Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg



CrossMark

Event time analysis of longitudinal neuroimage data

Mert R. Sabuncu ^{a,b,*,1}, Jorge L. Bernal-Rusiel ^{a,1}, Martin Reuter ^{a,c}, Douglas N. Greve ^a, Bruce Fischl ^{a,b}, For the Alzheimer's Disease Neuroimaging Initiative ²

^a Athinoula A. Martinos Center for Biomedical Imaging, Harvard Medical School/Massachusetts General Hospital, Charlestown, MA, USA

^b Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

^c Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

ARTICLE INFO

Article history: Accepted 4 April 2014 Available online 13 April 2014

Keywords: Longitudinal studies Linear mixed effects models Event time analysis Cox regression Survival analysis

ABSTRACT

This paper presents a method for the statistical analysis of the associations between longitudinal neuroimaging measurements, e.g., of cortical thickness, and the timing of a clinical event of interest, e.g., disease onset. The proposed approach consists of two steps, the first of which employs a linear mixed effects (LME) model to capture temporal variation in serial imaging data. The second step utilizes the extended Cox regression model to examine the relationship between *time-dependent* imaging measurements and the timing of the event of interest. We demonstrate the proposed method both for the univariate analysis of image-derived biomarkers, e.g., the volume of a structure of interest, and the exploratory mass-univariate method employs a recently developed spatial extension of the LME model. We applied our method to analyze structural measurements computed using FreeSurfer, a widely used brain Magnetic Resonance Image (MRI) analysis software package. We provide a quantitative and objective empirical evaluation of the statistical performance of the proposed method on longitudinal data from subjects suffering from Mild Cognitive Impairment (MCI) at baseline.

Introduction

Medical events, such as the onset of disease, represent major landmarks in the course of a patient's clinical history. A significant portion of biomedical research is dedicated to studying the risk factors associated with these events, aiming to predict, delay and ultimately prevent their occurrence.

In recent decades, neuroimaging has accelerated the study of brainrelated clinical conditions. A classical neuroimaging approach has been to contrast measurements obtained from those who have experienced the event (i.e., cases) with measurements from those who have not (i.e., controls). This methodology has yielded reliable markers of disease, e.g., (Jack et al., 2012), while providing insights about underlying biological mechanisms, e.g. (Buckner et al., 2005; Sabuncu et al., 2012).

Yet, the classical case-control approach treats the two groups as distinct entities and assumes a certain amount of within-group homogeneity. This approach can therefore be limited when the control group is a high-risk cohort, that is, when a significant proportion of subjects have not yet experienced the event of interest but are likely to do so in the not-too-distant future. Such "pre-event" cases, which, in the absence of other information will be treated as controls, typically fall in the gray area between a pure case and a pure control. Thus the within-group homogeneity assumption is violated, which will in turn impact statistical inference. Common examples for this are longitudinal studies of populations that are at high risk for disease, based on their genetic make-up (e.g., carriers of a faulty allele of the Huntingtin gene in a Huntington's study (Albin et al., 1990)), familial history (e.g., subjects who have a first-degree relative with schizophrenia (Whitfield-Gabrieli et al., 2009)) or clinical presentation (e.g., subjects with Mild Cognitive Impairment, or MCI, in an Alzheimer's study (Forsberg et al., 2008)). These examples are particularly relevant to drug trials focused on the pre-clinical or early phases of a disease and thus target high-risk populations. In such scenarios, an inappropriate statistical treatment of the group of subjects who have not been observed to experience the event (diagnosis or conversion to disease) during the follow-up period (sometimes referred to as "non-converters") can introduce bias into the analysis and/or reduce efficiency.

An alternative strategy that addresses this issue, directly models the timing of the event of interest, while accounting for finite follow-up or censoring. This is the event time (or survival) analysis approach (Kleinbaum and Klein, 2012), which includes classical models such as Cox proportional hazards regression (Cox, 1972). Standard event time



^{*} Corresponding author at: Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Building 149, 13th Street, Room 2301, Charlestown, MA 02129, USA. Fax: +1 617 726 7422.

E-mail address: msabuncu@nmr.mgh.harvard.edu (M.R. Sabuncu).

¹ Mert R. Sabuncu and Jorge L. Bernal-Rusiel contributed equally.

² Data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database. As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators is available at http://tinyurl.com/ADNI-main.

analysis models have been applied in prior neuroimaging studies (Desikan et al., 2009, 2010; Devanand et al., 2007; Geerlings et al., 2008; Marcus et al., 2007; Sabuncu, 2013; Stoub et al., 2005; Tintore et al., 2008; Vemuri et al., 2011) and have yielded novel insights about various clinical conditions. Most of these prior studies have analyzed associations between imaging measurements from a single baseline visit and the timing of the event of interest identified via follow-up clinical assessments. These analyses typically rely on survival models (e.g., the standard Cox model) that assume the explanatory variables are independent of time (gender, genetic marker, birth place, etc.). The employed models are useful for constructing individualized survival curves and making predictions about the timing of a future event. Furthermore, they offer insights about the relationships between independent variables and the event time. As such, survival models have been used to draw conclusions about associations between neuroimaging measurements (e.g., volume of a structure) and the clinical event (e.g., disease onset). This type of inference, however, suffers from two problems. Firstly, imaging measurements typically vary over time (e.g., due to anatomical changes). Yet, interpretation of the standard Cox model, for example, has to be done with respect to the baseline imaging measurements only and not with respect to the dynamically changing measurements. Secondly, in longitudinal designs that span an extended time period, imaging measurements are likely to vary substantially over time, making it harder to detect associations between baseline imaging markers and the clinical event.

Longitudinal neuroimaging (LNI) studies, where *multiple* serial images are acquired for each participant, provide a means to characterize the temporal trajectories of imaging measurements. Furthermore LNI studies can offer a substantial increase in statistical power for studying imaging markers (Bernal-Rusiel et al., 2013a,b), while opening up the possibility of examining the relationship between the temporal dynamics of imaging markers and clinical variables (Sabuncu et al., 2011). Today, the standard strategy for analyzing the association between LNI data and the occurrence of a clinical event, such as disease onset, is to perform a group comparison based on dichotomizing the subjects into, for example, "converters" versus "non-converters" (Borgwardt et al., 2011; Chetelat et al., 2005; Jack et al., 2008a; Morgan et al., 2011; Sun et al., 2009). However, as we discussed above, this approach can be sub-optimal, since the non-converter group likely includes subjects who might convert beyond the study follow-up.

The core goal of this paper is to propose a powerful method for the statistical analysis of the associations between longitudinal neuroimaging measurements, e.g., of gray matter density or cortical thickness, and the timing of a clinical event of interest, such as disease onset. The proposed approach combines a linear mixed effects (LME) model that captures the spatiotemporal correlation pattern in serial imaging data (Bernal-Rusiel et al., 2013a,b; Verbeke and Molenberghs, 2000) and an extended Cox regression model that allows the examination of associations between the *time-dependent* imaging measurements and the timing of a clinical event (Kleinbaum and Klein, 2012). Recent work showed that such a joint analysis can reduce bias and increase statistical efficiency by exploiting all available information (Tsiatis and Davidian, 2004).

We demonstrate the proposed method both for the univariate and mass-univariate analysis of imaging measurements automatically computed with FreeSurfer, a widely used brain Magnetic Resonance Image (MRI) data analysis software package (Dale et al., 1999; Fischl, 2012; Fischl and Dale, 2000; Fischl et al., 1999a,b). We include a quantitative and objective empirical evaluation of the statistical performance of the proposed method based on publicly available data (the Alzheimer's disease neuroimaging initiative, ADNI³) from a group of subjects with Mild Cognitive Impairment (MCI) (Gauthier et al.,

2006), a clinically defined condition associated with high-risk incipient dementia. Our experiments revealed that the proposed method offers a substantial increase in statistical efficiency relative to a "two-sample" benchmark method that compares those who convert from MCI to clinical AD against those who remain MCI through follow-up; and a classical Cox regression analysis that employs only baseline scans.

