized HP regions whose change accounts for the greatest amount of HP change variability across the population [17].

2 Methods

Fig 1 shows our framework for calculating individual hippocampus change. The cylinder-like HP surface is represented by a central axis, n_t slices along that axis, and n_θ surface points on each slice. The position of the central axis within each slice is calculated by taking the center of mass of all surface points in that slice. Each surface point is represented by its cylindrical coordinates (t, θ) , where t is the index of the slice and θ is the angle between a vector from the central axis to the surface point, and a vector in a canonical direction. The first step in the process is to align each individual’s follow-up HP to the corresponding baseline HP by shifting the follow-up HP along its central axis and cutting off slices at the ends that are not shared between the two segmentations (Fig 1a,b). Template alignment similarly shifts each baseline HP along its central axis to align it to the HP of a typical healthy individual (the template HP) and cuts off non-shared slices; it also deforms the baseline HP to correct for major differences in HP shape between individuals. For each surface point on each baseline HP, a set of weighted (soft) correspondences are then estimated between that point and points on the corresponding follow-up HP; the rate of HP change over time at each baseline surface point is calculated in terms of differences in local HP radius (*i.e.*, the length of the vector from the central axis to the point) between baseline and the follow-up surface points that it softly corresponds to (Fig 1b). Hard correspondences between baseline HP and the template HP allow for comparison of HP change rate across individuals at standardized surface locations ¹. At each location on the template HP we find individuals whose estimated radius change there is biologically implausible (Fig 1c); we then estimate a set of shape components, or spatially localized spatial patterns of coherent HP change (Fig 1d), while down-weighting the influence of such implausible change. The following sections describe each step in detail.

¹In our experiments, hard correspondences between baseline and template HP proved adequate, but subtle radius change measurements without weighted correspondences were extremely noisy

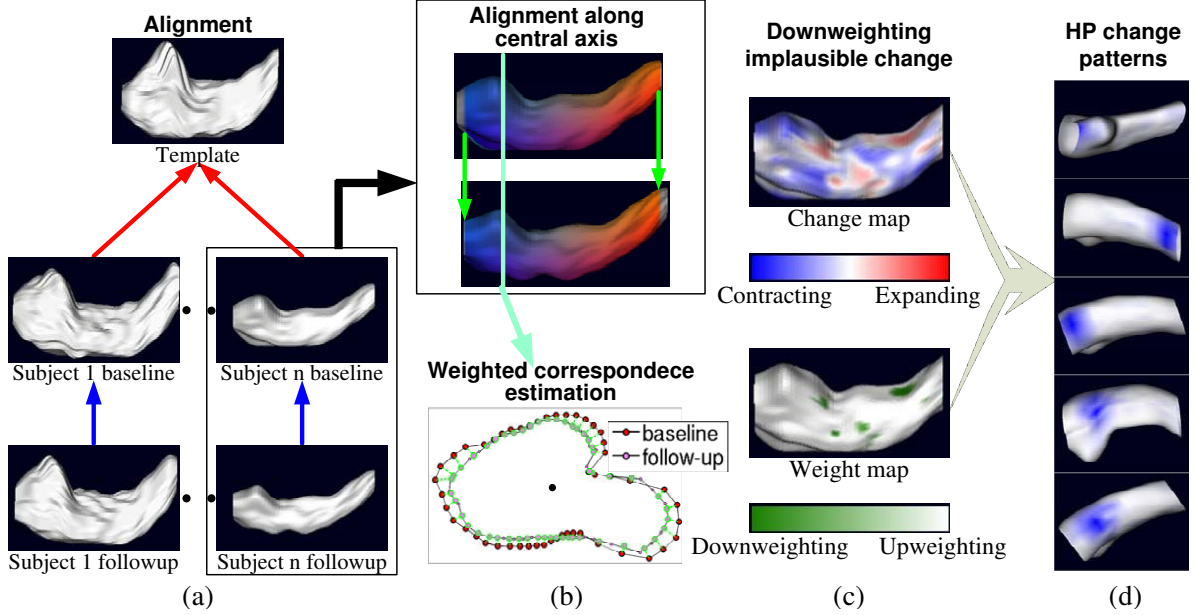


Figure 1: Framework for measuring longitudinal localized hippocampus change. (a) Follow-up HP are aligned to baseline HP (blue arrows) and baseline HP are aligned to a common template (red). (b) The alignment consists of deforming and shifting HP along their central axes, cutting off structure end slices not shared between the pair (top), and establishing weighted correspondences at each surface point (bottom). (c) Rates of local HP radius decrease are estimated (top) and points with implausible change rates are down-weighted. (d) Spatial patterns of coherent HP change (blue spots) are generated by a version of LoCA that accounts for the down-weighting.

