



OPEN ACCESS

EDITED BY

Loveleen Gaur,
Taylor's University, Malaysia

REVIEWED BY

Robert Zivadinov,
University at Buffalo, United States
Milon Biswas,
University of Alabama at Birmingham,
United States

*CORRESPONDENCE

Chenyu Wang
✉ chenyu.wang@sydney.edu.au

RECEIVED 29 March 2023

ACCEPTED 16 June 2023

PUBLISHED 06 July 2023

CITATION

Zhan G, Wang D, Cabezas M, Bai L, Kyle K,
Ouyang W, Barnett M and Wang C (2023)
Learning from pseudo-labels: deep networks
improve consistency in longitudinal brain
volume estimation.
Front. Neurosci. 17:1196087.
doi: 10.3389/fnins.2023.1196087

COPYRIGHT

© 2023 Zhan, Wang, Cabezas, Bai, Kyle,
Ouyang, Barnett and Wang. This is an
open-access article distributed under the terms
of the [Creative Commons Attribution License
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction
in other forums is permitted, provided the
original author(s) and the copyright owner(s)
are credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted which
does not comply with these terms.

Learning from pseudo-labels: deep networks improve consistency in longitudinal brain volume estimation

Geng Zhan^{1,2}, Dongang Wang^{1,2}, Mariano Cabezas¹, Lei Bai³,
Kain Kyle^{1,2}, Wanli Ouyang³, Michael Barnett^{1,2} and
Chenyu Wang^{1,2*}

¹Brain and Mind Center, The University of Sydney, Sydney, NSW, Australia, ²Sydney Neuroimaging
Analysis Center, Sydney, NSW, Australia, ³Shanghai AI Laboratory, Shanghai, China

Introduction: Brain atrophy is a critical biomarker of disease progression and treatment response in neurodegenerative diseases such as multiple sclerosis (MS). Confounding factors such as inconsistent imaging acquisitions hamper the accurate measurement of brain atrophy in the clinic. This study aims to develop and validate a robust deep learning model to overcome these challenges; and to evaluate its impact on the measurement of disease progression.

Methods: Voxel-wise pseudo-atrophy labels were generated using SIENA, a widely adopted tool for the measurement of brain atrophy in MS. Deformation maps were produced for 195 pairs of longitudinal 3D T1 scans from patients with MS. A 3D U-Net, namely DeepBVC, was specifically developed to overcome common variances in resolution, signal-to-noise ratio and contrast ratio between baseline and follow up scans. The performance of DeepBVC was compared against SIENA using McLaren test-retest dataset and 233 in-house MS subjects with MRI from multiple time points. Clinical evaluation included disability assessment with the Expanded Disability Status Scale (EDSS) and traditional imaging metrics such as lesion burden.

Results: For 3 subjects in test-retest experiments, the median percent brain volume change (PBVC) for DeepBVC and SIENA was 0.105 vs. 0.198% (subject 1), 0.061 vs. 0.084% (subject 2), 0.104 vs. 0.408% (subject 3). For testing consistency across multiple time points in individual MS subjects, the mean (\pm standard deviation) PBVC difference of DeepBVC and SIENA were 0.028% (\pm 0.145%) and 0.031% (\pm 0.154%), respectively. The linear correlation with baseline T2 lesion volume were $r = -0.288$ ($p < 0.05$) and $r = -0.249$ ($p < 0.05$) for DeepBVC and SIENA, respectively. There was no significant correlation of disability progression with PBVC as estimated by either method ($p = 0.86$, $p = 0.84$).

Discussion: DeepBVC is a deep learning powered brain volume change estimation method for assessing brain atrophy used T1-weighted images. Compared to SIENA, DeepBVC demonstrates superior performance in reproducibility and in the context of common clinical scan variances such as imaging contrast, voxel resolution, random bias field, and signal-to-noise ratio. Enhanced measurement robustness, automation, and processing speed of DeepBVC indicate its potential for utilisation in both research and clinical environments for monitoring disease progression and, potentially, evaluating treatment effectiveness.

KEYWORDS

brain volume change, longitudinal analysis, multiple sclerosis, atrophy measurement, deep learning, MRI

1. Introduction

Brain atrophy is a clinically relevant biomarker of disease progression in patients with multiple sclerosis (MS) that reflects irreversible tissue damage due to neuro-axonal destruction, demyelination and gliosis (Bermel and Bakshi, 2006). Accelerated brain tissue loss can be detected in MS cohorts compared to a healthy control population (De Stefano et al., 2016). Clinically, brain atrophy is a key predictor of future disease worsening and cognitive impairment in patients with MS (Popescu et al., 2013; De Stefano et al., 2014; Jacobsen et al., 2014); and has been used frequently in MS clinical trials as a secondary measure of treatment efficacy (Filippi et al., 2004; Cadavid et al., 2017).

The incorporation of brain atrophy into monitoring paradigms for individual patients requires significantly improved accuracy, precision and robustness. Over the past two decades, numerous methods (Hajnal et al., 1995; Rudick et al., 1999; Collins et al., 2001; Bermel et al., 2003; Friston, 2003; Horsfield et al., 2003; Prados et al., 2015; Smeets et al., 2016) have been developed for quantifying longitudinal brain volume change (BVC). Boundary Shift Integral (BSI) (Freeborough and Fox, 1997), gBSI (Prados et al., 2015), NeuroSTREAM (Dwyer et al., 2017), Jacobian integration (Nakamura et al., 2014), and SIENA (Smith et al., 2002) use different measurements to track the movement of the brain boundary. IPCA (Chen et al., 2004) uses an iterative principal component analysis method to find the outliers reflecting between-scan differences. MSmetrix (Smeets et al., 2016) adopts nonrigid registration and Jacobian integration of deformation fields to produce atrophy measures. Similarly, Quantitative Anatomical Regional Change (QUARC) (Holland et al., 2011) utilises non-rigid registration, but directly calculates hexahedral volumes to yield the fractional volume change. FreeSurfer-longitudinal (FS) (Reuter et al., 2012) is a segmentation-based method that performs tissue segmentation at each time point. Among those methods, SIENA (Smith et al., 2002) is arguably the most widely used algorithm in MS clinical trials.

Although continuous efforts have been made to develop new methods that improve the accuracy and reliability of BVC quantification, longitudinal MRI measurement is susceptible to inconsistency of imaging acquisition at baseline and follow-up (Lee et al., 2019; Medawar et al., 2021), particularly in routine clinical practise where scanner upgrades and protocol variations are normal. Common examples of inconsistencies that impact the reliability of quantitative BVC measurement, regardless of methodology, include image contrast differences (Preboske et al., 2006), intensity non-uniformity (Takao et al., 2010), noise, or different resolutions and voxel spacing (Vrenken et al., 2013). Additionally, there are few comparative studies that assess the performance of newer (versus older) measurement methods, such as SIENA and Boundary Shift Integral (Freeborough and Fox, 1997), respectively, in the presence of such acquisition inconsistencies.

However, several approaches have been proposed to ameliorate the influence of acquisition inconsistencies and thereby improve the accuracy of BVC measurement algorithms.

The first approach aims to reduce scan inhomogeneity during pre-processing (Smith et al., 2002; Learned-miller and Ahammad,

2004; Lewis and Fox, 2004; Vovk et al., 2004; Vemuri et al., 2005; Duffy et al., 2018; Higaki et al., 2019). By removing the bias field from the longitudinal input scans, these methods aim to reduce the variance in BVC estimates. Several other data harmonisation methods (Dewey et al., 2019, 2020; Beer et al., 2020; Garcia-Dias et al., 2020; Liu et al., 2021) aim to improve the qualitative and quantitative consistency of differently acquired MRI scans. In practise, these methods assign one scan as a reference and process images in the second scan to narrow the differences attributable to protocol inconsistency. Although many of these methods aim to improve the quantitative utility of MRI in long-term or multi-site studies, most are not specifically designed for consistent BVC measurement. Rather, they focus on harmonisation of images to a reference image or providing segmentation masks with greater consistency (as determined by Dice, Coefficient of Joint Variation or similar) with those derived from a reference image in a test-retest setting. Therefore, the effectiveness of these methods for producing consistent BVC measurements is not directly measured.

The second approach focuses on correcting BVC estimates during post-processing. Lee et al. (2019) estimates the fixed effect of scanner changes with a linear model and accounts for this factor during measurement of BVC. Sinnecker et al. (2022) estimates an additive fixed corrective term for scanner change by comparing BVC rates during scanner change and no scanner change for healthy control subjects. These methods are subject to scanner-specific variation and require group level test results to estimate the correction factor.

The addition of both pre-processing and post-processing steps to BVC measurement algorithms increase overall computational complexity and processing time. Additionally, while these methods may result in qualitative and quantitative improvements, most are confined to the research domain and their integration with (and effectiveness in) clinical workflows is not clear.

To improve the reliability of measurements while also considering computational cost, we introduce DeepBVC, a deep learning-based approach that addresses the impact of protocol inconsistency in longitudinal brain volume measurements. Extensive experimentation demonstrates its effectiveness in clinical settings, providing superior performance for potential large-scale clinical applications. Our DeepBVC combines a deep neural network with data augmentation to provide fully automated and robust BVC assessment. A deep neural network offers generalisation to common brain atrophy patterns; and comprehensive data augmentation (Shorten and Khoshgoftaar, 2019) provides robustness and mitigates protocol and other acquisition-related inconsistencies. Specifically, the deep neural network module is used to estimate shift at the brain boundary from baseline to follow-up; and the data augmentation module synthesises images that contain common image distortions and acquisition differences during the training stage of the deep neural network. We also propose a novel training regime for DeepBVC. In general, supervised neural network training requires a large-scale dataset with accurate labels. However, it is not practical to acquire accurate sub-voxel level atrophy estimations from images. Existing BVC measurement algorithms contain known or unknown bias and random variation factors for different scans (Thanellas and Pollari, 2010). Therefore, we used the brain

boundary shift produced by SIENA as the pseudo-label in training, noting precedents for the use of software-generated annotations as the label for tasks such as brain parcellation (Henschel et al., 2020). We then utilised the inherent regularisation properties of convolutional neural networks to tackle the underlying label noise (Goodfellow et al., 2014; Zhang H. et al., 2017; Zhang S. et al., 2017).

In summary, we have developed DeepBVC, a deep learning based method, for longitudinal BVC assessment with the following contributions:

- Data augmentation is introduced to cope with inconsistent imaging acquisitions that are largely unavoidable in clinical settings.
- Our deep learning model is trained with pseudo-labels from SIENA and the impact of label noise is ameliorated by a regularisation method.
- Experiments in two datasets demonstrate that DeepBVC has better accuracy, robustness, reliability, and consistency compared with SIENA, even though the intermediate outputs from SIENA were used for training.

2. Materials and methods

This section is organised as follows: We describe the data used in this study in Section 2.1 and the data preprocessing in Section 2.2. In Section 2.3, we illustrate the details of our method, including the model structure, training details, and the integration our model into the pipeline of BVC estimation. Next, we introduce the settings and metrics of evaluation experiments in Section 2.4.

2.1. Data acquisition

We use two data sources for this study: an in-house dataset (MS Clinical Dataset) comprising clinical data and matched MRI scans from patients with relapsing remitting multiple sclerosis (RRMS) subjects; and a public test-retest dataset (Maclaren test-retest dataset) (Maclaren et al., 2014) comprising scans from three healthy control subjects. The study was approved by the University of Sydney and followed the tenets of the Declaration of Helsinki.

2.1.1. MS Clinical Dataset

Written informed consent was obtained from all participants. In total, 2,457 T1-weighted MRI exams from 648 patients diagnosed with RRMS were included in this study (Table 1). All patients were recruited from the MS Clinic based at the Brain and Mind Centre, University of Sydney; and clinical MRI exams were acquired with one of three 3T MRI scanners (Table 2) between 2010 and 2020. Longitudinal scans for each patient were acquired with the same scanner using a consistent protocol. MRI acquisition parameters are summarised in Table 2.

TABLE 1 Demographic, clinical characteristics of the patients in the in-house dataset at baseline.

