



A Longitudinal Method for Simultaneous Whole-Brain and Lesion Segmentation in Multiple Sclerosis

Cerri, Stefano; Hoopes, Andrew; Greve, Douglas N. ; Mühlau, Mark; Van Leemput, Koen

Published in:

Machine Learning in Clinical Neuroimaging and Radiogenomics in Neuro-oncology

Link to article, DOI:

[10.1007/978-3-030-66843-3_12](https://doi.org/10.1007/978-3-030-66843-3_12)

Publication date:

2020

Document Version

Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):

Cerri, S., Hoopes, A., Greve, D. N., Mühlau, M., & Van Leemput, K. (2020). A Longitudinal Method for Simultaneous Whole-Brain and Lesion Segmentation in Multiple Sclerosis. In Machine Learning in Clinical Neuroimaging and Radiogenomics in Neuro-oncology: Third International Workshop, MLCN 2020, and Second International Workshop, RNO-AI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4–8, 2020, Proceedings (Vol. 12449, pp. 119-128). Springer. https://doi.org/10.1007/978-3-030-66843-3_12

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A Longitudinal Method for Simultaneous Whole-Brain and Lesion Segmentation in Multiple Sclerosis

Stefano Cerri¹, Andrew Hoopes², Douglas N. Greve^{2,3}, Mark Mühlau⁴, and Koen Van Leemput^{1,2}

¹ Department of Health Technology, Technical University of Denmark, Denmark

² Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, USA

³ Department of Radiology, Harvard Medical School, United States

⁴ Department of Neurology and TUM-Neuroimaging Center, School of Medicine, Technical University of Munich, Germany

Abstract. In this paper we propose a novel method for the segmentation of longitudinal brain MRI scans of patients suffering from Multiple Sclerosis. The method builds upon an existing cross-sectional method for simultaneous whole-brain and lesion segmentation, introducing subject-specific latent variables to encourage temporal consistency between longitudinal scans. It is very generally applicable, as it does not make any prior assumptions on the scanner, the MRI protocol, or the number and timing of longitudinal follow-up scans. Preliminary experiments on three longitudinal datasets indicate that the proposed method produces more reliable segmentations and detects disease effects better than the cross-sectional method it is based upon.

1 Introduction

Multiple Sclerosis (MS) is an inflammatory autoimmune disorder of the central nervous system. It is characterized by the formation of lesions in the white matter, as well as marked brain atrophy primarily in deep gray matter structures [1,2]. The increased availability of longitudinal magnetic resonance imaging (MRI) scans opens up the prospect of tracking lesion evolution and atrophy trajectories over time, enabling a better assessment of disease progression and treatment efficacy [3].

Despite high potential clinical impact, work on computational methods for quantifying longitudinal changes in MS has remained fairly limited to date (cf. [4] for an overview). The methods that do exist suffer from one or more of the following limitations: They only assess changes in white matter lesions [5,6,7,8] or in aggregate measures of brain atrophy such as global brain or gray matter volume [9,10], but not in individual brain structures; they can only compare between two consecutive time points [11,12,13,14] instead of characterizing entire temporal trajectories; or they are developed and tested in very specific settings

only, with degraded performance when applied to data from different scanners and acquisition protocols [15] which limits their usefulness in practice.

In order to address these limitations, here we propose a dedicated model for simultaneously segmenting anatomical brain structures and white matter lesion from longitudinal multi-contrast MRI scans. The proposed method builds upon a contrast-adaptive method for simultaneous whole-brain and lesion segmentation that we previously developed and validated [16]. Here we extend this approach to the longitudinal setting by additionally modeling the expected temporal consistency between repeated scans of the same subject, using latent variables that introduce a statistical dependency between the time points. By segmenting both white matter lesions and anatomical brain structures across time, the resulting method enables tracking deep gray matter atrophy trajectories and lesion evolution simultaneously. The model is fully adaptive to different MRI contrasts and scanners, and does not put any constraints on the number or the timing of longitudinal follow-up scans. To the best of our knowledge, no other method with these capabilities currently exists.