2.1 Alignment along central axis

For both the alignment of baseline to follow-up, and the alignment of baseline to template, the purpose of alignment along the central HP axis is to estimate a function f that maps slices t of one HP to slices $f(t)$ of the other. For both alignments, we allow $f(t)$ to be undefined for t at the ends of the HP, to account for the very common differences between HP (even of the same individual) in terms of the starting and stopping points of the segmentation (see Fig 2(e)(f)). Differences between baseline and follow-up HP in the region starting and stopping points represent inconsistencies in the segmentation algorithm; analogous differences between baseline and template HP may additionally represent real anatomical differences. For the alignment of baseline HP to the template, a deformable model is used to account for differences in HP morphology between individuals. For the alignment of follow-up HP to baseline, f is simply a translation designed to cut off non-shared HP slices from both segmentations; a deformable alignment between the two could in fact remove real HP change that had occurred over time.

For deformable alignment, we consider p slices spaced at uniform intervals along the central axis of each individual baseline HP. Let the slice positions of these p control slices along the baseline HP be denoted $\{c_1, c_2, \dots, c_p\}$. Moving these control points to new positions $\{c'_1, c'_2, \dots, c'_p\}$ defines a deformable mapping between baseline and template slices: $f(t) = \sum_{i=0}^p c'_i k_i(t)$ for kernel functions $k_i(t)$. The kernel functions insure that slice c_i in the baseline HP maps to slice c'_i in the template HP, and slices between c_i and c_{i+1} are smoothly mapped to slices between c'_i and c'_{i+1} . In our experiments we used linear kernel functions:

$$k_i(t) = \begin{cases} \frac{t - c_{i-1}}{c_i - c_{i-1}}, & c_{i-1} \leq t \leq c_i, i > 1 \\ \frac{c_{i+1} - t}{c_{i+1} - c_i}, & c_i \leq t \leq c_{i+1}, i < p \\ 0, & \text{otherwise} \end{cases}$$

Thus, the deformable alignment amounts to estimating $\{c'_1, c'_2, \dots, c'_p\}$, and the alignment of follow-up

to baseline is an estimation of a translation along the central axis. These transformations are optimized to maximize HP similarity. To calculate this similarity, we first define a function $d(t)$ along the length of the HP axis that gives the maximum local radius of all surface points on slice t —that is, the maximum distance from boundary points on slice t to the central axis (Fig 2(c)(d)). The similarity metric that drives alignment is the Pearson correlation coefficient between the $d(t)$ for the transformed HP and the $d(t)$ for the template one.² The summary curves provide a computationally efficient means for evaluating shape similarity while capturing the anatomically salient characteristics of HP morphology.