	Train	Test 1 (Section 2.4.3)	Test 2 (Section 2.4.4)	Test 3 (Section 2.4.5)
Patients, <i>n</i> (% female)	195 (74)	120 (81)	233 (77)	208 (33)
Age, (years)	41.6 (12.4)	40.3 (10.5)	41.5 (9.6)	39.7 (9.4)
Disease duration, (years)	9.4 (8.7)	8.7 (7.9)	10.4 (6.9)	8.5 (6.0)
EDSS ^a	2.1 (1.8)	1.9 (1.6)	1.6 (1.8)	2.0 (1.8)

The age, disease duration and EDSS are reported as mean (SD).

^aExpanded Disability Status Scale.

TABLE 2 MRI acquisition details for the MS Clinical Dataset.

Scanner	T1 parameters
GE discovery 3.0T	TE = 2.6 ms
	TR = 7 ms
	TI = 0.45 s
	0.93 × 0.93 × 1 mm ³
Philips Ingenia 3.0T	TE = 2.4 ms
	TR = 8 ms
	1 × 1 × 1 mm ³
SIEMENS Skyra 3.0T	TE = 2.5 ms
	TR = 2.2 s
	TI = 0.9 s
	0.90 × 0.90 × 0.90 mm ³

2.1.2. Test-retest data

We use the Maclaren test-retest dataset (Maclaren et al., 2014) to test measurement reliability. The dataset comprises 120 T1-weighted scans from three healthy subjects aged 26–31 years, acquired with a GE MR750 3T scanner. Each subject was scanned twice on 20 different days within a 31-day period. The acquisition protocol (TE: 3 ms, TI: 0.4 s, TR: 7.3 ms, 1.2 mm slice thickness) followed the recommendations of the Alzheimer's Disease Neuroimaging Initiative (ADNI) (Jack et al., 2008).

2.2. Data pre-processing

2.2.1. Format conversion and data selection

For the MS Clinical Dataset, images were originally stored in DICOM format. All acquired DICOM data were converted to NIFTI format by *dcm2nii* (Li et al., 2016). The quality of all MRI data was visually assessed by experienced neuroimaging analysts at the Sydney Neuroimaging Analysis Centre (Sydney, Australia). Images that failed quality assessment (incomplete brain coverage, severe imaging artifacts) were excluded from further study. N4 Bias Field Correction (Lowekamp et al., 2013) was applied to remove the bias field from all scans meeting quality criteria.

2.2.2. Brain and skull segmentation

Brain and skull segmentation was performed with BET (Jenkinson et al., 2005) for all eligible scans and the masks were manually refined by experienced neuroimaging analysts. Skull stripping was undertaken to generate images only containing brain tissues for subsequent analysis.

2.2.3. SIENA analysis

2.2.3.1. Co-registration

For the longitudinal scan pairs of each subject, T1-weighted scans at baseline and follow-up were aligned using two-step affine registration (Freeborough and Fox, 1997; Smith et al., 2002). First, the skull masks were used to optimise the scale and skew; then, the brain images were used to optimise image translation and rotation. All registration results were manually checked by trained neuroimaging analysts, and poorly aligned pairs excluded from further analysis.

2.2.3.2. Brain edge point segmentation

After brain alignment, FAST (Zhang et al., 2001) was used to segment longitudinal scans into the principal brain tissue compartments: white matter (WM), grey matter (GM), and cerebrospinal fluid (CSF). As whole brain volume change includes changes for both WM and GM, the probabilistic maps from FAST were binarised (using a threshold of 0.5) and defined the union of WM and GM as the foreground segmentation, and the remainder of the image as background. Consequently, brain edge points were defined by the edges of the foreground mask.

2.2.3.3. Change analysis

Voxel-wise atrophy/growth was estimated for all brain edge points from baseline to follow-up time points; and the mean edge point motion converted into PBVC (a single number) with a self-calibration step.

2.3. Model

2.3.1. Model structure

We used a 3D-Unet (Çiçek et al., 2016) as the backbone network for feature extraction, followed by a single convolutional layer as the prediction head. The main blocks of the networks use residual convolutional layers (He et al., 2016) with group normalisation (Wu and He, 2018). The model inputs comprise a pair of baseline and follow-up T1 images. The label is a voxel-wise estimation of brain boundary shift produced from SIENA. Because SIENA estimates the brain boundary shift for each voxel location in the input, the label and the model output are 3-D tensors with the same dimensions as the inputs. During both training and inference, the model output is masked with the brain boundary segmentation and only the outputs at the boundary locations are retained (outputs are set to 0 for non-boundary locations).

2.3.2. Training details

An overview of our method is shown in Figure 1. For each iteration, eight scan pairs were randomly selected from all training

scan pairs. Each pair was randomly cropped into a pair of patches of size $64 \times 64 \times 64$ due to GPU memory constraints. The patches were then fed into the model and the loss calculated accordingly to update the model weights.

We used mean squared error (MSE) as the loss function L_{MSE} to evaluate the model's deformation prediction when compared with the pseudo-labels from SIENA (edge point motion). Furthermore, to render the model insensitive to differences in imaging quality during training, we adopted a consistency regularisation loss. Specifically, the loss minimised the difference between the predictions derived from the original images and the augmented images, such that brain volume differences were maintained in the context of an isolated change in imaging acquisition conditions.

Formally, for a data point x , the regularisation loss term was defined as:

$$L_{reg} = \|p_{model}(y|x; \theta) - p_{model}(y|Augment(x); \theta)\|^2, \quad (1)$$

where $Augment(x)$ is a stochastic transformation, whose effect is not identical for each training sample, but rather follows a distribution. As the regularisation term requires the underlying brain volume change to remain constant, spatial transformations that deformed the original brains were not permitted. Therefore, we included both random spacing anisotropy re-sampling and random contrast alterations as possible augmentation steps. For clarity, only one type of augmentation was randomly selected for each sample.

Finally, the overall loss L was defined as:

$$L = L_{MSE} + \lambda L_{reg}, \quad (2)$$

where we set $\lambda = 0.01$.

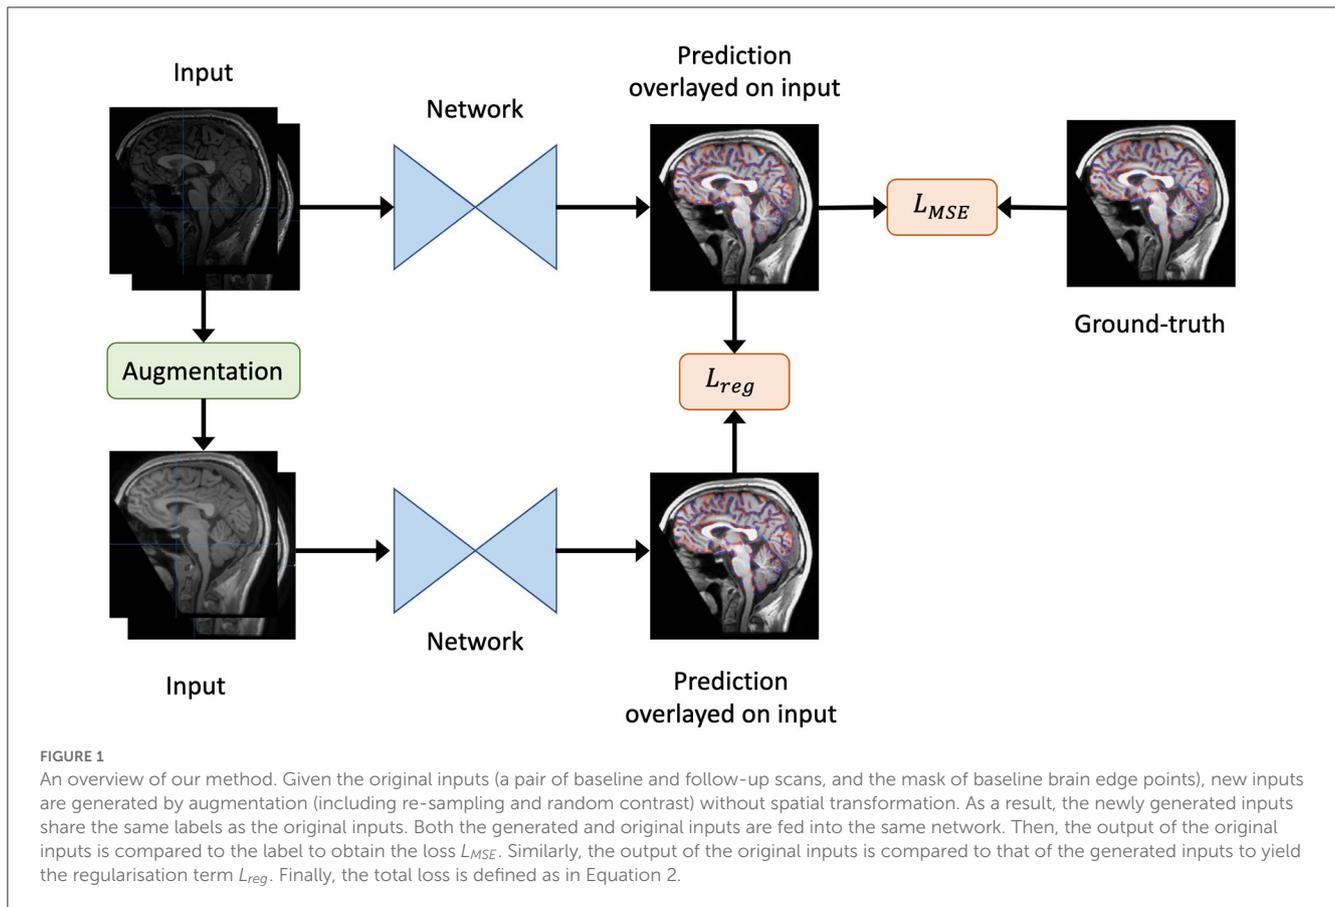
The final loss function was optimised using the Adam optimiser (Kingma and Ba, 2014) with an initial learning rate of 0.001 that was reduced by a factor of 0.5 when the loss did not drop for three consecutive epochs. The model was trained for 2,500 iterations per epoch and for 50 epochs in total. Model optimisation was performed with two NVIDIA V100 GPU cards on an NVIDIA DGX-1 station.

2.3.3. Training data

To train the model, we collected 195 pairs of scans (one pair per subject) from the MS Clinical Dataset (Section 2.1.1, Figure 2), from which 70% (137 pairs) were randomly selected for training and the remaining 30% (58 pairs) for validation. Two additional, independent imaging datasets were used for testing as described in Table 1. Testing datasets did not overlap with training or validation datasets; and all 195 subjects involved in model development were excluded from evaluation experiments.

2.4. Experimental setup

Five experiments were undertaken to evaluate the performance of DeepBVC with respect to test-retest consistency,



multi-step longitudinal consistency, protocol-agnostic test-retest consistency, relevance to T2 lesion and correlation with disability.

2.4.1. Consistency with test-retest data

The Maclaren test-retest data used for this experiment is described in Section 2.1.2. We grouped the data into baseline follow-up pairs as follows: the two scans from the same day and subject were used as a longitudinal pair with no atrophy. We used 60 pairs in total (20 pairs per subject).

For this experiment, we assumed that there would be no atrophy because the pairs were scanned back-to-back (i.e. the brain volume difference should be 0%). We ran SIENA and DeepBVC methods to measure the brain volume difference for each pair. We then grouped the results by subject and report each method's mean and standard deviation. We also report the mean absolute error for the results of each subject, where the error was acquired by comparing the BVC measurements against the 0% BVC.

2.4.2. Influence of the protocol inconsistency

To test the influence of various acquisition protocol inconsistencies, we used the Maclaren test-retest data (Section 2.1.2), coupled with imaging processing techniques to synthesise new image pairs with protocol pseudo-inconsistencies, as described below and shown in Figure 3. We then compared the experimental results with those derived from the original test-retest

data. Experiments followed the design described in Section 2.4.1, with the addition of a pre-processing step to include synthesised images as follows.

