Stochastic variational inference improves quantification of multiple timepoint arterial spin labelling perfusion MRI

Thomas F. Kirk[1,2]*, Georgia G. Kenyon[2,3], Martin S. Craig[1,2] and Michael A. Chappell[1,2]

[1]Quantified Imaging Limited, London, United Kingdom, [2]Sir Peter Mansfield Imaging Centre, School of Medicine, University of Nottingham, Nottingham, United Kingdom, [3]School of Computer and Mathematical Sciences, University of Adelaide, Adelaide, SA, Australia

Multiple-timepoint arterial spin labelling MRI is a non-invasive imaging technique that permits measurement of both cerebral blood flow and arterial transit time, the latter of which is an emerging biomarker of interest for cerebrovascular health. Quantification of arterial spin labelling data is challenging due to the low signal to noise ratio and non-linear tracer kinetics of this technique. In this work, we introduce a new quantification method called SSVB that addresses limitations in existing methods and demonstrate its performance using simulations and acquisition data. Simulations showed that the method is more accurate, particularly for estimating arterial transit time, and more robust to noise than existing techniques. On high spatial resolution data acquired at 3 T, the method produced less noisy parameter maps than the comparator method and captured greater variation in arterial transit time on a cross-sectional cohort.

KEYWORDS

perfusion, arterial transit time (ATT), arterial spin label (ASL) MRI, cerebral blood flow (CBF), quantification

Introduction

Arterial spin labelling (ASL) is a non-invasive magnetic resonance imaging (MRI) technique that permits the measurement of perfusion in the brain. Using strong magnetic gradient fields applied for a short time across the neck, inflowing blood-water is magnetically inverted shortly before an image of the brain is acquired, after which tracer kinetic modelling can be used to quantify the rate and speed of blood flow (Chappell et al., 2017; Buxton et al., 1998). By capturing the dynamics of inflowing label, multiple-timepoint ASL permits the measurement of both cerebral blood flow (CBF) and arterial transit time (ATT), the latter of which is an emerging marker of cerebrovascular health, for example associated with Alzheimer's disease and coronary artery disease (Nielsen et al., 2017; MacIntosh et al., 2015). Estimating ATT from multiple-timepoint arterial spin labelling (ASL) data is however challenging due to the non-linear signal model based on tracer kinetics and the inherently low signal-to-noise ratio (SNR) of this imaging technique (Chappell et al., 2017). This is particularly true in white matter, which has lower CBF and longer ATT than cerebral grey matter (thus reducing SNR), but is implicated in a range of neurological disorders such as schizophrenia and vascular dementia (Wright et al., 2014; Barker et al., 2014).

Existing methods for ATT estimation adopt a number of strategies to deal with this challenge. Signal-weighted delay (WD) approximates ATT using a first-order moment calculation that is simple to implement and robust to noise, though the accuracy of the

approximation is bounded by the choice of post label delays (PLD) used and it is unclear how to adapt the method to handle variable label durations, as used for example in Hadamard-encoded schemes in which multiple boluses are delivered (Dai et al., 2012; von Samson-Himmelstjerna et al., 2016). Non-linear least squares (NLLS) is widely available in software libraries but is vulnerable to noise, for which reason the data is commonly smoothed in pre-processing to increase SNR, or hard constraints on ATT may be implemented. FSL BASIL adopts a variational Bayesian (VB) approach that offers a more principled treatment of low SNR data via the use of priors, but this requires advance knowledge of the normative range of ATT which may not be available in certain populations (Chappell et al., 2023).

In this work we introduce a new method for perfusion estimation on multiple-timepoint ASL data and demonstrate its performance on both simulation and acquisition data. Compared to existing methods, Structured Stochastic Variational Bayes (SSVB) is a VB method that produces more accurate estimates, is more robust on low SNR data, and produces less noisy parameter maps.

Theory

Bayesian inference is an attractive strategy for model-fitting ASL data because it allows parameter constraints to be implemented in a principled manner (mitigating against low SNR) and quantifies the uncertainty of the resulting estimates (Chappell et al., 2009; Woolrich et al., 2009). Bayesian inference seeks to obtain the posterior distribution on model parameters conditioned on the modality specific signal model and the data that has been acquired. For all but the simplest problems, derivation of the true posterior is generally intractable due to the integrals that must be performed. Variational Bayesian (VB) techniques circumvent this problem by approximating the posterior using a simpler distribution with well-known properties, which is optimised to make the approximation as accurate as possible (Penny et al., 2003). Specifically, VB seeks to minimise the Kullback–Leibler divergence between the approximate and true posterior, which is equivalent to maximising the free energy

$$F(\theta) = \int q(\theta) \log \left(p(y|\theta) \frac{p(\theta)}{q(\theta)} \right) d\theta$$

Where θ represents the physiological parameters of interest, y the samples of data, $q(\theta)$ the approximating posterior distribution, $p(y|\theta)$ the likelihood, and $p(\theta)$ prior distribution for the parameters. Noise is explicitly modelled in the likelihood, typically as a zero-mean Gaussian distribution. Qualitatively, this expression can be seen as maximising the log-likelihood (a best fit to the data) whilst minimising the divergence of the approximating posterior from the priors which enforces the constraints on parameter estimates. Priors can be either distributional, for example a Gaussian of specified mean and variance; or spatial, which in the context of imaging encodes the belief that neighbouring voxels should not differ substantially from each other. Spatial priors have a similar effect to spatial smoothing which is commonly used in neuroimaging but have the key advantage of being able to determine the optimal amount of smoothing from the data itself rather than requiring the user to specify a value (Chappell et al., 2011; Penny et al., 2005).

FSL BASIL is an existing implementation of VB for ASL that uses an analytic formulation in which the signal model is linearised via a Taylor expansion, after which iterative update equations are derived via the calculus of variations (Chappell et al., 2009; Chappell et al., 2023). Though computationally fast, this approach has a number of drawbacks: firstly, in order to keep the derivation of update equations tractable, the parametric distributions that are used to represent the prior and posterior are generally restricted to the exponential family (e.g., Gaussian, often referred to as the 'conjugate-exponential' restriction), which may not be appropriate in a particular application. Secondly, should the analytic form of priors or noise model change, it is necessary to re-derive the update equations from scratch.