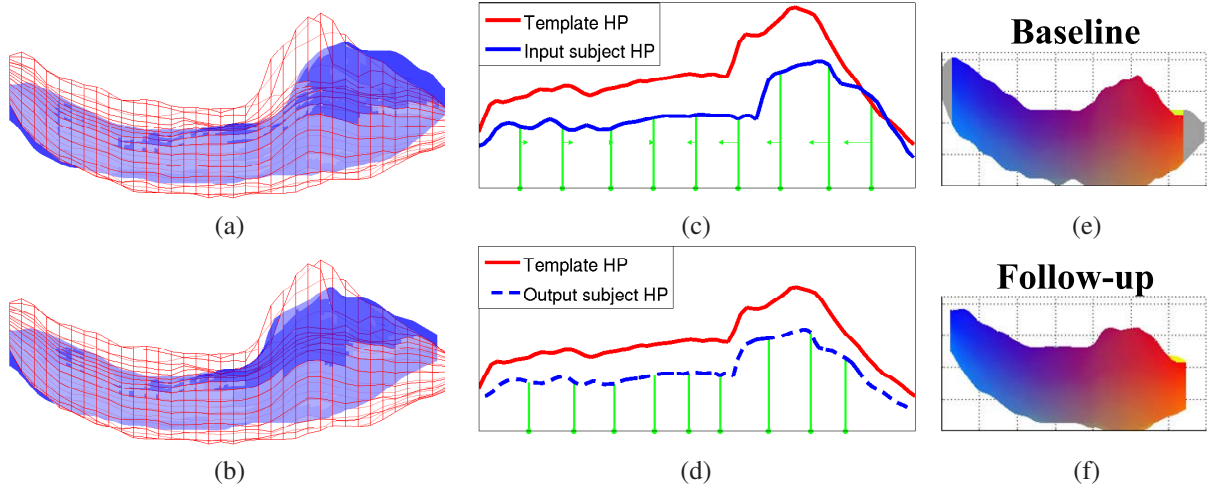


Figure 2: A baseline subject HP (blue surface) and a template HP (red wireframe) are shown overlaid before (a) and after (b) alignment from an inferior viewpoint. The curves plotting maximum distance from surface points to the central axis in each slice drive the alignment, and are more similar after alignment (d) than before (c). Green lines show positions of control slices before alignment (c) and after (d), with the direction and magnitude of control slice motion indicated by arrows. (e)(f) After baseline and follow-up HP are aligned, it is very common for end slices to occur in one HP that have no obvious counterpart in the other (gray). These slices are removed prior to change rate estimation.

The basic approach of modifying the transformation to maximize HP similarity is modified to account for additional factors. First, while we want to cut off slices of tissue that are not shared between HP, we also want to prevent the alignment from aggressively cutting off too many slices in the pursuit of higher similarity values. Therefore, the similarity metric is a linear combination of terms that try to maximize HP similarity, and the amount of HP that overlaps between the two, respectively. That is, we minimize $\lambda\rho + (1 - \lambda)r$, where ρ is the correlation between d curves; r is the percentage of HP slices on the fixed HP that overlap with the transformed one; and λ is a constant between 0 and 1. In addition, we encourage the transformations to stay numerically well conditioned by imposing lower and upper bounds on the distance $c'_i - c'_{i-1}$ between adjacent transformed control slices.

We use the `fminsearch` routine in Matlab to maximize this function with respect to the transformation parameters. See Fig 2(a)(b) for an example baseline HP whose similarity to the template increases substantially after alignment.

2.2 Weighted correspondence estimation

Given pairs of follow-up and baseline HP whose slices have been aligned, previous approaches have established correspondences between surface points with corresponding cylindrical coordinates [15]; in theory, the rate of HP change can be calculated by computing the difference in local HP radius at those corresponding surface points. However, because each HP slice is based on a relatively small number of image voxels, the number of non-redundant surface points is also small, thus causing many surface points to have no single natural correspondence on the other HP (see Fig 3). Therefore, we use weighted correspondences to assess the change

²The surface area and circumference, two alternative summary features of each slice, are possible to use for this similarity metric, but these did not provide high-quality alignment results.

rate in local HP radii: we map each surface point to a weighted average point on the other HP, where the weights are based on proximity to the surface point and proximity to the surface normal vector at that surface point.

More formally, we map each surface point in the baseline HP that has cylindrical coordinates (t, θ_0) to a weighted average point $\sum_{\theta} w(\theta_0, \theta) c(f(t), \theta)$, where $c(t, \theta)$ is the Cartesian coordinates of HP surface point (t, θ) .

The weighting function $w(\theta_0, \theta)$ is a sum of two terms that depend on the proximity of $c(f(t), \theta)$ and $c(t, \theta_0)$ and the proximity of vector from $c(t, \theta_0)$ to $c(f(t), \theta)$ to the surface normal vector at $c(t, \theta_0)$. For proximity, we use a logistic function $l(x) = 2 - \frac{2}{1 + e^{-\frac{x}{T}}}$ governed by a slope parameter T , and for similarity in surface normal orientation we calculate the inner product between the surface normal at $c(t, \theta_0)$ and the vector between $c(t, \theta_0)$ and $c(f(t), \theta)$. Let $i(\theta, \theta_0)$ denote this inner product, and let $d(\theta, \theta_0)$ denote the Euclidean distance between $c(t, \theta_0)$ and $c(f(t), \theta)$. The weighting function is then $w(\theta_0, \theta) = \alpha l(d(\theta, \theta_0)) + (1 - \alpha) i(\theta, \theta_0)$.