We assessed the segmentation performance of the proposed method on three longitudinal datasets. Preliminary results indicate that it produces more reliable segmentations and detects disease effects better than the cross-sectional method. An example result produced by the longitudinal method is shown in Fig.1.

2 Existing cross-sectional method

We first summarize the existing cross-sectional method for simultaneous whole-brain and lesion segmentation [16] the proposed method builds upon.

Let $\mathbf{D} = (\mathbf{d}_1, \dots, \mathbf{d}_I)$ be the image intensities of a multi contrast MRI scan with I voxels, where the vector $\mathbf{d}_i = (d_i^1, \dots, d_i^N)^T$ represents the log-transformed image intensity of voxel i for all the available N contrasts. Moreover, let $\mathbf{l} = (l_1, \dots, l_I)^T$ be corresponding segmentation labels, where $l_i \in \{1, \dots, K\}$ denotes one of the K possible anatomical structures assigned to voxel i . In order for the model to be capable of segmenting white matter lesions, a binary lesion map $\mathbf{z} = (z_1, \dots, z_I)$ is introduced, where $z_i \in \{0, 1\}$ indicates the presence of lesion in voxel i . We use a generative model, illustrated in black in Fig. 2, to estimate a joint segmentation $\{\mathbf{l}, \mathbf{z}\}$ from MRI data \mathbf{D} . The model consists of a segmentation prior $p(\mathbf{l}, \mathbf{z} | \mathbf{h}, \mathbf{x})$ with parameters \mathbf{h} and \mathbf{x} that encode shape information, and a likelihood function $p(\mathbf{D} | \mathbf{l}, \mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\theta}_z)$ with parameters $\boldsymbol{\theta}$ and $\boldsymbol{\theta}_z$ that govern intensity appearance. Below we briefly describe the segmentation prior and the likelihood function, as well as how the model is “inverted” to obtain automatic segmentations.

Segmentation prior: The segmentation prior is composed of two components $p(\mathbf{l} | \mathbf{x})$ and $p(\mathbf{z} | \mathbf{h}, \mathbf{x})$ that encode spatial information regarding the neuroanatomical labels \mathbf{l} and the lesion map \mathbf{z} respectively: $p(\mathbf{l}, \mathbf{z} | \mathbf{h}, \mathbf{x}) = p(\mathbf{l} | \mathbf{x})p(\mathbf{z} | \mathbf{h}, \mathbf{x})$. The first component is a deformable probabilistic atlas, encoded as a tetrahedral mesh [17] with node positions \mathbf{x} and with a deformation prior distribution




Fig.1: Example segmentation produced by the proposed method on a longitudinal scan with T1w and FLAIR contrast.

defined as:

$$p(\mathbf{x}) \propto \exp \left[-K \sum_d U_d(\mathbf{x}, \mathbf{x}_{ref}) \right].$$

Here K controls the stiffness of the mesh deformations, d loops over the tetrahedra in the mesh, and $U_d(\mathbf{x}, \mathbf{x}_{ref})$ is a cost [18] associated with deforming the d^{th} tetrahedron from its shape in the atlas's reference position \mathbf{x}_{ref} . Letting $p(l_i = k|\mathbf{x})$ denote the probability of observing label k at voxel i for a given deformation, assuming conditional independence of the labels between voxels yields $p(\mathbf{l}|\mathbf{x}) = \prod_{i=1}^I p(l_i|\mathbf{x})$.

The second component of the segmentation prior is a model of the form: $p(\mathbf{z}|\mathbf{h}, \mathbf{x}) = \prod_{i=1}^I p(z_i|\mathbf{h}, \mathbf{x})$, $p(\mathbf{h}) = \mathcal{N}(\mathbf{h}|\mathbf{0}, \mathbf{I})$, where $p(z_i = 1|\mathbf{h}, \mathbf{x})$ is the probability that voxel i is part of a lesion. This model takes into account both a voxel's spatial location within its neuroanatomical context (through \mathbf{x}), as well as lesion shape constraints through a variational autoencoder (VAE) [19] that "decodes" a low-dimensional latent code \mathbf{h} using a convolutional neural network.

