

RESEARCH ARTICLE

Subject-specific optimization of background suppression for arterial spin labeling magnetic resonance imaging using a feedback loop on the scanner

Kirsten Koolstra¹  | Marius Staring¹  | Paul de Bruin² | Matthias J. P. van Osch³ ¹Radiology, Division of Image Processing, Leiden University Medical Center, Leiden, The Netherlands²Philips, Best, The Netherlands³Radiology, C. J. Gorter MRI Center, Leiden University Medical Center, Leiden, The Netherlands**Correspondence**

Kirsten Koolstra, Leiden University Medical Center, Albinusdreef 2, 2333 ZA, Leiden, The Netherlands.

Email: k.koolstra@lumc.nl**Funding information**

HTSM; Innovational Research Incentives Scheme VICI, Grant/Award Number: 016.160.351; Universiteit Leiden

Background suppression (BGS) in arterial spin labeling (ASL) magnetic resonance imaging leads to a higher temporal signal-to-noise ratio (tSNR) of the perfusion images compared with ASL without BGS. The performance of the BGS, however, depends on the tissue relaxation times and on inhomogeneities of the scanner's magnetic fields, which differ between subjects and are unknown at the moment of scanning. Therefore, we developed a feedback loop (FBL) mechanism that optimizes the BGS for each subject in the scanner during acquisition. We implemented the FBL for 2D pseudo-continuous ASL scans with an echo-planar imaging readout. After each dynamic scan, the acquired ASL images were automatically sent to an external computer and processed with a Python processing tool. Inversion times were optimized on the fly using 80 iterations of the Nelder–Mead method, by minimizing the signal intensity in the label image while maximizing the signal intensity in the perfusion image. The performance of this method was first tested in a four-component phantom. The regularization parameter was then tuned in six healthy subjects (three males, three females, age 24–62 years) and set as $\lambda = 4$ for all other experiments. The resulting ASL images, perfusion images, and tSNR maps obtained from the last 20 iterations of the FBL scan were compared with those obtained without BGS and with standard BGS in 12 healthy volunteers (five males, seven females, age 24–62 years) (including the six volunteers used for tuning of λ). The FBL resulted in perfusion images with a statistically significantly higher tSNR (2.20) compared with standard BGS (1.96) ($p < 5 \times 10^{-3}$, two-sided paired t-test). Minimizing signal in the label image furthermore resulted in control images, from which approximate changes in perfusion signal can directly be appreciated. This could be relevant to ASL applications that require a high temporal resolution. Future work is needed to minimize the

Abbreviations used: ASL, arterial spin labeling; BGS, background suppression; EPI, echo-planar imaging; FBL, feedback loop; FLAIR, fluid-attenuated inversion recovery; TSE, turbo-spin-echo; tSNR, temporal signal-to-noise ratio; XTC, eXTernal Control.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2022 The Authors. *NMR in Biomedicine* published by John Wiley & Sons Ltd.

number of initial acquisitions during which the performance of BGS is reduced compared with standard BGS, and to extend the technique to 3D ASL.

KEYWORDS

adaptive MRI, ASL, background suppression, feedback loop, Nelder–Mead, protocol optimization

1 | INTRODUCTION

Suppression of background signal in arterial spin labeling (ASL) leads to perfusion images with a higher temporal signal-to-noise ratio (tSNR) compared with ASL without background suppression (BGS),¹ and is therefore commonly used in current ASL protocols.^{2–4} Using BGS is especially important for applications in which the ASL signal is subject to large temporal fluctuations. This is the case for most 3D ASL scans, where temporal fluctuations are introduced by the multishot nature of the acquisition process,^{5,6} as well as in abdominal imaging, where motion is much more dominant than in the brain.^{7,8}

BGS is obtained by applying multiple inversion pulses before and during the postlabel delay (PLD). If the pulses are applied at the right timings, the reduction in background signal can be of the order of 99%.⁹ However, the optimal inversion times, and therefore the quality of the BGS, depend on the T_1 relaxation times of the underlying tissue and on inhomogeneities of the scanner's magnetic fields (B_0 , B_1^+). These parameters differ between subjects¹⁰ and can therefore result in intersubject variability of image quality. Furthermore, close to optimal BGS settings, which can be obtained from an offline optimization step,⁹ may result in perfusion quantification errors as a result of magnitude subtraction artifacts. Nevertheless, current ASL protocols rely on one set of predefined inversion times for all subjects, primarily because the interscan variations are not known at the moment of scanning. This means that the quality of the resulting perfusion images is not optimal for all subjects.

In this work, we designed and implemented a feedback mechanism that optimized the quality of single-slice ASL perfusion images on the scanner during acquisition. Specifically, we optimized the timings of four BGS pulses by minimizing the signal intensity in the label image, while maximizing the perfusion signal. By doing this, we aimed to improve the performance of the BGS and therefore the tSNR of the perfusion images. We validated the BGS performance of the designed feedback loop (FBL) with a phantom experiment, in which we also tested the precision of our approach regarding inversion time updates. We then acquired the FBL scan in 12 healthy volunteers, comparing the individually optimized BGS with standard BGS and with an ASL scan without BGS, using an experimentally tuned regularization parameter. We finally demonstrated how such a FBL can be used to track the neuronal activation in stimulus scans that require a high temporal resolution.

2 | METHODS

2.1 | Feedback mechanism

Directly after each dynamic scan, the acquired label and control images were sent to an external computer via the remote connection software eXTernal Control (XTC) (Philips, The Netherlands).¹¹ We used an inhouse-developed Python (Python Software Foundation, <https://www.python.org>) processing framework that ran on this external computer and received and processed the images during acquisition. The updated inversion times were sent back to the scanner computer, where we implemented a parameter import functionality that reads the updates during scanning, such that the new inversion times were taken into account at 10^{-3} ms precision for the next dynamic scan. This closed the FBL (Figure 1A).

2.2 | Optimization of inversion times

Four inversion times in the ASL sequence were optimized during acquisition, such that the signal in the label image was minimized while maximizing the perfusion signal to avoid magnitude subtraction errors for near-optimal BGS, that is,

$$\hat{\mathbf{T}}_I = \underset{\mathbf{T}_I}{\operatorname{argmin}} \{ \|\mathbf{u}_{\text{label}}(\mathbf{T}_I)\|_2 - \lambda \|\mathbf{u}_{\text{control}}(\mathbf{T}_I) - \mathbf{u}_{\text{label}}(\mathbf{T}_I)\|_2 \}, \quad (1)$$

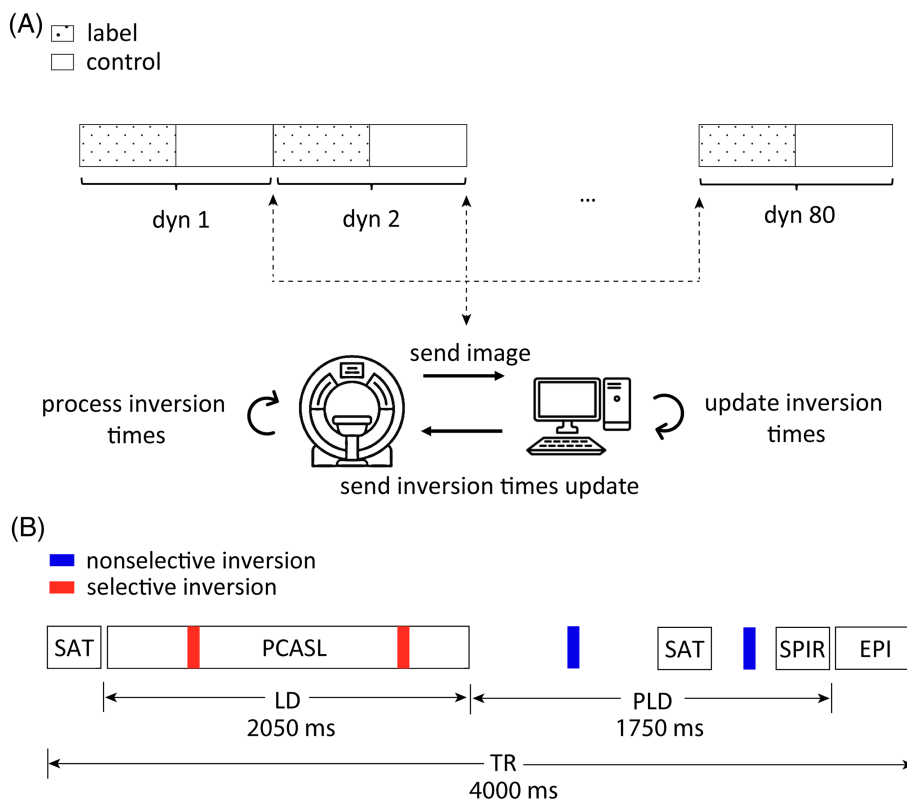


FIGURE 1 Schematic overview of the feedback loop (FBL) and the arterial spin labeling (ASL) sequence. (A) The FBL scan consists of 80 dynamics, in each of which one label and one control image is acquired. After each dynamic, the acquired label and control images are sent to an external computer via the remote connection software eXTernal Control (XTC). At the external computer, the arrival of new images is automatically detected and processing using a Python tool is commenced. The updated inversion times are stored in a text file, which is read by the scanner before acquiring the new dynamic. (B) ASL images are acquired with pseudo-continuous ASL (PCASL) and a single-shot echo-planar imaging (EPI) readout, preceded by spectral presaturation with inversion recovery (SPIR) fat suppression. Two selective saturation modules (SAT) are used to suppress (1) residual magnetization at the beginning of each dynamic, and (2) fresh inflowing spins that arrive in the imaging plane later than the labeled blood. Two inversion/background suppression (BGS) pulses are applied during the label duration (LD) (in red), while two are applied during the postlabel delay (PLD) (in blue). The timings of these four inversion pulses are adjusted between dynamics to obtain optimal background suppression at the end of the FBL scan