After weight estimation, for every baseline point, the follow-up boundary points with the highest N weights are picked. The weights of these follow-up points are normalized so that they sum up to 1. The image of a baseline point is the weighted average of the Cartesian coordinates of these N points.

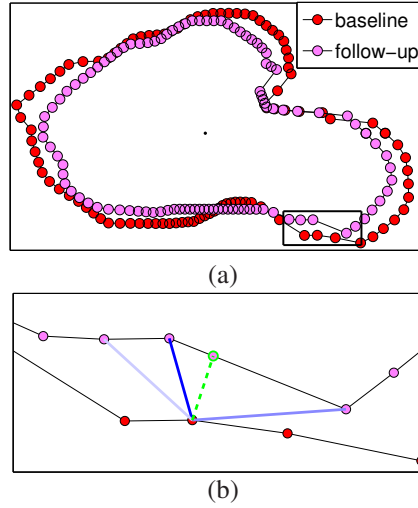


Figure 3: (a) Corresponding slices from a baseline and follow-up HP are shown overlaid. (b) Expanded view of a sub-region of these slices illustrating weighted correspondence. A point on the baseline HP is mapped to the green point on the follow-up HP, which was found by a weighted average of the positions of follow-up points (pink) by their distance to the baseline point and how they deviate from the surface normal direction at the baseline point. Higher weights are indicated by darker blue lines to the baseline point.

2.3 Downweighting

The weighted correspondences allow us to calculate the difference in local radius between each baseline HP surface point and its weighted correspondence point on the follow-up HP; dividing this difference by the interval between baseline and follow-up scans provides a rate of local HP change. The alignment of baseline HP to the template allows us to assemble measurements of rates of local radius change across all individuals at the template surface points. However, because noise in HP boundary measurements is high, some individuals are likely to have spurious weighted correspondences at some surface points, even after alignment and weighted correspondence estimation, leading to spurious estimates of local radius change. Therefore, for each surface point we discount rates of change that are outliers over the population, and also use biological knowledge about the aging and AD processes to discount the influence of such points on estimated spatial patterns of change. Specifically, aging and AD pathology has been universally reported to decrease the size of the HP by promoting injury and death of neurons, without increasing its size in any way; thus, at each HP surface location, positive rates of radius change, suggesting HP growth, and considered implausible. Let $r_{i,t,\theta}$ denote the radius change rate of subject i at template HP cylindrical coordinate (t, θ) . We calculate the mean and standard deviation of

$r_{i,t,\theta}$ over i , $\mu_{t,\theta}$ and $\sigma_{t,\theta}$, and for each i calculate a weight $w_{i,t,\theta}$ for point (t, θ) that tends toward zero as change rates become greater than the mean plus one standard deviation:

$$w_{i,t,\theta} = \begin{cases} l\left(\frac{r_{i,t,\theta} - \mu_{t,\theta}}{\sigma_{t,\theta}}\right), & r_{i,t,\theta} > \mu_{t,\theta} + \sigma_{t,\theta} \\ 1, & \text{otherwise} \end{cases}$$

The logistic function $l()$ is defined as above. In our experiments we set the logistic slope T to .25. These weights are incorporated into the following step that estimates spatially-localized patterns of coherent local HP radius change. The algorithm makes less of an attempt to account for HP change in individual points that have low $w_{i,t,\theta}$, *i.e.* surface points whose change rates are implausible high outliers.