In recent years, the increasing use of neural networks has been accompanied by improvements in optimisation techniques based on fast and efficient gradient descent which have been released to the public in libraries such as TensorFlow and PyTorch (Abadi et al., 2016; Paszke et al., 2019). The availability of these alternative optimisation techniques offers a new possibility for implementing VB, namely, direct optimisation of the expression for free energy. Specifically, using a Monte Carlo sampling technique, an approximation for $F(\theta)$ may be obtained by drawing L samples of θ, the physiological parameters of interest, from the posterior $q(\theta)$, where * denotes a sample.

$$F(\theta) \approx \frac{1}{L} \sum_l^L \left[\log \left(p\left(y|\theta^{*l}\right) \right) - \log \left(\frac{q\left(\theta^{*l}\right)}{p\left(\theta^{*l}\right)} \right) \right]$$

This is the objective function for SSVB that is directly optimised via gradient descent [a more detailed derivation is given in (Chappell et al., 2020)]. The first term approximates the log-likelihood, encouraging a good model fit to the data, and the second term penalises divergence of the model parameters from the prior. In this work, a multi-variate normal distribution over CBF and ATT has been used without covariance, and zero-mean Gaussian noise is assumed (the noise variance is a parameter of the model). Under these assumptions, the log-likelihood may be expressed as

$$\log p(y|\theta) = \frac{N}{2} \log \phi - \frac{1}{2} \phi (y - M(\theta))^T (y - M(\theta))$$

where N is the number of observations (number of ASL images that have been acquired), ϕ is the noise precision (inverse variance) and the latter term represents the sum of squared differences between the model prediction for the current parameters θ and the acquired data y.

Within SSVB, a spatial prior is used to implement adaptive (data-driven) smoothing of the parameter maps, which is beneficial to mitigate against the low SNR of ASL. A first-order graph Laplacian operator is used, represented by the v x v matrix D where v is the number of voxels, off-diagonal elements are 1 if the corresponding row and column voxels are adjacent, and the diagonal elements are the negative sum of that voxel's number of direct neighbours (−6 in the case of a fully-connected voxel). The prior for a single model parameter θ_i may then be expressed as

$$\log p(\theta_i) = \frac{1}{2} \log \alpha - \frac{\alpha}{2} \theta_i^T D \theta_i$$

Where α represents the spatial precision, the weight given to the spatial prior, and is a target of the optimisation itself. This means the amount of smoothing is automatically determined from the data and need not be specified by the user, which should better preserve spatial detail and generalise well across datasets without requiring hand-tuning, as has been observed in previous work (Zhao et al., 2017; Chappell et al., 2011).

The stochastic approach offers a number of advantages over analytic VB. For example, prior and posterior distributions no longer need to be drawn from the same family which removes the conjugate-exponential restriction. Secondly, because the free energy is optimised directly, it is no longer necessary to derive update equations which permits greater flexibility in modifying priors or noise models between different applications.

The ASL model used in this work is the Buxton single-compartment pseudo-continuous ASL (PCASL) model (Buxton et al., 1998), in which the magnetisation signal arriving in a voxel may be expressed as piecewise function

$$
M(t) = \begin{cases}
0 & \text{for } t < \Delta t \\[2mm]
2M_{0a}fT_{1app}\left(1 - e^{-\frac{t - \Delta t}{T_{1app}}}\right) & \text{for } \Delta t \le t < \Delta t + \tau \\[2mm]
2M_{0a}fT_{1app}e^{-\frac{\Delta t}{T_{1b}}}e^{-\frac{t - \Delta t - \tau}{T_{1app}}}\left(1 - e^{-\frac{\tau}{T_{1app}}}\right) & \text{for } \Delta t + \tau \le t
\end{cases}
$$

Where $\frac{1}{T_{1app}} = \frac{1}{T_1} + \frac{f}{\lambda}$, t is the time since of the start of label creation (magnetic inversion of blood), f is perfusion, $\ddot{A}t$ is arterial transit time, M_{0a} is the equilibrium magnetisation of arterial blood,

T_1 and T_{1b} are the longitudinal relaxation times for tissue and blood water respectively, and τ is the label duration.

Methods and materials

Datasets

Simulation data

Two simulation datasets were used. The first of these ("grey paper") used the optimal sampling scheme proposed by Woods et al. (2024) for perfusion measurement when ATT is expected to be in the range of 0.5 to 2.5 s specifically a 2.05 s label duration, 9 PLDs of 0.200, 0.775, 0.775, 0.775, 1.800, 2.275, 2.475, 2.675, 2.800 s with 4 repeats at each PLD. The second ("HCP ASL") used the sampling scheme of the HCP Lifespan ASL dataset detailed in the following paragraph, specifically a 1.5 s label duration at 5 PLDs of 0.2, 0.7, 1.2, 1.7 and 2.2 s with 6, 6, 6, 10 and 15 repeats, respectively. For both datasets, data was simulated on an isotropic voxel grid of 5 voxels in each dimension. Ground truth CBF was held constant at 60 units whilst ATT was varied from 0.5 to 3.0 s inclusive in steps of 0.25 s. Zero-mean Gaussian noise with standard deviations (SD) of 10 to 40 inclusive in steps of 10 was added to each dataset to represent varying SNR. CBF was not varied because the PCASL signal scales linearly with CBF and would thus be equivalent to varying SNR. Noiseless signal curves for a range ground truth ATT values from each sequence are shown in Figure 1.

HCP ASL

The Human Connectome Project (HCP) Lifespan ASL dataset contains high spatial resolution multiple-timepoint ASL for over 3,000 subjects aged between 5 and 100+ years; a subset of 90 subjects (51 female, mean 31 years old, minimum 8 and maximum 88) was used in this work. This dataset was used because it used a sampling schedule specifically designed to capture ATT, should contain a wide range of
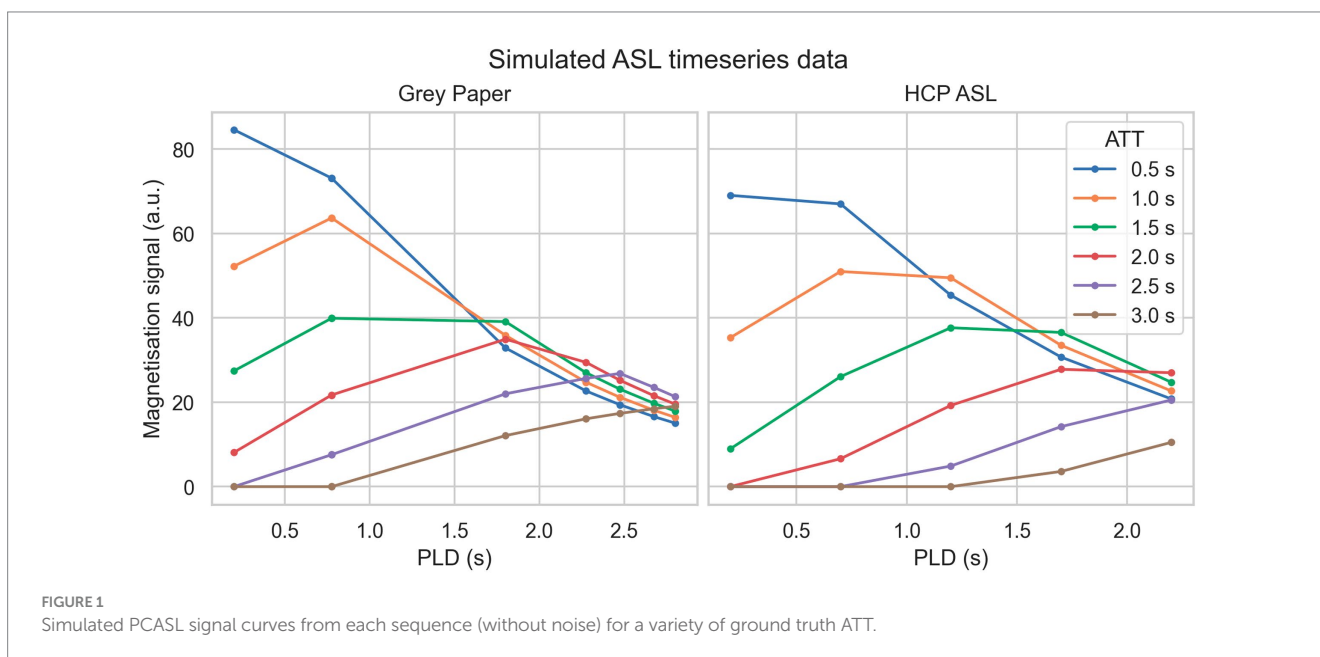

FIGURE 1
Simulated PCASL signal curves from each sequence (without noise) for a variety of ground truth ATT.