The paper is organized as follows. The Cox proportional hazards model and its extension section and the Linear mixed effects models for longitudinal data section review the Cox proportional hazards and linear mixed effects models, respectively. The Proposed strategy for joint analysis of event time and LNI data section presents the proposed method that unifies these two frameworks. The Alternative methods section describes the alternative analysis strategies that we will use to benchmark our experimental results. The ADNI data section offers a description of the data used in the experiments and the Statistical models section details the statistical analyses conducted on these data. In the Experimental results section, we present experimental results that illustrate the proposed joint modeling approach and compare it against benchmarks. Finally, the Discussion section provides a discussion of the main experimental findings and the Conclusions section closes with concluding remarks.

Material and methods

The Cox proportional hazards model and its extension

In this section, we provide a brief overview of the classical Cox proportional hazards model (Cox, 1972) and its extension for time-varying explanatory (independent) variables. For a detailed treatment, the reader is referred to dedicated texts, such as Kleinbaum and Klein (2012).

A core component of event time models is the so-called hazard function h(t), which is the instantaneous probability of experiencing the event of interest (e.g., disease onset), given no event up to time t. The hazard function is mathematically defined as:

$$h(t) = \lim_{dt \to 0} \frac{P(t \le T < t + dt | T \ge t)}{dt},$$

where *T* is the random variable that represents the time of event and *p*(. |.) denotes conditional probability. The classical Cox model assumes that the hazard function of a sample with *p* time-independent explanatory variables $X = (X_1, X_2, ..., X_p)$ can be expressed as:

$$h(t,X) = h_0(t) \exp\left(\sum_{i=1}^p \alpha_i X_i\right),\tag{1}$$

where $h_0(t)$ is the so-called baseline hazard function and $\alpha = (\alpha_1, ..., \alpha_p)$ is the coefficients associated with the explanatory variables. This model assumes that the hazard function can be written as a product of two factors: one that varies with time but is independent of *X*, and another that is a function of the *time-independent* explanatory variables *X* and thus is fixed over time. The foundation of the classical Cox model is the proportional hazards assumption, i.e., the proportion of the hazard

functions of two samples is constant over time: $\frac{h(t,X^1)}{h(t,X^2)} = \exp(t)$

 $\left(\sum_{i=1}^{p} \alpha \left(X_{i}^{1}-X_{i}^{2}\right)\right) = \text{const.}$, where X^{j} is the independent variables of the *j*'th sample.

A popular strategy to estimate the coefficients α is the so-called *partial* likelihood maximization method (Cox, 1972, 1975). In particular, the partial likelihood is expressed as a product of *K* terms, each corresponding to the likelihood of an observed event computed based on the time of occurrence (i.e., there are *K* events observed during the

³ http://tinyurl.com/ADNI-main.

study and $K \le N$, where *N* is the total number of samples⁴). Note that, although only observed events are considered, their likelihood values depend on those samples that haven't experienced the event yet and still remain in the study (i.e., have not dropped out or are not "censored" before the particular event time⁵).

For readability, let us index the samples such that the first *K* are those that we have event timing information on. All remaining samples are thus censored. That is, they either drop out of the study before experiencing the event or do not experience the event during their follow-up. Then, mathematically, the partial likelihood is computed as:

$$L = \prod_{k=1}^{K} L^{k}, \text{ where } L^{k} = \frac{h\left(t^{k}, X^{k}\right)}{\sum_{r \in \mathbb{R}^{k}} h\left(t^{k}, X^{r}\right)}.$$
(2)

In Eq. (2), t^k and X^k denote the event time and exploratory variables of the *k*'th sample, respectively. R^k is the so-called risk set at t^k , i.e., the set of samples that are *known* to have not experienced the event at t^k . The partial likelihood function of Eq. (2) is then maximized with respect to the unknown model parameters α via a numerical optimization strategy, such as Newton–Raphson.

As mentioned above, the original Cox model requires that the explanatory variables be constant over time (e.g., a genetic marker). Hence, this model is not appropriate for analyzing longitudinal imaging measurements that typically vary over time. We note that baseline imaging measurements, however, have been employed in prior neuro-imaging studies with standard proportional hazard models (Desikan et al., 2009, 2010; Devanand et al., 2007; Geerlings et al., 2008; Marcus et al., 2007; Tintore et al., 2008; Vemuri et al., 2011). Although these analyses offer predictive models, their use for inferring associations is restricted to baseline measurements.

The Cox model can easily be extended to handle time-varying variables (Kleinbaum and Klein, 2012). The hazard function is then expressed as:

$$h(t, X, Y(t)) = h_0(t) \exp\left(\sum_{i=1}^p \alpha_i X_i + \sum_{j=1}^q \gamma_j Y_j(t)\right),\tag{3}$$

where the second term in the exponential includes the effects of q time-varying variables $Y(t) = (Y_1(t), ..., Y_q(t))$ with associated coefficients $\gamma = (\gamma_1, ..., \gamma_q)$. Partial likelihood maximization can also be employed to solve the extended Cox model of Eq. (3).

Here, we make a crucial observation. To evaluate the partial likelihood function of the extended Cox model, at each event time we need to be able to compute or observe the time-dependent variable for all samples in the risk set, that is those samples *under observation and have not experienced the event yet*. Note that, this means that the value of a time-dependent variable of a subject needs to be identifiable not only for the event time of that particular subject, but also for other relevant event times that are prior to the subject's own event/censor time as well.

Longitudinal neuroimaging (LNI) studies offer us the opportunity for identifying the time-varying imaging measurements at different time points. However, in a typical LNI study, image data are acquired at certain intervals (e.g., every six months). Furthermore, these serial scans are usually unbalanced across subjects, i.e., their number and timing can vary between subjects. Finally, the visits for clinical assessments and image acquisitions might not coincide. *Hence, the fundamental challenge* of the application of the (extended) Cox model to the analysis of LNI data is the computation (or estimation) of the image data at the times of the clinical events. Note that, even if we have image data coinciding with a subject's own event time (e.g. disease onset), we typically will not have the corresponding imaging measurements for that subject for other relevant event times observed in the study. We propose to use a linear mixed effects (LME) model, which captures the temporal trajectories of each individual's imaging measurements, to estimate image data at all observed event times. The following section provides a brief description of the LME approach and its recently introduced spatial extension.

Linear mixed effects models for longitudinal data

In two recent papers (Bernal-Rusiel et al., 2013a,b), we illustrated the use of linear mixed effects (LME) models for the analysis of longitudinal neuroimage data. The classical LME model can handle unbalanced data with high inter-subject variability in scan times and missing data points, while offering a parsimonious yet effective strategy to model the mean and covariance structure in longitudinal data (Fitzmaurice et al., 2011; Verbeke and Molenberghs, 2000). The central idea in LME is to allow a subset of the regression parameters to vary randomly across subjects. Hence, the mean trajectory is modeled as a combination of population-level "fixed" effects and subject-specific "random" effects (together, they are called mixed effects).