Likelihood: For the likelihood, which links segmentations $\{\mathbf{l}, \mathbf{z}\}$ to intensities \mathbf{D} , we use a multivariate Gaussian intensity model for each structure, and model the MRI bias field artifact as a linear combination of spatially smooth basis functions that is added to the local voxel intensities [20,21]. Letting $\boldsymbol{\theta}_z = \{\boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z\}$ denote the mean and variance of lesion intensities, and $\boldsymbol{\theta}$ the collection of bias field parameters and intensity means and variances $\{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}$ of all K anatomical structures, the likelihood is defined as $p(\mathbf{D}|\mathbf{l}, \mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\theta}_z) = \prod_{i=1}^I p(\mathbf{d}_i|l_i, z_i, \boldsymbol{\theta}, \boldsymbol{\theta}_z)$, where

$$p(\mathbf{d}_i|l_i = k, z_i, \boldsymbol{\theta}, \boldsymbol{\theta}_z) = \begin{cases} \mathcal{N}(\mathbf{d}_i|\boldsymbol{\mu}_z + \mathbf{C}^T \boldsymbol{\phi}_i, \boldsymbol{\Sigma}_z) & \text{if } z_i = 1, \\ \mathcal{N}(\mathbf{d}_i|\boldsymbol{\mu}_k + \mathbf{C}^T \boldsymbol{\phi}_i, \boldsymbol{\Sigma}_k) & \text{otherwise.} \end{cases}$$

Here $\boldsymbol{\phi}_i$ evaluates the bias field basis functions at the i^{th} voxel, and $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_N)$, where \mathbf{c}_n denotes the parameters of the bias field model for the

n^{th} contrast. The model is completed by a flat prior on θ , and a weak conditional prior $p(\theta_z|\theta)$ that ensures that the method can be robustly applied to scans with no or very small lesion loads [16].

Segmentation: Given an MRI scan \mathbf{D} , segmentation proceeds by approximating the segmentation posterior using point estimates of the parameters \mathbf{x} and θ :

$$p(\mathbf{l}, \mathbf{z}|\mathbf{D}) \simeq p(\mathbf{l}, \mathbf{z}|\mathbf{D}, \hat{\theta}, \hat{\mathbf{x}}), \quad (1)$$

and Markov chain Monte Carlo sampling to marginalize over the remaining, lesion-specific parameters θ_z and \mathbf{h} . For the purpose of finding the point estimates $\hat{\mathbf{x}}$ and $\hat{\theta}$, a simplified model is fitted to the data:

$$\hat{\Omega} = \operatorname{argmax}_{\Omega} p(\Omega|\mathbf{D}) \quad \text{with} \quad \Omega = \{\mathbf{x}, \theta, \theta_z\}, \quad (2)$$

where the lesion-shape encoding VAE and its parameters \mathbf{h} are temporarily removed to simplify the optimization process. More details can be found in [16].

3 Longitudinal extension

In the longitudinal setting we are given T scans with image intensities $\{\mathbf{D}_t\}_{t=1}^T$, and we wish to compute for each time point t the corresponding segmentation $\{\mathbf{l}_t, \mathbf{z}_t\}$. In contrast to the cross-sectional setting where each image is treated independently, here we can exploit the fact that all images belong to the same subject to produce more consistent (and potentially more accurate) segmentations. Towards this end, we introduce subject-specific latent variables \mathbf{x}_0 and θ_0 in the segmentation prior and likelihood, respectively, imposing a statistical dependency between the time points that encourages the segmentations to be similar to one another. The augmented generative model is depicted in Fig. 2, where the parameters $\mathbf{x}_t, \mathbf{h}_t, \theta_t$ and $\theta_{t,z}$ denote the model parameters of time point t , and the blue parts indicate the additional components compared to the cross-sectional model.