2.4.2.1. Random contrast adjustments

Image intensities were adjusted by the parameter γ . Each voxel intensity x was updated as

$$x = \left(\frac{x - x_{min}}{x_{range}} \right)^\gamma \cdot x_{range} + x_{min}, \quad (3)$$

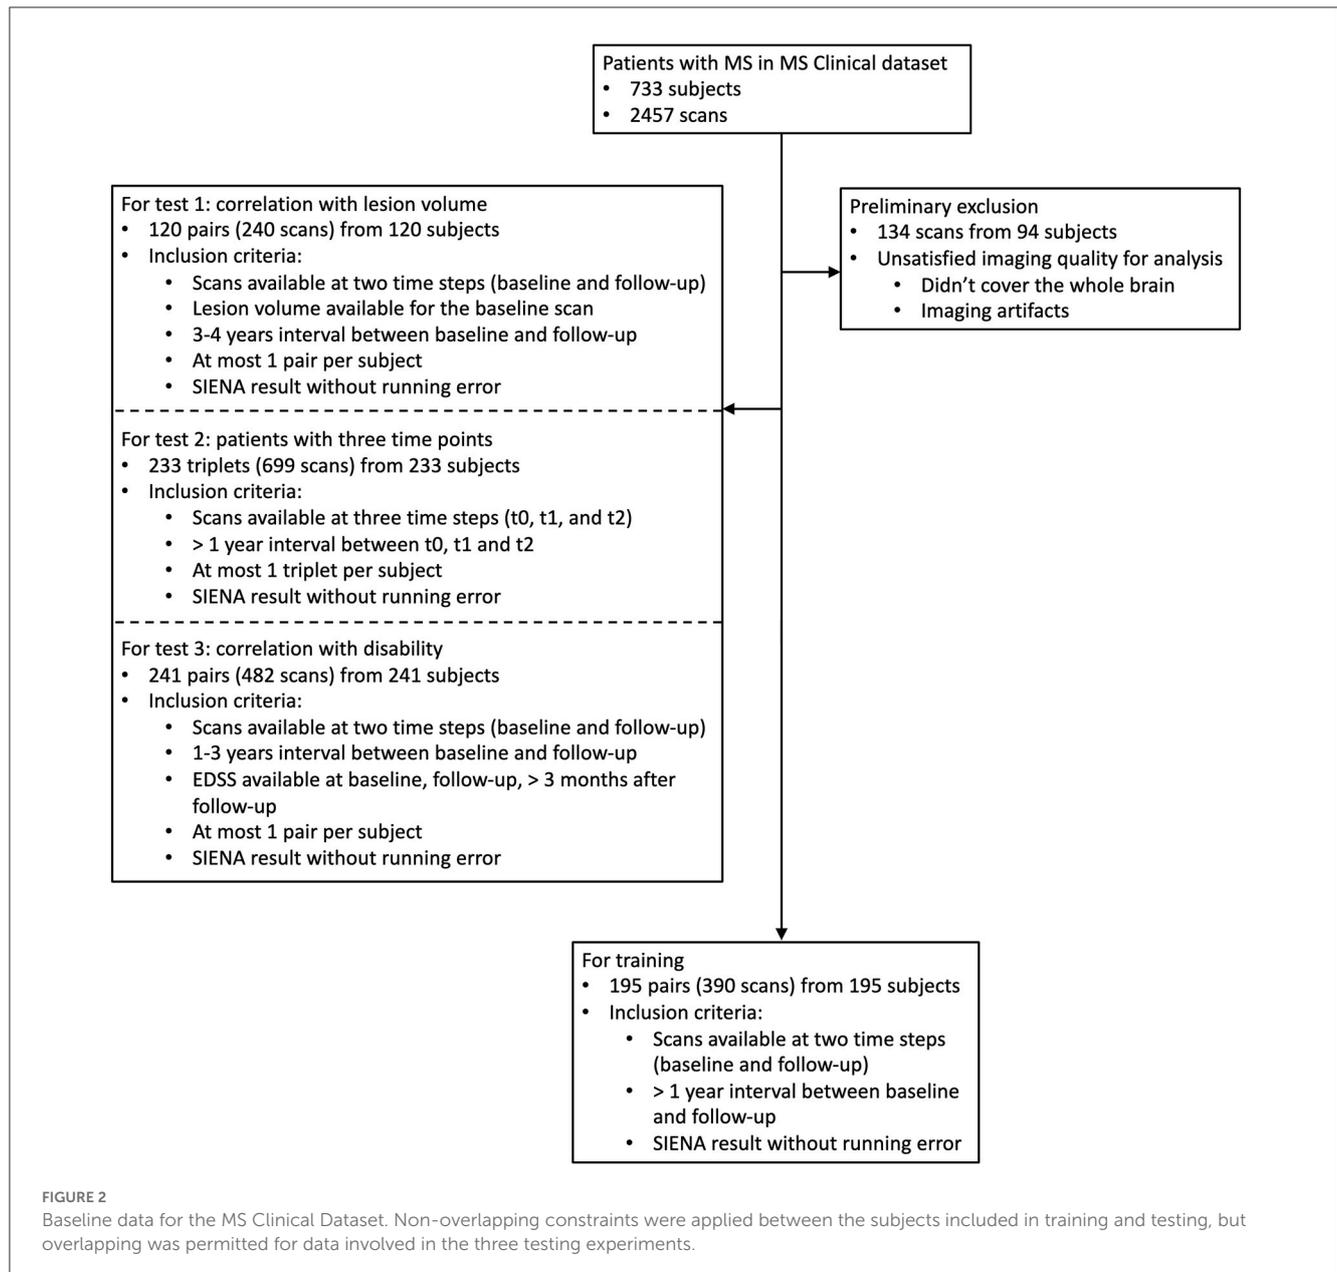
where x_{min} is the minimal voxel value in the original brain image, and x_{range} is the difference between the maximal and minimal voxel value in the scan, excluding the background.

2.4.2.2. Random bias field

The bias field was generated from a linear combination of smoothly varying basis functions, as described in Van Leemput et al. (1999). In practise, we observed that intensity inhomogeneity was more likely to occur along one of three axes (sagittal, coronal, and axial). Therefore, we synthesised the random bias field along one of the (randomly selected) axes.

2.4.2.3. Random spatial anisotropy

Spacing anisotropy was introduced by down sampling an image along an axis and then re-sampling back to its original spacing. In our experiment, we simultaneously performed this random transformation along all three axes (sagittal, coronal, and axial).



2.4.2.4. Gaussian noise

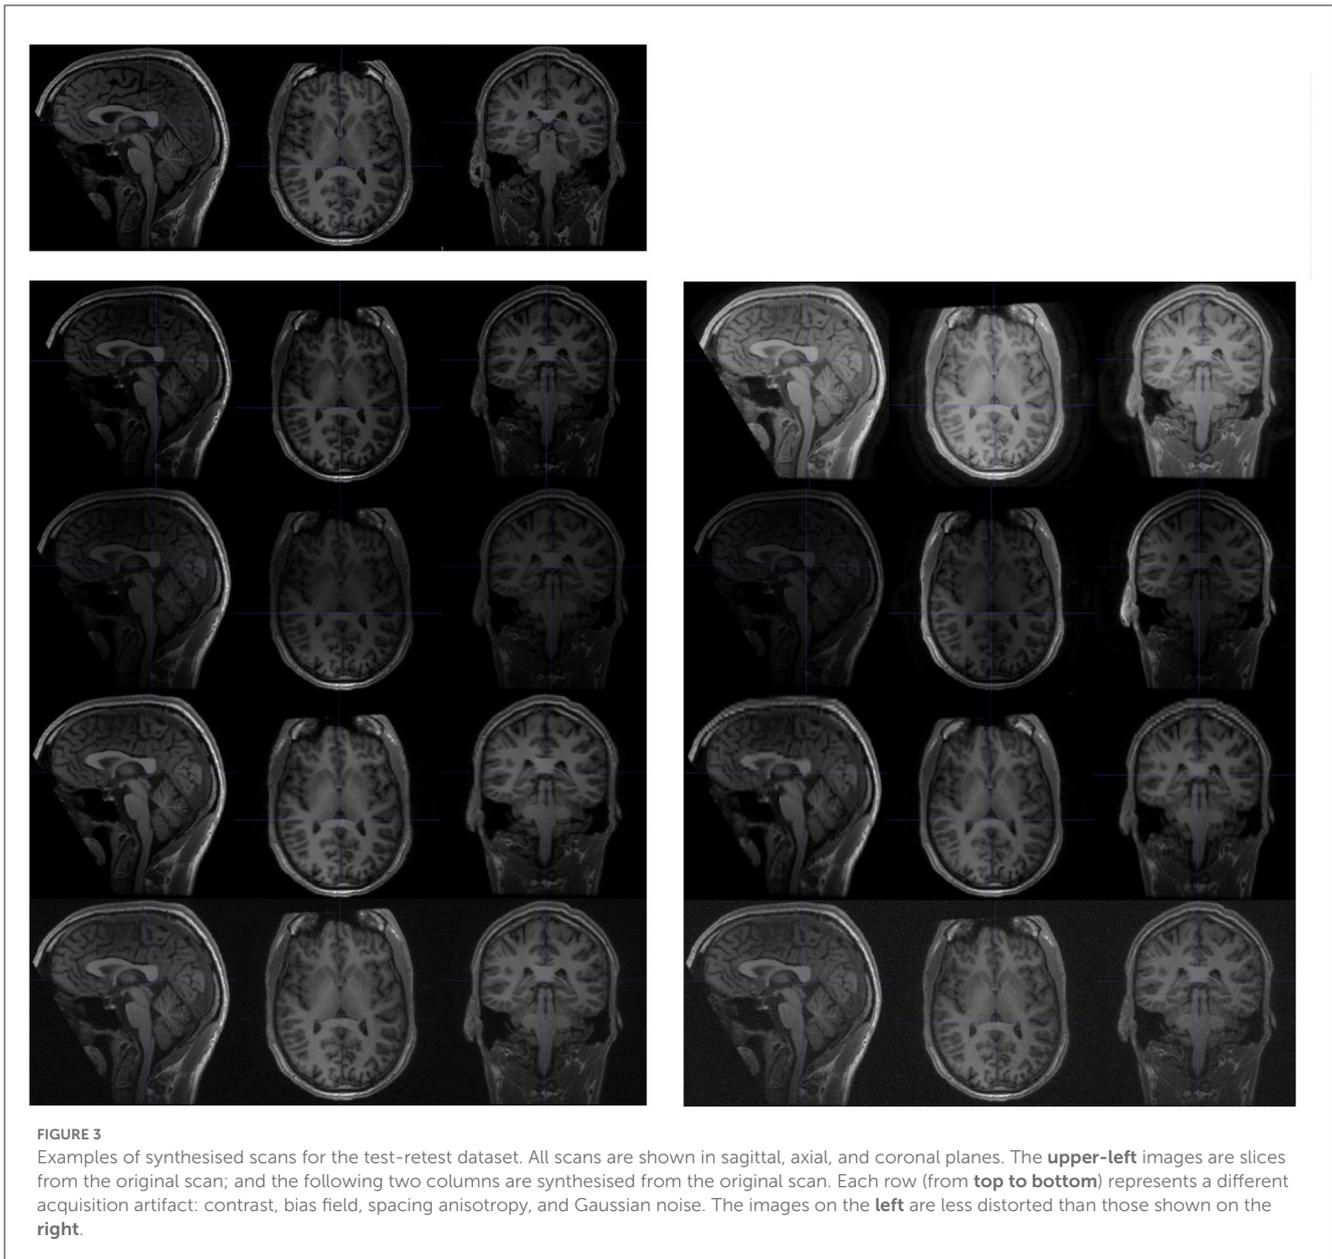
The noise for each voxel was sampled from a normal distribution with random parameters, and added to the original image.