with $\mathbf{u}_{\text{label/control}} \in \mathbb{R}^{256^2 \times 1}$ the label/control image obtained with inversion times $\mathbf{TI} = [T_{I1}, T_{I2}, T_{I3}, T_{I4}]^T$. The regularization parameter λ determines how much the perfusion signal should be taken into account in the minimization process and was tuned experimentally. Because there was no accurate signal model available at the time of scanning, the optimization was performed using the Nelder–Mead method.¹² This optimization method is a value-based method, requiring only function values itself (i.e., values of the cost function), and no additional information about gradients. The reflection coefficient, the expansion coefficient, and the contraction coefficient were set to the default parameters of $\alpha = 1, \beta = \frac{1}{2}$, and $\gamma = 2$. The initial simplex was defined around the initial timings (described in section 2.4), placing one additional simplex point in each unit direction, using a 3% increase relative to the initial timings. Before optimization, the scan without BGS was used to compute a background mask based on thresholding, which was then used to exclude background noise in the optimization process. The optimization was terminated after 80 dynamics, covering a total duration of approximately 11 min.

2.3 | MR data acquisition

Experiments were performed in one phantom and in 12 healthy volunteers (five males, seven females, age 24–62 years) after informed consent was obtained, conforming to the local ethical regulations. Scanning was performed on an Ingenia 3-T dual transmit MR system (Philips, The Netherlands) equipped with a 32-channel head coil.

Phantom: The phantom contained four tubes with mixtures of water and different agar concentrations corresponding to varying T_1 values of 248, 522, 862, and 3500 ms. Pseudo-continuous ASL (PCASL) data were acquired in this phantom with the same scan parameters as the in vivo

scans, for which details are provided in the next paragraph, except using different initial timings, as described in section 2.4. In this experiment, a varying number of BGS pulses (the last two, the last three or all four) were optimized, while the total number of BGS pulses was always four. In a second experiment, the first inversion pulse (during labeling) and the last inversion pulse (during the PLD) were updated in steps of 0.01 and 0.001 ms between dynamics in the FBL framework, using the same initial timings as for the first experiment. This was carried out to test the precision of updating the inversion timings in the FBL mechanism by examining the smoothness of the resulting signal change.

In vivo: PCASL data were acquired with a 2D single-shot echo-planar imaging (EPI) readout and the following scan parameters: field of view (FOV) = 240 × 240 mm², in-plane resolution = 2.75 × 2.75 mm², slice thickness = 7 mm, label duration/PLD = 2050/1750 ms, TE/TR = 17/4000 ms, spectral presaturation with inversion recovery (SPIR) fat suppression, SENSE factor of 2.3, and a scan time of 10 n 48 min. Labeling was performed using balanced gradients with a maximum gradient strength of 6 mT/m and 0.6-ms long RF pulses, obtaining a 21° flip angle. The imaging plane was positioned through the ventricles, parallel to the caudal part of the corpus callosum. The labeling plane was positioned perpendicular to the internal carotid artery approximately 60 mm below the imaging plane. One selective saturation module was applied at the beginning of each dynamic to suppress any residual magnetization. A second selective saturation module was applied 420 ms before the excitation pulse, using a 60-mm thick saturation slab at 10 mm below the imaging slice, to suppress fresh inflowing spins traveling from the neck to the imaging plane that arrive later than the labeled blood. Four inversion pulses were used for BGS. The timings were initialized as described in section 2.4, and optimized using the FBL with 80 dynamics. PCASL scans were also acquired with constant inversion times (equal to the initial timings for the FBL scan), hereafter referred to as standard BGS, and without BGS (20 dynamics), for comparison. Figure 1B shows a schematic overview of the PCASL sequence.

A 2D T₁-weighted turbo spin echo (TSE) image, used to generate a gray matter mask for postprocessing of the perfusion images, was acquired with the following scan parameters and the same geometry settings as the ASL imaging slice: FOV = 240 × 240 mm², in-plane resolution = 2.75 × 2.75 mm², slice thickness = 7 mm, flip angle = 90°, TSE factor = 4, TE/TR = 14/500 ms, SPIR fat suppression, and a scan time of 5 s.

In six of the 12 volunteers, the scan with individually optimized BGS using the FBL was repeated for different amounts of regularization: λ = 0, 2, 4, and 6. For all other volunteers the regularization parameter was set to λ = 4 based upon the initial results.

In three of the 12 volunteers, the scan with individually optimized BGS using the FBL was repeated twice with the same initial timings to test the repeatability of the optimization procedure. The FBL scan was also repeated with initial timings 100 ms shorter than the optimized initial timings to test the stability of the optimization procedure. To test whether individual optimization of inversion times is necessary, PCASL scans were also acquired with the individually optimized inversion timings from the previously scanned subject in three of the 12 volunteers. Finally, in the last experiment of this work, the FBL was used as a preparation step for a stimulus scan in three of the 12 volunteers, to demonstrate that this enables tracking of the neuronal activation at twice the temporal resolution, similar to what was performed in Ref.¹³ To this end, the inversion times were optimized during the first 80 dynamics, after which they were fixed for another 20 dynamics. During these 20 dynamics, an 8-Hz flickering checkerboard pattern was presented twice (five dynamics each), each time followed by a period of rest. During these activation scans (dynamics 81–100), the label image was not acquired anymore, allowing measuring of the control image twice in each dynamic (twice the temporal resolution). Because the signal in the label images was minimized, these control images were expected to directly show approximate perfusion changes. The results were compared with a separate scan, in which both label and control were acquired in each dynamic (original scan).

2.4 | Simulation of initial timings

Optimizing four parameters at once results in a minimization problem with multiple local minima. Therefore, the initial timings need to be chosen close to the optimal inversion times, which were estimated via simulation. First, a T₁ map was generated after segmenting a publically available brain image¹⁴ into four different tissues and assigning T₁ values of gray matter (T₁ = 1331 ms), white matter (T₁ = 832 ms), cerebrospinal fluid (CSF) (T₁ = 4000 ms), and corpus callosum (T₁ = 690 ms) to each of the compartments.¹⁰ Second, the longitudinal magnetization left in all the voxels of an ASL image, $M_z \in \mathbb{R}^{256^2 \times 1}$, at the time of the ASL readout, was minimized according to

$$\hat{\mathbf{T}}\mathbf{l} = \operatorname{argmin}_{\mathbf{T}\mathbf{l}} \|M_z(\mathbf{T}\mathbf{l})\|_2. \quad (2)$$

Each element of M_z (representing voxel \mathbf{x}) is given by^{15,16}

$$M_z(\mathbf{x}, T_1(\mathbf{x}), \mathbf{T}\mathbf{l}) = 1 + (-1)^{n+1} e^{-\frac{LD+PLD}{T_1(\mathbf{x})}} + 2 \sum_{i=1}^n (-1)^i e^{-\frac{LD+PLD-T_i}{T_1(\mathbf{x})}}. \quad (3)$$

Equation 2 was minimized using 350 iterations of the same simplex algorithm as was used to optimize the inversion times on the scanner, resulting in the following initial timings for in vivo experiments: $\hat{\mathbf{T}}\mathbf{l} = [683, 1948, 2980, 3597]^T$. This distribution of timings means that the algorithm

optimized two FOCI pulses during labeling¹⁷ and two pulses during the PLD. For the phantom experiment (details below), T_1 values of 248, 522, 862, and 3500 ms resulted in initial timings of $\hat{T}_1 = [1022, 2431, 3267, 3679]^T$. In this case, the algorithm optimized one pulse during labeling and three pulses during the PLD. To simplify the modifications in the scanner software, care was taken during the optimization that none of the pulses moved inside or outside of the labeling block.