Segmentation prior: In order to obtain temporal consistency in the segmentation prior, we use the concept of a “subject-specific atlas” [22]: a deformation of the cross-sectional atlas to represent the average subject-specific anatomy across all time points. In particular,

$$p(\{\mathbf{x}_t\}_{t=1}^T|\mathbf{x}_0) = \prod_{t=1}^T p(\mathbf{x}_t|\mathbf{x}_0), \quad p(\mathbf{x}_t|\mathbf{x}_0) \propto \exp \left[-K_1 \sum_d U_d(\mathbf{x}_t, \mathbf{x}_0) \right],$$

where \mathbf{x}_0 are latent atlas node positions encoding subject-specific brain shape, with prior $p(\mathbf{x}_0) \propto \exp[-K_0 \sum_d U_d(\mathbf{x}_0, \mathbf{x}_{ref})]$. Here the mesh stiffnesses K_0 and K_1 are hyperparameters of the model; by choosing $K_0 = \infty$ and $K_1 = K$ the model devolves into the cross-sectional segmentation prior for each time point separately.

Likelihood: In a similar vein, we also introduce subject-specific latent variables to encourage temporal consistency in the Gaussian intensity models. For each anatomical structure, we condition the Gaussian parameters $\{\boldsymbol{\mu}_{t,k}, \boldsymbol{\Sigma}_{t,k}\}$ on latent variables $\{\boldsymbol{\mu}_{0,k}, \boldsymbol{\Sigma}_{0,k}\}$ using a normal-inverse-Wishart (NIW) distribution: $p(\{\boldsymbol{\theta}_t\}_{t=1}^T | \boldsymbol{\theta}_0) = \prod_{t=1}^T p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_0)$ with

$$p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_0) \propto \prod_{k=1}^K \mathcal{N}(\boldsymbol{\mu}_{t,k} | \boldsymbol{\mu}_{0,k}, P_{0,k} \boldsymbol{\Sigma}_{0,k}) \text{IW}(\boldsymbol{\Sigma}_{t,k} | P_{0,k} \boldsymbol{\Sigma}_{0,k}, P_{0,k} - N - 2).$$

Here $\boldsymbol{\theta}_0 = \{\boldsymbol{\mu}_{0,k}, \boldsymbol{\Sigma}_{0,k}\}_{k=1}^K$ with prior $p(\boldsymbol{\theta}_0) \propto 1$, and $P_{0,k} \geq 0$ is a hyperparameter that governs the strength of the regularization across time for label k . Note that choosing $P_{0,k} = 0, \forall k$ yields the cross-sectional likelihood for each time point independently.

Segmentation: We follow the same overall segmentation strategy as in the cross-sectional setting: we first compute point estimates $\{\hat{\boldsymbol{\theta}}_t, \hat{\mathbf{x}}_t\}_{t=1}^T$ using a simplified model in which the lesion shape codes $\{\mathbf{h}_t\}_{t=1}^T$ are removed, and subsequently obtain segmentations as described in the cross-sectional setting, i.e., by using (1) for each time point separately. As in the cross-sectional case, we obtain the required point estimates by fitting the longitudinal model to the data:

$$\{\hat{\boldsymbol{\Omega}}_1, \dots, \hat{\boldsymbol{\Omega}}_T, \hat{\boldsymbol{\theta}}_0, \hat{\mathbf{x}}_0\} = \underset{\{\boldsymbol{\Omega}_1, \dots, \boldsymbol{\Omega}_T, \boldsymbol{\theta}_0, \mathbf{x}_0\}}{\operatorname{argmax}} p(\boldsymbol{\Omega}_1, \dots, \boldsymbol{\Omega}_T, \boldsymbol{\theta}_0, \mathbf{x}_0 | \mathbf{D}_1, \dots, \mathbf{D}_T), \quad (3)$$