2.4 Localized Change Components

After downweighting, the HP of subject i can be represented by a vector \mathbf{v}_i of the local radius change rates $r_{i,t,\theta}$, and a weight vector \mathbf{w}_i of the corresponding $w_{i,t,\theta}$. We then use a modified version of Localized Components Analysis (LoCA) to estimate a linear subspace for the set of \mathbf{v} , \mathbf{w} pairs [1]. Briefly, LoCA attempts to find an orthogonal set of basis vectors $\mathbf{e}_1, \dots, \mathbf{e}_n$ that provide a concise and spatially localized explanation for the set of \mathbf{v}_i vectors. Because the basis vectors capture a dominant pattern of local HP radius change, we refer to them as *HP change patterns*. After estimation of the \mathbf{e} vectors, individual HP are represented by several *change pattern coefficients* $\alpha_{i,j}$, where $\mathbf{v}_i \approx \mathbf{e}_0 + \sum_j \alpha_{i,j} * \mathbf{e}_j$. Each $\alpha_{i,j}$ quantifies the degree to which a particular HP change pattern \mathbf{e}_j is indicated in the HP local radius change vector \mathbf{v}_i . LoCA generates the initial basis vectors from the covariance matrix of the \mathbf{v}_i , and iteratively optimizes them to minimize an energy function that is a linear combination of 2 terms: one term, like PCA, attempts to make the smallest number of \mathbf{e} vectors account for the largest amount of variance in the \mathbf{v}_i as possible, by analyzing the change pattern coefficients $\alpha_{i,j}$; the other term penalizes basis vectors that are not spatially localized. The addition of the weight vectors $w_{i,t,\theta}$ modifies this optimization in two ways. First, the calculation of \mathbf{e}_0 , the population average HP change rate vector, becomes a weighted average. That is, the entry $\mathbf{e}_0[k]$ representing the average change rate in template HP point k is calculated by: $\mathbf{e}_0[k] = \frac{\sum_i \mathbf{w}_i[k] \mathbf{v}_i[k]}{\sum_i \mathbf{w}_i[k]}$. Similarly, the covariance matrix becomes a weighted

covariance matrix whose entries $C[j, k]$ are: $C[j, k] = \frac{\sum_i \mathbf{w}_i[j] \mathbf{v}_i[j] \mathbf{w}_i[k] \mathbf{v}_i[k]}{\sum_i \mathbf{w}_i[k] \mathbf{w}_i[j] - 1}$.

These modifications effectively decrease the influence of low-weight points on the change pattern calculation. Note that if the elements of weight vector are all 1, the computations of the average vector and covariance matrix are the same as their canonical definitions.

3 Experiments

We tested our methods on a set of 79 AD patients; 146 patients with amnesic mild cognitive impairment (MCI), an AD precursor condition; and 95 normal controls from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) who had baseline and 12-month follow-up MRIs, and semi-automated HP segmentations [11][9].

Average HP change. As expected, local HP radii decreased on average faster among AD patients than MCI patients, and faster among MCI patients than controls (Fig 4). However, if none of the processing steps described above are applied to the segmentations, many local regions appear to grow over time (red), even in the AD group, in clear violation of the conventional wisdom that HP radii should be either stable or decreasing in each group. Each additional step of processing had a visible effect in reducing the magnitude and extent of local HP growth regions, thus increasing the plausibility of the change measurements. The final maps of change calculated after all processing steps among MCI and AD groups suggest prominent atrophy on average in the anterior medial head and lateral body sub-regions, both of which appear to correspond to the CA1 HP sub-field that is preferentially damaged by early AD pathology [6]. The relatively spared (white) HP region in MCI and AD appears to lie within the CA2-4 or dentate gyrus sub-fields, which are relatively preserved by early AD. In addition, the final mean change maps agree with prior studies by suggesting that even among cognitively normal individuals the HP atrophies mildly over time [14].