ATT [a parameter that is known to increase with age (Feron et al., 2024)] and has a relatively challenging SNR (a consequence of the high spatial resolution), both of which make it difficult to accurately estimate perfusion. ASL data was acquired with the aforementioned sampling scheme at 2.5 mm isotropic resolution on a Siemens Prisma 3 T scanner. A calibration image was acquired at the same resolution with 8 s TR. T1 anatomical scans were also acquired at 0.8 mm isotropic. Further details of the acquisition are given in (Harms et al., 2018).

HCP ASL data was pre-processed using the HCP ASL minimal pre-processing pipeline version 0.1.5.post14 (stages 0 to 6 inclusive) to correct acquisition-related geometric and intensity artefacts and register the ASL/calibration images to the anatomical images (Kirk et al., 2024). Anatomical images were pre-processed using the HCP structural minimal processing pipelines, during which structural segmentation was performed using FreeSurfer (Glasser et al., 2013; Fischl, 2012). After perfusion estimation, reference region calibration was performed using the mean M0 value of ventricular CSF (Pinto et al., 2020).
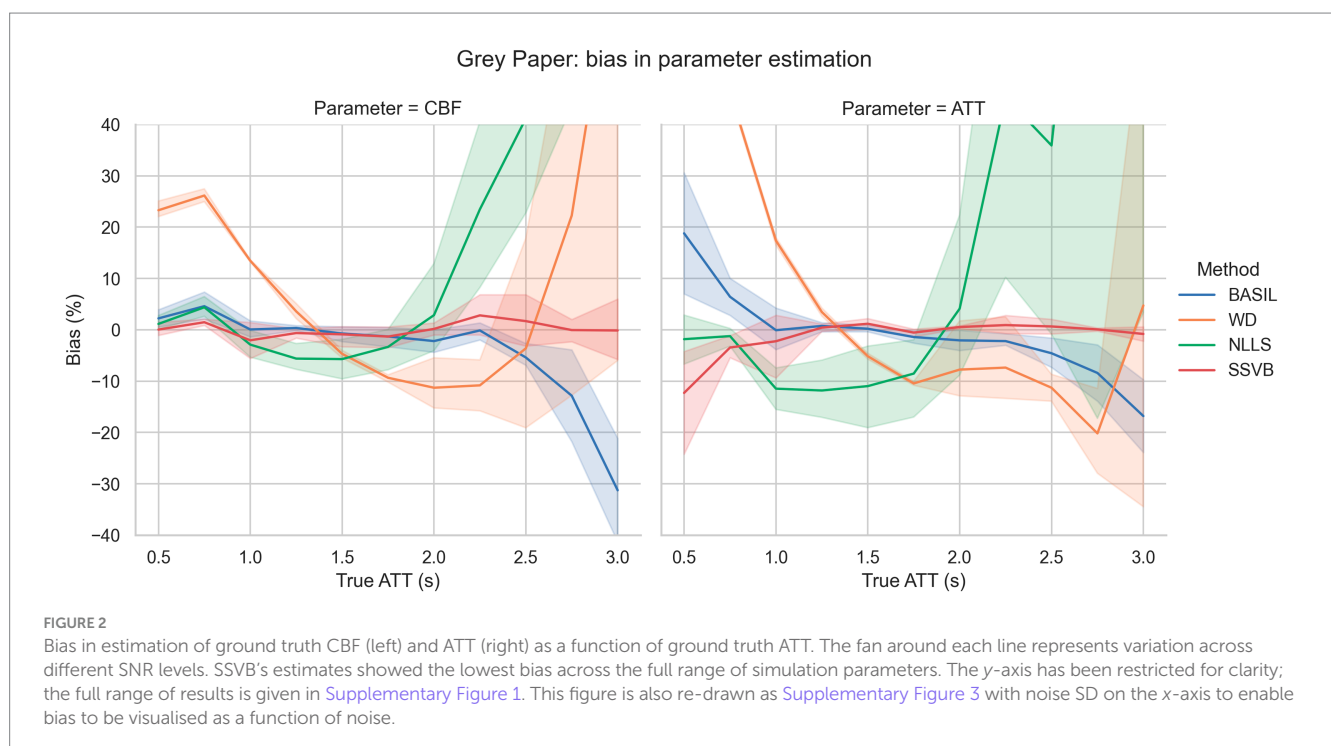
Inference methods

SSVB version 0.2.0 running on Python 3.11 using TensorFlow 2.16 was used for stochastic inference. A spatial prior was used on both CBF and ATT. RMSProp with a learning rate of 0.1 was used for optimisation (Ruder, 2017). After 50 successive steps without a decrease in cost, optimisation was reverted to the previous best state and the sample size increased by 1 from an initial starting size of 2; optimisation was run until 5 such reversions had taken place (typically around 1,000 epochs on a complete image, taking about 5 min with 8GB of RAM, or approximately twice as long as FSL BASIL).

Three comparator methods were used. FSL BASIL ("BASIL") was used for conventional analytic VB inference, with a spatial prior for CBF whereas a normal distribution prior with mean 1.3 s and standard deviation 0.5 s was used on ATT (Chappell et al., 2009; Chappell et al., 2023). Signal weighted decay ("WD") was implemented in two stages: first, using the first moment calculation to obtain ATT estimates and then providing these as fixed values to FSL BASIL to estimate CBF (Dai et al., 2012). Non-linear least squares (NLLS) was implemented using the FSL FABBER tool with –*method* = *nlls*. No constraints were implemented on either CBF or ATT, nor was any data pre-processing used (for example spatial smoothing). Like SSVB, BASIL explicitly models acquisition noise using a zero-mean Gaussian, whilst NLLS and WD make this assumption implicitly. On the basis of simulation results, only SSVB and BASIL were run on the HCP ASL acquisition data.

Evaluation methods

For the simulation datasets, the mean across voxels of CBF and ATT estimates produced by each method was calculated, and the bias with respect to ground truth expressed as a percentage of ground truth.