where $\boldsymbol{\Omega}_t = \{\mathbf{x}_t, \boldsymbol{\theta}_t, \boldsymbol{\theta}_{t,z}\}$. For optimizing (3) we use coordinate ascent, updating one variable at a time in an iterative fashion. Because $p(\mathbf{x}_t | \mathbf{x}_0)$ is of the same form as the cross-sectional deformation prior, and the NIW distribution used in $p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_0)$ is the conjugate prior for the mean and variance of a Gaussian distribution, estimating $\boldsymbol{\Omega}_t$ from \mathbf{D}_t for given values of $\boldsymbol{\theta}_0$ and \mathbf{x}_0 simply involves performing an optimization of the form of (2) for each time point t separately. Conversely, for given values $\{\boldsymbol{\Omega}_t\}_{t=1}^T$ the update for $\boldsymbol{\theta}_0$ is given in closed form:

$$\boldsymbol{\mu}_{0,k} \leftarrow \left(\sum_{t=1}^T \boldsymbol{\Sigma}_{t,k}^{-1} \right)^{-1} \sum_{t=1}^T \boldsymbol{\Sigma}_{t,k}^{-1} \boldsymbol{\mu}_{t,k}, \quad \boldsymbol{\Sigma}_{0,k}^{-1} \leftarrow \left(\frac{1}{T} \sum_{t=1}^T \boldsymbol{\Sigma}_{t,k}^{-1} \right) \frac{P_{0,k}}{P_{0,k} - N - 2},$$

whereas updating \mathbf{x}_0 involves the optimization (cf. [22])

$$\underset{\mathbf{x}_0}{\operatorname{argmin}} \sum_d \left[K_0 U_d(\mathbf{x}_0, \mathbf{x}_{ref}) + K_1 \sum_{t=1}^T U_d(\mathbf{x}_0, \mathbf{x}_t) \right],$$

which we solve numerically using a limited-memory BFGS algorithm.

Implementation: In order to avoid longitudinal processing biases resulting from not treated all time points in exactly the same way, we first compute an unbiased within-subject template using an inverse consistent registration

method [23]. This template is a robust representation of the average subject anatomy over time, and we use it as an unbiased reference to register all time points to in a preprocessing step. We also use it to start the proposed iterative algorithm optimizing (3): we apply the cross-sectional method to the template, and use the estimated model parameters $\boldsymbol{\Omega}$ to initialize $\boldsymbol{\Omega}_t, t = 1, \dots, T$. The proposed algorithm, which interleaves updating the latent variables $\boldsymbol{\theta}_0$ and \mathbf{x}_0 with updating the parameters $\{\boldsymbol{\Omega}_t\}_{t=1}^T$, is then run for five iterations, which we have found to be sufficient to reach convergence.

Based on initial pilot experiments on scans from the ADNI project⁵ (distinct from the ones used in the experiments below), we set the method’s hyperparameter values to $K_1 = 14K$ and $K_0 = K$, where K is the mesh stiffness in the existing cross-sectional method, and $P_{0,k}$ to the number of voxels assigned to class k in the segmentation of the within-subject template.

Our implementation builds upon the C++ and Python code of [24,16], and is publicly available from FreeSurfer⁶. Segmenting one subject takes approximately 15 minutes per time point on an Intel 12-core i7-8700K processor with a GeForce GTX 1060 graphics card.