Formally, the LME model for longitudinal data can be expressed as:

$$Y(t) = \sum_{i=1}^{f} \beta_i F_i(t) + \sum_{j=1}^{r} b_j R_j(t) + e,$$
(4)

where Y is the time-dependent outcome measurement of a subject, F(t)denotes the value of *F* at time *t*, $F = (F_1, ..., F_f)$ is the *f* so-called "fixed" effects that can include subject-level constant variables (e.g., gender, genotype) or time-varying variables (clinical status, measurement time, etc.), $R = (R_1, ..., R_r)$ is the *r* "random" effects, which can include a constant bias term and/or a subset of the time-varying fixed effect variables (e.g., measurement time). $\beta = (\beta_1, ..., \beta_f)$ and $b = (b_1, ..., \beta_f)$ b_r) are the unknown fixed and random effect coefficients, respectively and *e* is independent and identically distributed zero-mean Gaussian measurement noise with an unknown variance σ^2 . We further assume that the random effect vector *b* is sampled for each subject from a zero-mean Gaussian with an unknown, non-diagonal $r \times r$ covariance matrix D. The unknown model parameters are the measurement variance σ^2 and random effect covariance matrix *D*. The traditional LME approach solves for the model parameters via an iterative restricted maximum likelihood (ReML) procedure.

Given estimates \hat{D} and $\hat{\sigma}$, we have a closed-form solution for the maximum likelihood (ML) estimate of the fixed effect coefficients $\hat{\beta}$. We can further use these estimates to compute a prediction for the values of the subject-specific random effect coefficients \hat{b}_s , where the subscript *s* denotes subject index (Fitzmaurice et al., 2011). One can then easily compute an unbiased prediction for the temporal trajectory for subject *s* as:

$$\hat{Y}_{s}(t) = \sum_{i=1}^{f} \hat{\beta}_{i} F_{si}(t) + \sum_{j=1}^{r} \hat{b}_{sj} R_{sj}(t).$$
(5)

We recently extended the classical LME model to handle spatial data, such as image-wide measurements in a mass-univariate analysis (Bernal-Rusiel et al., 2013b). This approach, called ST-LME, essentially builds on the LME framework but modifies it to model spatial correlations via a parametric spatial covariance matrix. The parameters associated with the spatial matrix are added to the list of unknown model parameters and estimated via ReML. The prediction of subject-level trajectories can then be computed exactly the same way using the estimated model coefficients and Eq. (5). Our prior experiments have demonstrated

⁴ In this paper, we assume that each sample can experience the event of interest at most once. There are extended treatments that relax this assumption. We consider these outside the scope of this manuscript.

⁵ We note that the Cox model does not assume that the event will definitely occur for each individual during his or her lifetime, since T can theoretically approach infinity.

that the ST-LME model offers superior statistical efficiency since it exploits the spatial structure in image data.

Proposed strategy for joint analysis of event time and LNI data

As discussed above, the extended Cox model requires that the value for the time-dependent variables be specified for each subject in the risk set at each observed event time. In general, many of these imaging measurements are unavailable. However, one can estimate these data based on serial measurements available in a study. We propose to use the LME model to compute these estimates. The LME approach (and its mass-univariate, spatial extension) provides a way to parsimoniously model the spatial and temporal correlation structure in serial imaging data, while accounting for unbalanced longitudinal data collection.

The proposed strategy consists of two steps. In the first step, we fit an LME model to the longitudinal imaging data. The details of model selection, the determination of random effects, and parameter estimation are provided in Bernal-Rusiel et al. (2013a,b). Once the model parameters are estimated, one then computes and saves the population-level fixed-effect coefficients and subject-level random-effect coefficients.

In the second step, we fit an extended Cox regression model (Eq. (3)) to the event time data, where for each observed event time we compute the imaging measurements of each subject based on the LME model fit in the previous step (Eq. (5)). Statistical inference (i.e., hypothesis testing) on the association between imaging measurements and the event can then be conducted based on the Wald statistic of the approximately standard normal distribution of $\frac{\hat{\gamma}_l}{\sqrt{Var}(\hat{\gamma}_l)}$ (a *Z* score),

which yields a p-value for the effect of interest (Kleinbaum and Klein, 2012). The estimate for $V\hat{a}r(\hat{\gamma}_i)$ is computed as the negative of the inverse of the Hessian matrix associated with the ML solution of the extended Cox model.

Here, we would like to make several remarks. First, the LME model of the first step and the extended Cox model of the second step can contain different sets of time-independent variables, although in our experiments we chose to use the same set for both models. This decision was motivated by recent joint frameworks, where models for the event time distribution and longitudinal data are taken to depend on a common set of effects (Tsiatis and Davidian, 2004). We note that the selection of the explanatory variables to be included in a model is a general problem in regression and should be made based on domain knowledge and study constraints. Second, the time-dependent variables of the LME model have to be identifiable for arbitrary timepoints, since we need to compute predictions for all relevant observed event times, which in general do not coincide with the imaging times. Common examples for such time-varying variables are time elapsed from baseline, subject age, and functions of these, e.g., time squared. Finally, the predicted image measurements computed at the exact imaging times, in general, will not be equal to the actual measurements themselves. Rather, they can be considered as "denoised" measurements, where the error in longitudinal observations is estimated and discounted via the LME model. This is similar in spirit to recent joint longitudinal and survival models, e.g. (Kleinbaum and Klein, 2012).

Alternative methods

To date, the most common approach to perform an analysis between neuroimage data and a clinical event of interest, such as disease onset, relies on a two-group comparison (Borgwardt et al., 2011; Chetelat et al., 2005; Jack et al., 2008a; Julkunen et al., 2009; Morgan et al., 2011; Risacher et al., 2009; Sun et al., 2009). In this method, the subjects are divided into two groups: "converters", i.e., those who experience the event during follow-up, and "non-converters", i.e. those who remain clinically stable for a certain amount of follow-up time. In our experiments, we used the two-group approach with a LME model that has been demonstrated to offer excellent statistical power for the analysis of longitudinal neuroimage data (Bernal-Rusiel et al., 2013a, b). Here, the output variables are the imaging measurements and the analyses examine the differences between the intercepts *and* slope coefficients of the two groups.

As a second alternative, we consider employing the extended Cox model but with a simpler method to estimate the longitudinal trajectories of the imaging data. Recall that the proposed approach fits an LMEbased statistical model to longitudinal neuroimage data, which is then utilized to estimate the imaging measurements for all observed event times in the second step. The LME-based statistical model examines the entire data, in order to characterize and estimate the individuallevel temporal trajectories. A more basic strategy would be to fit a linear function independently to each individual's serial imaging data. Note that this alternative approach, while computationally very efficient, ignores the spatial structure in the images. Furthermore, when estimating the individual temporal trajectories, it does not pool information across the population, the way the LME approach does. Hence, we expect the estimates of the temporal trajectories of the imaging measurements to be noisier and thus the inference of the extended Cox model to be statistically less efficient.

Finally, to our knowledge, all prior neuroimaging studies that conducted a Cox regression analysis, simply utilized the baseline scans of each subject. In our experiments, we considered this method as a benchmark as well. However, as we discuss below, this analysis tests a slightly different association.

The ADNI data

MRI processing

We analyzed serial brain MRI data (T1-weighted, 1.5 T), which were acquired and made public by the Alzheimer Disease Neuroimaging Initiative (ADNI). We processed all MRI scans automatically using FreeSurfer (Fischl, 2012) (version 5.1.0, http://surfer.nmr.mgh.harvard. edu, specifically its longitudinal processing pipeline (http://surfer.nmr.mgh.harvard.edu/fswiki/LongitudinalProcessing) (Reuter and Fischl, 2011; Reuter et al., 2010, 2012)).