For the HCP ASL dataset, the FreeSurfer DKT segmentation labels were grouped into five major cortical structures (frontal, temporal, parietal, occipital and cingulate) and the median CBF/ATT estimate reported within each by resampling the respective parameter maps onto the same voxel grid using trilinear interpolation. Methodological differences in parameter estimates were analysed using paired *t*-tests. Linear regression was used to investigate the relationship between age and CBF or ATT. Statistical significance for all tests was set at $p = 0.05$ and correcting for multiple comparisons between the five cortical structures reduced the threshold to $p = 0.01$.

FIGURE 2
Bias in estimation of ground truth CBF (left) and ATT (right) as a function of ground truth ATT. The fan around each line represents variation across different SNR levels. SSVB's estimates showed the lowest bias across the full range of simulation parameters. The *y*-axis has been restricted for clarity; the full range of results is given in Supplementary Figure 1. This figure is also re-drawn as Supplementary Figure 3 with noise SD on the *x*-axis to enable bias to be visualised as a function of noise.

Results

Simulation data

Figure 2 shows bias in estimation of ground truth parameters from simulated data using the grey paper sequence. Across the full range of true ATT and SNR levels, SSVB's estimates showed consistently low bias (always less than 12% averaged across noise levels) in both CBF and ATT. The most substantial bias of −12% was observed in ATT for true ATT = 0.5 s. BASIL produced similar results to SSVB for true ATT between 1.0 s and 2.5 s but showed greater bias outside this range (19% in ATT for true ATT = 0.5 s and − 30% in CBF for true ATT = 3.0 s). WD showed large positive bias for true ATT < 1.5 s (e.g., 20% in ATT at true ATT = 1.0 s) and NLLS showed large positive bias for true ATT > 2 s (e.g., 30% in ATT at true ATT = 2 s).

FIGURE 3
Bias in estimation of ground truth CBF (left) and ATT (right) as a function of ground truth ATT. The fan around each line represents variation across different SNR levels. SSVB's estimates showed the lowest bias across the full range of simulation parameters, though performance deteriorated for true ATT > 2.5 s. The y-axis has been restricted for clarity; the full range of results is given in Supplementary Figure 2. This figure is also re-drawn as Supplementary Figure 4 with noise SD on the x-axis to enable bias to be visualised as a function of noise.

FIGURE 4
CBF maps (calibrated to ml/100 g/min) for subject HCA6007044 demonstrated the ability for high spatial resolution ASL to reveal fine anatomical detail in perfusion. There was broadly high agreement between SSVB and BASIL, though SSVB estimated slightly higher perfusion in posterior regions, visible for example in the sagittal view of the occipital lobe.
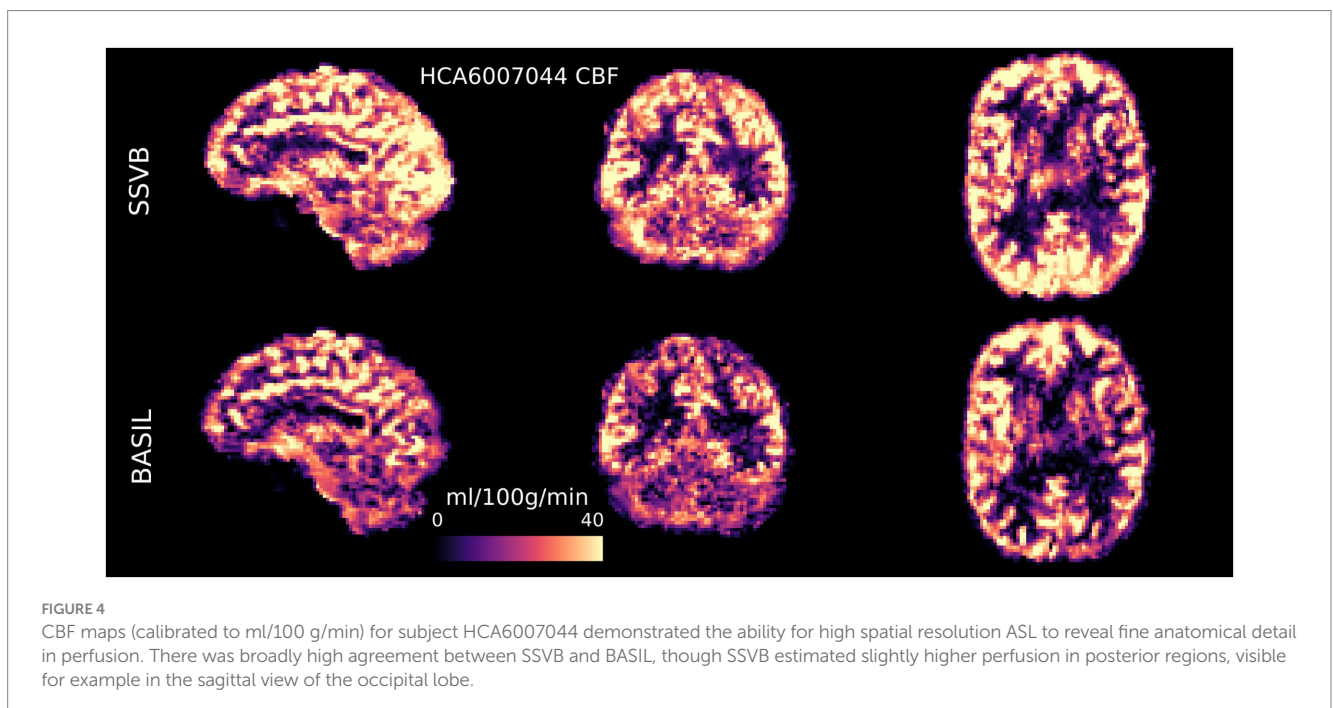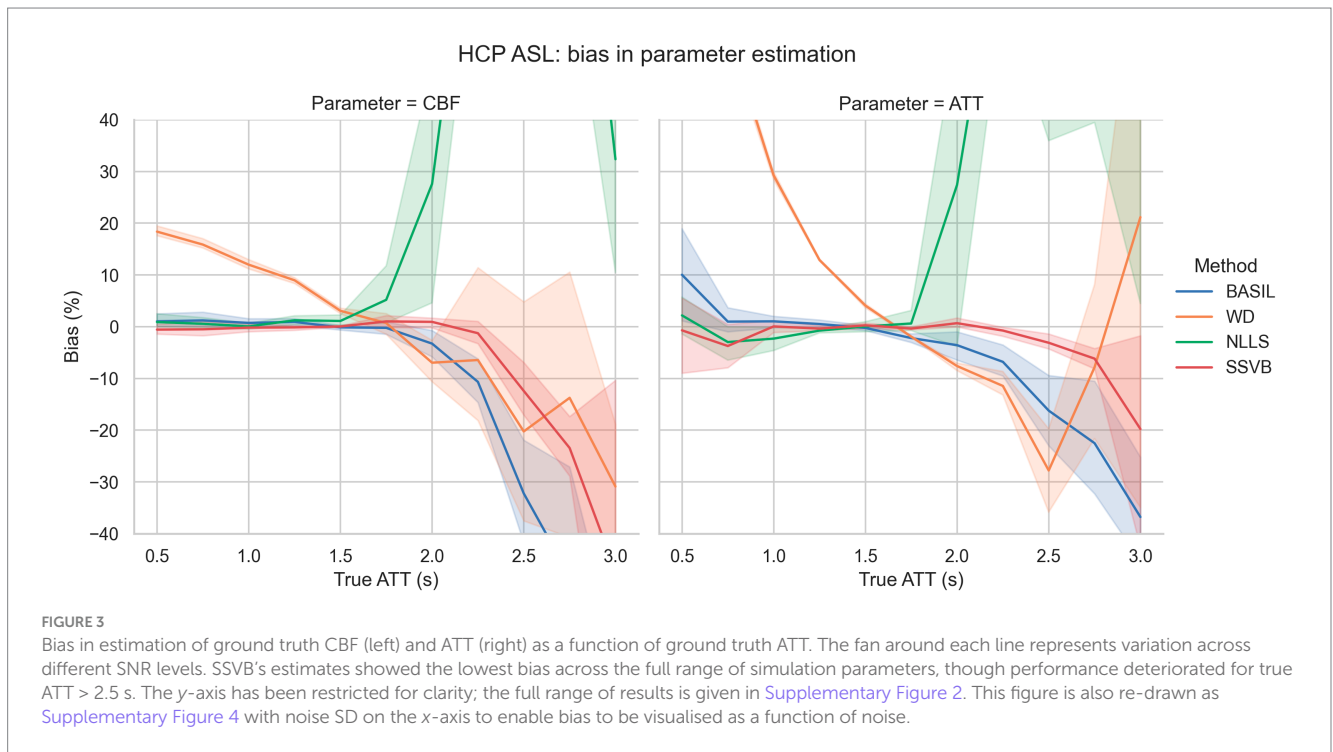
Figure 3 shows bias in estimation of ground truth parameters from simulated data using the HCP ASL sequence. For true ATT < 2.5 s, SSVB's estimates showed consistently low bias (less than 13% in both CBF and ATT); BASIL performed similarly for true