To systematically explore the impact of protocol inconsistency, we analysed the performance of both DeepBVC and SIENA for different levels of the four types of inconsistency, controlled by different parameters during synthesis. For protocol inconsistency in Gaussian noise, spatial resolution anisotropy and bias field, the value of the parameters is positively correlated to the degree of inconsistency. For the inconsistency in image contrast, the measurement follows a different relationship (Equation 3), namely image contrast is controlled by γ . The larger the difference between γ and 1, the higher the inconsistency between synthesised and original images.

2.4.3. Multi-step consistency with three time points

Inspired by the experiments of Smith et al. (2002), we included data from three time points in our analysis, enabling both single-step and a multi-step measurement strategies. For the single step approach, we estimated brain atrophy between the first (t0) and last (t2) time points, while for the multi-step approach, we combined the estimated intermediate atrophy between the t0 and the second timepoint (t1), and t1 and t2. A smaller difference between the single-step and multi-step approaches indicates a more consistent measurement.

For these experiments, we use data selected from the MS Clinical Dataset (Section 2.1.1), restricting inclusion to those subjects with three available brain scans with an interval of at least



1 year between successive time points. In total, 233 subjects (three scans each subject) were used for this experiment.

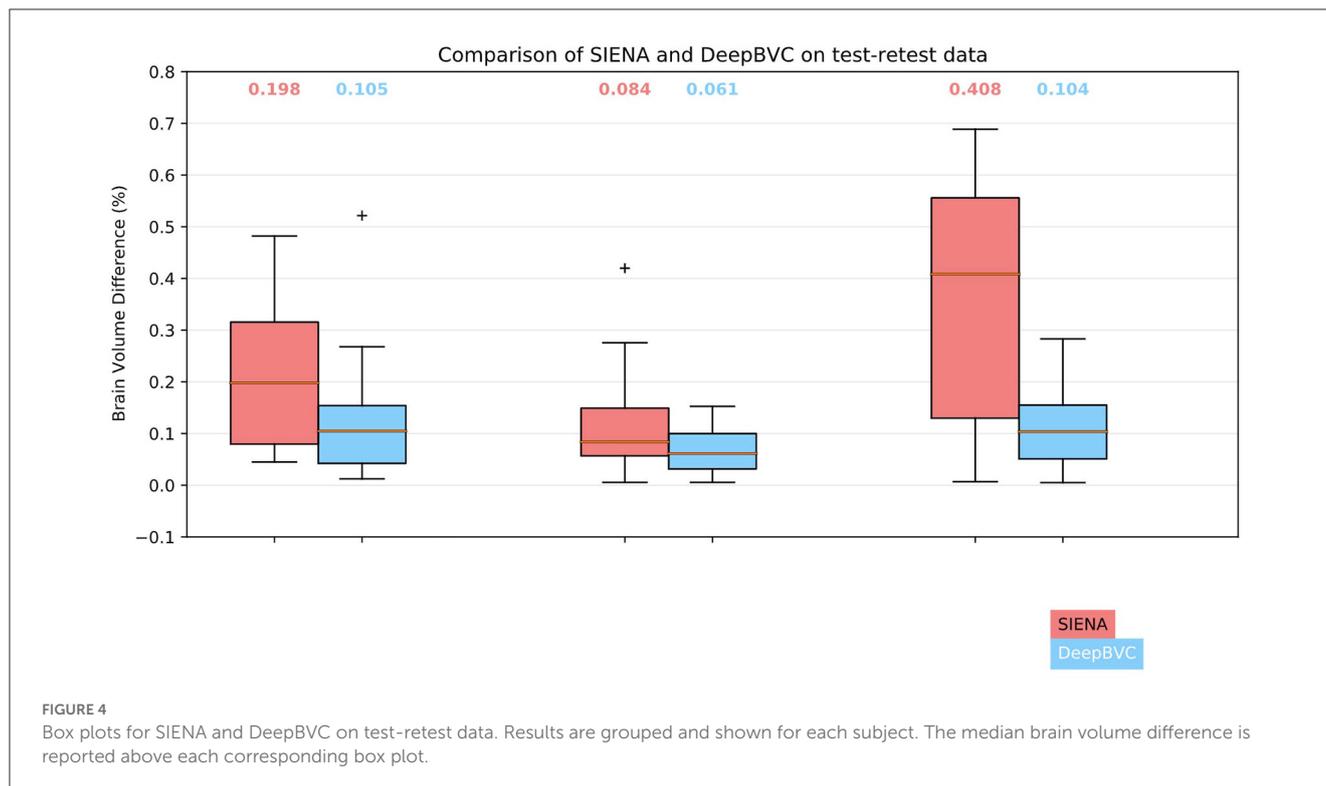
2.4.4. Brain T2 lesion volume and brain volume change

To further investigate the predictions of our method and their clinical impact, we analysed the correlation between brain volume change and baseline lesion volume as the rate of brain volume loss has been found to correlate with baseline T2 lesion volume (Tedeschi et al., 2005). We assumed that improved accuracy of BVC measurement would enhance this correlation. For this experiment, scans were selected from the MS Clinical Dataset (Section 2.1.1), based on availability of T1-w and FLAIR sequences at baseline and T1-w images at follow up (3–4 years subsequent to the baseline time point). The baseline lesion volume

was derived from an in-house deep learning lesion segmentation algorithm (Liu et al., 2022) followed by manual review (and, if required, correction) of lesion masks by trained neuroimaging analysts at the Sydney Neuroimaging Analysis Centre. We report linear correlation and partial correlation (controlled for age and disease duration) for DeepBVC and SIENA. SPSS (Field, 2013) was used for this analysis.

2.4.5. Disability and brain volume change

Sustained progression of the expanded disability status scale (EDSS) score has reported to correlate with higher rates of brain atrophy (Rudick et al., 2000; Bermel and Bakshi, 2006). To investigate the clinical relevance of the BVC, we therefore compared the annualised BVC with EDSS progression over 1–3 years in subjects with both T1-w scans and clinical data available at



two time points with an interval of at least 12 months. Subjects were firstly grouped into EDSS progressors and non-progressors, where EDSS progression was determined by three strata as previously described by Kalincik et al. (2015) and confirmed over 3 months. BVC was then determined by both DeepBVC and SIENA and reported for each group. We also analysed BVC for matched subjects from each group, using 1-to-multiple propensity score matching based on age and disease duration.

3. Results

For each validation experiment involving the MS Clinical data, we included all subjects that met the relevant inclusion criteria (Sections 2.3, 2.4 and Figure 2). The demographic and clinical characteristics of the subjects eligible for training and each validation are listed in Table 1. Among 2,457 scans from 648 subjects, 134 scans from 94 subjects were first excluded from all experiments due to poor imaging quality. For multi-step consistency experiments (Section 2.4.3), 233 eligible subjects (77% female) with three available scan timepoints were used. At the time of baseline imaging, mean age and disease duration was 41.5 and 10.4 years, respectively; and mean EDSS was 1.6, in this group. For the lesion experiment, 120 pairs of scans were available. In this group, 81% of the patients were female, with a mean age and disease duration of 40.3 and 8.7 years, respectively; and a mean EDSS of 1.9. The remaining subjects (195 subjects/scan pairs, 74% female) were used for training. In this group, mean age and disease duration was 41.6 and 9.4 years, respectively; and mean EDSS was 2.1.

TABLE 3 Brain volume difference in percentage for three subjects in the test-retest dataset.

Subject	SIENA (%)	DeepBVC (%)
Sub1	0.212 (\pm 0.139)	0.126 (\pm 0.116)
Sub2	0.118 (\pm 0.103)	0.066 (\pm 0.042)
Sub3	0.351 (\pm 0.228)	0.111 (\pm 0.077)

The numbers are reported as average (\pm standard deviation) across all the sessions for each subject.

3.1. Consistency with test-retest data

We illustrate the performance of SIENA and DeepBVC in Figure 4 and Table 3. The subject-wise means and standard deviations of PBVC measured by DeepBVC were smaller than by SIENA (Table 3). For all three subjects, the BVC measured by DeepBVC was less dispersed and was closer to 0 (Figure 4). The median PBVC for DeepBVC smaller than the equivalent plot for SIENA for all subjects (0.105 vs. 0.198, 0.061 vs. 0.084, 0.104 vs. 0.408, respectively). One outlier with a large BVC was found for DeepBVC (subject 1) and another for SIENA (subject 2).

3.2. Influence of the protocol inconsistency

The expected PBVC for both methods was 0 for pairs of identical scans.

For contrast inconsistencies (Figure 5), the box plot for DeepBVC showed a much lower distribution of errors than SIENA; and, unlike SIENA, no significant increase as gamma

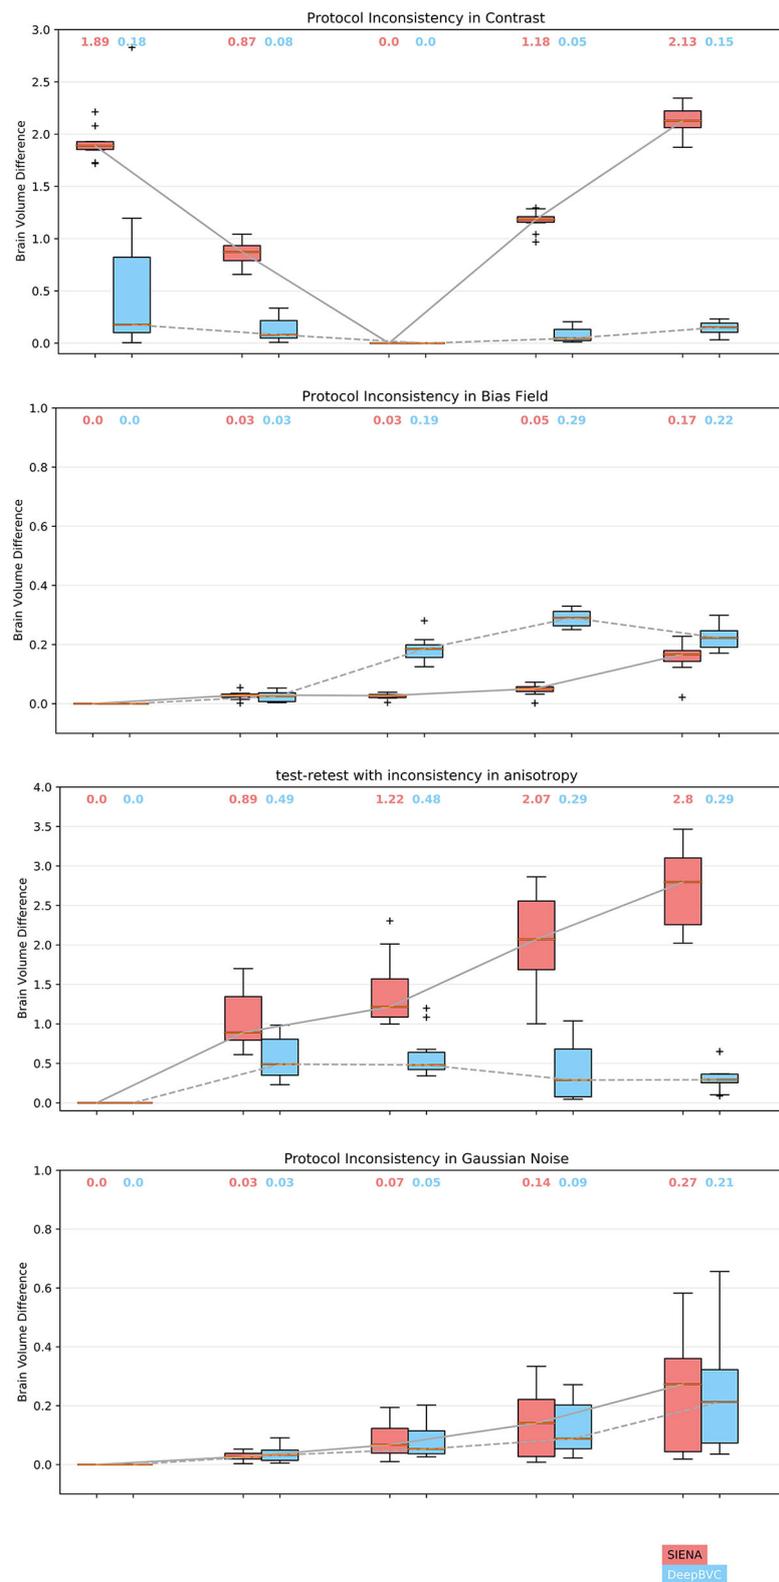


FIGURE 5 Brain volume difference estimation for stratified protocol inconsistency. For random contrast, the imaging quality is unchanged when the parameter is 1. The larger the difference between the parameter and 1, the lower the image quality and more inconsistent the protocol. The x-values are parameters determining the image processing step and the output image quality. For random Gaussian noise, bias field and anisotropy, parameters are negatively correlated to the quality of the output image (the higher the value, the lower the image quality, the more inconsistent the protocol). The box plots at the left-most of each figure represent the imaging quality is unchanged.

changed (up to ± 0.5). Furthermore, the median brain volume difference of DeepBVC was lower than SIENA by one order of magnitude. Finally, the variance in the error was similar for both methods when λ was 0.75, 1.25, and 1.5. However, when $\lambda = 0.5$, the errors for DeepBVC showed a wider distribution than SIENA.