FreeSurfer computes volume estimates for a wide range of brain structures such as the hippocampus, and estimates of the intra-cranial volume (ICV). In all subsequent analyses, we summed the volumes of the left and right hippocampi to obtain the total hippocampal volume (HV). Additionally, for each MRI scan, FreeSurfer automatically computes subject-specific thickness measurements across the entire cortical surface of each cerebral hemisphere. These measurements are further spatially re-sampled onto a standard surface-based template (*fsaverage*), which represents an average brain.

In our experiments we performed both univariate and massunivariate analyses. Our goal was to detect the association between the longitudinal measurements of neuroimaging biomarkers of AD and clinical progression from MCI to AD. Total hippocampal volume (HV) was the imaging variable of interest in the univariate analyses. The mass-univariate analyses were conducted on cortical thickness data computed across the entire cortex. These two types of measurements were chosen since they have been shown to be strongly associated with progression from MCI to AD (Dickerson et al., 2001, 2009; Jack et al., 1997; Lerch et al., 2005). Cortical thickness maps were smoothed by applying an iterative nearest neighbor averaging procedure that approximates Gaussian kernel smoothing on the high resolution surface of FreeSurfer's *fsaverage* template subject (full-width at half max = 15 mm) (Hagler et al., 2006; Han et al., 2006). For computational efficiency, the mass-univariate analyses were conducted on the left hemisphere of *fsaverage6*, which is a lower resolution version of fsaverage (FreeSurfer's average template surface) and has about 35k vertices.

Longitudinal data from MCI subjects

In our experiments, we analyzed longitudinal imaging and clinical data from the ADNI subjects with MCI (N = 374, 75.7 \pm 6.7 years, 32.6% female). As can be appreciated from Table 1, there is substantial variation between the timing and number of longitudinal visits across these subjects. In many prior studies, the MCI subjects were subdivided into two categories: *progressor MCIs* (or converters), i.e., those who convert to clinical AD during follow-up; and *stable MCIs* (or non-converters), i.e., those who do not progress. However, as we have emphasized above, many so-called stable MCIs, in fact, drop out from the study prematurely or might convert beyond the study follow-up. Hence, a better illustration of the MCI-to-AD conversion data is the non-parametric estimate of the cumulative AD diagnosis probability (see Fig. 1).

In the analyzed data, there were 160 observed event times (i.e., diagnosis of clinical AD), which was on average 1.32 years from baseline with a standard deviation of 0.77 years. For the remaining 214 MCI subjects, the average censor time, i.e., the final follow-up visit, was on average 2.35 years from baseline with a standard deviation of 1.16 years.

Statistical models

Proposed method

The first step of the proposed method fits an LME model to the longitudinal neuroimage data. Here, two important design decisions need to be made: (1) the specification of time-dependent variables that model the mean temporal trajectory, and (2) the selection of the intercept and/or time-dependent variables that will determine the temporal covariance structure. For further detail, the reader is referred to Bernal-Rusiel et al. (2013a). In the mass-univariate setting, these model specification/selection questions are particularly challenging due to the large number of tests that need to be conducted. In our previous analyses of the ADNI data (Bernal-Rusiel et al., 2013a,b), we found that a clinical group-specific linear trajectory was an appropriate model for Alzheimer-associated hippocampal atrophy and cortical thinning during the 4.5-year follow-up period of the ADNI study.

In all reported analyses with the proposed method, the following variables were included as independent (fixed effect) variables: time from baseline, baseline age, sex, APOE genotype status (one if e4 carrier or zero otherwise), interaction between time and APOE genotype status (note that this variable was included based on the evidence that e4 accelerates atrophy during the prodromal phases of AD (Jack et al., 2008b)), and education (in years). For hippocampal volume (HV), we further added ICV to normalize for the confounding effect of head size. Intercept and time from baseline were the only random effect variables in the LME models. For the mass-univariate analyses, we used the recently developed spatial extension of LME, namely ST-LME, with the parameter setting recommended in Bernal-Rusiel et al. (2013b).

We report the results of the statistical tests that examine the association between LNI data and the timing of MCI-to-AD progression. This test employs an extended Cox model with the aforementioned

 Table 1

 The ADNI MCI cohort: Number and timing of analyzed longitudinal MRI scans per visit.

Visit	Number of individual scans	Time from baseline (years)	
Baseline	374	0	
Year 0.5 (month 6)	354	0.58 ± 0.07 [0.32-0.94]	
Year 1	319	1.08 ± 0.07 [0.82-1.35]	
Year 1.5	281	1.59 ± 0.07 [1.26-1.92]	
Year 2	210	2.08 ± 0.08 [1.70-2.52]	
Year 3	100	3.09 ± 0.08 [2.98-3.45]	
Year 4	13	$4.10 \pm 0.10 [3.98 - 4.38]$	
Total	1651		

Time from baseline (in years) is in mean \pm standard deviation; ranges are listed in square brackets.



Fig. 1. Non-parametric estimate of the AD diagnosis cumulative distribution function for the ADNI MCI subjects. The dashed lines show the 95% confidence interval.

explanatory variables and the time-dependent imaging measurements (variable of interest) estimated using the LME model fit in the first step. For all the mass-univariate analyses, multiple comparisons were corrected by employing a powerful two-stage adaptive False Discovery Rate (FDR) procedure at q-level = 0.05 (Benjamini et al., 2006). We note that the proposed method tests for the hypothesis that there is an association between the imaging measurements and concurrent AD onset.

Benchmark two-group method

Our two-sample LME-based analyses used the same independent fixed and random effect variables as the proposed method. In addition, these models included a binary clinical group membership variable (1 if the subject converted to from MCI to AD or 0 if subject was stable during follow-up⁶), and the interaction between group membership and time as additional fixed effects. Under the null hypothesis of the extended Cox analysis there is no association between the imaging measurements and MCI-to-AD conversion. The corresponding null hypothesis for a two-group analysis is that the coefficients associated with group membership are zero. In other words, under the null, the progressor and stable MCIs have the *same* intercept *and* slope coefficient. All the reported results for the benchmark two-group method employed an F-test that was based on this null hypothesis.

Alternative Cox models

We conducted two alternative event time analyses using the Cox proportional hazards model. The first method replaces the first step of the proposed approach with a simple line-fitting scheme. So, instead of fitting an LME based (or ST-LME based) model to the longitudinal imaging data, we fit the best line (in the least squares sense) to the serial measurements of each individual. In the mass-univariate setting, each spatial location was treated independently. We then computed the imaging measurements for each subject and at each observed event time based on these estimated linear trajectories. Given these estimates, the extended Cox regression model and statistical test were identical to the proposed method. The tested hypothesis was identical to that of the proposed model.

In the second method, we ignored the temporal trajectories in the imaging data and simply treated *baseline* measurements as time-independent variables in a classical Cox regression analysis. All

⁶ Stable MCIs were those subjects who were categorized as MCI at baseline and remained so during a clinical follow-up of at least 1 year. Converter MCIs were those MCI subjects who were diagnosed with clinical AD at follow-up.

remaining explanatory variables were identical to the extended Cox models. The reported results were for the test of association between the baseline image variables and time of event. Here, we caution the reader that the tested hypothesis is in fact somewhat different from the one tested with the longitudinal neuroimage data. Relying on baseline measurements allow us only to test associations between the baseline measurement and the timing of a future event. In contrast, the extended Cox analysis we propose in this paper can be used to test associations between the (time-dependent) imaging marker and clinical event. We return to this issue in the Discussion section, where we emphasize the distinction between the two approaches.

Experimental results

Univariate analysis of hippocampal volume

In our first experiment, we compared the statistical performance of the proposed method with the alternative benchmarks based on detecting the known association between MCI-to-AD conversion and total hippocampal volume. To achieve this, we utilized an empirical strategy inspired by Thirion et al. (2007).