ATT < 2.0 s. WD showed large positive bias for true ATT < 1.5 s (e.g., 30% in ATT at true ATT = 1.0 s) and NLLS showed large positive bias for true ATT > 2.0 s (e.g., 25% in ATT at true ATT = 2.0 s).
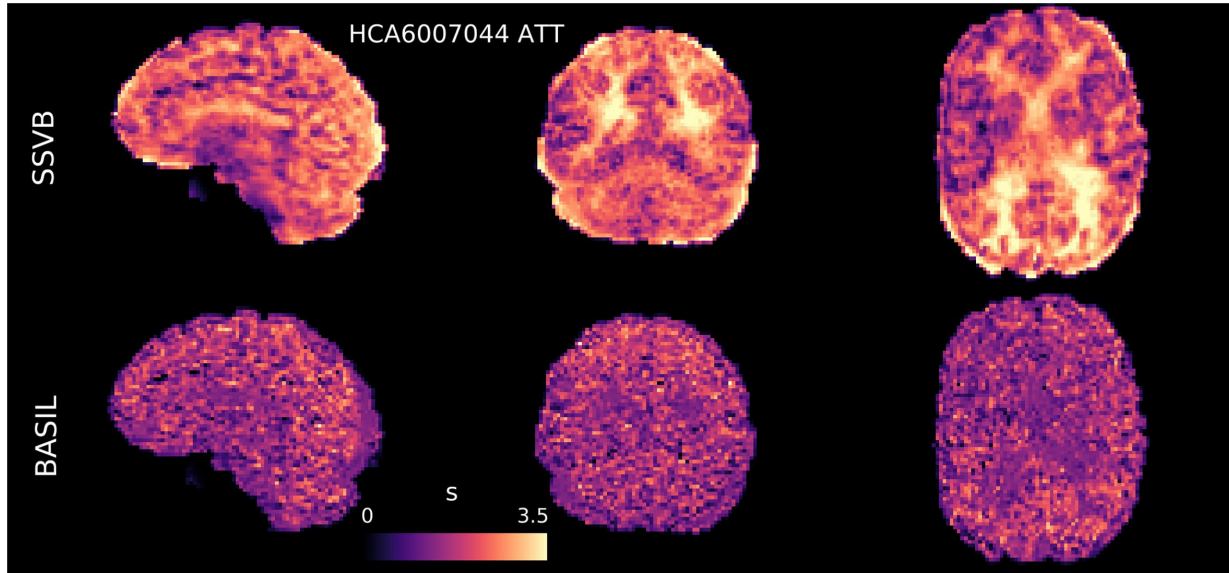

FIGURE 5
ATT maps (s) for subject HCA6007044 demonstrated substantial differences between the two methods: SSVB's estimates were both less noisy and covered a wider range of ATT values. ATT estimated by SSVB in posterior WM was substantially higher than BASIL.
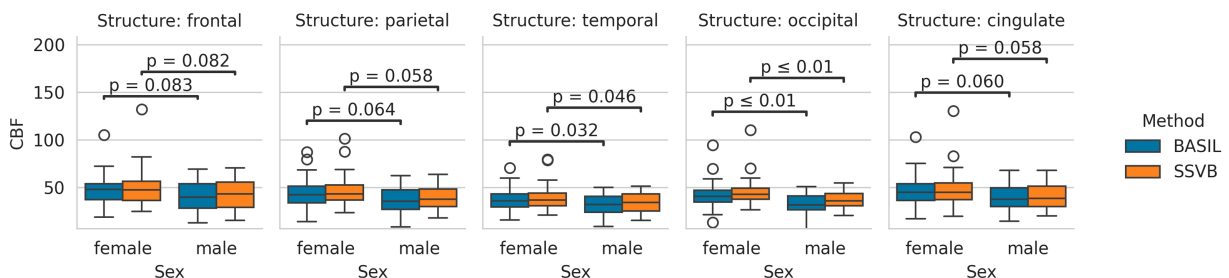

FIGURE 6
Distributions of cortical CBF estimates produced by each method, separated by sex. Both methods produced broadly similar estimates in all lobes, and both methods produced higher CBF estimates in females than males.
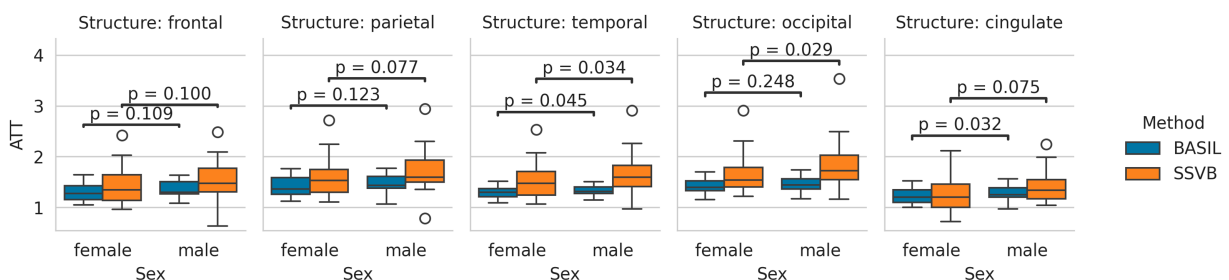

FIGURE 7
Distributions of cortical ATT estimates produced by each method, separated by sex. In all cortical lobes, the estimates produced by SSVB covered a wider range of values than BASIL. For both methods, males tended to have longer ATT than females; the difference was more pronounced for SSVB.