For bias field inconsistencies, the medians and interquartile ranges for DeepBVC were higher than SIENA, though brain volume difference was low for both techniques (range 0.00–0.29 and 0.00–0.17, respectively).

For spatial resolution anisotropy, the differences measured by SIENA became greater as the inconsistency level increased. As shown in Figure 5, the median difference gradually increased from 0.89% ($\gamma = 2$) to 2.8% ($\gamma = 5$). For DeepBVC, the median difference remained low at the different levels of inconsistency tested. For the highest level of inconsistency, the median difference measured by SIENA reached 2.8%, while DeepBVC remained as low as 0.29%. The error distribution of DeepBVC was less scattered and closer to 0 than the equivalent of SIENA for all levels of inconsistency tested.

For Gaussian noise, the error distribution of DeepBVC was less dispersed and closer to 0 than the equivalent of SIENA when the Gaussian noise level was 0.2, 0.3 and 0.4. The error distribution of DeepBVC was wider than SIENA and the median was 0.3% for both methods when Gaussian noise level was 0.1.

3.3. Multi-step consistency with three time points

Among the 233 subjects involved in this evaluation experiment, the mean (\pm SD) time interval for $t_0 \rightarrow t_1$ was 1.6 (± 0.9) years (range 1.0–6.9 years); for $t_1 \rightarrow t_2$, 1.6 (± 0.8) years (range 1.0–6.7 years); and for $t_0 \rightarrow t_2$, 3.2 (± 1.2) years (range 2.0–9.4 years).

We report the difference in single-step and two-step BVC measurements for the two methods in Table 4 and Figure 6.

The mean (\pm standard deviation) for SIENA and DeepBVC was 0.031% ($\pm 0.154\%$) and 0.028% ($\pm 0.145\%$), respectively. The mean absolute error (\pm standard deviation) was 0.123% ($\pm 0.097\%$) and 0.120% ($\pm 0.087\%$), respectively. The difference between direct ($t_0 \rightarrow t_2$) and two-step measurements ($t_0 \rightarrow t_1$ and $t_1 \rightarrow t_2$) was relatively smaller for our method ($p = 0.78$).

When comparing the direct and indirect measurement for each method using Bland-Altman plots, the points are scattered randomly above and below 0 for both SIENA and DeepBVC (Figure 6). The 1.96 SD and -1.96 SD for SIENA is 0.27 and -0.38 , respectively (approximating the values reported by Smith et al., 2002), whereas the 1.96 SD and -1.96 SD for DeepBVC is 0.26 and -0.31 , respectively. DeepBVC was less likely than SIENA to generate an output that suggested brain volume growth (biologically less likely) over time; and there was no obvious bias between two methods. For points with large atrophy (brain loss $> 2\%$), most points are between SD 1.96 = 0.53 and $-SD$ 1.96 = -0.45 with only one point outside this range.

TABLE 4 The difference between direct measurements ($t_0 \rightarrow t_2$) and two-step measurement ($t_0 \rightarrow t_1$ and $t_1 \rightarrow t_2$).

Method	Mean (\pm SD) %	Mean absolute (\pm SD) %
SIENA	0.031 (\pm 0.154)	0.123 (\pm 0.097)
DeepBVC	0.028 (\pm 0.145)	0.120 (\pm 0.087)

The numbers are reported in percentage and as mean atrophy (\pm standard deviation) across all triplets described in Section 2.4.3. The mean absolute error (\pm standard deviation) of BVC is also reported.

3.4. Brain T2 lesion volume and brain volume change

We report the correlation between the annualised brain atrophy rate and the total lesion volume at baseline. As shown in Figure 7, our method had a slightly stronger linear correlation with baseline lesion volume ($r_s = -0.288$, $p < 0.05$) than PBVC-SIENA ($r_s = -0.249$, $p < 0.05$). A similar trend was observed on partial correlation controlled for age and disease duration (Table 5), with $r_s = -0.373$ ($p < 0.05$) for the deep learning model, and $r_s = -0.339$ ($p < 0.05$) for SIENA.

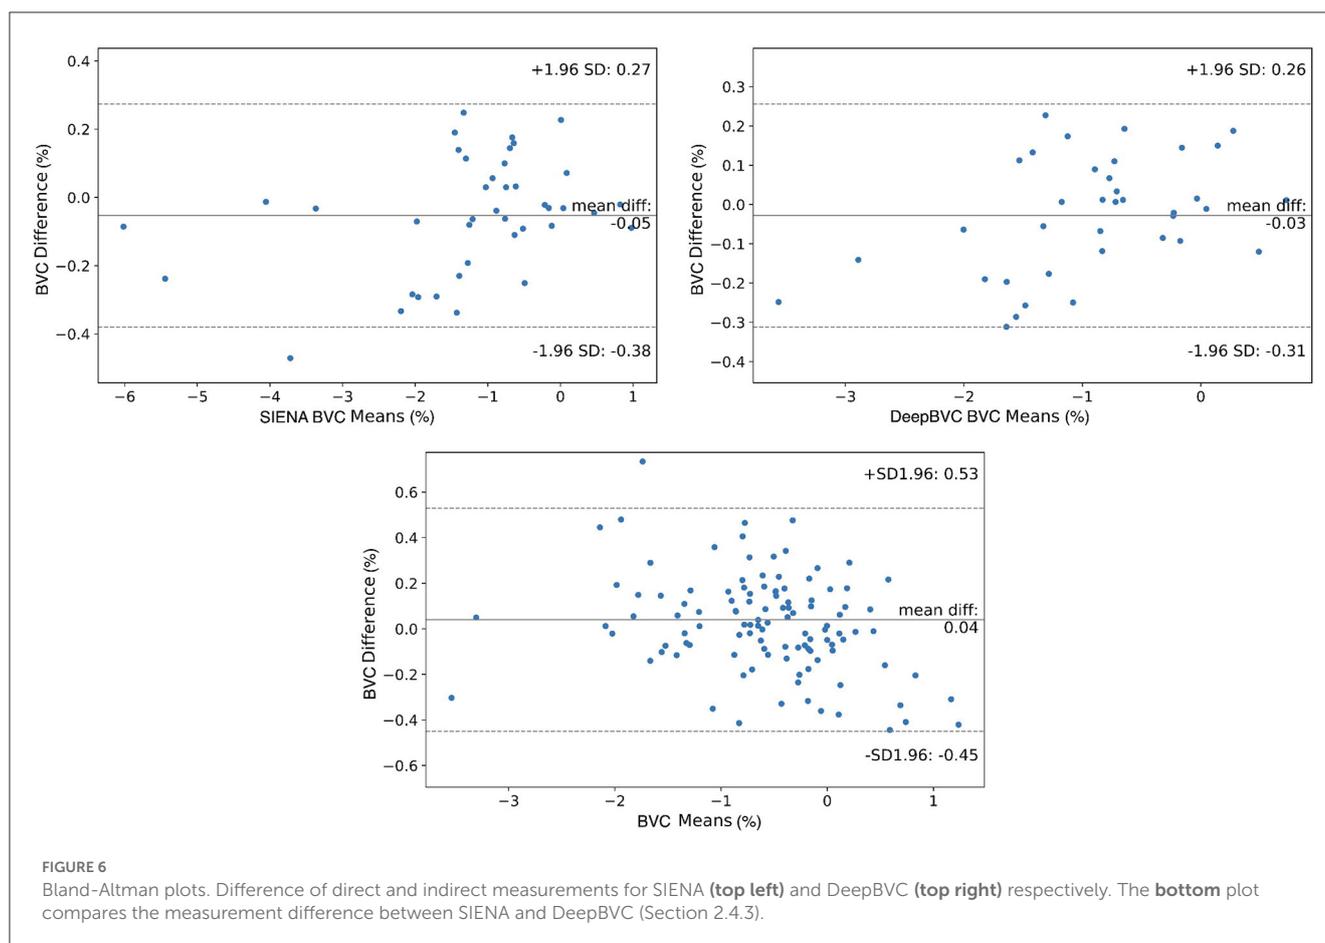
3.5. Disability and brain volume change

The annualised BVC distribution is shown in Figure 8 for subjects with and without sustained EDSS progression. For both DeepBVC and SIENA, the average annualised BVC is slightly larger for the group with sustained EDSS progression (for DeepBVC $p = 0.86$, for SIENA $p = 0.84$). The annualised BVC for matched subjects with and without sustained EDSS progression is shown in Figure 9. For both DeepBVC and SIENA, the average annualised BVC for the subjects with progression was larger (for DeepBVC $p = 0.31$, for SIENA $p = 0.35$). For both experiments, there was no significant correlation of disability progression with BVC as estimated by either method.

4. Discussion

Deep learning with pseudo labels has been previously used in the field of neuroimaging. For example, FastSurfer's deep learning algorithms were trained on outputs generated by the conventional neuroimaging pipelines that underpin FreeSurfer (Henschel et al., 2020). The reliability, sensitivity, and time efficiency of FastSurfer is proven to be superior to FreeSurfer (Fischl, 2012). While FastSurfer and DeepBVC share the concept of using pseudo-labels for training, there are two significant differences. First, the output of FastSurfer is a segmentation mask and is generated by classification, while the output of our method is produced by regression. Second, FastSurfer focuses on designing a novel deep learning network to improve efficiency, whereas our method uses a simple but effective network and regularisation technique to reduce the impact of noise in the pseudo-label.

The data augmentation and consistency regularisation of our method are only used during training; therefore the efficiency of the model is not affected by those techniques. Data augmentation



improves the model's generalisability to unseen data (Zhou et al., 2022). In our case, augmentation simulates inconsistent protocols during the acquisition of scan pairs. Therefore, the DeepBVC model is adapted to those types of protocol inconsistency and potentially generalises well to similar data. Though we only included random contrast and spacing anisotropy in training, the model demonstrated improved reliability on test-retest data with noise-related inconsistencies. Consistency regularisation renders the predictions invariant to noise applied to the input, and is widely used in semi-supervised learning and learning from noisy labels (Sajjadi et al., 2016; Clark et al., 2018; Miyato et al., 2018). Specifically, the incorporation of regularisation into our model maintains an identical brain volume change between scan pairs acquired with either a consistent protocol or a (synthetically generated) inconsistent protocol. These two techniques enable the model to learn clean predictions from noisy pseudo-labels; and improve the estimation reliability, especially for longitudinal scan pairs with an inconsistent acquisition protocol.