For different sample size values (N = 30-100), we randomly selected two sets of independent MCI samples (i.e., two independent samples of size N), from the MCI subjects in the ADNI sample (Total N = 374). Note that there was no overlap between the two independent samples. We repeated this procedure 1000 times to obtain 1000 random pairs of independent MCI samples of a certain size (that is, for a given size we had 2000 random MCI samples in total).

To assess sensitivity, we computed the detection (true positive) rate across the 2000 samples for a range of p-value thresholds and sample sizes (N = 30–100). Here we assumed that the underlying ground truth was that there is an association between hippocampal volume and MCI-to-AD conversion. Instances where the p-value was less than a threshold were considered a "detection" and remaining cases were treated as a false negative. The true positive rate (or sensitivity) was quantified as the fraction of detections.

Fig. 2 shows empirical sensitivity as a function of the p-value threshold and sample size. Fig. 3 plots repeatability, defined as the rate of detection in both of the independent samples. These results demonstrate that the proposed method offers the highest statistical sensitivity and repeatability for detecting associations between imaging measurements and the clinical event of interest. Yet, as we emphasized above, technically, each method is testing a slightly different hypothesis. In particular, the classical Cox analyses of the baseline measurements are testing associations between these values and the timing of the future diagnosis of AD. The two-group method is testing for differences in the trajectories of imaging measurements between those who convert to AD and those who remain MCI during follow-up. Finally, the extended Cox analyses directly test for the associations between the value of imaging measurements and concurrent risk of AD diagnosis.

The boost in statistical efficiency with respect to the common twogroup method is substantial: for a given sample size (e.g. 50) and p-value threshold (e.g. 0.05), the increase in statistical power can be over 10%. These results further illustrate that employing longitudinal imaging data can improve the statistical efficiency of a Cox regression model. Finally, these univariate analyses revealed that the proposed LME-based two-step approach can offer a modest, yet consistent improvement over a more basic extended Cox strategy that uses a simple line-fitting scheme to interpolate the imaging data. This result suggests that the LME-based model, which pools data across subjects, provides improved estimates of the longitudinal trajectories of the imaging measurements.

Finally, to examine the type I error control offered by the Cox models, we conducted an additional analysis. Here, we randomly permuted the event time information within each sample (1000 times),



Fig. 2. Empirical sensitivity (statistical power) as a function of (a) the p-value threshold (N = 50) and (b) sample size (p-value threshold = 0.05) for detecting the association between total hippocampal volume, a *univariate marker*, and MCI-to-AD conversion. The proposed method (Ext. Cox with LME) yields the most statistical power. Ext. Cox (line) replaces the LME-based first step of the proposed method with a simple line fit. Two-class LME is the popular approach of comparing converter MCIs with stable MCIs. Cox baseline uses only the baseline scans and treats imaging measurements as time-independent exploratory variables in a classical Cox regression. See text for further details.

and repeated the analyses with the three Cox-based methods: (i) a classical proportional hazards model that uses the baseline scans only, extended Cox models, with (ii) a line-based interpolation scheme, and (iii) the proposed LME-based estimation strategy. The random permutations simulate a null hypothesis, where the imaging measurements and the event of MCI-to-AD conversion were statistically independent. Since the random permutations broke down the relationship between the image data and the timing of MCI-to-AD conversion, we considered these data as samples from the null hypothesis (Good, 2000; Nichols and Holmes, 2002). Then, for each p-value threshold, we computed the detection rate of the association between the event time and imaging measurement, with each method. Under these random permutations, a method with good type I error control should achieve a detection rate that is close to the used theoretical threshold. Table 2 shows that this is indeed the case. All three Cox-based methods achieve type I error rates that are very close to each other and to the theoretical threshold.



Fig. 3. Repeatability (the frequency at which a method detects an association between total hippocampal volume, a *univariate* marker, and MCI-to-AD conversion in two independent samples) as a function of: (a) the p-value threshold (N = 50), (b) sample size (p-value threshold = 0.05). See caption of Fig. 2 for further details.

Mass-univariate analysis of cortical thickness

In our second experiment, we exploited the known association between regional cortical thinning and MCI-to-AD conversion (Dickerson et al., 2009) and used an empirical strategy similar to the one in the first experiment. Out of 374 ADNI MCI subjects, we randomly drew independent pairs of samples of 80 subjects. Each independent pair shared no common subject. This way, we generated 1000 pairs of independent samples (or 2000 samples in total). For each sample, we used the aforementioned methods to compute significance maps for the association between cortical thickness values and MCI-to-AD conversion. We used the two-stage adaptive FDR procedure with an

Table 2

Empirical type I error rate for the three Cox-based event time models and different p-value thresholds. The null hypothesis was simulated by randomly permuting the event time data.

Theoretical p-value threshold 0.01 0.03 0.05 0.10 Cox with baseline scans 0.00 0.03 0.06 0.11 Extended Cox with line-based interp. 0.01 0.03 0.06 0.12 Extended Cox with LME 0.01 0.03 0.06 0.11					
Cox with baseline scans 0.00 0.03 0.06 0.11 Extended Cox with line-based interp. 0.01 0.03 0.06 0.12 Extended Cox with LME 0.01 0.03 0.06 0.11	Theoretical p-value threshold	0.01	0.03	0.05	0.10
	Cox with baseline scans Extended Cox with line-based interp. Extended Cox with LME	0.00 0.01 0.01	0.03 0.03 0.03	0.06 0.06 0.06	0.11 0.12 0.11



Fig. 4. Empirical statistical power as a function of FDR q-value for detecting the association between *mass-univariate* cortical thickness measurements and MCI-to-AD conversion in groups of 80 MCI subjects. The proposed method (Ext. Cox with LME) yields the most statistical power. Ext. Cox (line) replaces the LME-based first step of the proposed method with a simple line fit. Two-class LME is the popular approach of comparing converter MCIs with stable MCIs. Cox baseline uses only the baseline scans and treats imaging measurements as time-independent exploratory variables in a classical Cox regression. See text for further details.

array of q-values (Benjamini et al., 2006) to control for multiple comparisons. Thus, for each sample, each method and each q-value threshold, we obtained a map of significant associations.

We computed the statistical power (sensitivity) at the sample-level as the fraction of instances (out of the 2000) where a statistical method detected a significant association at a given FDR q-value (see Fig. 4). Next, we assessed repeatability via the overlap area between the two independent MCI samples of size 80 (for FDR q-value = 0.05). Fig. 5 shows the means and standard errors across the 1000 random draws for the four methods that were compared in this study. These results demonstrate that the benchmark two-group method offers the least repeatability, while the extended Cox strategy yields a dramatic improvement in the repeatability of the results. The proposed method



Fig. 5. Average overlap area (in mm²) between the detection maps obtained from two independent samples consisting of 80 MCI subjects. The detection maps were binary masks, where the FDR-corrected statistical association (q-value = 0.05) between cortical thickness and MCI-to-AD conversion was recorded. Vertices that exhibited a significant association in both independent samples were included in the overlap area. See caption of Fig. 4 and text for further details. Errorbars show the standard error of the mean.

of using an LME-based first step to capture the spatiotemporal patterns in the neuroimaging data provides a subtle increase in repeatability.

Analyzing the entire ADNI MCI sample

Finally, we employed the proposed strategy to analyze the association between longitudinal cortical thickness measurements and MCIto-AD conversion in the entire ADNI MCI sample. Fig. 6 presents the map of significant associations. We note that these maps, which include regions associated with the default-mode network, show a striking resemblance to previously reported regions where cortical atrophy was correlated with AD dementia (Buckner et al., 2005; Dickerson et al., 2009). The strong agreement between these results and prior studies increases our confidence in the validity of the proposed method.