4 Experiments and Results

In order to assess whether introducing subject-specific latent variables leads to better longitudinal performance, we compared the proposed method and the cross-sectional method on three different datasets:

- **Test-retest** [25]: This dataset consists of longitudinal T1-weighted (T1w) and FLuid Attenuation Inversion Recovery (FLAIR) scans of 2 MS subjects. For each subject 6 repeated scans were acquired from 3 different 3T scanners (Philips Achieva; Siemens Verio; GE Signa MR750) within 3 weeks.
- **Achieva**: This dataset consists of longitudinal T1w and FLAIR scans of 86 MS subjects. The subjects were scanned between 3 and 6 times (time between scans between 6 and 12 months) with a 3T Philips Achieva scanner at the Department of Neurology, School of Medicine, at the Technical University of Munich in the context of the in-house cohort study on MS named TUM-MS. All the subjects were diagnosed as relapsing-remitting MS.
- **ADNI**: This dataset consists of longitudinal T1w scans of 135 subjects randomly selected from the ADNI project. Scanners from multiple sites were used to acquired the scans, and subjects were scanned between 2 and 6 times, with 6 or 12 months between scans. The subjects were divided into 3 groups: cognitively normal (CN, n=45), mild cognitive impairment (MCI, n=54), and Alzheimer disease (AD, n=36).

We report results on the estimated volumes of the following 26 regions: left and right cerebral white matter, cerebellum white matter, cerebral cortex, cerebellum

⁵ <http://adni.loni.usc.edu/>

⁶ <http://freesurfer.net/>

cortex, lateral ventricle, hippocampus, thalamus, putamen, pallidum, caudate, amygdala and nucleus accumbens, as well as brain stem and lesions. To avoid cluttering, results for left and right structures are averaged.

Temporal consistency: We wished to assess whether the proposed method is able to reduce non-biological variations in longitudinal volume measurements, both within the short (< 3 weeks) and longer (< 6 years) time intervals of the test-retest and the Achieva datasets, respectively. For the test-retest dataset one can expect true biological changes to be minimal, and we therefore computed the coefficient of variation (ratio of the standard deviation to the mean) for each brain structure. The results, shown in Table 1, indicate that the longitudinal method indeed performs better in this respect than the cross-sectional one for almost all the structures.

For the Achieva dataset, one may assume that the true change in volume of a structure over the span of a few years is approximately linear, except for lesions whose temporal trajectory is affected by disease effects, with growing and shrinking lesions occurring at the same time. We therefore fitted, for each structure and for each subject, a linear regression model to the longitudinal volumes estimated by each method, and computed the ratio of the standard deviation of the residuals to the intersect (time of the first scan is taken as zero). The results are summarized in Table 2 and indicate that the proposed model indeed yields generally better results in this respect.

Detecting disease effects: In order to ensure that the proposed method is not simply over-regularizing, we also assessed whether it can capture known group differences in the temporal evolution of certain brain structures better than the cross-sectional method. Towards this end, we compared the annualized percentage change (the slope of a linear regression model divided by its intersect) in the volume of the hippocampus between the CN, MCI and AD groups in the ADNI dataset. The results, shown in Fig. 3, indicate that the longitudinal method can indeed detect group differences better this way.

5 Discussion and conclusion

In this paper we have proposed a novel method for the segmentation of longitudinal brain MRI scans of patients suffering from MS. The method is based on an existing cross-sectional method for simultaneous whole-brain and lesion segmentation, and leverages subject-specific latent variables to encourage segmentations across time points to be similar to each other. Preliminary results indicate that it is able to produce more consistent and reliable segmentations compared to the cross-sectional version, while being more sensitive to group differences. Future work will involve an extensive analysis of disease progression in different MS patient groups, as well as a more careful tuning of the hyperparameters of the model.




Fig. 2: Graphical representation of the proposed model. In black the existing cross-sectional method of [16] for each time point t ; in blue the proposed additional latent variables for modeling temporal consistency between longitudinal scans with T time points.