HCP ASL data

Figures 4, 5 show CBF and ATT maps produced by SSVB and BASIL for a single subject of the HCP ASL dataset. Whilst there was high similarity in the CBF maps (both exhibiting anatomical detail consistent with the high spatial resolution of the data), SSVB's ATT map showed a much wider range of values, with notably longer ATT in white matter (WM) or watershed regions.

Figure 6 shows the distribution of cortical CBF estimates produced by each method across the dataset. No significant differences in CBF were observed between methods (comparisons given in Supplementary Table 1). Both methods found females to have higher CBF by approximately 5–8 mL/100 g/min, though this was significant only in the occipital lobe (comparisons are given in Supplementary Table 2).
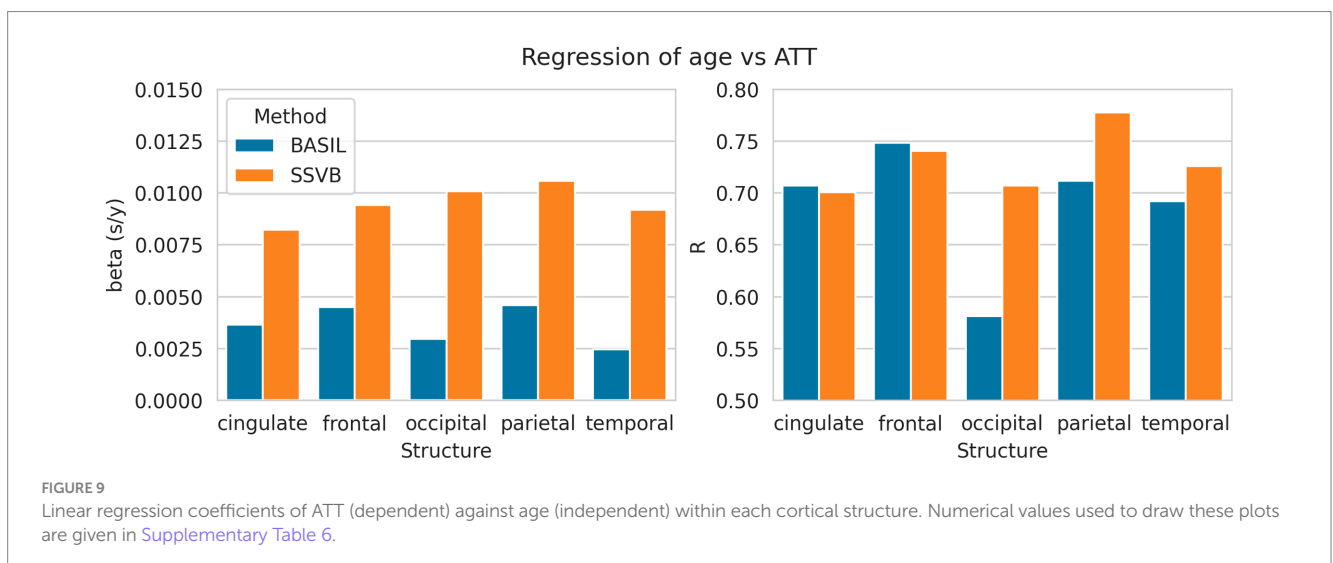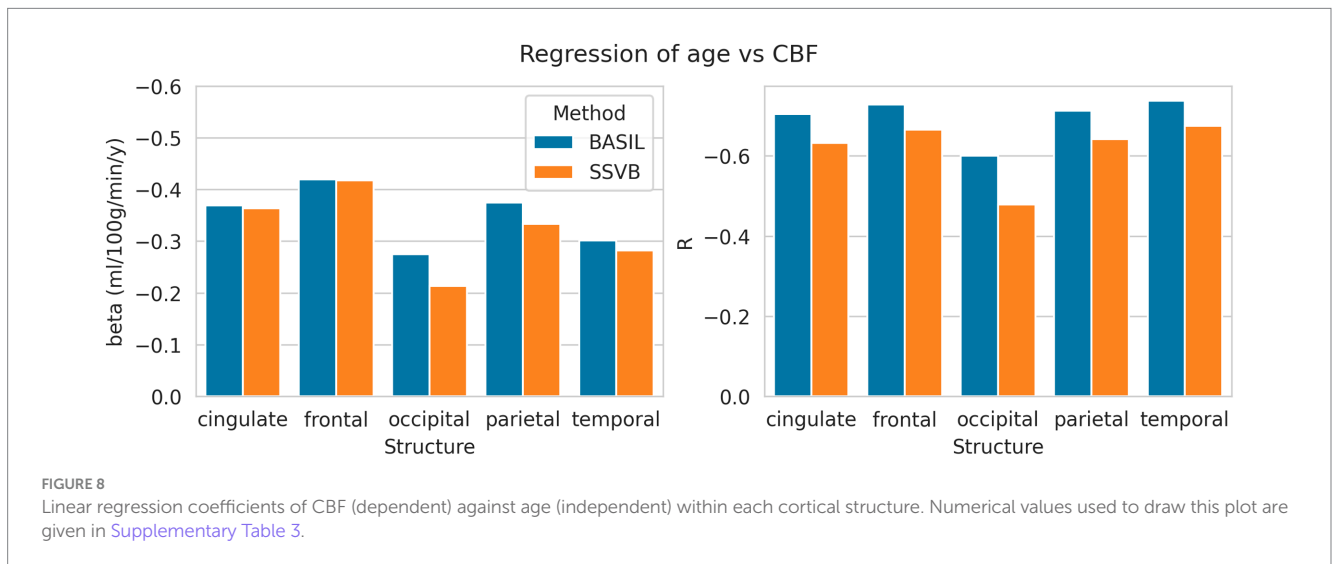
Figure 7 shows the distribution of cortical ATT estimates produced by each method. Across all cortical lobes, SSVB produced estimates covering a wider range of values than BASIL. Comparing methods within individuals, SSVB's ATT estimate was significantly higher in all lobes except the cingulate

(comparisons are given in Supplementary Table 4). For both methods, males were observed to have longer ATT than females but these differences were not significant (comparisons are given in Supplementary Table 5).

Figure 8 shows linear regression coefficients for CBF against age (all regressions were significant). In all cortical lobes, SSVB revealed a shallower (less negative) gradient and weaker correlation coefficient than BASIL, though the differences were not very substantial (for example, −0.27 vs. −0.21 in the occipital lobe). Numerical values are reported in Supplementary Table 3.