Discussion

As shown by recent studies, the event time analysis framework can provide a substantial increase in statistical efficiency when examining associations between imaging biomarkers and a clinical event of interest in a longitudinal study design (Vemuri et al., 2011). In contrast with the more popular two-group comparison method that compares converters and non-converters (i.e., those who experience the event and those who do not), the event time analysis method exploits the variation in the event timing data and, crucially, accounts for finite follow-up. Hence non-converters, i.e., those who are not observed to experience the event of interest, are not treated as a uniform group, separate from the converters, since some of these subjects might eventually experience the event beyond the study follow-up.

The Cox proportional hazards model, which has been employed in neuroimage analysis before (Desikan et al., 2009, 2010; Devanand et al., 2007; Geerlings et al., 2008; Marcus et al., 2007; Stoub et al., 2005; Tintore et al., 2008; Vemuri et al., 2011), is a very flexible and powerful method for conducting event time analyses. However, this method has only been employed to analyze baseline measurements with respect to the timing of a future event, which restricts the associations we can examine and detect. An alternative approach is to model the temporal trajectory in image data and use an extended Cox model that handles time-dependent variables. Longitudinal neuroimaging studies offer us this opportunity. However, the fundamental challenge is that in the extended Cox modeling approach, time-varying variables have to be identifiable at all relevant observed event times, and not just the time of event for the corresponding sample. In this paper, we proposed to model the spatiotemporal patterns in neuroimaging measurements using an LME-based approach (Bernal-Rusiel et al., 2013a,b). The LME model is then used to estimate the image data at observed event times.

We conducted an empirical evaluation of the proposed method, along with three alternative strategies. Each comparison with a benchmark allowed us to quantify the effect of a different factor. The comparison with the popular two-group analysis, which simply contrasts the longitudinal data of converters versus non-converters, reveals the influence of explicit event time modeling achieved with the Cox approach. The comparison with the baseline Cox model, which ignores the longitudinal neuroimage data and simply uses the baseline scans (as done in prior neruoimaging studies), gives us information on the effect of exploiting longitudinal imaging. Finally, the comparison with an alternative extended Cox method that uses a simple line-fitting step to model the temporal trajectory in the imaging measurements, uncovers the impact of the LME-based first step in the proposed approach.

In agreement with prior studies, we found that the Cox proportional hazards approach offers a significant improvement in statistical power and repeatability. For example, the agreement between the maps of two independent MCI samples of 80 subjects was about 50 times greater for the proposed method compared to the two-group benchmark. Similar gains were consistently observed in other univariate and mass-univariate analyses.

Secondly, the proposed method was substantially more powerful and reliable than a classical Cox analysis of the baseline scans. In the mass-univariate experiment, this gain was slightly less than the improvement with respect to the two-group method, yet it was consistent across all analyses. These results highlight the advantage of utilizing and modeling longitudinal imaging data in a Cox analysis. Obviously, this might be impractical in experimental designs with no serial imaging



Fig. 6. Brain regions where there is an association between cortical thickness and MCI-to-AD conversion, as determined by the proposed method and based on analyzing the entire longitudinal ADNI MCI sample. Results are overlaid on an inflated representation of the cortex of an average human brain (fsaverage). The gray pattern reflects the underlying cortical folding. Colored regions represent significant association (uncorrected p < 0.01). The color reflects the significance of the association, which was computed as $-\log_10(p-value)$ (see color bar). The sign encodes the direction of association, where a positive value suggests that thinner cortex at that particular location is associated with an earlier date of onset of AD dementia. Top and bottom rows show views of the lateral and medial surfaces, respectively. Left and right panels correspond to the left and right hemispheres, respectively.

follow-up. However, to our knowledge, we offer the first discussion of this issue in neuroimaging and propose a strategy that might boost statistical efficiency for detecting associations of interest.

Finally, we can quantify the effect of the LME-based first step in the proposed approach. The LME model attempts to capture the spatiotemporal structure in the neuroimaging data by examining the entire longitudinal sample. In contrast, a simpler strategy would be to estimate the temporal trajectory of each imaging measurement in isolation, without considering other subjects or spatial locations in a mass-univariate analysis. From a theoretical standpoint, we expect the proposed method to yield a more accurate model for the longitudinal data, and hence improve the power of the Cox regression, which relies on the estimates of the imaging data at time-points with no observation. Empirically, we find this is indeed the case. Our experiments suggest that the improvement in statistical power and reliability due to the LME step is subtle, yet consistent. In the mass-univariate setting, the agreement between the maps of two independent samples can increase by over 5% as a result of the LME step.

The extended Cox analysis we advocate in this article has several drawbacks and technical subtleties, as discussed in prior work (Fisher and Lin, 1999). Firstly, the model parameters have no straightforward interpretation as in the classical Cox model, where under the proportional hazards assumption, the estimated coefficients can be interpreted as a constant hazard ratio. Therefore, we are largely constrained to testing the statistical strength of associations, rather than presenting interpretable hazard ratios. Secondly, we need to underscore the difference between internal versus external time-dependent variables, as distinguished by Kalbfleisch and Prentice (2011). Internal variables are those that are generated by the studied subject directly (e.g., blood pressure) and are directly related to the event (i.e., the event is defined via this variable or these variables cannot be measured after the event e.g., blood pressure after death). When dealing with such variables and events, the relationship between the conditional hazard function and survival function breaks down. For example, a measurable value of blood pressure is indicative of the subject being still alive and hence survival at that time point is known. However, in the scenarios we considered in this article, longitudinal imaging measurements, which might be considered as internal, do not generally suffer from this technical problem. This is because scans can be and are acquired after the event of interest and the event is not directly defined via imaging measurements. Another issue with time-dependent variables is that usually they do not yield individual predictions of event risk curves. This is because these curves depend on the typically unknown temporal trajectories of the time-dependent variables. Finally, our analysis interrogated the relationship between imaging measurements and concurrent risk of event. Alternative functional forms for this relationship can also be used within the extended Cox analysis. For instance, one could consider the cumulative history of the variable (e.g., history of high blood pressure). One alternative promising approach in neuroimaging might be to consider the concurrent or past slope (rate of change) in imaging measurements as a potential biomarker of the event. This would require reliable estimates of derivatives of the longitudinal imaging trajectory, which could also be obtained from the LME step. We leave the exploration of this direction to future work.

Conclusions

We presented a statistical method for the event time analysis of longitudinal neuroimage data. We have implemented and validated the proposed method for mapping longitudinal brain atrophy effects within the FreeSurfer framework, yet its adaptation to other types of spatial data is straightforward. Our results suggest that the proposed method can offer excellent statistical power for detecting associations between longitudinal imaging measurements and a clinical event, such as disease onset.