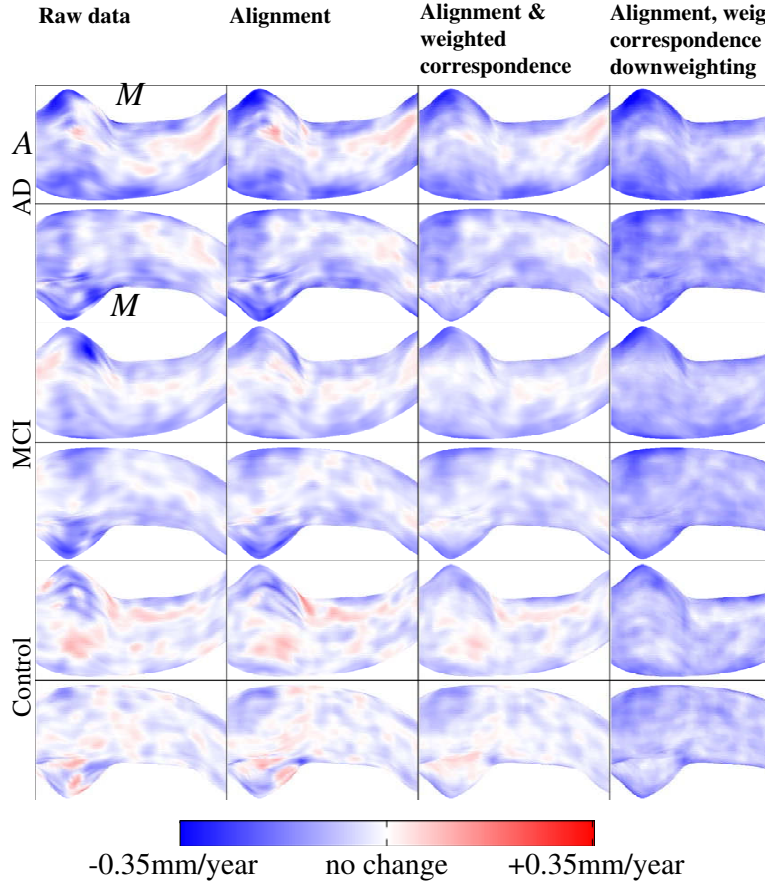


Figure 4: Color-coded maps of average local radius change rates among ADNI control, MCI, and AD groups calculated after application of varying sets of algorithm steps. Red points indicate positive change, *i.e.* HP growth over time, while blue points indicate local atrophy. Odd and even rows render the HP from superior and inferior viewpoints respectively (anterior and medial directions are marked with *A*, *M*).

Localized HP change patterns. We ran LoCA on three versions of HP change maps that represented differing levels of processing to assess whether the various steps ease the estimation of plausible, spatially localized HP change patterns. LoCA was applied to: 1. the raw data, 2. the data after applying only the down-weighting step in Section 2.4, and 3. the data after applying alignment, weighted correspondence estimation, and down-weighting. The individual radius change rate map of each subject can be expressed as the sum of a population average change rate map plus a linear combination of spatial change patterns estimated by LoCA. The five LoCA change patterns that account for the highest percentage of variability in HP change rate for each method are shown in Figure 5, with blue areas representing a spatial pattern of coherent HP change that was indicated by the data set. The renderings suggest that the processing steps eased the estimation of biologically plausible, localized change patterns. In particular, some of the change patterns based on the raw data and downweighting only are not spatially localized, including blue areas at both the top and bottom of the structure simultaneously (raw data columns 1 and 3; downweighting columns 3 and 5). The goal of LoCA is to exclude such patterns, which are more difficult to interpret anatomically. In addition, the change patterns based on raw data and downweighting only suggest that substantial HP radial change is only occurring in the HP head (top of each sub-figure), while change patterns covering a broader extent of the structure, especially the CA1 and subiculum regions preferentially damaged by early AD pathology, are found from the aligned and downweighted data. Together, these observations suggest that biologically-plausible local change patterns may be easier to discover from the data due to the processing steps.

Relation between HP change and cognition. Given LoCA change patterns, the HP change pattern of each individual can be expressed in terms of change pattern coefficients: these quantify the degree to which any




Figure 5: For 3 differing algorithm variants, the 5 LoCA change patterns accounting for the greatest amount of HP change rate variability are shown. For each rendering, the change pattern accounts for a coherent pattern of change in the region shown in blue.