2.5 | Processing of MR data

The individual perfusion images were defined as the absolute difference between the magnitude of the individual label and the control images, and were normalized by element-wise dividing the perfusion signal by an estimate of the M_0 map, obtained from the scan without BGS. Complex subtraction was avoided because of the unstable phase behavior for low SNR scans. Averaged label, control, and perfusion images were calculated by averaging the individual images over the last 20 dynamics of the scan, unless mentioned otherwise. Corresponding tSNR maps of the perfusion signal were defined as the temporal mean of the normalized perfusion signal divided by its standard deviation in each voxel. For each volunteer, the average tSNR was reported in the gray matter region, for which a mask was derived from the T_1 -weighted image using k-means clustering with five clusters. The average tSNR in the gray matter was compared between standard BGS and individually optimized BGS using the FBL, and tested for statistically significant differences using a two-sided paired t-test. A Shapiro–Wilk test was first performed to confirm sufficient

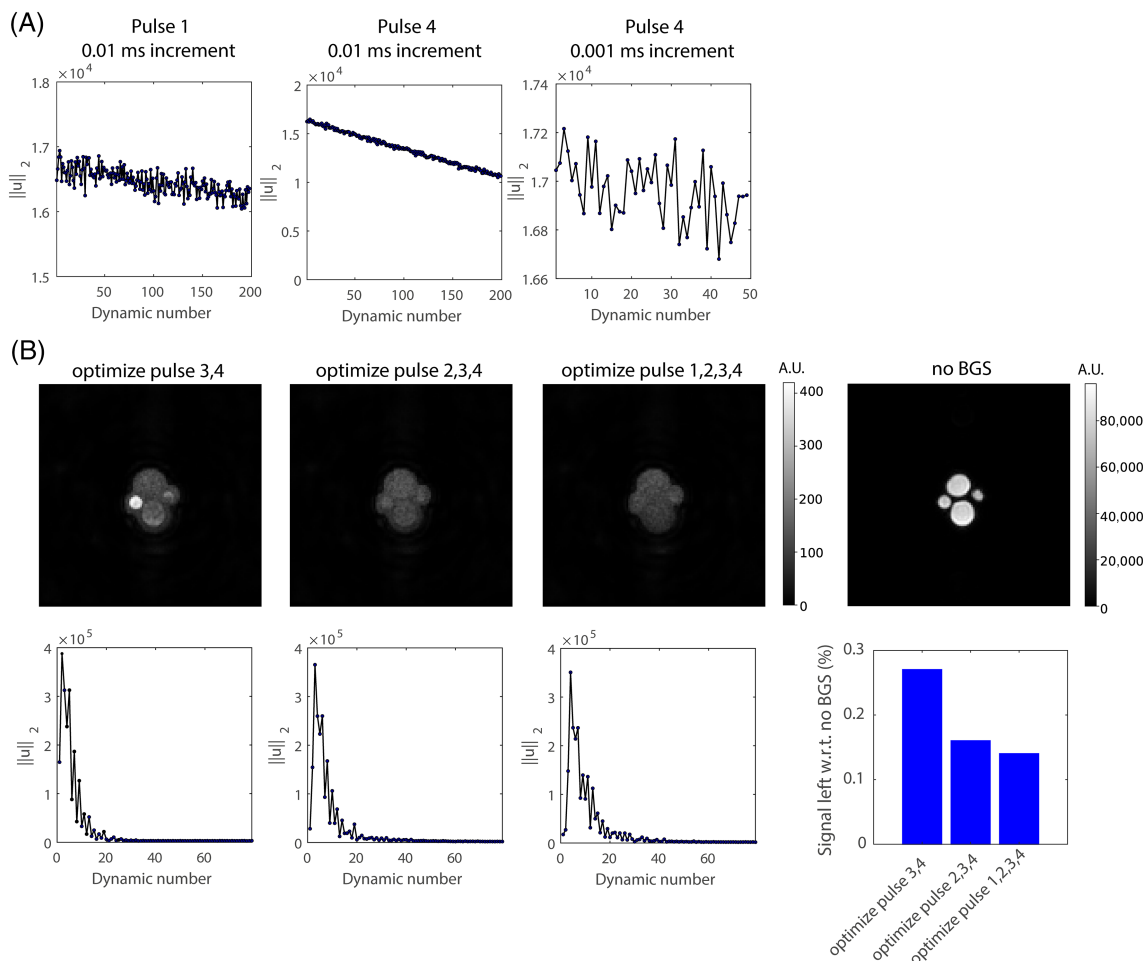


FIGURE 2 Validation of the feedback loop (FBL) scan in a phantom. (A) Updating the first and the fourth inversion pulses in steps of 0.01 ms in time resulted in a linear, smoothly decreasing signal over 200 dynamics, corresponding to a 2-ms time interval. Shifting the fourth inversion pulse in much smaller time steps of 0.001 ms also resulted in a linearly decreasing signal over 50 dynamics (0.05 ms time interval). (B) Averaged label images (u) over the last 20 dynamics of the FBL scan (top row) and corresponding convergence plots of the cost as a function of dynamic scan number (bottom row). Fixing the timings of the first two inversion pulses, while optimizing the other two timings with the FBL scan, resulted in a background suppression (BGS) of 0.27% with respect to using no BGS. Fixing only the timing of the first inversion pulse, while optimizing the other three pulses with the FBL scan, resulted in a slightly better BGS of 0.16%. Optimizing the timings of all four inversion pulses resulted in the best BGS: 0.14%. The convergence was slower as the number of inversion timings to optimize increased

normality of the SNR values ($p = 0.75$ for standard BGS, $p = 0.80$ for FBL BGS), considering the analysis on magnitude images and low SNR data. The significance level was set at p less than 0.05 for both tests.

3 | RESULTS

Figure 2A shows a linearly decreasing signal norm when the first inversion pulse (left plot) or the last inversion pulse (middle plot) is shifted in time with steps of 0.01 ms over 200 dynamics, corresponding to a total time interval of 2 ms. Shifting the last inversion pulse in even smaller time steps of 0.001 ms (right plot) still results in a linearly decaying signal norm over 50 dynamics, corresponding to a total time interval of 0.05 ms. This linear behavior and the absence of sharp signal jumps (larger than the observed noise level) suggest that there are no unexpected errors or discontinuities in the inversion time updates, even when the inversion pulse is surrounded by PCASL pulses. Note that the differences in the noise level between the plots in Figure 2A are introduced by the differences in the scaling of the y-axes. The results therefore show that the inversion times (during and after labeling) can be updated in the scanner at high precision. Figure 2B shows the performance of the FBL scan in the four-tube phantom for a varying number of optimized inversion timings. Only optimizing the timings of the last two inversion pulses, while fixing the timings of the first two inversion pulses, resulted in a BGS with 0.27% signal left compared with the scan without BGS. Optimizing the timings of the last three inversion pulses, while fixing the timings of the first inversion pulse, resulted in a slightly better BGS, with 0.16% signal left. Optimizing all four pulses resulted in the best BGS, leaving 0.14% signal compared with without BGS. The plots of the cost as a function of the dynamic number show that the convergence is slightly slower when more timings are being optimized.

Figure 3A shows the performance of the FBL scan for different regularization parameters in one volunteer. Without regularization ($\lambda = 0$), the FBL does not guarantee that label and control signal have the same signs, leading to magnitude subtraction errors in the averaged perfusion images. This effect is circumvented using regularization ($\lambda = 2, 4, 6$), resulting in a higher tSNR and signal intensity of the perfusion signal in the gray matter region compared with the FBL without regularization. Figure 3B shows the average tSNR as a function of λ in the gray matter for all six volunteers. For some volunteers, $\lambda = 6$ resulted in a decrease in tSNR compared with $\lambda = 4$, whereas $\lambda = 4$ resulted in a larger or approximately equal tSNR compared with $\lambda = 2$ in most volunteers. Therefore, all further experiments were performed with $\lambda = 4$.