In general, smaller BVC on the test-retest data indicates better reliability. As in Table 3, for each of the three subjects, our method estimated a smaller test-retest brain volume difference (range 0.066–0.126%) and smaller standard deviations than SIENA (range 0.118–0.351%) across sessions ($p = 0.03, p = 0.08, p < 0.001$ for subject 1, 2, and 3, respectively). The reliability of DeepBVC also appears to be less sensitive to subject-related factors, based on the results of subjects 1 and 3 from the test-retest dataset. For

example, the mean BVC estimated by SIENA for subject 3 was greater than three times that of subject 1, whereas estimation by our method was less than two times. Based on these observations, the reliability of DeepBVC is superior to SIENA for the test-retest data with a consistent protocol. We simulated a limited array of protocol inconsistencies commonly observed in clinical practise in the test-retest dataset. Specifically, we investigated the impact of contrast variance, which may be introduced by changes in head coil or sequence parameters such as TE (Constable et al., 1992); bias field, which may relate to spatial variance in coil sensitivity and the interaction between the scanner and the subject (Kim et al., 2011); and image resolution, which is determined by scanner/sequence settings. DeepBVC demonstrated superior or at least equivalent performance when compared to SIENA in all scenarios other than inconsistency in the context of bias field. Among the four types of inconsistency tested, contrast and spacing anisotropy variance had the greatest impact on SIENA measurements, followed by Gaussian noise and bias field. DeepBVC showed significantly better performance in the context of contrast and spacing anisotropy pseudo-inconsistencies, despite the lack of spacing anisotropy variation and far less extreme contrast variation in training data augmentation, indicating generalisability of the method to unseen scenarios. Surprisingly, Gaussian noise and random bias field did not significantly impact the measurement for SIENA. DeepBVC demonstrates better reliability against most levels of Gaussian noise. However, potentially reflecting the fact that SIENA

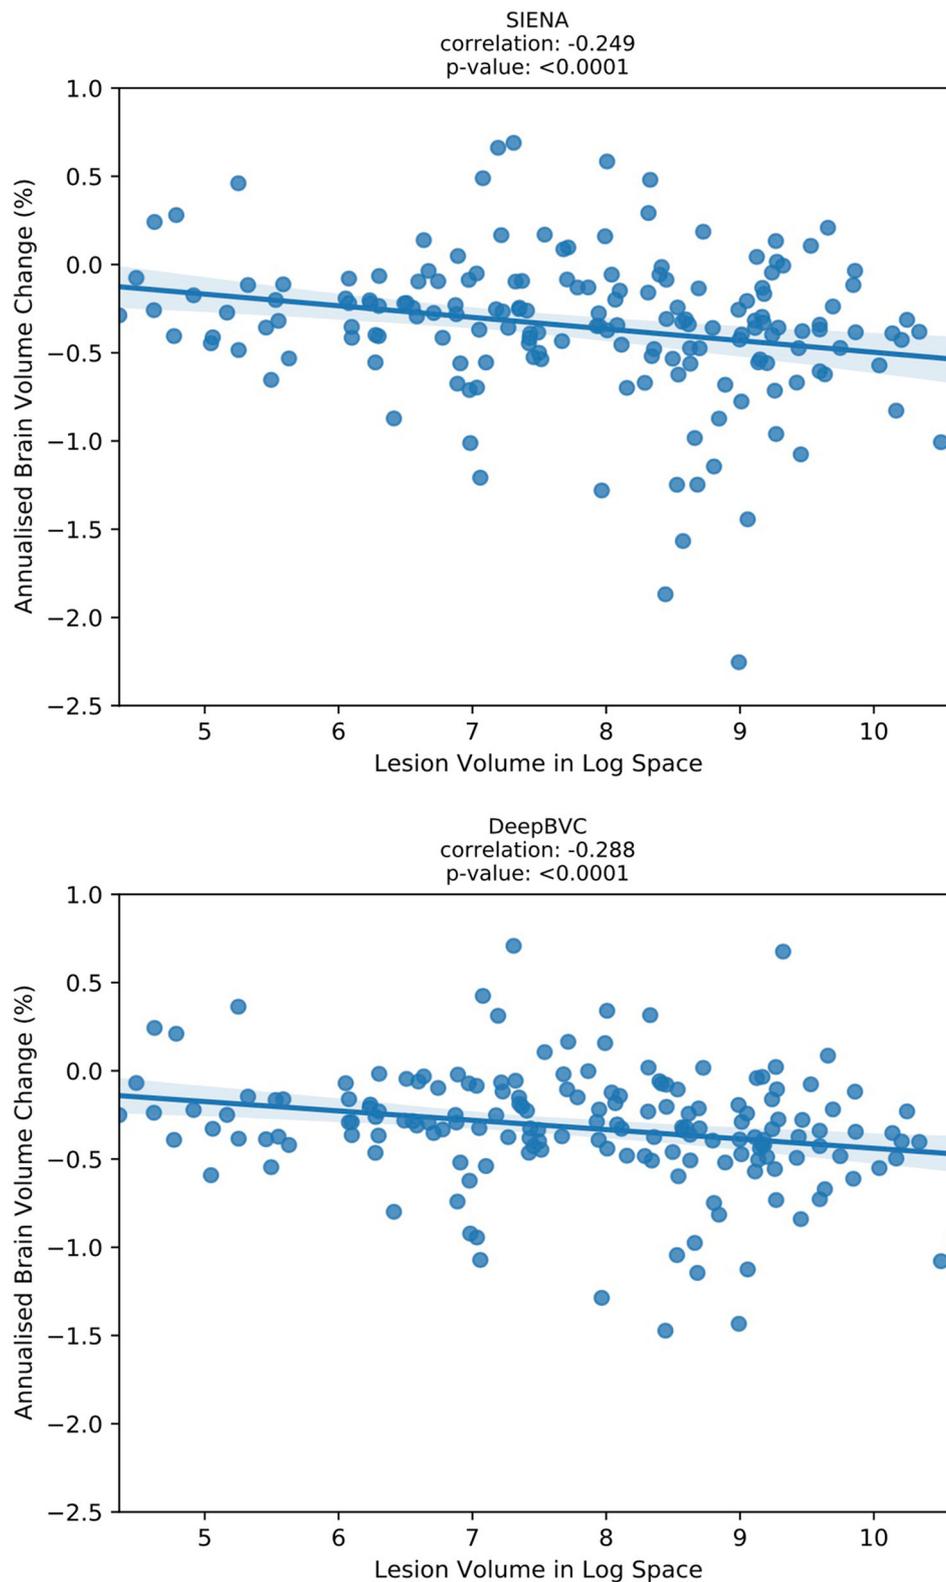


FIGURE 7

The correlation between annualised brain atrophy rate (over 3–4 years) and baseline lesion volume.

relies on edge-enhanced image profiles, which are robust to local intensity differences between images. For DeepBVC, input images are pre-processed with voxel intensity normalisation and bias

field alters the distribution of the input voxel intensities after the normalisation step. Although SIENA estimates are more robust to random bias field fluctuations, the largest median BVC for

TABLE 5 Correlation (controlled for age and disease duration) between annualised brain atrophy rate (over 3–4 years) and baseline lesion volume.

Method	Correlation r_s	p -value
SIENA	−0.339	0.005
Ours	−0.373	0.002

DeepBVC was only 0.29%, a fraction of the magnitude of error introduced into SIENA estimates by other types of inconsistency. In general, DeepBVC is therefore more reliable and less sensitive to protocol inconsistency.

Using patient data from three time points, a smaller difference between one-step and multi-step measurements indicates better consistency. The application of both SIENA and DeepBVC in this experimental paradigm yielded a small difference ($\approx 0.03\%$) between one-step ($t_0 \rightarrow t_2$) and two-step ($t_0 \rightarrow t_1$ and $t_1 \rightarrow t_2$) measurements, indicating high consistency for both methods, marginally in favour of DeepBVC. The difference between one-step and multi-step measurements can also reveal systematic errors; a smaller difference therefore indicates that less effort is required for calibration for studies that involve multiple (≥ 3) data timepoints. For both SIENA and DeepBVC, the differences between the two measurements randomly scattered above and below 0 in the Bland Altman comparison (Figure 6), suggesting that there were no significant accumulative or systematic measurement errors. For both SIENA and DeepBVC, a high agreement between the direct and indirect measurements was observed (+1.96 SD = 0.27 and −1.96 SD = −0.38 for SIENA; +1.96 SD = 0.26 and −1.96 SD = −0.31 for DeepBVC), especially for the subjects with large mean atrophy.

Correlation experiments (Figure 7) illustrated that DeepBVC estimates of BVC had a marginally stronger linear correlation with baseline brain lesion volume (DeepBVC: $r = 0.288$; SIENA: $r = 0.249$), an advantage that was retained when confounding variables (age and disease duration) were controlled (DeepBVC: $r = 0.373$; SIENA: $r = 0.339$). These findings suggest fewer subjects may be required to power group-level studies that use our tool to estimate BVC as an endpoint.

EDSS experiments demonstrate that a higher annualised BVC was observed amongst EDSS progressors compared with non-progressors group, but the difference was not significant for either DeepBVC ($p = 0.86$) or SIENA ($p = 0.84$). Similarly, analysis of propensity score matched subjects showed higher, but not significant, annualised BVC amongst EDSS progressors for both methods (DeepBVC: $p = 0.31$; SIENA: $p = 0.35$). The magnitude of group level differences in annualised BVC between EDSS progressors and non-progressors were less than in previous reports (Rudick et al., 2000; Bermel and Bakshi, 2006). Patient populations in these earlier studies differed from the modern MS cohort, in which the majority of patients are treated with high efficacy disease modifying agents (that are known to reduce BVC), studied in the present work. Additionally, these studies employed different measures to determine brain volume (change), such as brain parenchymal fraction and normalised whole brain grey-matter volume.

4.1. Study limitations and future directions

While our experiments suggest that DeepBVC more consistently and reliably estimates BVC than the classical tool, SIENA, in several scenarios, there are a number of limitations.

First, model training requires pseudo-labels from SIENA. While the use of pseudo-labels generates improved performance, the overall framework and concept follow the principles of the classical method, namely the requirements for co-registration of baseline and follow-up scans, segmentation of the brain to find edge points, and a calibration step for the final volume change estimate. As a consequence, our experimental results may be confounded by errors propagated from each pre-processing step. For example, the registration step potentially changes the scale and skew of the brain image, which can in turn affect the final BVC estimation. Additionally, lesion inpainting changes the image context and affects brain edge point segmentation, impacting the edge locations at which voxel-wise atrophy/growth is subsequently estimated. Future pipeline optimisation, in which each step is integrated as a component of the learning process, may mitigate this cascading effect and enhance the performance of deep learning based solutions for BVC.

Second, we simulated four protocol inconsistencies and independently tested their impact on BVC estimates using DeepBVC. In real-world clinical imaging acquisitions, protocol inconsistencies are more complex, overlapping and generated by different and, at times, multiple sources. Decomposing these inconsistencies into isolated categories is challenging. Similarly, the availability of inconsistently acquired scans, acquired back-to-back on the same subject, would be required to complete the test-retest study in the absence of simulated data.

Third, it is challenging to validate the actual measurement accuracy of tools such as SIENA or DeepBVC, because the ground truth BVC is unknown, other than for test-retest subjects in whom brain volume should essentially remain static when scans are acquired back to back, thereby avoiding changes in hydration or diurnal factors that could impact brain volume. Although DeepBVC did not inappropriately detect BVC in the test-retest cohort, this does not necessarily demonstrate capacity to accurately determine atrophy estimates in cases with true brain tissue volume differences. In this regard, BVC estimates from SIENA are noisy: they can be used to produce pseudo-labels for the training phase, but they cannot be used as ground truth during the validation phase, particularly without rigorous manual quality control checks. We therefore resorted to validating DeepBVC using proxy techniques (Section 2.4) in the absence of ‘clean’ labels for testing model performance. While these methods were comprehensive and approximated real world imaging scenarios, multiple time-consuming steps in the experimental pipeline hindered rapid model development.

Lastly, the limitations of this study also include the lack of a “true” multi-centre dataset. While the study integrated data from three different scanners, as depicted in Table 2, all images were obtained from a single MS clinic. This approach, to a certain extent, normalised the scan quality while prescribing the MRI exams. Future investigations may find value in validating DeepBVC using a variety of scanners from different clinical centers.