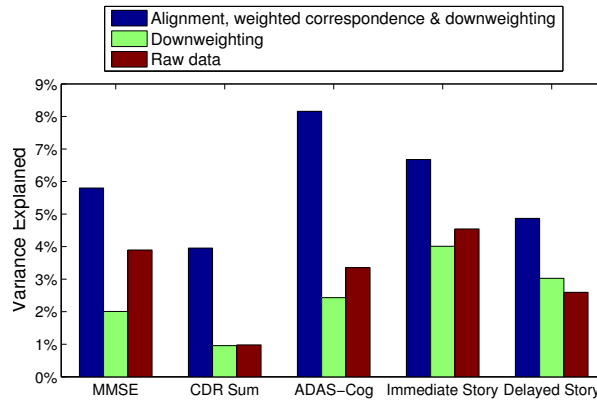


Figure 6: For 3 different algorithm variants, the amount of variance in the ADNI cognitive measures that is explained by the top 5 LoCA change patterns is plotted. See text for details on variance explanation calculation. particular change pattern (such as the ones shown in Figure 5) is evident in that individual. To test whether the processing steps encourage the discovery of clinically relevant HP change patterns, we analyzed associations between cognitive measures collected by ADNI at baseline, and HP change pattern coefficients. Based on the natural history of AD, we would expect diminished cognitive scores to indicate greater levels of AD pathology and therefore greater subsequent HP atrophy; thus, HP change patterns that more strongly associate with baseline cognition may be more clinically relevant as indicators of AD-related HP injury. For each cognitive measure and algorithm variant, we estimated multiple linear regression models with the coefficients for the 5 change patterns shown in Figure 5 as predictors, and the cognitive measures as the outcome. We use the total percent of variance explained by the predictors (the total R^2 of the model) to quantify the strength of association between the change patterns and the cognitive measure. For each of the 5 cognitive measures, the R^2 for the data based on alignment and downweighting was substantially higher than the R^2 based on downweighting only or the R^2 based on raw data (Fig 6), suggesting that the full algorithm enhances the discovery of HP change patterns that are more relevant in terms of stronger associations with real-world outcome measures.

4 Discussion

The hippocampus is a very small brain structure relative to conventional MRI voxel sizes, it is difficult to delineate in a reliable way, and its changes can be subtle even in degenerative disease. Our results suggest that a precise sequence of HP-specific algorithmic steps, including alignment, weighted correspondence estimation, and downweighting, may be required to overcome these issues and estimate rates of localized HP change that

are both biologically plausible and clinically relevant. Future work should extend the technique to longer series of longitudinal scans, for which groupwise alignment and reweighting of points may be required to assess localized change. In addition, HP change from other modalities that best represent the structure as a filled volume, including fMRI and diffusion MRI, is an increasingly important problem that the current approach may be extended to.