Figure 4 shows a comparison between the ASL performance without BGS, with standard BGS, and with optimized BGS using the FBL mechanism, with and without regularization, in one volunteer. Without regularization ($\lambda = 0$), the optimized BGS using the FBL resulted in label images containing much less signal (0.57%) compared with standard BGS (0.77%), suggesting improved BGS performance. The quality of the perfusion images is lower, however, because of the occurrence of magnitude subtraction errors. With regularization ($\lambda = 4$), the optimized BGS using the FBL resulted in label images with a slightly higher signal norm (0.90%) compared with without regularization and standard BGS (due to a slightly less effective CSF suppression), leading to averaged perfusion images with the highest tSNR in the gray matter region. This shows that the regularized FBL can avoid suboptimal inversion time settings, which the offline optimization simulation calculated for standard BGS. Furthermore, the




FIGURE 3 Tuning the regularization parameter in vivo. (A) Averaged perfusion images (top row) over the last 20 dynamics of the feedback loop (FBL) scan in one volunteer, using different regularization parameters ($\lambda = 0, 2, 4, 6$) and corresponding temporal signal-to-noise ratio (tSNR) maps masked by the gray matter region (bottom row). Without regularization ($\lambda = 0$), the FBL scan converged to label and control signal with opposite signs in a majority of voxels, leading to underestimation of perfusion signal. This effect was circumvented using regularization ($\lambda = 2, 4, 6$), resulting in perfusion images of higher quality and with higher tSNR. (B) For all six volunteers, the tSNR in the gray matter was larger when using regularization ($\lambda = 2, 4, 6$) compared with the tSNR without using regularization ($\lambda = 0$). In some volunteers, $\lambda = 6$ resulted in a reduction of tSNR, while in most volunteers $\lambda = 4$ resulted in a larger or equal tSNR than $\lambda = 2$. Therefore, the regularization parameter was set to $\lambda = 4$ for all other experiments. Note that the same regularization parameter could be used both for female (solid lines) and for male subjects (dashed lines), although the tSNR is in general larger for female than for male subjects

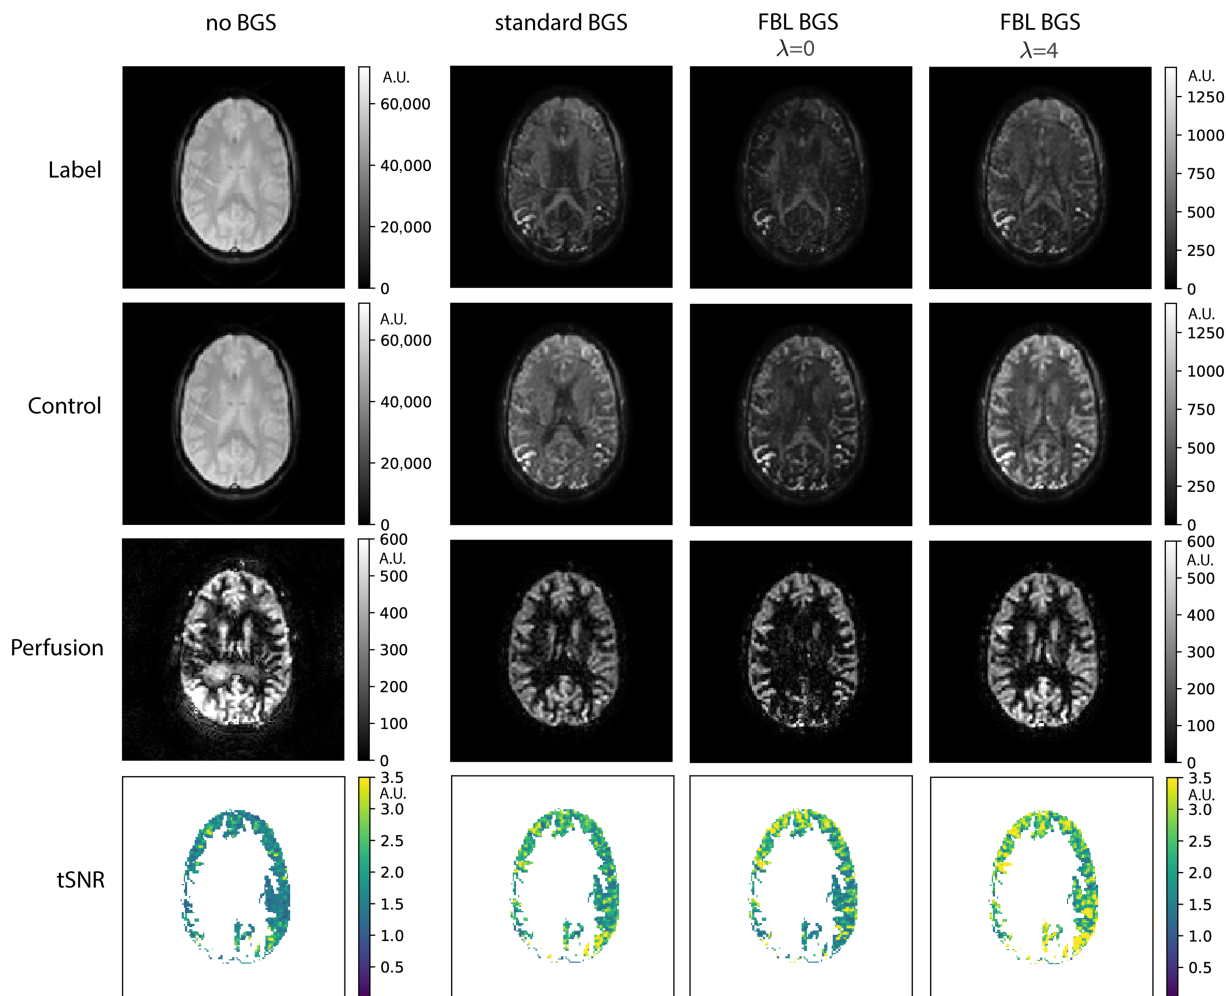


FIGURE 4 Comparison of arterial spin labeling (ASL) images acquired without background suppression (BGS), with standard BGS, and with individually optimized BGS using the feedback loop (FBL) scan. Using BGS resulted in a much lower signal in the unsubtracted label and control images and a higher quality of the corresponding perfusion images, both with standard BGS and with individually optimized BGS using the FBL. Optimizing the BGS with the FBL without regularization ($\lambda = 0$) resulted in a much darker label image, suggesting improved BGS, but the quality of the perfusion image is lower due to magnitude subtraction artifacts. Using regularization ($\lambda = 4$) when optimizing the BGS with the FBL resulted in slightly more signal in the label image compared with without regularization, but the quality of the perfusion image was best. This was confirmed by the temporal signal-to-noise ratio (tSNR) map (masked to only show the gray matter region). Furthermore, the FBL resulted in an averaged control image, from which the perfusion signal can directly be appreciated. Note that the reduction in perfusion signal for scans with BGS compared with scans without BGS is likely caused by an imperfect inversion efficiency (i.e., a loss of ASL-signal due to applying the inversion pulses). Note that magnitude subtraction errors for standard BGS and for the FBL without regularization not only result in a lower tSNR, but also in a lower perfusion signal intensity compared with the FBL with regularization

FBL mechanism resulted in an averaged control image in which the perfusion signal can directly be appreciated, possibly mixed with a small amount of gray matter residue. This is much less the case for standard BGS, where the control image contains a larger amount of white matter signal. The perfusion signal intensity in the FBL scan was also slightly higher than that obtained with standard BGS, which is an indication that standard BGS also slightly suffered from magnitude subtraction errors. Note that the reduction in perfusion signal for standard BGS and for FBL BGS compared with not using BGS is likely caused by an imperfect inversion efficiency.¹⁸

Figure S1 shows the label, control, and perfusion images for all the 80 dynamics. The same comparison is summarized in Figure 5A,B for all 12 volunteers: the FBL improved the tSNR of the perfusion signal in all volunteers. The mean tSNR of the perfusion signal in the gray matter was statistically significantly higher ($p = 2 \times 10^{-4}$) with the individually optimized BGS using the FBL compared with standard BGS. The FBL leads to a percentage increase of 10%–60% in perfusion tSNR with respect to no BGS, whereas the tSNR improvement for standard BGS ranged from 1% to 30%. The standard deviation of the tSNR in the gray matter region increases with the mean tSNR, showing that the tSNR increase is non-uniform: some regions benefit more from the optimized BGS than other regions. Figure 5E furthermore shows in three volunteers that the individually optimized BGS still outperformed the standard BGS in terms of tSNR when more dynamics were used to calculate the averaged perfusion