Structure	Subject 1						Subject 2						Avg	
	GE	Philips	Siemens	GE	Philips	Siemens	GE	Philips	Siemens	GE	Philips	Siemens	Cross	Long
Lesions	3.79	3.90	4.78	3.79	6.27	3.39	2.84	4.04	2.57	1.95	2.35	1.22	3.77	3.05
Cerebral white matter	0.38	0.46	0.14	0.12	1.70	1.24	0.93	0.79	0.78	0.17	1.16	1.00	0.85	0.63
Cerebellum white matter	0.82	0.56	0.28	0.19	2.03	1.96	1.28	0.91	0.71	0.41	1.14	1.00	1.04	0.84
Cerebral cortex	0.53	0.50	0.50	0.44	1.50	1.84	0.60	1.08	0.30	0.40	0.74	1.10	0.70	0.89
Cerebellum cortex	0.38	0.34	0.31	0.36	1.13	1.25	0.61	0.63	0.42	0.36	0.86	0.57	0.62	0.58
Lateral ventricles	1.10	1.03	1.37	0.38	1.11	0.57	3.73	2.95	1.62	1.10	2.71	1.24	1.94	1.21
Hippocampus	1.03	0.67	0.55	0.43	0.84	1.46	1.88	1.39	0.55	0.57	1.37	0.91	1.04	0.90
Thalamus	0.44	0.44	0.52	0.31	1.21	0.54	0.92	0.79	0.50	0.35	1.02	0.36	0.77	0.46
Putamen	0.58	0.14	0.84	0.76	0.70	0.71	0.51	0.31	1.10	0.52	1.44	1.08	0.86	0.59
Pallidum	2.25	1.35	2.04	1.58	3.94	1.31	2.77	2.39	1.12	0.77	4.35	1.84	2.74	1.54
Caudate	0.80	1.09	0.96	0.90	1.00	0.84	1.79	1.04	0.92	0.54	0.61	0.41	1.01	0.80
Amygdala	1.51	0.47	0.42	0.20	1.85	1.04	1.17	0.64	1.05	0.41	0.60	0.42	1.10	0.53
Accumbens	1.84	1.27	1.77	1.40	1.80	1.39	1.96	0.69	0.80	0.80	1.73	0.80	1.65	1.06
Brain stem	0.70	0.53	0.32	0.19	1.10	0.93	0.78	0.59	0.54	0.34	0.61	0.63	0.67	0.53
Intracranial volume	0.30	0.06	0.13	0.18	0.53	0.56	0.46	0.06	0.22	0.08	0.52	0.36	0.36	0.22

Table 1: Coefficients of variation in [%] on the test-retest dataset, both for the proposed longitudinal (“Long”) and the cross-sectional (“Cross”) method.

Structure	Cross	Long	Structure	Cross	Long
Lesions	9.41	9.27	Cerebral white matter	0.56	0.31
Cerebellum white matter	0.59	0.45	Cerebral cortex	0.54	0.59
Cerebellum cortex	0.45	0.42	Lateral ventricles	2.04	1.85
Hippocampus	0.69	0.55	Thalamus	0.51	0.49
Putamen	0.70	0.42	Pallidum	1.23	0.78
Caudate	1.01	0.80	Amygdala	0.70	0.40
Accumbens	1.39	0.82	Brain stem	0.62	0.46
Intracranial volume	0.28	0.08			

Table 2: Average deviation from a linear trajectory in [%] for volumetric measurements in the Achieva dataset.




Fig. 3: Annualized percentage change (APC) in the volume of the hippocampus for the three groups of the ADNI dataset. Statistical significance was computed with a Welch’s t-test and effect size with Cohen’s d.

Acknowledgments: This project has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie project TRABIT agreement No 765148, the German Research Foundation No 428223038, and the NIH NINDS No R01NS112161.