References

- [1] D. Alcantara, O. Carmichael, E. Delson, W. Harcourt-Smith, K. Sterner, S. Frost, R. Dutton, P. Thompson, H. Aizenstein, O. Lopez, J. Becker, and N. Amenta. Localized components analysis. In *Proc. IPMI*, 2007. 6
- [2] M. Bobinski, M. D. Leon, J. Wegiel, S. Desanti, A. Convit, L. S. Louis, H. Rusinek, and H. Wisniewski. The histological validation of post mortem magnetic resonance imaging-determined hippocampal volume in alzheimer's disease. *J Neurosci*, 95(3):721–5, 1999. 1
- [3] M. Boccardi, M. Bocchetta, R. Ganzola, N. Robitaille, A. Redolfi, G. Bartzokis, R. Camicioli, J. Csernansky, M. de Leon, L. deToledo Morrell, R. Killiany, S. Lehricy, J. Pantel, J. Pruessner, H. Soininen, C. Watson, S. Duchesne, C. Jack, and G. Frisoni. Estimating the impact of differences among protocols for manual hippocampal segmentation on alzheimers disease-related atrophy: Preparatory phase for a harmonized protocol. *Alzheimer's and Dementia*, 7, 2011. 2
- [4] O. Carmichael, J. Xie, E. Fletcher, B. Singh, C. DeCarli, , and the Alzheimers Disease Neuroimaging Initiative. Localized hippocampus measures are associated with Alzheimer pathology and cognition independent of total hippocampal volume. *Neurobiol. Aging*, 2011. 1
- [5] J. Csernansky, J. Hamstra, L. Wang, D. McKeel, M. Gado, and J. Morris. Correlations between antemortem hippocampal volume and postmortem neuropathology in ad subjects. *Alzheimer Dis Assoc Disord*, 18, 2004. 1
- [6] J. Csernansky, L. Wang, J. Swank, J. Miller, M. Gado, D. McKeel, M. Miller, and J. Morris. Preclinical detection of alzheimer's disease: hippocampal shape and volume predict dementia onset in the elderly. *NeuroImage*, 25:783–792, 2005. 1, 6
- [7] E. Geuze, E. Vermetten, and J. Bremner. Mr-based in vivo hippocampal volumetrics: 1. review of methodologies currently employed. *Mol. Psych.*, pages 1–13, August 31 2004. 2
- [8] B. Gutman, Y. Wang, J. Morra, A. W. Toga, and P. M. Thompson1. Disease classification with hippocampal shape invariants. *Hippocampus*, 19(6):572–578, June 2009. 1
- [9] Y.-Y. Hsu, N. Schuff, A.-T. Du, K. Mark, X. Zhu, D. Hardin, and M. W. Weiner. Comparison of automated and manual mri volumetry of hippocampus in normal aging and dementia. *J Magn Reson*, 16(3):305–310, 2002. 6
- [10] C. Jack, F. Barkhof, M. Bernstein, M. Cantillon, P. Cole, C. DeCarli, B. Dubois, S. Duchesne, N. Fox, G. Frisoni, H. Hampel, D. Hill, K. Johnson, J.-F. Mangin, P. Scheltens, A. Schwarz, R. Sperling, J. Suhy, P. Thompson, M. Weiner, and N. Foster. Steps to standardization and validation of hippocampal volumetry as a biomarker in clinical trials and diagnostic criterion for alzheimer's disease. *Alzheimer's and Dementia*, 7, 2011. 2
- [11] C. Jack, M. Bernstein, N. Fox, P. Thompson, G. Alexander, D. Harvey, B. Borowski, P. Britson, J. Whitwell, C. Ward, A. Dale, J. Felmlee, J. Gunter, D. Hill, R. Killiany, N. Schuff, S. Fox-Bosetti, C. Lin, C. Studholme, C. DeCarli, G. Krueger, H. Ward, G. Metzger, K. Scott, R. Mallozzi, D. Blezek, J. Levy, J. Debbins, A. Fleisher, M. Albert, R. Green, G. Bartzokis, G. Glover, J. Mugler, and M. Weiner. The alzheimer's disease neuroimaging initiative (adni): Mri methods. *J Magn Reson*, 27(4):685–691, 2008. 6
- [12] C. Jack, D. Dickson, J. Parisi, Y. Xu, R. Cha, P. O'Brien, S. Edland, G. Smith, B. Boeve, E. Tangalos, E. Kokmen, and R. Petersen. Antemortem MRI findings correlate with hippocampal neuropathology in typical aging and dementia. *Neurology*, 58(5):750–757, March 2002. 1
- [13] J.-M. Peyrat, H. Delingette, M. Sermesant, X. Pennec, C. Xu, and N. Ayache. Registration of 4d time-series of cardiac images with multichannel diffeomorphic demons. In *Proc. MICCAI*, pages 972–979, 2008. 2
- [14] N. Raz, U. Lindenberger, K. M. Rodrigue, K. M. Kennedy, D. Head, A. Williamson, C. Dahle, D. Gerstorf, and J. D. Acker. Regional brain changes in aging healthy adults: General trends, individual differences and modifiers. *Cereb. Cortex*, 15:1676–1689, November 2005. 6
- [15] P. Thompson, K. Hayashi, G. de Zubicaray, A. Janke, S. Rose, J. Semple, M. Hong, D. Herman, D. Gravano, D. Dordrell, and A. Toga. Mapping hippocampal and ventricular change in Alzheimer's disease. *NeuroImage*, 22(4):1754–1766, August 2004. 1, 4
- [16] G. Wu, Q. Wang, H. Jia, and D. Shen. Registration of longitudinal image sequences with implicit template and spatial-temporal heuristics. In *Proc. MICCAI*, pages 618–625, 2010. 2
- [17] J. Xie, D. Alcantara, N. Amenta, E. Fletcher, O. Martinez, M. Persianinova, C. DeCarli, and O. Carmichael. Spatially localized hippocampal shape analysis in late-life cognitive decline. *Hippocampus*, 19(6):526–532, June 2009. 1, 2