Figure 9 shows linear regression coefficients for ATT against age (all regressions were significant). In all cortical lobes, SSVB revealed a steeper relationship between age and ATT (the gradient was 2–3 times larger than for BASIL) with an approximately equal or larger correlation coefficient (particularly in the occipital, parietal and temporal lobes). Numerical values are reported in Supplementary Table 6.

Figure 10 shows the interaction between age and methodological differences for ATT and CBF, restricted to the
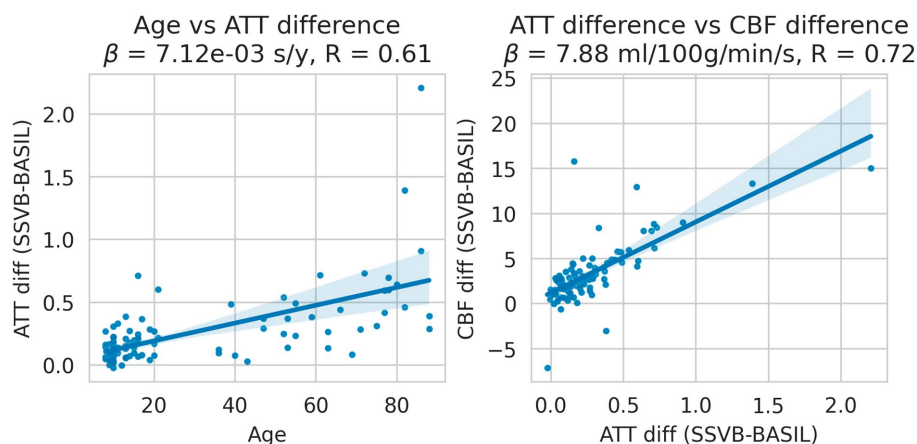

FIGURE 8
Linear regression coefficients of CBF (dependent) against age (independent) within each cortical structure. Numerical values used to draw this plot are given in Supplementary Table 3.

FIGURE 9
Linear regression coefficients of ATT (dependent) against age (independent) within each cortical structure. Numerical values used to draw these plots are given in Supplementary Table 6.

FIGURE 10
Linear regressions of methodological ATT differences (SSVB-BASIL) against age (left panel); and methodological ATT differences against CBF differences (right panel). Both panels are drawn for the occipital lobe only and both regressions were significant. ATT differences were observed to increase with age with reasonable correlation (0.61), such that SSVB estimated approximately 0.6 s higher ATT in 80-year-old subjects. These ATT differences correlated with CBF differences between the methods (R = 0.72), where are longer SSVB ATT estimated also implied a higher CBF estimate.

occipital lobe (which is most posterior and therefore typically has extended ATT). Methodological ATT differences were found to increase with age, such that SSVB estimated longer ATT in elderly subjects with a correlation of 0.61. Subject-wise ATT differences were in turn found to correlate (R = 0.72) with CBF differences such that longer ATT estimates were associated with higher CBF estimates.

Discussion

Simulation results showed that SSVB was the most accurate and robust inference method over the range of simulation parameters investigated (Figures 2, 3). Particularly notable was the extra accuracy for ground truth ATT > 2.0 s, and consistency across SNR levels. These results indicate that SSVB is better suited to subjects with long ATT and low CBF (equivalent to low SNR), which is often the case for elderly cohorts or in neurodegenerative disease.

Results from HCP ASL data demonstrated the benefits of SSVB's capabilities in ATT estimation. SSVB's ATT maps were visibly less noisy than BASIL's (Figure 5) and revealed expected anatomical features such as watershed regions or elongated ATT in WM. In agreement with previously reported results, both SSVB and BASIL's ATT estimates increased with age, though the slope of the relationship was 2–3 times greater with SSVB than BASIL (Figure 9). Coupled with the smaller regression slope observed for SSVB between CBF and age (Figure 8), and the mapping of methodological ATT differences to CBF differences (Figure 10), this suggests that the extent of CBF decrease with age which has been previously reported is at risk of being overestimated if ATT is not accurately estimated.

Addressing between-method differences specifically, differences in ATT were most pronounced in elderly HCP ASL subjects (Figure 10),

which is consistent with prior knowledge that ATT is elongated in elderly subjects, and consistent with the simulation results which revealed greater methodological differences for true ATT > 2 s (Figure 3). Potential sex differences in ATT were observed (Figure 7; Supplementary Table 4), though the number of subjects in the dataset was insufficient to reach statistical significance. If confirmed on a larger dataset, these would be an interesting addition to previously-reported sex differences in CBF (Leidhin et al., 2021).

Finally, the visual quality of the HCP ASL parameter maps demonstrated the potential of high spatial resolution ASL. Although ASL is often acquired at low spatial resolution to mitigate low SNR, the quality of SSVB's maps showed that sophisticated perfusion estimation methods can "look through" noise to recover anatomical details that would otherwise be lost at low spatial resolution. This observation is consistent with the robust performance of SSVB on low SNR simulation data, and echoes previous work investigating the feasibility of high-resolution ASL (Kashyap et al., 2024).

Future work could investigate whether the extra flexibility afforded by SSVB can be exploited advantageously. For example, though MRI acquisition noise is routinely approximated as a zero-mean Gaussian distribution, a Rician distribution is theoretically more appropriate, and could be implemented in SSVB.

Conclusion

SSVB is a new method for estimating perfusion from multiple-timepoint ASL data that offers increased accuracy and robustness compared to existing methods. Simulations show that the gains in performance over existing methods are greatest for true ATT > 2 s, which is pertinent for the study of elderly or potentially diseased populations. When run on high spatial resolution HCP ASL data,

SSVB produced ATT maps with low noise and expected anatomical features such as watershed regions.

Ethics statement

Ethical approval was not required for the study involving humans in accordance with the local legislation and institutional requirements. Written informed consent to participate in this study was not required from the participants or the participants' legal guardians/next of kin in accordance with the national legislation and the institutional requirements.

Author contributions

TK: Investigation, Software, Writing – original draft. GK: Investigation, Writing – review & editing. MCr: Software, Writing – review & editing. MCh: Conceptualization, Funding acquisition, Writing – review & editing.

Conflict of interest

TK, MCr, MCh were employed by Quantified Imaging Limited.

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

Generative AI statement

The authors declare that no generative AI was used in the creation of this manuscript.

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.

Supplementary material

The Supplementary material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fnins.2025.1536752/full#supplementary-material

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., et al. (2016). Tensorflow: large-scale machine learning on heterogeneous distributed systems. *arXiv*.

Barker, R., Ashby, E. L., Wellington, D., Barrow, V. M., Palmer, J. C., Kehoe, P. G., et al. (2014). Pathophysiology of White matter perfusion in Alzheimer's disease and vascular dementia. *Brain* 137, 1524–1532. doi: 10.1093/brain/awu040

Buxton, R. B., Frank, L. R., Wong, E. C., Siewert, B., Warach, S., and Edelman, R. R. (1998). A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magn. Reson. Med.* 40, 383–396. doi: 10.1002/mrm.1910400308

Chappell, M. A., Craig, M. S., and Woolrich, M. W. (2020). Stochastic Variational Bayesian inference for a nonlinear forward model. *arXiv*.