Acknowledgments

Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI; Principal Investigator: Michael Weiner; NIH grant U01 AG024904). The ADNI is funded by the National Institute on Aging (NIA), the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: Abbott Laboratories, AstraZeneca AB, Bayer Schering Pharma AG, Bristol-Myers Squibb, Eisai Global Clinical Development, Elan Corporation Plc, Genentech Inc., GE Healthcare, GlaxoSmithKline, Innogenetics, Johnson and Johnson Services Inc., Eli Lilly and Company, Medpace Inc., Merck and Co. Inc., Novartis International AG, Pfizer Inc., F. Hoffmann-La Roche Ltd., Schering-Plough Corporation, CCBR-SYNARC Inc., and Wyeth Pharmaceuticals, as well as nonprofit partners the Alzheimer's Association and Alzheimer's Drug Discovery Foundation, with participation from the US Food and Drug Administration. Private sector contributions to the ADNI are facilitated by the Foundation for the NIH. The grantee organization is the Northern California Institute for Research and Education Inc., and the study is coordinated by the Alzheimer's Disease Cooperative Study at the University of California, San Diego. The ADNI data are disseminated by the Laboratory for NeuroImaging at the University of Southern California

This research was carried out in whole or in part at the Athinoula A. Martinos Center for Biomedical Imaging at the Massachusetts General Hospital, using resources provided by the Center for Functional Neuroimaging Technologies, P41EB015896, a P41 Biotechnology Resource Grant supported by the National Institute of Biomedical Imaging and Bioengineering (NIBIB), National Institutes of Health. Further support for this research was provided in part by the National Center for Research Resources (P41-RR14075), the National Institute for Biomedical Imaging and Bioengineering (R01EB006758), the National Institute on Aging (AG022381), the National Center for Alternative Medicine (RC1 AT005728-01), the National Institute for Neurological Disorders and Stroke (R01 NS052585-01, 1R21NS072652-01, 1R01NS070963, 2R01NS042861-06A1, 5P01NS058793-03), the National Institute of Child Health and Human Development (R01-HD071664), and was made possible by the resources provided by Shared Instrumentation Grants 1S10RR023401, 1S10RR019307, and 1S10RR023043. Additional support was provided by The Autism & Dyslexia Project funded by the Ellison Medical Foundation, by the NIH 5R01AG008122-22 grant, and by the NIH Blueprint for Neuroscience Research (5U01-MH093765), part of the multi-institutional Human Connectome Project. Dr. Sabuncu received support from an AHAF (BrightFocus) Alzheimer's Disease pilot grant (AHAF A2012333), and an NIH K25 grant (NIBIB 1K25EB013649-01).

Finally, the authors would like to thank Nick Schmansky and Louis Vinke for their efforts in downloading and processing the ADNI MRI scans.

References

- Albin, R.L., Young, A.B., Penney, J.B., Handelin, B., Balfour, R., Anderson, K.D., Markel, D.S., Tourtellotte, W.W., Reiner, A., 1990. Abnormalities of striatal projection neurons and N-methyl-D-aspartate receptors in presymptomatic Huntington's disease. N. Engl. J. Med. 322, 1293–1298.
- Benjamini, Y., Krieger, A.M., Yekutieli, D., 2006. Adaptive linear step-up procedures that control the false discovery rate. Biometrika 93, 491–507.
- Bernal-Rusiel, J.L., Greve, D.N., Reuter, M., Fischl, B., Sabuncu, M.R., 2013a. Statistical analysis of longitudinal neuroimage data with linear mixed effects models. Neuroimage 66, 249–260.
- Bernal-Rusiel, J.L., Reuter, M., Greve, D.N., Fischl, B., Sabuncu, M.R., 2013b. Spatiotemporal linear mixed effects modeling for the mass-univariate analysis of longitudinal neuroimage data. Neuroimage 81.
- Borgwardt, S., McGuire, P., Fusar-Poli, P., 2011. Gray matters: mapping the transition to psychosis. Schizophr. Res. 133, 63–67.
- Buckner, R.L., Snyder, A.Z., Shannon, B.J., LaRossa, G., Sachs, R., Fotenos, A.F., Sheline, Y.I., Klunk, W.E., Mathis, C.A., Morris, J.C., 2005. Molecular, structural, and functional characterization of Alzheimer's disease: evidence for a relationship between default activity, amyloid, and memory. J. Neurosci. 25, 7709–7717.

- Chetelat, G., Landeau, B., Eustache, F., Mezenge, F., Viader, F., de La Sayette, V., Desgranges, B., Baron, J.C., 2005. Using voxel-based morphometry to map the structural changes associated with rapid conversion in MCI: a longitudinal MRI study. Neuroimage 27, 934–946.
- Cox, D.R., 1972. Regression models and life-tables. J. R. Stat. Soc. Ser. B Methodol. 187–220.
- Cox, D.R., 1975. Partial likelihood. Biometrika 62, 269–276.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis* 1: I. Segmentation and surface reconstruction. Neuroimage 9, 179–194.
- Desikan, R.S., Cabral, H.J., Fischl, B., Guttmann, C.R., Blacker, D., Hyman, B.T., Albert, M.S., Killiany, R.J., 2009. Temporoparietal MR imaging measures of atrophy in subjects with mild cognitive impairment that predict subsequent diagnosis of Alzheimer disease. Am. J. Neuroradiol. 30, 532–538.
- Desikan, R.S., Cabral, H.J., Settecase, F., Hess, C.P., Dillon, W.P., Glastonbury, C.M., Weiner, M.W., Schmansky, N.J., Salat, D.H., Fischl, B., 2010. Automated MRI measures predict progression to Alzheimer's disease. Neurobiol. Aging 31, 1364–1374.
- Devanand, D., Pradhaban, G., Liu, X., Khandji, A., De Santi, S., Segal, S., Rusinek, H., Pelton, G., Honig, L., Mayeux, R., 2007. Hippocampal and entorhinal atrophy in mild cognitive impairment: prediction of Alzheimer disease. Neurology 68, 828–836.
- Dickerson, B.C., Goncharova, I., Sullivan, M., Forchetti, C., Wilson, R., Bennett, D., Beckett, L., deToledo-Morrell, L., 2001. MRI-derived entorhinal and hippocampal atrophy in incipient and very mild Alzheimer's disease. Neurobiol. Aging 22, 747–754.
- Dickerson, B.C., Bakkour, A., Salat, D.H., Feczko, E., Pacheco, J., Greve, D.N., Grodstein, F., Wright, C.I., Blacker, D., Rosas, H.D., 2009. The cortical signature of Alzheimer's disease: regionally specific cortical thinning relates to symptom severity in very mild to mild AD dementia and is detectable in asymptomatic amyloid-positive individuals. Cereb. Cortex 19, 497–510.
- Fischl, B., 2012. FreeSurfer. Neuroimage 62 (2), 774–781.
- Fischl, B., Dale, A.M., 2000. Measuring the thickness of the human cerebral cortex from magnetic resonance images. Proc. Natl. Acad. Sci. 97, 11050.
- Fischl, B., Sereno, M.I., Dale, A.M., 1999a. Cortical surface-based analysis* 1: II: inflation, flattening, and a surface-based coordinate system. Neuroimage 9, 195–207.
- Fischl, B., Sereno, M.I., Tootell, R.B.H., Dale, A.M., 1999b. High-resolution intersubject averaging and a coordinate system for the cortical surface. Hum. Brain Mapp. 8, 272–284.
- Fisher, L.D., Lin, D., 1999. Time-dependent covariates in the Cox proportional-hazards regression model. Annu. Rev. Public Health 20, 145–157.
- Fitzmaurice, G.M., Laird, N.M., Ware, J.H., 2011. Applied Longitudinal Analysis. Wiley.
- Forsberg, A., Engler, H., Almkvist, O., Blomquist, G., Hagman, G., Wall, A., Ringheim, A., Langstrom, B., Nordberg, A., 2008. PET imaging of amyloid deposition in patients with mild cognitive impairment. Neurobiol. Aging 29, 1456–1465.
- Gauthier, S., Reisberg, B., Zaudig, M., Petersen, R.C., Ritchie, K., Broich, K., Belleville, S., Brodaty, H., Bennett, D., Chertkow, H., 2006. Mild cognitive impairment. Lancet 367, 1262–1270.
- Geerlings, M., Den Heijer, T., Koudstaal, P., Hofman, A., Breteler, M., 2008. History of depression, depressive symptoms, and medial temporal lobe atrophy and the risk of Alzheimer disease. Neurology 70, 1258–1264.
- Good, P.I., 2000. Permutation Tests. Wiley Online Library.
- Hagler Jr., D.J., Saygin, A.P., Sereno, M.I., 2006. Smoothing and cluster thresholding for cortical surface-based group analysis of fMRI data. Neuroimage 33, 1093–1103.
- Han, X., Jovicich, J., Salat, D., van der Kouwe, A., Quinn, B., Czanner, S., Busa, E., Pacheco, J., Albert, M., Killiany, R., 2006. Reliability of MRI-derived measurements of human cerebral cortical thickness: the effects of field strength, scanner upgrade and manufacturer. Neuroimage 32, 180–194.
- Jack Jr., C.R., Petersen, R.C., Xu, Y.C., Waring, S.C., O'Brien, P.C., Tangalos, E.G., Smith, G.E., Ivnik, R.J., Kokmen, E., 1997. Medial temporal atrophy on MRI in normal aging and very mild Alzheimer's disease. Neurology 49, 786–794.
- Jack Jr., C.R., Petersen, R.C., Grundman, M., Jin, S., Gamst, A., Ward, C.P., Sencakova, D., Doody, R.S., Thal, L.J., 2008a. Longitudinal MRI findings from the vitamin E and donepezil treatment study for MCI. Neurobiol. Aging 29, 1285–1295.
- Jack Jr., C.R., Weigand, S.D., Shiung, M.M., Przybelski, S.A., O'Brien, P.C., Gunter, J.L., Knopman, D.S., Boeve, B.F., Smith, G.E., Petersen, R.C., 2008b. Atrophy rates accelerate in amnestic mild cognitive impairment. Neurology 70, 1740–1752.