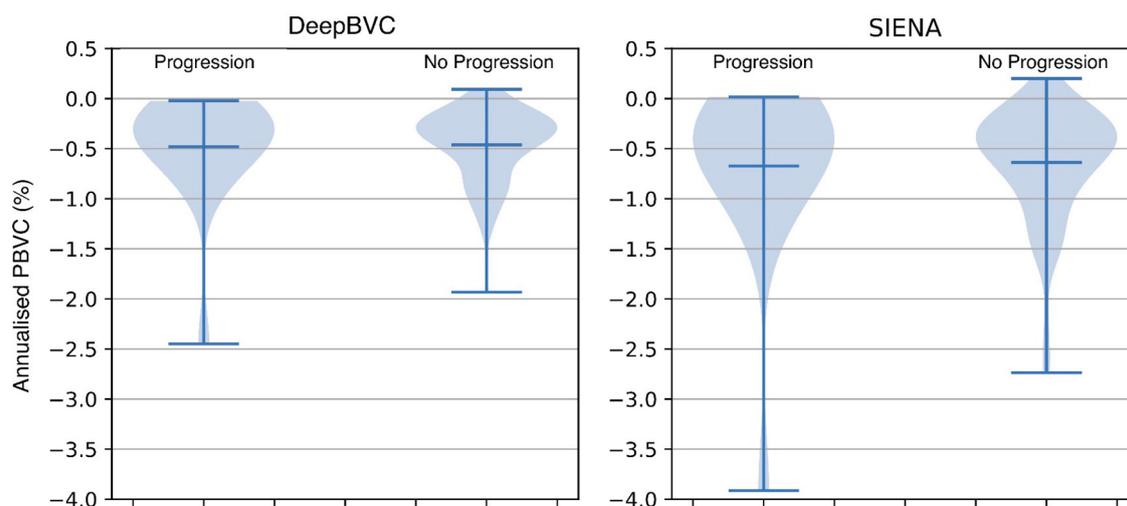


FIGURE 8

The violin plots of annualised BVC for two groups: sustained EDSS progression and no sustained EDSS progression.

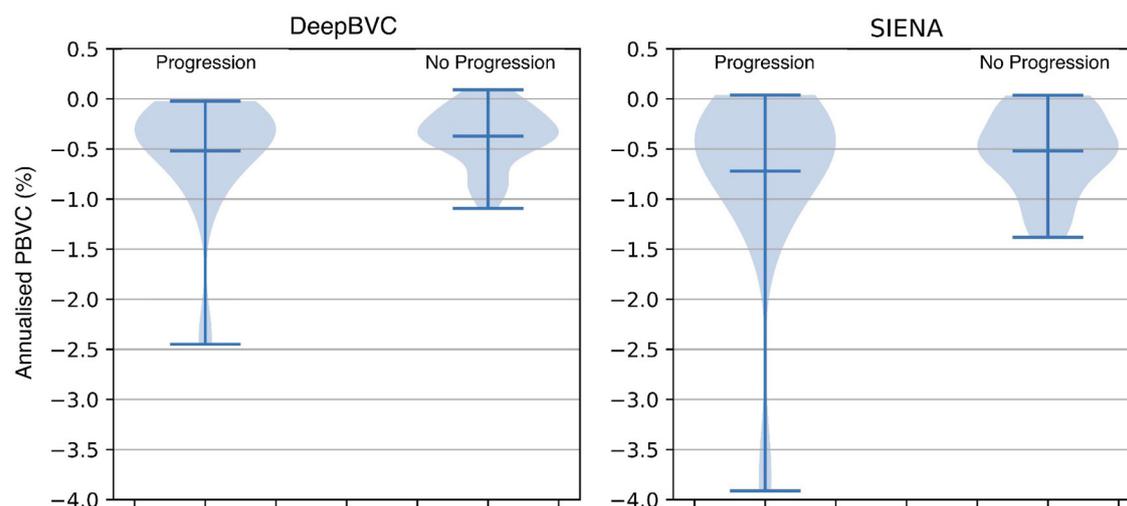


FIGURE 9

The violin plots of annualised BVC for matched subjects from two groups: sustained EDSS progression and no sustained EDSS progression. The matching is one-to-multiple propensity score matching on age and disease duration.

5. Conclusions

We demonstrate that a deep learning model trained for whole BVC estimation with pseudo-labels derived from SIENA can achieve better performance in terms of consistency, invariance to protocol change, and correlation between BVC and baseline lesion volume in a cohort of subjects with MS. Brain atrophy is a common endpoint in MS clinical trials that will become more relevant as neuroprotective and pro-reparative therapies are developed. Similarly, there is a need for robust monitoring of brain atrophy in neurodegenerative disease, the imperative

for which has been heightened by the recent advent of disease modifying therapies in this patient population. DeepBVC is a fast and robust method for estimating brain atrophy that may have particular application in both clinical trials and precision medicine.

Data availability statement

The in-house datasets presented in this article are availability upon request in written and relevant institutional approval.

Ethics statement

The studies involving human participants were reviewed and approved by the University of Sydney. The patients/participants provided their written informed consent to participate in this study.

Author contributions

GZ was responsible for processing data, conducting experiments, and drafting the paper. DW was responsible for processing data, conducting experiments, and revising the paper. MC was responsible for discussing the idea and implementation and revising the paper. LB and WO were responsible for discussing the idea and experimental design. KK was responsible for preparing and processing the data. MB and CW were responsible for discussing the idea and revising the paper. All authors contributed to the article and approved the submitted version.

Funding

CW would like to acknowledge research support from Australia Medical Research Future Fund (MRFFAI000085),

Australia Cooperative Research Centres Projects Grants (CRCPFIVE000141), and Multiple Sclerosis Research Australia (18-0461).

Conflict of interest

GZ, DW, KK, and CW are a part-time employee at the Sydney Neuroimaging Analysis Centre. MB has received institutional support for research, speaking and/or participation in advisory boards for Biogen, Merck, Novartis, Roche, and Sanofi Genzyme, and is a research consultant to RxPx and research director for the Sydney Neuroimaging Analysis Centre.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Beer, J. C., Tustison, N. J., Cook, P. A., Davatzikos, C., Sheline, Y. I., Shinohara, R. T., et al. (2020). Longitudinal ComBat: a method for harmonizing longitudinal multi-scanner imaging data. *Neuroimage* 220, 117129. doi: 10.1016/j.neuroimage.2020.117129
- Bermel, R. A., and Bakshi, R. (2006). The measurement and clinical relevance of brain atrophy in multiple sclerosis. *Lancet Neurol.* 5, 158–170. doi: 10.1016/S1474-4422(06)70349-0
- Bermel, R. A., Sharma, J., Tjoa, C. W., Puli, S. R., and Bakshi, R. (2003). A semiautomated measure of whole-brain atrophy in multiple sclerosis. *J. Neurol. Sci.* 208, 57–65. doi: 10.1016/S0022-510X(02)00425-2
- Cadavid, D., Balcer, L., Galetta, S., Aktas, O., Ziemssen, T., Vanopdenbosch, L., et al. (2017). Safety and efficacy of opicinumab in acute optic neuritis (renew): a randomised, placebo-controlled, phase 2 trial. *Lancet Neurol.* 16, 189–199. doi: 10.1016/S1474-4422(16)30377-5
- Chen, K., Reiman, E. M., Alexander, G. E., Bandy, D., Renaut, R., Crum, W. R., et al. (2004). An automated algorithm for the computation of brain volume change from sequential MRIS using an iterative principal component analysis and its evaluation for the assessment of whole-brain atrophy rates in patients with probable Alzheimer's disease. *Neuroimage* 22, 134–143. doi: 10.1016/j.neuroimage.2004.01.002
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T., and Ronneberger, O. (2016). "3d u-net: learning dense volumetric segmentation from sparse annotation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Cham: Springer), 424–432. doi: 10.1007/978-3-319-46723-8_49
- Clark, K., Luong, M.-T., Manning, C. D., and Le, Q. (2018). "Semi-supervised sequence modeling with cross-view training," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (Brussels: Association for Computational Linguistics), 1914–1925. doi: 10.18653/v1/D18-1217
- Collins, D. L., Montagnat, J., Zijdenbos, A. P., Evans, A. C., and Arnold, D. L. (2001). "Automated estimation of brain volume in multiple sclerosis with BICCR," in *Biennial International Conference on Information Processing in Medical Imaging* (Cham: Springer), 141–147. doi: 10.1007/3-540-45729-1_12
- Constable, R., Anderson, A., Zhong, J., and Gore, J. (1992). Factors influencing contrast in fast spin-echo MR imaging. *Magn. Reson. Imaging* 10, 497–511. doi: 10.1016/0730-725X(92)90001-G
- De Stefano, N., Airas, L., Grigoriadis, N., Mattle, H. P., O'Riordan, J., Oreja-Guevara, C., et al. (2014). Clinical relevance of brain volume measures in multiple sclerosis. *CNS Drugs* 28, 147–156. doi: 10.1007/s40263-014-0140-z
- De Stefano, N., Stromillo, M. L., Giorgio, A., Bartolozzi, M. L., Battaglini, M., Baldini, M., et al. (2016). Establishing pathological cut-offs of brain atrophy rates in multiple sclerosis. *J. Neurol. Neurosurg. Psychiatry* 87, 93–99. doi: 10.1136/jnnp-2014-309903
- Dewey, B. E., Zhao, C., Reinhold, J. C., Carass, A., Fitzgerald, K. C., Sotirchos, E. S., et al. (2019). Deepharmony: a deep learning approach to contrast harmonization across scanner changes. *Magn. Reson. Imaging* 64, 160–170. doi: 10.1016/j.mri.2019.05.041
- Dewey, B. E., Zuo, L., Carass, A., He, Y., Liu, Y., Mowry, E. M., et al. (2020). "A disentangled latent space for cross-site MRI harmonization," in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Cham: Springer), 720–729. doi: 10.1007/978-3-030-59728-3_70
- Duffy, B. A., Zhang, W., Tang, H., Zhao, L., Law, M., Toga, A. W., et al. (2018). Retrospective correction of motion artifact affected structural MRI images using deep learning of simulated motion. *Med Imaging Deep Learn.*
- Dwyer, M. G., Silva, D., Bergsland, N., Horakova, D., Ramasamy, D., Durfee, J., et al. (2017). Neurological software tool for reliable atrophy measurement (neurostream) of the lateral ventricles on clinical-quality t2-flair MRI scans in multiple sclerosis. *Neuroimage Clin.* 15, 769–779. doi: 10.1016/j.nicl.2017.06.022
- Field, A. (2013). *Discovering Statistics using IBM SPSS Statistics*. London: Sage.
- Filippi, M., Rovaris, M., Inglesse, M., Barkhof, F., De Stefano, N., Smith, E. S., et al. (2004). Interferon beta-1a for brain tissue loss in patients at presentation with syndromes suggestive of multiple sclerosis: a randomised, double-blind, placebo-controlled trial. *Lancet* 364, 1489–1496. doi: 10.1016/S0140-6736(04)17271-1
- Fischl, B. (2012). Freesurfer. *Neuroimage* 62, 774–781. doi: 10.1016/j.neuroimage.2012.01.021
- Freeborough, P. A., and Fox, N. C. (1997). The boundary shift integral: an accurate and robust measure of cerebral volume changes from registered repeat MRI. *IEEE Trans. Med. Imaging* 16, 623–629. doi: 10.1109/42.640753
- Friston, K. J. (2003). "Statistical parametric mapping," in *Neuroscience Databases* (Cham: Springer), 237–250. doi: 10.1007/978-1-4615-1079-6_16