Chappell, M. A., Groves, A. R., MacIntosh, B. J., Donahue, M. J., Jezzard, P., and Woolrich, M. W. (2011). Partial volume correction of multiple inversion time arterial spin labeling MRI data. *Magn. Reson. Med.* 65, 1173–1183. doi: 10.1002/mrm.22641

Chappell, M. A., Groves, A. R., Whitcher, B., and Woolrich, M. W. (2009). Variational Bayesian inference for a nonlinear forward model. *IEEE Trans. Signal Process.* 57, 223–236. doi: 10.1109/TSP.2008.2005752

Chappell, M., Jenkinson, M., and Beckmann, C. (2017). Introduction to perfusion quantification using arterial spin labeling. Oxford, UK: Oxford University Press.

Chappell, M. A., Kirk, T. F., Craig, M. S., Kennedy, F. A., McConnell, M. Y., Zhao, B. J., et al. (2023). BASIL: a toolbox for perfusion quantification using arterial spin labelling. *Imaging Neurosci.* 1, 1–16. doi: 10.1162/imag_a_00041

Dai, W., Robson, P. M., Shankaranarayanan, A., and Alsop, D. C. (2012). Reduced resolution transit delay Prescan for quantitative continuous arterial spin labeling perfusion imaging. *Magn. Reson. Med.* 67, 1252–1265. doi: 10.1002/mrm.23103

Feron, J., Segaert, K., Rahman, F., Fosstveit, S. H., Joyce, K. E., Lohne-Seiler, H., et al. (2024). Determinants of cerebral blood flow and arterial transit time in healthy older adults. *Aging (Albany NY)* 16, 12473–12497. doi: 10.18632/aging.206112

Fischl, B. (2012). FreeSurfer. *NeuroImage* 62, 774–781. doi: 10.1016/j.neuroimage.2012.01.021

Glasser, M. F., Sotiropoulos, S. N., Anthony Wilson, J., Coalson, T. S., Fischl, B., Andersson, J. L., et al. (2013). The minimal preprocessing pipelines for the human connectome project. *NeuroImage* 80, 105–124. doi: 10.1016/j.neuroimage.2013.04.127

Harms, M. P., Somerville, L. H., Ances, B. M., Andersson, J., Barch, D. M., Bastiani, M., et al. (2018). Extending the human connectome project across ages: imaging protocols for the lifespan development and aging projects. *NeuroImage* 183, 972–984. doi: 10.1016/j.neuroimage.2018.09.060

Kashyap, S., Oliveira, Í. A. F., and Uludağ, K. (2024). Feasibility of high-resolution perfusion imaging using arterial spin labeling MRI at 3 tesla. *Front. Physiol.* 14:1271254. doi: 10.3389/fphys.2023.1271254

Kirk, T. F., Kennedy McConnell, F. A., Toner, J., Craig, M. S., Carone, D., Li, X., et al. (2024). Arterial spin labelling perfusion MRI analysis for the human connectome project lifespan ageing and development studies. *Imaging Neurosci.* doi: 10.1101/2024.09.19.613907

Leidhin, C. N., McMorrow, J., Carey, D., Newman, L., Williamson, W., Fagan, A. J., et al. (2021). Age-related normative changes in cerebral perfusion: data from the Irish longitudinal study on ageing (TILDA). *NeuroImage* 229:117741. doi: 10.1016/j.neuroimage.2021.117741

MacIntosh, B. J., Swardfager, W., Robertson, A. D., Tchistiakova, E., Saleem, M., Oh, P. I., et al. (2015). Regional cerebral arterial transit time hemodynamics correlate with vascular risk factors and cognitive function in men with coronary artery disease. *Am. J. Neuroradiol.* 36, 295–301. doi: 10.3174/ajnr.A4094

Nielsen, R. B., Egefjord, L., Angleys, H., Mouridsen, K., Gejl, M., Møller, A., et al. (2017). Capillary dysfunction is associated with symptom severity and neurodegeneration in Alzheimer's disease. *Alzheimers Dement.* 13, 1143–1153. doi: 10.1016/j.jalz.2017.02.007

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). PyTorch: an imperative style, High-Performance Deep Learning Library. *arXiv.* doi: 10.48550/arXiv.1912.01703

Penny, W., Kiebel, S., and Friston, K. (2003). Variational Bayesian inference for fMRI time series. *NeuroImage* 19, 727–741. doi: 10.1016/S1053-8119(03)00071-5

Penny, W. D., Trujillo-Barreto, N. J., and Friston, K. J. (2005). Bayesian fMRI time series analysis with spatial priors. *NeuroImage* 24, 350–362. doi: 10.1016/j.neuroimage.2004.08.034

Pinto, J., Chappell, M. A., Okell, T. W., Mezue, M., Segerdahl, A. R., Tracey, I., et al. (2020). Calibration of arterial spin labeling data—potential pitfalls in post-processing. *Magn. Reson. Med.* 83, 1222–1234. doi: 10.1002/mrm.28000

Ruder, S. (2017). An overview of gradient descent optimization algorithms. *arXiv.* doi: 10.48550/arXiv.1609.04747

von Samson-Himmelstjerna, F., Madai, V. I., Sobesky, J., and Guenther, M. (2016). Walsh-ordered Hadamard time-encoded Pseudocontinuous ASL (WH pCASL). *Magn. Reson. Med.* 76, 1814–1824. doi: 10.1002/mrm.26078

Woods, J. G., Achten, E., Asllani, I., Bolar, D. S., Dai, W., Detre, J. A., et al. (2024). Recommendations for quantitative cerebral perfusion MRI using multi-Timepoint arterial spin labeling: acquisition, quantification, and clinical applications. *Magn. Reson. Med.* 92, 469–495. doi: 10.1002/mrm.30091

Woolrich, M. W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., et al. (2009). Bayesian analysis of neuroimaging data in FSL. *NeuroImage* 45, S173–S186. doi: 10.1016/j.neuroimage.2008.10.055

Wright, S. N., Kochunov, P., Chiappelli, J., McMahon, R. P., Florian Muellerklein, S., Wijtenburg, A., et al. (2014). Accelerated White matter aging in schizophrenia: role of White matter blood perfusion. *Neurobiol. Aging* 35, 2411–2418. doi: 10.1016/j.neurobiolaging.2014.02.016

Zhao, M. Y., Mezue, M., Segerdahl, A. R., Okell, T. W., Tracey, I., Xiao, Y., et al. (2017). A systematic study of the sensitivity of partial volume correction methods for the quantification of perfusion from Pseudo-continuous arterial spin labeling MRI. *NeuroImage* 162, 384–397. doi: 10.1016/J.NEUROIMAGE.2017.08.072