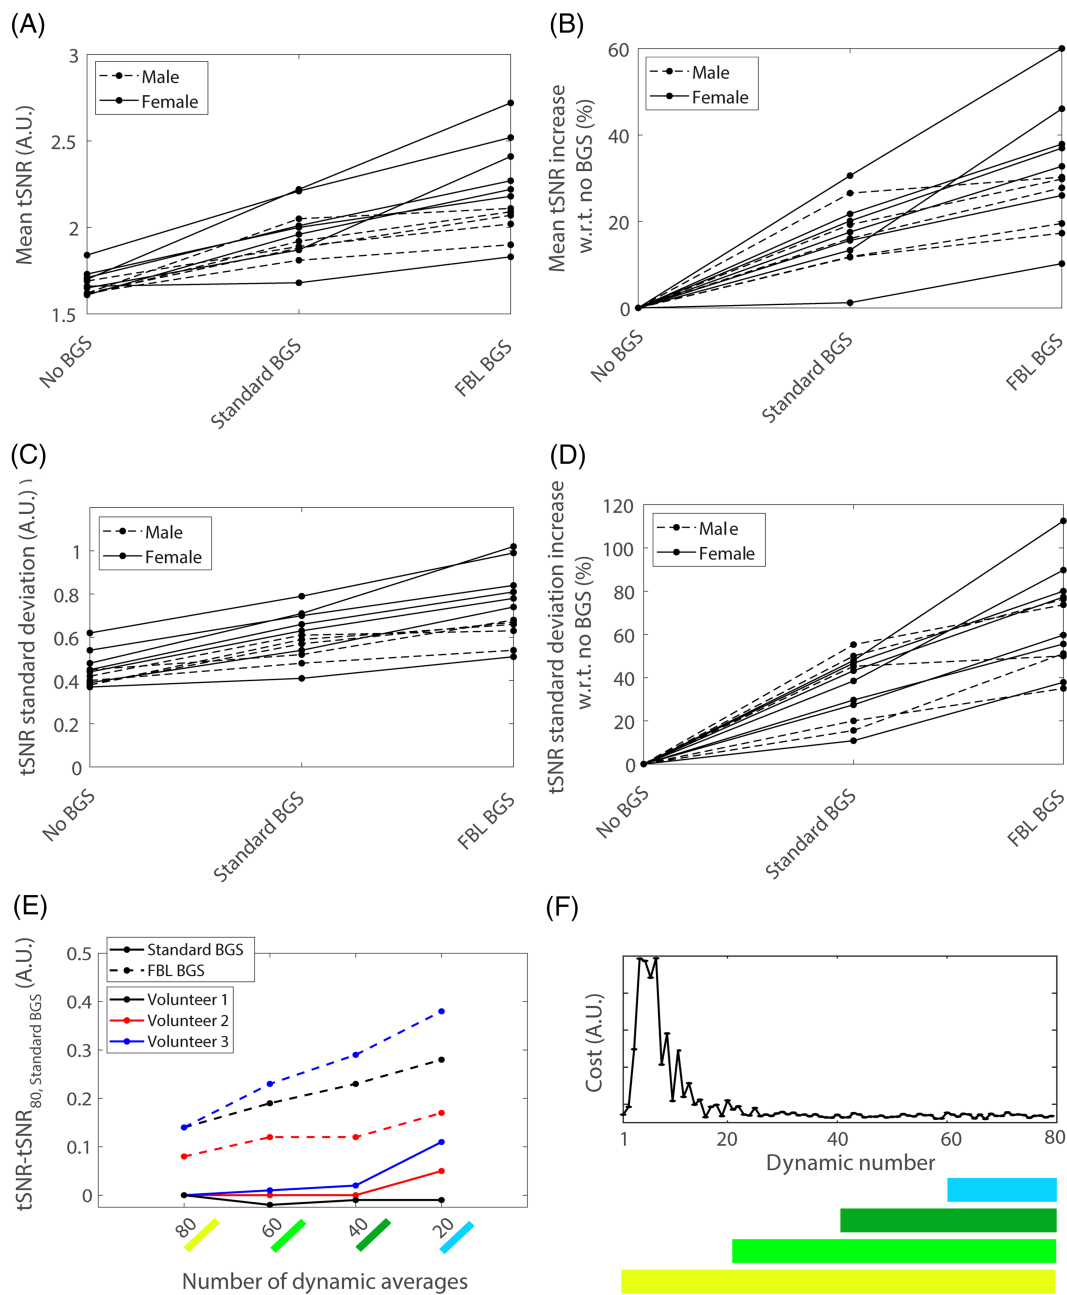


FIGURE 5 Comparison of temporal signal-to-noise ratio (tSNR) in all volunteers. (A) The individually optimized background suppression (BGS) resulted in an increased tSNR in the gray matter compared with standard BGS in all 12 volunteers. (B) The percentage increase for the individually optimized BGS with the feedback loop (FBL) compared with a scan without BGS ranged from 10% to 60%. This is a much larger range than for the standard BGS (1%–30%). The tSNR of the optimized BGS with the FBL was statistically significantly higher ($p = 2 \times 10^{-4}$) compared with standard BGS. (C) and (D) The standard deviation of the tSNR in the gray matter region increases with the mean tSNR, showing that the tSNR increase is nonuniform: some regions benefit more from the optimized BGS than other regions. (E) The FBL scan also resulted in a higher tSNR than standard BGS when more (and earlier) dynamics were used (yellow, light green, dark green) to compute the perfusion images, as shown in all three subjects. This is the case even although earlier dynamics of the FBL correspond with a higher cost, and thus a higher variance of the perfusion signal compared with later dynamics, as shown in (F). This increasing variance for earlier dynamics leads to a decreasing tSNR for a larger number of averages. The graphs for standard BGS confirm that in this case the tSNR is independent of the number of averages used, except for small dynamic numbers, indicating that more than 20 averages are needed for accurate tSNR computation. Note that the tSNR values are shown as difference with respect to the tSNR value obtained with 80 dynamic averages and standard BGS ($tSNR - tSNR_{80, Standard\ BGS}$), to equalize the starting point of the curves for standard BGS. Further note that the mean perfusion signal was always computed over 80 dynamics, to minimize the effect of random sampling on the tSNR curves

image, even although earlier dynamics of the FBL scan correspond with a poor BGS performance and therefore a suboptimal variance of the perfusion signal, as shown in Figure 5F. This suboptimal variance explains the decreasing trend in tSNR with an increasing number of averages for the FBL scan. For standard BGS the tSNR is constant, except for a small number of averages, suggesting that more than 20 dynamics are required for accurate tSNR calculation. The averaged perfusion images for 20 and 80 dynamics can be found in Figure S2, as well as a sliding window analysis of the tSNR during the FBL.

The FBL scan converged to a different set of inversion times for each volunteer. The final inversion times, normalized with respect to the group mean inversion times, are shown in Figure 6A. All inversion times differ by a maximum of 45 (−15 to +30) ms. The largest differences between volunteers are observed for the second inversion pulse, which could be due to its initial timing being relatively close to the end of the labeling block (~100 ms), and is therefore more restricted in its possible update steps compared with the other pulses. The final inversion time updates at the end of the FBL with respect to the initial timings, as plotted in Figure 6B, show a maximum update of 33 ms. Figure 6C furthermore shows that these relatively small updates are necessary to achieve an increase in perfusion tSNR: reusing the individually optimized timings from a previously scanned subject resulted in a lower tSNR compared with individually optimizing the BGS in each of the three volunteers.

Repeating the FBL scan twice with the same initial timings resulted in very similar inversion time updates over the 80 dynamics, as shown in Figure 7A,B, indicating a good repeatability of the FBL scan. The FBL scan converged to different final inversion times when the initial timings were chosen 100 ms shorter than the optimized initial timings for each of the four pulses, as shown in Figure 7C,D.

Figure 8 shows how individually optimizing the BGS with the FBL can be used as a preparation step for functional ASL experiments. After completing the optimization of the inversion times during the first 80 dynamics, the signal intensity in the control image shows a clear increase following the start of the visual stimulus. Acquiring only the control image during the activation scans enables tracking an estimate of the fast changing perfusion signal during the stimulus at twice the temporal resolution. Figure 8C shows that these changes can directly be appreciated in the control images, while yielding similar relative changes in perfusion signal compared with a traditional label-control approach.

4 | DISCUSSION

The results in this study have shown that individually optimizing the BGS using a FBL improves the quality of the perfusion images compared with standard BGS. Even although the updates of the inversion times (~15 ms) with respect to the initial timings were relatively small, they resulted in a statistically significant increase in tSNR. Reusing the optimized timings from a different volunteer resulted in a lower tSNR compared with individually optimizing the BGS for each of the volunteers. This suggests that optimal perfusion quality cannot be achieved with one single protocol. The increase in tSNR could furthermore potentially result in a more precise cerebral blood flow quantification.^{1,17} The individually optimized timings resulted in control images from which an estimate of the perfusion signal can directly be appreciated. This was used to track the neuronal activation in stimulus scans at twice the temporal resolution compared with the original ASL scans.

