

RESEARCH ARTICLE

Shim optimization with region of interest-specific Tikhonov regularization: Application to second-order slice-wise shimming of the brain

Yuhang Shi  | Stuart Clare | Signe Johanna Vannesjo 

Wellcome Centre for Integrative Neuroimaging, FMRIB, University of Oxford, John Radcliffe Hospital, Oxford, United Kingdom

Correspondence

Johanna Vannesjo, Department of Physics, NTNU, Høgskoleringen 5, 7491 Trondheim, Norway.
Email: johanna.vannesjo@gmail.com

Funding information

Dunhill Medical Trust; H2020 Marie Skłodowska-Curie Actions, Grant/Award Number: 659263; Wellcome Trust, Grant/Award Number: 202788/Z/16/Z and 203139/Z/16/Z

Abstract

Purpose: Slice-wise shimming can improve field homogeneity, but suffers from large noise propagation in the shim calculation. Here, we propose a robust shim current optimization for higher-order dynamic shim updating, based on Tikhonov regularization with a variable regularization parameter, λ .

Theory and Methods: λ was selected for each slice separately in a fully automatic procedure based on a combination of boundary constraints and an L-curve search algorithm. Shimming performance was evaluated for second order slice-wise shimming of the brain at 7T, by simulation on a database of field maps from 143 subjects, and by direct measurements in 8 subjects.

Results: Simulations yielded on average 36% reduction in the shim current norm for just 0.4 Hz increase in residual field SD as compared to unconstrained unregularized optimization. In vivo results yielded on average 34.0 Hz residual field SD as compared to 34.3 Hz with a constrained unregularized optimization, while simultaneously reducing the shim current norm to 2.8 A from 3.9 A. The proposed regularization also reduced the average step in the shim current between slices.

Conclusion: Slice-wise variable Tikhonov regularization yielded reduced current norm and current steps to a negligible cost in field inhomogeneity. The method holds promise to increase the robustness, and thereby the utility, of higher-order shim updating.

KEYWORDS

dynamic B_0 shimming, higher-order shimming, regularization, shim optimization

1 | INTRODUCTION

Most MR acquisitions benefit from a homogeneous background magnetic field, especially imaging techniques relying on fast echo planar imaging (EPI) readouts, such as

functional MRI and diffusion-weighted imaging. The B_0 field is, however, distorted by differences in magnetic susceptibility between air and tissue, causing image artifacts, such as signal loss, blurring, and geometric distortion. To counteract the B_0 field inhomogeneity, modern MR

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Magnetic Resonance in Medicine* published by Wiley Periodicals LLC on behalf of International Society for Magnetic Resonance in Medicine

scanners are equipped with shim coils that generate low-order spherical harmonic magnetic fields.^{1,2} The currents through the shim coils are adjusted to improve magnetic field homogeneity within a given region of interest (ROI). Typically, the ROI is defined to be the full imaging volume, such as the whole brain. However, optimizing the shims for smaller volumes can generally improve field homogeneity further, as the field in smaller regions can be better approximated by low-order spherical harmonics.³ Hence, it has been suggested to perform shimming on a slice-wise basis instead of over the full imaging volume in 2D imaging. Slice-wise shimming has been shown to improve field homogeneity and geometric accuracy of EPI acquisitions in high-field brain imaging.⁴⁻⁸

Optimal shim currents are typically obtained by minimizing the residual field SD within a given ROI based on an acquired B_0 field map.^{2,9} As the ROI gets smaller, however, fewer data points are included in the optimization, which consequently becomes more susceptible to noise and measurement errors in the B_0 field map. Additionally, reduced ROIs are more likely to degrade the orthogonality of the spherical harmonic shim terms, especially for off-center ROIs.¹⁰ The reduced orthogonality makes the shim calculation increasingly ill-conditioned, thereby exacerbating the noise propagation. Taken together, this leads to increasingly unstable shim calculations for smaller shim volumes, depending on the shape and location of the volume. Unstable shim calculation has various negative effects, including less robust field homogenization, larger variability between slices/ROIs, and excessively high shim currents.