- Garcia-Dias, R., Scarpazza, C., Baecker, L., Vieira, S., Pinaya, W. H., Corvin, A., et al. (2020). Neuroharmony: a new tool for harmonizing volumetric MRI data from unseen scanners. *Neuroimage* 220, 117127. doi: 10.1016/j.neuroimage.2020.117127
- Goodfellow, I. J., Shlens, J., and Szegedy, C. (2014). Explaining and harnessing adversarial examples. *arXiv*. [preprint]. doi: 10.48550/arXiv.1412.6572
- Hajnal, J. V., Saeed, N., Oatridge, A., Williams, E. J., Young, I. R., Bydder, G. M., et al. (1995). Detection of subtle brain changes using subvoxel registration and subtraction of serial MR images. *J. Comput. Assist. Tomogr.* 19, 677–691. doi: 10.1097/00004728-199509000-00001
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). “Deep residual learning for image recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV: IEEE), 770–778. doi: 10.1109/CVPR.2016.90
- Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M., et al. (2020). Fastsurfer—a fast and accurate deep learning based neuroimaging pipeline. *Neuroimage* 219, 117012. doi: 10.1016/j.neuroimage.2020.117012
- Higaki, T., Nakamura, Y., Tatsugami, F., Nakaura, T., and Awai, K. (2019). Improvement of image quality at CT and MRI using deep learning. *Jpn. J. Radiol.* 37, 73–80. doi: 10.1007/s11604-018-0796-2
- Holland, D., Dale, A. M., and Alzheimer’s Disease Neuroimaging Initiative (2011). Nonlinear registration of longitudinal images and measurement of change in regions of interest. *Med. Image Anal.* 15, 489–497. doi: 10.1016/j.media.2011.02.005
- Horsfield, M., Rovaris, M., Rocca, M., Rossi, P., Benedict, R., Filippi, M., et al. (2003). Whole-brain atrophy in multiple sclerosis measured by two segmentation processes from various MRI sequences. *J. Neurol. Sci.* 216, 169–177. doi: 10.1016/j.jns.2003.07.003
- Jack Jr, C. R., Bernstein, M. A., Fox, N. C., Thompson, P., Alexander, G., Harvey, D., et al. C., et al. (2008). The Alzheimer’s disease neuroimaging initiative (ADNI): MRI methods. *J. Magn. Reson. Imaging* 27, 685–691. doi: 10.1002/jmri.21049
- Jacobsen, C., Hagemeyer, J., Myhr, K.-M., Nyland, H., Lode, K., Bergsland, N., et al. (2014). Brain atrophy and disability progression in multiple sclerosis patients: a 10-year follow-up study. *J. Neurol. Neurosurg. Psychiatry* 85, 1109–1115. doi: 10.1136/jnnp-2013-306906
- Jenkinson, M., Pechaud, M., Smith, S., et al. (2005). “Bet2: MR-based estimation of brain, skull and scalp surfaces,” in *Eleventh Annual Meeting of the Organization for Human Brain Mapping*, volume 17 (Toronto, ON), 167.
- Kalincik, T., Cutter, G., Spelman, T., Jokubaitis, V., Havrdova, E., Horakova, D., et al. (2015). Defining reliable disability outcomes in multiple sclerosis. *Brain* 138, 3287–3298. doi: 10.1093/brain/awv258
- Kim, K., Habas, P. A., Rajagopalan, V., Scott, J. A., Corbett-Detig, J. M., Rousseau, F., et al. (2011). Bias field inconsistency correction of motion-scattered multislice MRI for improved 3d image reconstruction. *IEEE Trans. Med. Imaging* 30, 1704–1712. doi: 10.1109/TMI.2011.2143724
- Kingma, D. P., and Ba, J. (2014). “Adam: a method for stochastic optimization,” in *International Conference on Learning Representations 2015* (San Diego).
- Learned-miller, E., and Ahammad, P. (2004). “Joint MRI bias removal using entropy minimization across images,” in *Advances in Neural Information Processing Systems 17*. Neural Information Processing Systems, NIPS 2004, December 13–18, 2004 (Vancouver, BC).
- Lee, H., Nakamura, K., Narayanan, S., Brown, R. A., Arnold, D. L., Initiative, A. D. N., et al. (2019). Estimating and accounting for the effect of MRI scanner changes on longitudinal whole-brain volume change measurements. *Neuroimage* 184, 555–565. doi: 10.1016/j.neuroimage.2018.09.062
- Lewis, E. B., and Fox, N. C. (2004). Correction of differential intensity inhomogeneity in longitudinal MR images. *Neuroimage* 23, 75–83. doi: 10.1016/j.neuroimage.2004.04.030
- Li, X., Morgan, P. S., Ashburner, J., Smith, J., and Rorden, C. (2016). The first step for neuroimaging data analysis: DICOM to NIFTI conversion. *J. Neurosci. Methods* 264, 47–56. doi: 10.1016/j.jneumeth.2016.03.001
- Liu, D., Cabezas, M., Zhan, G., Wang, D., Ly, L., Kyle, K., et al. (2022). DAMS-Net: a domain adaptive lesion segmentation framework in patients with multiple sclerosis from multiple imaging centers (p18-4.001). *Neurology* 98.
- Liu, M., Maiti, P., Thomopoulos, S., Zhu, A., Chai, Y., Kim, H., et al. (2021). “Style transfer using generative adversarial networks for multi-site MRI harmonization,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Cham: Springer), 313–322. doi: 10.1007/978-3-030-87199-4_30
- Lowekamp, B. C., Chen, D. T., Ibáñez, L., and Blezek, D. (2013). The design of simpleitk. *Front. Neuroinform.* 7, 45. doi: 10.3389/fninf.2013.00045
- Maclaren, J., Han, Z., Vos, S. B., Fischbein, N., and Bammer, R. (2014). Reliability of brain volume measurements: a test-retest dataset. *Sci. Data* 1, 1–9. doi: 10.1038/sdata.2014.37
- Medawar, E., Thieleking, R., Manuilova, I., Paerisch, M., Villringer, A., Witte, A. V., et al. (2021). Estimating the effect of a scanner upgrade on measures of grey matter structure for longitudinal designs. *PLoS ONE* 16, e0239021. doi: 10.1371/journal.pone.0239021
- Miyato, T., Maeda, S.-i., Koyama, M., and Ishii, S. (2018). Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Trans. Pattern Anal. Mach. Intell.* 41, 1979–1993. doi: 10.1109/TPAMI.2018.2858821
- Nakamura, K., Guizard, N., Fonov, V. S., Narayanan, S., Collins, D. L., Arnold, D. L., et al. (2014). Jacobian integration method increases the statistical power to measure gray matter atrophy in multiple sclerosis. *Neuroimage Clin.* 4, 10–17. doi: 10.1016/j.nicl.2013.10.015
- Popescu, V., Agosta, F., Hulst, H. E., Sluimer, I. C., Knol, D. L., Sormani, M. P., et al. (2013). Brain atrophy and lesion load predict long term disability in multiple sclerosis. *J. Neurol. Neurosurg. Psychiatry* 84, 1082–1091. doi: 10.1136/jnnp-2012-304094
- Prados, F., Cardoso, M. J., Leung, K. K., Cash, D. M., Modat, M., Fox, N. C., et al. (2015). Measuring brain atrophy with a generalized formulation of the boundary shift integral. *Neurobiol. Aging* 36, S81–S90. doi: 10.1016/j.neurobiolaging.2014.04.035
- Preboske, G. M., Gunter, J. L., Ward, C. P., and Jack Jr, C. R. (2006). Common MRI acquisition non-idealities significantly impact the output of the boundary shift integral method of measuring brain atrophy on serial MRI. *Neuroimage* 30, 1196–1202. doi: 10.1016/j.neuroimage.2005.10.049
- Reuter, M., Schmansky, N. J., Rosas, H. D., and Fischl, B. (2012). Within-subject template estimation for unbiased longitudinal image analysis. *Neuroimage* 61, 1402–1418. doi: 10.1016/j.neuroimage.2012.02.084
- Rudick, R., Fisher, E., Lee, J.-C., Simon, J., Jacobs, L., Group, M. S. C. R., et al. (1999). Use of the brain parenchymal fraction to measure whole brain atrophy in relapsing-remitting MS. *Neurology* 53, 1698–1698. doi: 10.1212/WNL.53.8.1698
- Rudick, R. A., Fisher, E., Lee, J.-C., Duda, J. T., and Simon, J. (2000). Brain atrophy in relapsing multiple sclerosis: relationship to relapses, edss, and treatment with interferon β -1a. *Mult. Scler. J.* 6, 365–372. doi: 10.1177/13524585000600601
- Sajjadi, M., Javanmardi, M., and Tasdizen, T. (2016). Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Adv. Neural Inf. Process. Syst.* 29.
- Shorten, C., and Khoshgoftar, T. M. (2019). A survey on image data augmentation for deep learning. *J. Big Data* 6, 1–48. doi: 10.1186/s40537-019-0197-0
- Sinnecker, T., Schädelin, S., Benkert, P., Ruberte, E., Amann, M., Lieb, J. M., et al. (2022). Brain atrophy measurement over a MRI scanner change in multiple sclerosis. *Neuroimage Clin.* 36, 103148. doi: 10.1016/j.nicl.2022.103148
- Smeets, D., Ribbens, A., Sima, D. M., Cambron, M., Horakova, D., Jain, S., et al. (2016). Reliable measurements of brain atrophy in individual patients with multiple sclerosis. *Brain Behav.* 6, e00518. doi: 10.1002/brb3.518
- Smith, S. M., Zhang, Y., Jenkinson, M., Chen, J., Matthews, P. M., Federico, A., et al. N. (2002). Accurate, robust, and automated longitudinal and cross-sectional brain change analysis. *Neuroimage* 17, 479–489. doi: 10.1006/nimg.2002.1040
- Takao, H., Abe, O., Hayashi, N., Kabasawa, H., and Ohtomo, K. (2010). Effects of gradient non-linearity correction and intensity non-uniformity correction in longitudinal studies using structural image evaluation using normalization of atrophy (SIENA). *J. Magn. Reson. Imaging* 32, 489–492. doi: 10.1002/jmri.22237
- Tedeschi, G., Lavorgna, L., Russo, P., Prinster, A., Dinacci, D., Savettieri, G., et al. (2005). Brain atrophy and lesion load in a large population of patients with multiple sclerosis. *Neurology* 65, 280–285. doi: 10.1212/01.wnl.0000168837.87351.1f
- Thanellas, A.-K., and Pollari, M. (2010). “Sensitivity of volumetric brain analysis to systematic and random errors,” in *2010 IEEE 23rd International Symposium on Computer-Based Medical Systems (CBMS)* (Bentley, WA: IEEE), 238–242. doi: 10.1109/CBMS.2010.6042648
- Van Leemput, K., Maes, F., Vandermeulen, D., and Suetens, P. (1999). Automated model-based tissue classification of MR images of the brain. *IEEE Trans. Med. Imaging* 18, 897–908. doi: 10.1109/42.811270
- Vemuri, P., Kholmovski, E. G., Parker, D. L., and Chapman, B. E. (2005). “Coil sensitivity estimation for optimal snr reconstruction and intensity inhomogeneity correction in phased array MR imaging,” in *Biennial International Conference on Information Processing in Medical Imaging* (Cham: Springer), 603–614. doi: 10.1007/11505730_50
- Vovk, U., Pernuš, F., and Likar, B. (2004). MRI intensity inhomogeneity correction by combining intensity and spatial information. *Phys. Med. Biol.* 49, 4119. doi: 10.1088/0031-9155/49/17/020
- Vrenken, H., Jenkinson, M., Horsfield, M., Battaglini, M., Van Schijndel, R., Rostrup, E., et al. (2013). Recommendations to improve imaging and analysis of brain lesion load and atrophy in longitudinal studies of multiple sclerosis. *J. Neurol.* 260, 2458–2471. doi: 10.1007/s00415-012-6762-5
- Wu, Y., and He, K. (2018). “Group normalization,” in *Proceedings of the European Conference on Computer Vision (ECCV)* (Cham: Springer), 3–19. doi: 10.1007/978-3-030-01261-8_1

Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. (2017). "mixup: beyond empirical risk minimization," in *International Conference on Learning Representations* (Toulon). doi: 10.48550/arXiv.1710.09412

Zhang, S., Hou, Y., Wang, B., and Song, D. (2017). Regularizing neural networks via retaining confident connections. *Entropy* 19, 313. doi: 10.3390/e19070313

Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* 20, 45–57. doi: 10.1109/42.906424

Zhou, K., Liu, Z., Qiao, Y., Xiang, T., and Loy, C. C. (2022). Domain generalization: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* 01, 1–20. doi: 10.1109/TPAMI.2022.3195549