ASL images obtained with standard BGS sometimes showed a lower perfusion signal intensity (besides lower SNR) than those obtained with the FBL scan. This can be explained by magnitude subtraction errors in the case of standard BGS. Although the level of BGS obtained with an off-line optimization of inversion times (such as performed for standard BGS) can be very high (>99%), there is no guarantee that the offline optimized timings will lead to label and control images with equal sign in the perfused region for all subjects. Even although the inversion time updates obtained with the FBL were small (~15 ms), the FBL was able to steer the initial BGS away from the regime in which magnitude subtraction errors




FIGURE 6 Comparison of the final inversion times (TI) at the end of the feedback loop (FBL) in all volunteers. (A) The final inversion times differed by a maximum of 45 (−15 to +30) ms for all the volunteers. (B) The maximum final inversion time update at the end of the FBL with respect to the initial timings is 33 ms. Different colors represent different volunteers. (C) Reusing the individually optimized inversion times from a previously scanned subject results in a lower temporal signal-to-noise ratio (tSNR) compared with individually optimizing the background suppression (BGS) for each volunteer




FIGURE 7 Repeatability and stability of the feedback loop (FBL) in two volunteers. (A) and (B) Repeating the FBL scan with the same initial timings twice resulted in very similar inversion time updates for the different dynamics: both scans converged to the same final inversion times. (C) and (D) Adding a 100-ms perturbation to each of the four pulses of the initial timings resulted in different convergence behavior of the FBL scan, especially for the timing of the first two pulses. Note that all updates in (B) and (C) are given with respect to the optimized initialization values

occur. Perfusion regularization, which indirectly penalizes opposite signs in the label and control images, was the key driver in this process. This was demonstrated by the improved performance of the FBL scan with $\lambda = 4$ compared with $\lambda = 0$. In practice this means that the FBL resulted in improved tSNR of the perfusion images, while the signal norm of the corresponding label images was sometimes slightly larger compared with that of standard BGS. This last effect can also be interpreted as an improved quantification.

The individually optimized BGS resulted in control images from which the perfusion signal can directly be appreciated. Although it might sound counterintuitive that the control acquisition contains perfusion information, this is explained by the fact that (1) the label image was minimized to the target of full suppression, serving as a “zero-image” in the subtraction, and (2) magnitude subtraction errors were avoided via regularization. In this case, we can write $|\text{perfusion}| = |\text{control-label}| \approx |\text{control-0}| = |\text{control}|$, and as such, we can appreciate an estimate of the perfusion signal in the control image. It should be noted that the control image gives a good estimate of the perfusion signal, but cannot be stated to be equal to the perfusion signal. This is because local variations of BGS efficiency can lead to minor static signal residues in the corresponding regions of the optimally suppressed label image. Furthermore, the effective BGS of the label image that was reached after optimization of the inversion times could be disrupted by the physiological changes upon neuronal stimulation. Therefore, ASL acquisitions at twice the temporal resolution (in which the label images are not acquired) are most interesting in applications where the temporal resolution is more important than accuracy of the perfusion estimate, such as in perfusion functional magnetic resonance imaging with feedback to the participant during the acquisition. Future experiments should investigate how temporal variations of perfusion affect the efficiency of the BGS over time and therefore the accuracy of the perfusion estimate. We speculate that this may depend on the ratio between the gray matter tissue and the blood tissue, which determines the effective overall T_1 for each voxel in the perfused area. Considering this, the benefit of our FBL approach should be compared

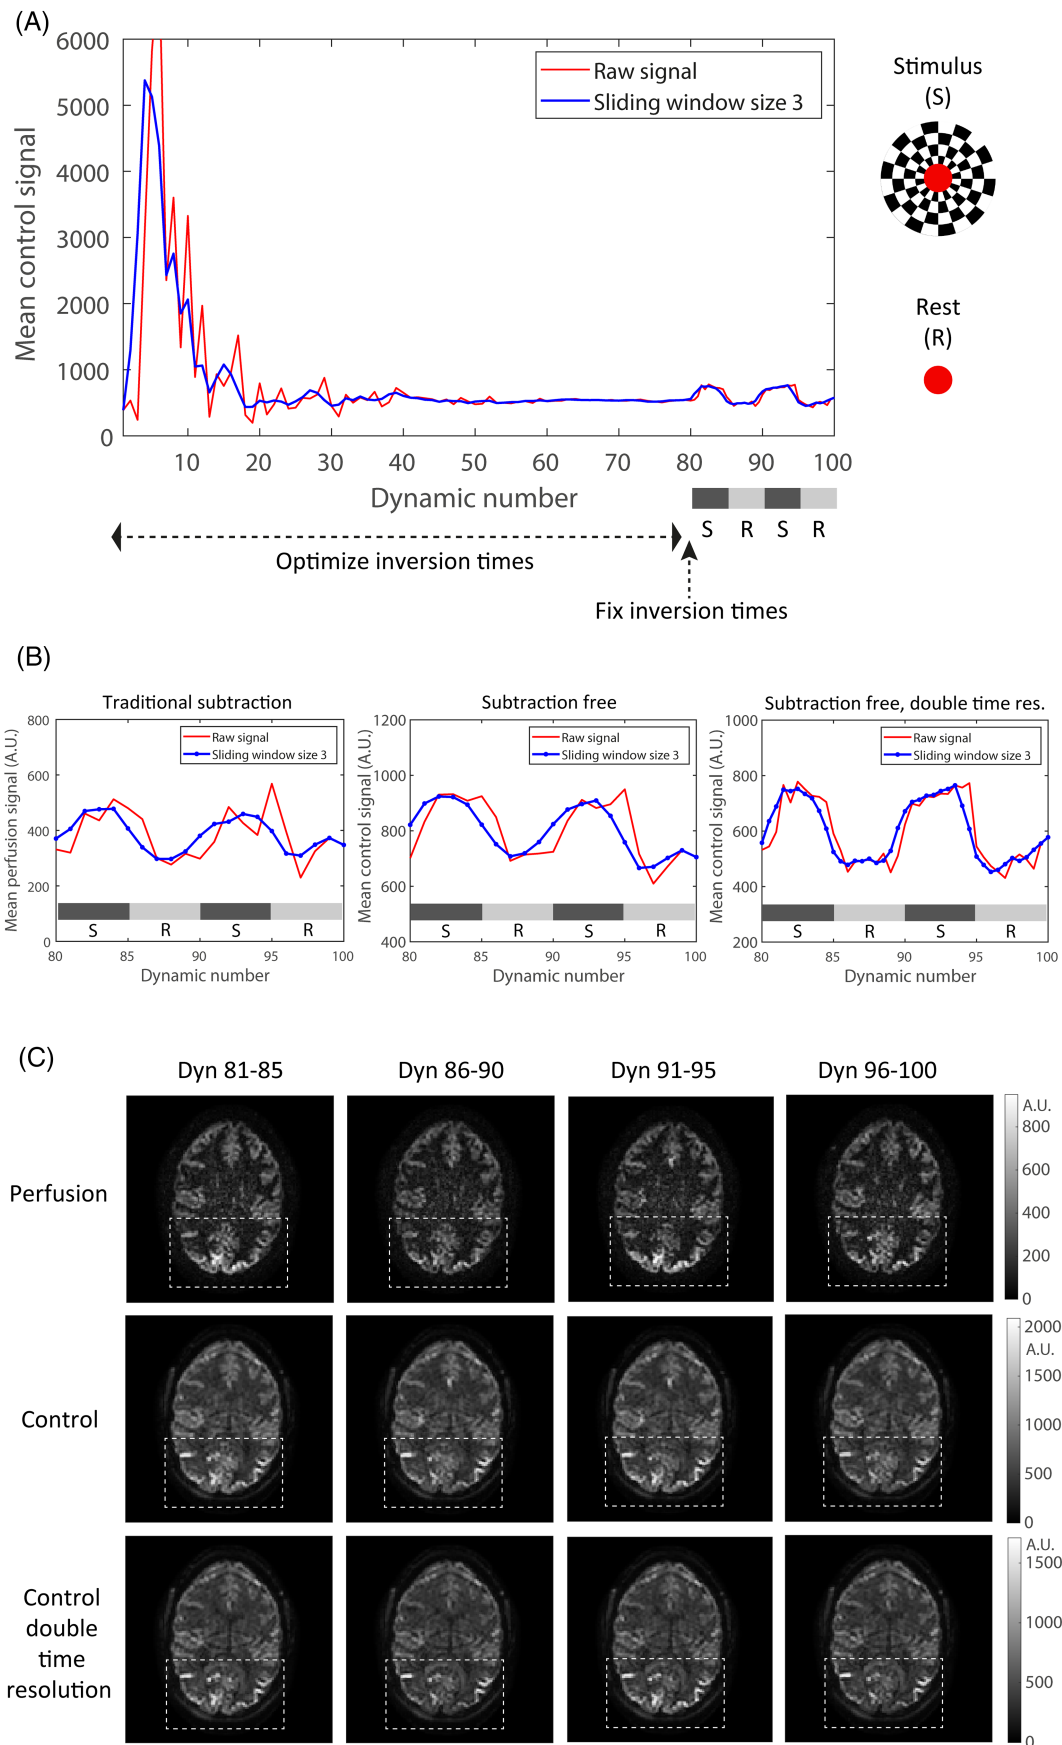


FIGURE 8 Legend on next page.