High shim currents may exceed hardware constraints of the shim amplifiers, especially for the higher-order shims. To ensure hardware compliance, it is common to use constrained optimization methods for slice-wise shimming.¹¹⁻¹⁵ Constrained optimization, however, does not alleviate the degree of ill-conditioning of the optimization problem. The calculated solution is therefore still susceptible to noise, and may result in unnecessarily high shim currents, although within hardware limits. High currents combined with large variability leads to large shim steps between slices, which increases eddy currents induced after shim switching. Well-calibrated eddy current compensation is crucial for higher-order dynamic shimming applications.⁶⁻⁸ Any residual uncompensated eddy current fields may degrade image quality, with the effect scaling with the size of the shim steps. Eddy current compensation requires overdriving the shim currents at each shim step, thereby limiting the range of shim currents available for field homogenization.^{7,8} Reducing the size of shim steps is therefore beneficial, also with eddy current compensation. To this end, it has been suggested to directly include the shim steps as an additional term in the cost function of the shim optimization.¹⁶

To improve the orthogonality of the shim optimization, it has been proposed to omit degenerate shim terms in the shim calculation.^{6,17} However, the degeneracy analysis only works for specific volumes, eg thin slices. Furthermore, truncation of through-slice shim terms may degrade the achievable shimming quality, as signal loss in 2D imaging is largely related to field variations in the through-slice direction. To avoid shim degeneracy in slice-wise shimming, the shim volume can be extended to include additional slices centered at the target slice.⁸ However, even when covering a few slices, the orthogonality of the shim fields within a thin shim volume is likely to be poor, and the shim optimization correspondingly ill-conditioned.

A common approach to deal with ill-conditioned problems in various areas of engineering is to use regularized optimization. For shim optimization, Kim et al demonstrated that truncated singular value decomposition (TSVD) improved the numerical stability of whole-brain static shimming.¹⁸ A fundamental question in any regularization approach is how to choose the degree of regularization. Too little regularization yields a solution that is sensitive to perturbations in the data, and results in high shim currents. Too much regularization, on the other hand, yields a robust solution and low shim currents, but at the cost of suboptimal field homogeneity. In TSVD the degree of regularization is controlled by the number of truncated singular values. Nassirpour et al proposed a TSVD-based optimization approach where the number of truncated singular values is adjusted for each ROI separately to ensure that the shim currents do not exceed hardware limits.¹⁹ However, like the constrained optimization methods, the hardware-limited TSVD only regularizes solutions that otherwise would have exceeded hardware constraints.

In this work, we present a novel approach to regularized shim optimization based on Tikhonov regularization, where the degree of regularization is individually tuned for each ROI (eg, per slice). The variable selection of the regularization parameter serves to balance the trade-off between shim current norm and residual field homogeneity depending on the conditioning of the problem in each ROI. The selection is performed in a fully automated procedure designed to satisfy three aims: (i) ensuring compliance with hardware constraints, (ii) including a small degree of regularization in all ROIs to improve robustness to noise, while (iii) avoiding excessively increased residual field SD where possible. The performance of the proposed shim calculation was evaluated for second-order slice-wise shimming of the brain at 7T in 143 subjects in simulation and 8 subjects in experiment. The proposed method was compared to previously suggested shim optimization methods, including constrained linear least-squares (CLS) optimization and hardware-constrained TSVD.

2 | THEORY

2.1 | Linear least-squares shim optimization

Shim optimization is typically performed by minimizing the L2 norm of the residual magnetic field:

$$\min \|\Delta B - S \cdot I\|_2^2, \quad (1)$$

where ΔB is a $N_V \times 1$ vector with the measured magnetic field offset in N_V voxels inside the ROI, S is a $N_V \times N_C$ matrix representing the magnetic field profile of each shim channel within the ROI, and I is a $N_C \times 1$ vector with the unknown shim currents. The number of shim channels, N_C , is typically much lower than N_V , leading to an overdetermined linear least-squares problem, with a solution given by the pseudo-inverse of the matrix S , denoted by S^+ :

$$I_{LS} = S^+ \cdot \Delta B.$$

The pseudoinverse yields the theoretically optimal solution, given unlimited current supply. In reality, however, the shim hardware imposes limits on the maximum shim currents, leading to a constrained minimization problem:

$$\min \|\Delta B - S \cdot I\|_2^2, \quad |I_c| \leq I_{\max}.$$

There is generally no analytic solution to the CLS problem, but it can be obtained numerically with various iterative solvers. Noise and measurement errors in ΔB propagate into the shim current solution, I_{LS} , at a rate that is inversely proportional to the singular values of S , meaning that small singular values amplify perturbations in corresponding subspaces. Low orthogonality between the shim terms, as frequently encountered for small shim volumes, generally degrade the conditioning of the shim matrix S , such that small perturbations in the measured field map lead to large deviations in the calculated shim currents.