References

1. F. Barkhof *et al.* Imaging outcomes for neuroprotection and repair in multiple sclerosis trials. *Nature Reviews Neurology*, 5(5):256–266, may 2009.
2. C.J. Azevedo *et al.* Thalamic atrophy in multiple sclerosis: A magnetic resonance imaging marker of neurodegeneration throughout disease. *Annals of Neurology*, 83(2):223–234, 2018.
3. A.J. Thompson *et al.* Diagnosis of multiple sclerosis: 2017 revisions of the McDonald criteria. *The Lancet Neurology*, 17(2):162–173, 2018.
4. A. Carass *et al.* Longitudinal multiple sclerosis lesion segmentation data resource. *Data in Brief*, 12:346–350, 2017.
5. C.R.G. Guttman *et al.* Quantitative follow-up of patients with multiple sclerosis using mri: Reproducibility. *Journal of Magnetic Resonance Imaging*, 9(4):509–518, 1999.
6. G. Gerig *et al.* Exploring the discrimination power of the time domain for segmentation and characterization of active lesions in serial MR data. *Medical Image Analysis*, 4(1):31–42, 2000.
7. P. Schmidt *et al.* Automated segmentation of changes in FLAIR-hyperintense white matter lesions in multiple sclerosis on serial magnetic resonance imaging. *NeuroImage: Clinical*, 23, 2019.
8. R. McKinley *et al.* Automatic detection of lesion load change in Multiple Sclerosis using convolutional neural networks with segmentation confidence. *NeuroImage: Clinical*, 25, 2020.
9. S.M. Smith *et al.* Accurate, Robust, and Automated Longitudinal and Cross-Sectional Brain Change Analysis. *NeuroImage*, 17(1):479–489, 2002.
10. D. Smeets *et al.* Reliable measurements of brain atrophy in individual patients with multiple sclerosis. *Brain and Behavior*, 6(9), 2016.
11. S.M. Smith *et al.* Normalized accurate measurement of longitudinal brain change. *Journal of Computer Assisted Tomography*, 25(3):466–475, 2001.
12. D. Rey *et al.* Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis. *Medical Image Analysis*, 6(2):163–179, 2002.
13. M. Battaglini *et al.* Automated identification of brain new lesions in multiple sclerosis using subtraction images. *Journal of Magnetic Resonance Imaging*, 39(6):1543–1549, 2014.
14. S. Jain *et al.* Two time point MS lesion segmentation in brain MRI: An expectation-maximization framework. *Frontiers in Neuroscience*, 10, 2016.
15. D. García-Lorenzo *et al.* Review of automatic segmentation methods of multiple sclerosis white matter lesions on conventional magnetic resonance imaging. *Medical Image Analysis*, 17(1):1–18, 2013.
16. S. Cerri *et al.* A Contrast-Adaptive Method for Simultaneous Whole-Brain and Lesion Segmentation in Multiple Sclerosis, 2020.
17. K. Van Leemput. Encoding probabilistic brain atlases using Bayesian inference. *IEEE Transactions on Medical Imaging*, 28(6):822–837, 2009.

18. J. Ashburner *et al.* Image registration using a symmetric prior - In three dimensions. *Human Brain Mapping*, 9(4):212–225, 2000.
19. D.P. Kingma and M. Welling. Auto-encoding variational bayes, 2013.
20. W.M. Wells *et al.* Adaptive segmentation of MRI data. *IEEE Transactions on Medical Imaging*, 15(4):429–442, 1996.
21. K. Van Leemput *et al.* Automated model-based bias field correction of MR images of the brain. *IEEE Transactions on Medical Imaging*, 18(10):885–896, 1999.
22. J.E. Iglesias *et al.* Bayesian longitudinal segmentation of hippocampal substructures in brain MRI using subject-specific atlases. *NeuroImage*, 141:542–555, 2016.
23. M. Reuter *et al.* Within-subject template estimation for unbiased longitudinal image analysis. *NeuroImage*, 61(4):1402–1418, 2012.
24. O. Puonti *et al.* Fast and sequence-adaptive whole-brain segmentation using parametric Bayesian modeling. *NeuroImage*, 143:235–249, 2016.
25. V. Biberacher *et al.* Intra- and interscanner variability of magnetic resonance imaging based volumetry in multiple sclerosis. *NeuroImage*, 142:188–197, 2016.