FIGURE 8 Neuronal activation measured at low and high temporal resolution. (A) The first 80 dynamics of the activation scan were used to optimize the background suppression (BGS) inversion times. During the following 20 dynamics, two blocks of visual stimulation (S) were imposed by showing an 8-Hz flickering checkerboard. Shortly after starting the stimulation, an increased mean control signal (acquired with optimal inversion times) in a region of interest of the visual cortex could be observed. (B) With the feedback loop, the change in control signal (middle) is of similar amplitude compared with the perfusion signal (left). Acquiring only control images during the last 20 dynamics of the scan resulted in activation signal (right) at twice the temporal resolution compared with the control signal in the traditional arterial spin labeling (ASL) scan, showing much sharper signal changes between stimulus blocks. Note that the first two subplots were computed from the same scan, and that the range of the axes is the same, allowing comparison of the amplitude of the signal changes. The fact that we observed similar control curves without subtraction and with traditional ASL subtraction seems to point to marginal influence of a physiology-induced disturbance of the BGS during activation. (C) The corresponding perfusion and control images averaged over five neighboring dynamics showed the signal increase in the visual cortex (dashed rectangle) during block 1 (dynamic 81–85) and block 3 (dynamic 91–95)

with using the first control image as baseline in the subtraction for all other label images, or by minimizing the control image with the FBL and only acquiring the label images. These three approaches will double the temporal resolution at the cost of reduced quantification accuracy. Furthermore, comparison with a less well-optimized BGS case in combination with subtraction from a control image acquired before activation could be an interesting direction for further investigation.

Although the final individually optimized inversion times resulted in improved quality of the perfusion images, the convergence of the FBL was rather slow: signal intensity and inversion times converged after approximately 60 dynamics, with a duration of 8 s for each dynamic. This means that the total scan time for the individually optimized protocol (~11 min for 80 dynamics) is in its current form roughly twice as long compared with that of a standard ASL protocol, which typically lasts about 5 min. To minimize the overhead time of the FBL scan, the optimization process could potentially be accelerated with the help of a more efficient optimization scheme. In our current protocol, both label and control images were acquired in each dynamic of the FBL scan to provide perfusion signal (control minus label) needed for regularization. In the first half of the FBL scan, however, the regularization has negligible effect as a result of the relatively large signal intensity in the label image compared with the signal intensity in the perfusion image. Therefore, the first half of the FBL scan could potentially also be performed without acquiring the control image, which could lead to a reduction in scan time of 25%, leaving a total scan time of 8 min 15 s. The scan time could be reduced further by optimizing the size of the initial simplex, which determines the amount of signal increase during the first iterations. Furthermore, the phantom experiment in Figure 2 showed faster convergence when only the timings of the last two inversion pulses were optimized, while fixing the timings of the first two pulses. Even although optimizing all four pulses resulted in the best BGS in a phantom, it is worth investigating whether fixing two of the four pulses still provides good BGS in vivo, where more different tissues and stronger field inhomogeneities are present. Finally, reinforcement learning could potentially be used to accelerate the convergence even further,¹⁹ in this way bringing the total scan time closer to that of the standard ASL protocol. This would be especially important for 3D ASL acquisitions, which often require a much longer scan time due to their multishot nature.

It should be noted that the current FBL scan may take longer than offline optimization using a preacquired T_1 map. However, to obtain optimal BGS under the specific conditions of a certain subject, the efficiency of the inversion should be known, for example, by acquiring T_2 , B_0 , and B_1^+ maps. The FBL, however, adapts the inversion times on the fly to achieve optimal BGS with imperfect inversion efficiencies, without explicitly requiring knowledge of B_0 or B_1^+ . This is an advantage of using an online FBL scan compared with an offline optimization approach. One could consider estimating a T_1 map from the early iterations of the FBL scan and using this to steer the optimization in the right direction. Note that the current FBL will not correct for signal loss due to noninstantaneous inversion and saturation pulses, although the BGS will be optimal.

In this work we chose to minimize the signal intensity of the label image instead of the control image. This choice was motivated by the fact that the magnitude of the control signal is larger than that of the label signal. This means that, when starting the optimization with positive real signals, the real parts of both the label and the control images do not have to cross the x-axis when minimizing the signal intensity in the label image, whereas this would be the case when minimizing the signal intensity in the control image. If the real part of the label or the control signal needs to cross the x-axis at a later time point in the optimization process, this leads to a less stable FBL scan. This effect is shown in Figure S3, where only in one of three volunteers did minimizing the control image result in a darker (i.e., less negative) control image than the label image. This can be understood, because for near-optimal BGS, the regularization will start to dominate the optimization before both label and control signals have fully crossed the x-axis.

The performance of the FBL is sensitive to some of the hyperparameters of the optimization algorithm. First of all, due to the local minima of the minimization problem, it is important to choose the initial timings close to the global minimum. Our results showed that initial timings 100 ms shorter than the optimal initial timings resulted in different final inversion times. This is especially true for the first two pulses. We speculate that this is because small errors in the first two timings can partially be corrected by changing the timings of the last two pulses, whereas errors in the last two timings directly influence the final BGS, being close to the actual readout. Subject-specific initial timings could be obtained via simulation by acquiring a T_1 map prior to the FBL scan, although this would add to the total scan time. Second, the size of the initial simplex influences how far the optimization algorithm moves away from the minimum before converging. This parameter can affect both the final inversion times and the

convergence speed: the size of the simplex should be chosen small enough to ensure rapid convergence, but large enough to make sure that the algorithm does not get stuck in a local minimum during the first iterations. Third, the regularization parameter λ balances the minimization of the label signal and the maximization of the perfusion signal. If λ is chosen too small, the algorithm can converge to label and control signal having opposite signs, leading to magnitude subtraction artifacts in the perfusion images. If λ is chosen too large, the algorithm will become sensitive to temporal fluctuations in the perfusion signal, as well as to motion occurring between the label and control acquisitions, which will spoil the ranking of function values in the simplex, and lead to remaining signal in the label and control images. In our experiments, $\lambda = 4$ resulted in a good trade-off between eliminating magnitude subtraction artifacts and minimizing the label signal.

One should also note that, because the two-norm of the perfusion signal (used as the regularization term) sums all the pixels together, larger compartments in the image are automatically weighted stronger than smaller compartments. The regularization term should, at the end of the optimization process, also prevent opposite signs in the label and control images for relatively smaller compartments, finding a set of inversion times that satisfies the maximized perfusion constraint for all tissues. However, if the structures are too small, such as for example a small tumor, the changes in perfusion signal in the corresponding pixels may be of the order of the physiological perfusion fluctuations, making them essentially invisible to the optimization scheme. This could, despite achieving a good overall BGS performance, for example result in quantification errors in smaller tumors. In this case, on-the-fly segmentation and assigning weights to each of the compartments based on their respective sizes may circumvent this problem.

A limitation of this study is the fact that the performance of the FBL scan was only demonstrated in single-slice 2D acquisitions, while the largest benefit is expected for volumetric acquisitions. Moving to multislice 2D acquisitions is expected to have little benefit over presetting the BGS timings, because only the first slice would have perfect BGS. The use of single-shot 3D approaches should be a relatively straightforward extension of our proof-of-concept study, although blurring in the z-direction is still a challenge for these readout approaches.^{20,21} However, extension to multishot 3D acquisitions would be more challenging, because the update time would become too long for a FBL approach within a reasonable scan time. This is worth addressing, because the benefit of optimized BGS is expected to be largest for multishot approaches due to artifacts arising from signal variations between shots. Further research is needed to investigate whether the use of timings obtained from a single-slice 2D optimization would also result in improved 3D perfusion images subsequently acquired with the fixed, optimal 2D BGS timings. This would be the case if optimal timings are more different between subjects than between different locations in the brain. Using simultaneous multislice imaging in the optimization process, and thus optimization over a larger coverage of the brain, could in principle result in a better BGS performance over the whole brain.

Although, as discussed, more research is needed to investigate whether a single control acquisition can represent the same relevant perfusion information as the label-control approach, the results in this study have already shown that changes in perfusion signal in a functional, subtractionless FBL scan were similar to those in a traditional label-control scan. BOLD signal may also have contributed to the changes in control signal that we observed, but likely to a much smaller extent compared with perfusion.²² Although the accuracy of perfusion quantification is much reduced and integrated motion compensation becomes difficult with low static signals, the perfusion changes were easily detected (visually and spatially) in the subtraction-free approach. Such an approach may, in the future, open opportunities for functional neurofeedback ASL, where perfusion changes can thus be tracked at much higher temporal resolution to visualize fast changes in the neuronal response. Early detection of such changes could enable effectively adjusting the stimulus presented during scanning based on the perfusion response, or to automatically optimize the neuronal response using neurofeedback.^{23,24} The approach that was taken in this study could furthermore be used to optimize other types of image contrasts, such as that obtained with (double) inversion recovery or fluid-attenuated inversion recovery (FLAIR). Finally, the FBL mechanism offers a flexible framework that could be used to optimize a wide variety of scan parameters, besides inversion times, on the fly.^{25–28} This could help to improve the image quality for scan protocols that are sensitive to system imperfections and other intersubject differences,²⁹ such as in the case of diffusion-weighted imaging.