Regularization is a common approach to improve the conditioning of least-squares problems. TSVD, is designed to directly remove the contribution of subspaces with small singular values, where the truncation factor, k , determines the degree of regularization. Tikhonov regularization instead adds the solution norm to the cost function of the optimization problem:

$$\min \left\{ \|\Delta B - S \cdot I_\lambda\|_2^2 + \lambda^2 \|I_\lambda\|_2^2 \right\},$$

corresponding to dampening the rate of decay of the singular values.²⁰ In the context of shim optimization, Tikhonov optimization thus allows for a direct trade-off between minimizing the residual field inhomogeneity and




FIGURE 1 Comparison of the shim current norm versus the residual field SD of the TSVD (circles and diamonds) and the Tikhonov (lines) regularization when varying the regularization level. The blue diamonds/line show one example of whole-brain shimming, and the red circles/line show shim optimization for one example slice. The left end point of each curve corresponds to the results with no regularization ($\lambda = 0$), which is equivalent to the unconstrained method

reducing the shim current norm. While larger regularization reduces errors due to noise, it also introduces a penalty in the form of regularization errors. A suitable choice of the regularization factor, thus, becomes crucial. The relation between the residual and the solution norm while varying the regularization parameter using Tikhonov regularization and TSVD is visualized in Figure 1, for whole-brain static shimming and for an example slice in the middle of the brain. The current norm is attenuated with increasing regularization parameters at the expense of increased residual field norm. It can be noted that the Tikhonov regularization factor λ is a continuous variable, in contrast with the discrete TSVD truncation factor k . Thus, Tikhonov regularization with varying λ can be sampled on an arbitrarily fine grid, producing a curve commonly referred to as the L-curve, due to its characteristic L-shape. Note also that the L-curve defines the lower boundary of possible solutions to a given optimization problem, ie any solution with or without regularization will lie on or above the L-curve.

2.2 | Slice-wise variable Tikhonov regularization

We here propose a Tikhonov-based regularization with variable λ , such that the degree of regularization can be tuned to the condition of each specific ROI. One approach to systematic selection of the regularization parameter is to pick λ at the corner of the L-curve where the curvature

reaches its maximum.²¹ We propose to combine the L-curve search method with boundaries on the search range, and term the method Tikh-LC. The procedure is illustrated schematically in Figure 2, and outlined step-by-step below:

1. First, the L curve (blue line in Figure 2B,C) is calculated for a wide range of regularization parameters (including the unconstrained solution, $\lambda = 0$).
2. In the second step, a set of boundaries are defined that limit the range of acceptable solutions:
 - the maximum current norm complying with hardware limits (pink line),
 - the maximum allowed increase in residual field SD (green line) in order to avoid over-regularization, and
 - the minimum improvement in the conditioning of the problem (black line) in order to always ensure some degree of regularization.
3. The set of boundaries create two cases that need to be handled separately:
 - If solutions exist within the boundaries (Figure 2B), the L-curve method is used to select the optimal regularization parameter within that range (dashed red line).²⁰
 - If no solution fulfills all three boundary conditions (Figure 2C), ie when the maximum allowed increase in the residual field SD is exceeded due to shim current constraints, the solution with the lowest residual field SD within shim current constraints is selected (dashed red line).

3 | METHODS

3.1 | Shim setup

Experiments were performed on a Siemens Magnetom whole-body 7T MRI scanner (Siemens Healthineers, Erlangen, Germany), equipped with actively shielded 70 mT/m gradients and second-order spherical harmonic shim coils without active shielding. Imaging data were acquired with a 32-channel head coil (Nova Medical, Wilmington, MA). All in vivo acquisitions were performed in compliance with the local Institutional Review Board.

Dynamic updating of the second-order shims was performed using shim amplifiers from Resonance Research Inc (RRI, Billerica, MA, USA). The shim amplifiers were controlled via auxiliary measurement control units (MPCU—measurement and physiologic control unit) of the parallel transmit (pTx) system of the scanner. Digital signal from the MPCUs was converted to analogue and connected to the analogue voltage input of the shim amplifiers. Each MPCU exerted real-time control over three virtual gradient channels, which were mapped to the shim channels. Hence two auxiliary MPCUs were needed to perform shim updating on all five second-order shim channels. Control over the shim waveforms via the MPCUs was accessible through the sequence programming of the pTx system. First-order shim updating was achieved using the gradients, by setting slice-wise gradient offsets in the ordinary sequence programming, and the zeroth-order shim term was