- Jack Jr., C.R., Vemuri, P., Wiste, H.J., Weigand, S.D., Lesnick, T.G., Lowe, V., Kantarci, K., Bernstein, M.A., Senjem, M.L., Gunter, J.L., 2012. Shapes of the trajectories of 5 major biomarkers of Alzheimer disease. Arch. Neurol. 69 (7), 856–867. http://dx.doi.org/ 10.1001/archneurol.2011.3405 (archneurol. 2011.3405 v2011).
- Julkunen, V., Niskanen, E., Muehlboeck, S., Pihlajamaki, M., Kononen, M., Hallikainen, M., Kivipelto, M., Tervo, S., Vanninen, R., Evans, A., 2009. Cortical thickness analysis to detect progressive mild cognitive impairment: a reference to Alzheimer's disease. Dement. Geriatr. Cogn. Disord. 28, 404–412.
- Kalbfleisch, J.D., Prentice, R.L., 2011. The Statistical Analysis of Failure Time Data. John Wiley & Sons.
- Kleinbaum, D.G., Klein, M., 2012. Survival Analysis. Springer.
- Lerch, J.P., Pruessner, J.C., Zijdenbos, A., Hampel, H., Teipel, S.J., Evans, A.C., 2005. Focal decline of cortical thickness in Alzheimer's disease identified by computational neuroanatomy. Cereb. Cortex 15, 995–1001.
- Marcus, K.J., Astrakas, L.G., Zurakowski, D., Zarifi, M.K., Mintzopoulos, D., Poussaint, T.Y., Anthony, D.C., De Girolami, U., Black, P.M., Tarbell, N.J., 2007. Predicting survival of children with CNS tumors using proton magnetic resonance spectroscopic imaging biomarkers. Int. J. Oncol. 30, 651.
- Morgan, V., Riches, S., Thomas, K., Vanas, N., Parker, C., Giles, S., Desouza, N., 2011. Diffusion-weighted magnetic resonance imaging for monitoring prostate cancer progression in patients managed by active surveillance. Br. J. Radiol. 84, 31–37.
- Nichols, T.E., Holmes, A.P., 2002. Nonparametric permutation tests for functional neuroimaging: a primer with examples. Hum. Brain Mapp. 15, 1–25.
- Reuter, M., Fischl, B., 2011. Avoiding asymmetry-induced bias in longitudinal image processing. Neuroimage 57, 19–21.
- Reuter, M., Rosas, H.D., Fischl, B., 2010. Highly accurate inverse consistent registration: a robust approach. Neuroimage 53, 1181–1196.
- Reuter, M., Schmansky, N.J., Rosas, H.D., Fischl, B., 2012. Within-subject template estimation for unbiased longitudinal image analysis. Neuroimage. http://dx.doi.org/10. 1016/j. neuroimage .2012.02.084.
- Risacher, S.L., Saykin, A.J., West, J.D., Shen, L., Firpi, H.A., McDonald, B.C., 2009. Baseline MRI predictors of conversion from MCI to probable AD in the ADNI cohort. Curr. Alzheimer Res. 6, 347.
- Sabuncu, M.R., 2013. A Bayesian algorithm for image-based time-to-event Prediction. MLMI Workshop, MICCAI 2013.
- Sabuncu, M.R., Desikan, R.S., Sepulcre, J., Yeo, B.T.T., Liu, H., Schmansky, N.J., Reuter, M., Weiner, M.W., Buckner, R.L., Sperling, R.A., 2011. The dynamics of cortical and hippocampal atrophy in Alzheimer disease. Arch. Neurol. 68, 1040.
- Sabuncu, M.R., Buckner, R.L., Smoller, J.W., Lee, P.H., Fischl, B., Sperling, R.A., 2012. The association between a polygenic Alzheimer score and cortical thickness in clinically normal subjects. Cereb. Cortex 22, 2653–2661.
- Stoub, T., Bulgakova, M., Leurgans, S., Bennett, D., Fleischman, D., Turner, D., 2005. MRI predictors of risk of incident Alzheimer disease: a longitudinal study. Neurology 64, 1520–1524.
- Sun, D., Phillips, L., Velakoulis, D., Yung, A., McGorry, P.D., Wood, S.J., van Erp, T.G., Thompson, P.M., Toga, A.W., Cannon, T.D., 2009. Progressive brain structural changes mapped as psychosis develops in at risk individuals. Schizophr. Res. 108, 85–92.
- Thirion, B., Pinel, P., Mériaux, S., Roche, A., Dehaene, S., Poline, J.B., 2007. Analysis of a large fMRI cohort: statistical and methodological issues for group analyses. Neuroimage 35, 105–120.
- Tintore, M., Rovira, A., Rio, J., Tur, C., Pelayo, R., Nos, C., Tellez, N., Perkal, H., Comabella, M., Sastre-Garriga, J., 2008. Do oligoclonal bands add information to MRI in first attacks of multiple sclerosis? Neurology 70, 1079–1083.
- Tsiatis, A.A., Davidian, M., 2004. Joint modeling of longitudinal and time-to-event data: an overview. Stat. Sin. 14, 809–834.
- Vemuri, P., Weigand, S.D., Knopman, D.S., Kantarci, K., Boeve, B.F., Petersen, R.C., Jack Jr., C. R., 2011. Time-to-event voxel-based techniques to assess regional atrophy associated with MCI risk of progression to AD. Neuroimage 54, 985–991.
- Verbeke, G., Molenberghs, G., 2000. Linear Mixed Models for Longitudinal Data. Springer, N.Y.
- Whitfield-Gabrieli, S., Thermenos, H.W., Milanovic, S., Tsuang, M.T., Faraone, S.V., McCarley, R.W., Shenton, M.E., Green, A.I., Nieto-Castanon, A., LaViolette, P., 2009. Hyperactivity and hyperconnectivity of the default network in schizophrenia and in first-degree relatives of persons with schizophrenia. Proc. Natl. Acad. Sci. 106, 1279–1284.