In conclusion, a FBL mechanism can be used to optimize the performance of BGS in ASL scans. This leads to optimal perfusion quality for each subject in the scanner.

ACKNOWLEDGMENTS

This publication is part of the project HTSM (with project number 17104), which is financed by the Dutch Research Council (NWO). MJPvO is funded by the research program Innovational Research Incentives Scheme VICI with project number 016.160.351, which is financed by the Netherlands Organization for Scientific Research (NWO). This research was in part funded by the Academy Van Leersum grant of the Academy Medical Sciences Fund, Royal Netherlands Academy of Arts & Sciences.

CONFLICT OF INTEREST

Kirsten Koolstra, Marius Staring, and Matthias J. P. van Osch declare that they have no conflicts of interest. Paul de Bruin is an employee of Philips Best.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

ORCID

Kirsten Koolstra  <https://orcid.org/0000-0002-7873-1511>

Marius Staring  <https://orcid.org/0000-0003-2885-5812>

Matthias J. P. van Osch  <https://orcid.org/0000-0001-7034-8959>

REFERENCES

1. Ye FQ, Frank JA, Weinberger DR, McLaughlin AC. Noise reduction in 3D perfusion imaging by attenuating the static signal in arterial spin tagging (ASSIST). *Magn Reson Med*. 2000;44(1):92-100. doi:10.1002/1522-2594(200007)44:1%3C92::AID-MRM14%3E3.0.CO;2-M
2. Shen Q, Duong T. Background suppression in arterial spin labeling MRI with a separate neck labeling coil. *NMR Biomed*. 2011;24(9):1111-1118. doi:10.1002/nbm.1666
3. Jao T, Nayak K. Demonstration of velocity selective myocardial arterial spin labeling perfusion imaging in humans. *Magn Reson Med*. 2018;80(1):272-278. doi:10.1002/mrm.26994
4. Spann SM, Shao X, Wang DJ, et al. Robust single-shot acquisition of high resolution whole brain ASL images by combining time-dependent 2D CAP-IRINHA sampling with spatio-temporal TGV reconstruction. *Neuroimage*. 2020;206:116337. doi:10.1016/j.neuroimage.2019.116337
5. Amukotuwa SA, Yu C, Zaharchuk G. 3D pseudocontinuous arterial spin labeling in routine clinical practice: A review of clinically significant artifacts. *J Magn Reson Imaging*. 2016;43(1):11-27. doi:10.1002/jmri.24873
6. Gevers S, van Osch MJ, Bokkers RPH, et al. Intra- and multicenter reproducibility of pulsed, continuous and pseudo-continuous arterial spin labeling methods for measuring cerebral perfusion. *J Cereb Blood Flow Metab*. 2011;31(8):1706-1715. doi:10.1038/jcbfm.2011.10
7. Madhuranthakam AJ, Yuan Q, Pedrosa I. Quantitative Methods in Abdominal MRI: Perfusion Imaging. *Top Magn Reson Imaging*. 2017;26(6):251.
8. Taso M, Zhao L, Guidon A, Litwiller DV, Alsop DC. Volumetric abdominal perfusion measurement using a pseudo-randomly sampled 3D fast-spin-echo (FSE) arterial spin labeling (ASL) sequence and compressed sensing reconstruction. *Magn Reson Med*. 2019;82(2):680-692. doi:10.1002/mrm.27761
9. Maleki N, Dai W, Alsop DC. Optimization of background suppression for arterial spin labeling perfusion imaging. *Magn Reson Mater Phys Biol Med*. 2012;25(2):127-133. doi:10.1007/s10334-011-0286-3
10. Wansapura JP, Holland SK, Dunn RS, Ball WS. NMR relaxation times in the human brain at 3.0 Tesla. *J Magn Reson Imaging*. 1999;9(4):531-538. doi:10.1002/(SICI)1522-2586(199904)9:4%3C531::AID-JMRI4%3E3.0.CO;2-L
11. Smink J, Häkkinen M, Holthuizen R, et al. eXternal Control (XTC): a flexible, real-time, low-latency, bi-directional scanner interface. Proceedings of the 19th Annual Meeting of ISMRM, Montréal, Canada, 7-13 May 2011. 2011; 1755.
12. Nelder JA, Mead R. A simplex method for function minimization. *Comput J*. 1965;7(4):308-313. doi:10.1093/comjnl/7.4.308
13. Duyn JH, Tan CX, Van Gelderen P, et al. High-sensitivity single-shot perfusion-weighted fMRI. *Magn Reson Med*. 2001;94:88-94. doi:10.1002/mrm.1163
14. Guerquin-Kern M, Lejeune L, Pruessmann KP, Unser M. Realistic analytical phantoms for parallel magnetic resonance imaging. *IEEE Trans Med Imaging*. 2012;31(3):626-636. doi:10.1109/TMI.2011.2174158
15. Mani S, Pauly J, Conolly S, Meyer C, Nishimura D. Background suppression with multiple inversion recovery nulling: Applications to projective angiography. *Magn Reson Med*. 1997;37(6):898-905. doi:10.1002/mrm.1910370615
16. Dixon WT, Sardashti M, Castillo M, Stomp GP. Multiple inversion recovery reduces static tissue signal in angiograms. *Magn Reson Med*. 1991;18(2):257-268. doi:10.1002/mrm.1910180202
17. van Osch MJP, Teeuwisse WM, van Walderveen MA, Hendrikse J, Kies DA, van Buchem MA. Can arterial spin labeling detect white matter perfusion signal? *Magn Reson Med*. 2009;62(1):165-173. doi:10.1002/mrm.22002
18. Garcia DM, Duhamel G, Alsop DC. Efficiency of inversion pulses for background suppressed arterial spin labeling. *Magn Reson Med*. 2005;54(2):366-372. doi:10.1002/mrm.20556
19. Pineda L, Basu S, Romero A, Calandra R, Drozdal M. Active MR k-space sampling with reinforcement learning. International Conference on Medical Image Computing and Computer-Assisted Intervention, Lima, Peru (virtual), Oct 4-8 2020. 2020; 23-33.
20. Paschoal AM, Leoni RF, Pastorello BF, Osch MJP. Three-dimensional gradient and spin-echo readout for time-encoded pseudo-continuous arterial spin labeling: influence of segmentation factor and flow compensation. *Magn Reson Med*. 2021;86(3):1454-1462. doi:10.1002/mrm.28807
21. Günther M, Oshio K, Feinberg DA. Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements. *Magn Reson Med*. 2005;54(2):491-498. doi:10.1002/mrm.20580
22. Liu HL, Pu Y, Nickerson LD, et al. Comparison of BOLD and perfusion based event-related functional MRI: Annual International Conference of the IEEE Engineering in Medicine and Biology – Proceedings, Chicago, USA, 23-28 July. 2000;3:2012-2016.
23. Kim S, Birbaumer N. Real-time functional MRI neurofeedback: A tool for psychiatry. *Curr Opin Psychiatry*. 2014;27(5):332-336. doi:10.1097/YCO.000000000000087
24. Dewiputri WI, Auer T. Functional magnetic resonance imaging (fMRI) neurofeedback: Implementations and applications. *Malaysian J Med Sci*. 2013;20(5):5-15.
25. Schiano JL, Blaich AJ, Ginsberg MD. Optimization of NQR pulse parameters using feedback control. *Z Nat Forsch A J Phys Sci*. 2000;55(1-2):67-73. doi:10.1515/zna-2000-1-213
26. Schiano JL, Routhier T, Blaich AJ, Ginsberg MD. Feedback optimization of pulse width in the SORC sequence. *J Magn Reson*. 1999;140(1):84-90. doi:10.1006/jmre.1999.1824
27. Navkar NV, Deng Z, Shah DJ, et al. Visual and force-feedback guidance for robot-assisted interventions in the beating heart with real-time MRI. IEEE International Conference on Robotics and Automation, St Paul, USA, 14-18 May 2012. 2012;689-694.
28. Christoforou E, Akbudak E, Ozcan A, Karanikolas M, Tsekos NV. Performance of interventions with manipulator-driven real-time MR guidance: implementation and initial in vitro tests. *Magn Reson Imaging*. 2007;25(1):69-77. doi:10.1016/j.mri.2006.08.016

29. Gagoski B, Xu J, Wighton P, et al. Automatic detection and reacquisition of motion degraded images in fetal HASTE imaging at 3T. 28th Annual Meeting of the International Society for Magnetic Resonance in Medicine, Sydney, Australia (virtual), 18-23 April 2020. 2020.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Koolstra K, Staring M, de Bruin P, van Osch MJP. Subject-specific optimization of background suppression for arterial spin labeling magnetic resonance imaging using a feedback loop on the scanner. *NMR in Biomedicine*. 2022;35(9):e4746. doi:[10.1002/nbm.4746](https://doi.org/10.1002/nbm.4746)