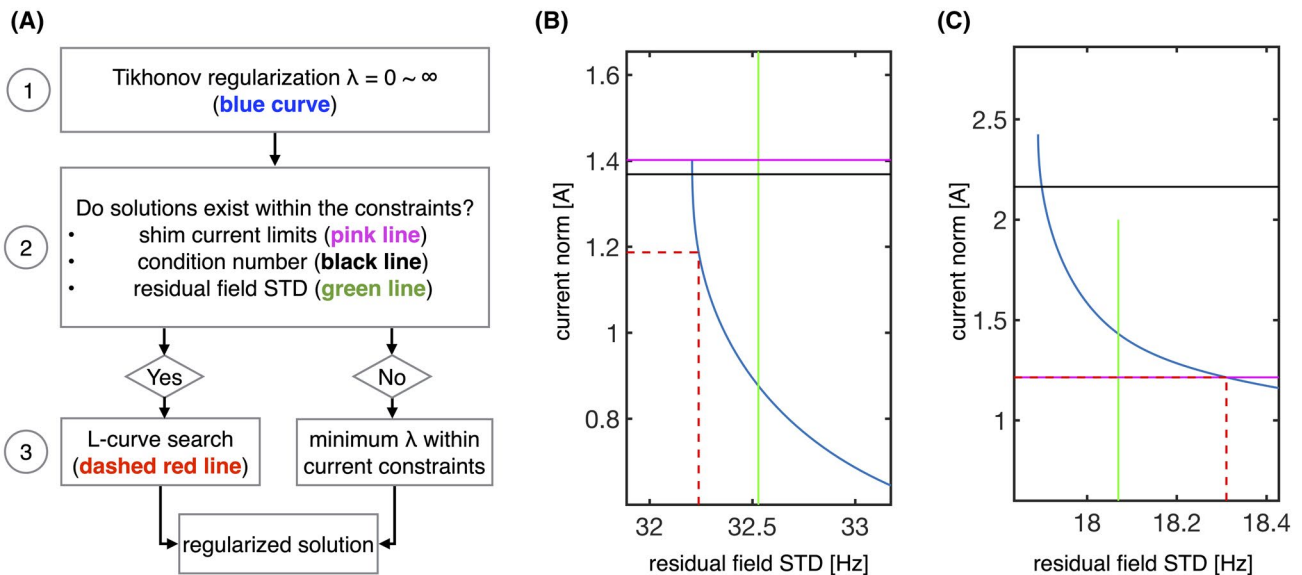


FIGURE 2 A, Flow chart of the Tikh-LC regularization for shim optimization, showing how the regularization parameter is selected. B,C, The selected shim solution (dashed red line) is determined from a range of regularized solutions (blue line) using the L-curve method along with boundary constraints on the minimum regularization level (black line), maximum current norm (pink line) and maximum residual field SD (green line)

controlled by adjusting the transmit/receive frequency offset. Thus it was possible to synchronize the slice-wise updating of all zero- to second-order shim terms with the imaging sequence via the sequence programming software.

Higher-order shims typically couple strongly with the cryostat of the magnet and thereby induce long-living eddy currents that hinder fast shim updating. Multi-exponential eddy current compensation with up to five terms was, therefore, implemented for all second-order shims, including cross-terms to B_0 and selected second-order shims.^{22,23} Cross-term eddy current compensation between second-order shims was possible for shim channels controlled by the same auxiliary MPCU. An image-based calibration scan was used to measure the eddy currents induced by each shim channel.²⁴ Eddy current amplitudes and decay time constants were fitted to the measured eddy currents, and were used to calculate the pre-emphasis parameters.^{22,23} The resulting eddy current compensation parameters are given in Table S1, which is available online. In order to minimize effects of residual short-lived eddy currents after pre-emphasis, a delay time of 10 ms after each shim update was included in the imaging sequences.

The spatial field distribution of each individual shim coil was calibrated as detailed in the Supporting Information. The resulting field distributions were subsequently used to create the shim matrix, S , for each imaging volume. To comply with power and cooling constraints of the shim amplifiers, the shim currents for the second-order shim channels were limited to a maximum of 5 A per channel, and a total of 20 A summed over all second-order channels. After pre-emphasis, the output currents can temporarily reach over 150% of the input amplitude. Thus, in the slice-wise shim calculations, the limits of the second-order shims were set to 3 A to prevent output currents with pre-emphasis from exceeding the current constraints of the amplifiers.^{7,8}

3.2 | Simulations

Shim simulations were conducted on a database with field maps from 143 healthy subjects (age range 20–66). The field maps were acquired as part of various studies running at the center, and were shared anonymously in compliance with approved ethics protocols. For each subject, a B_0 field map was acquired using a 2D dual gradient-echo sequence (2 mm isotropic resolution, 64 angulated axial slices, TE1 4.08 ms, Δ TE 1.02 ms, flip angle 39°, TR 620 ms, field of view [FOV] 220×220 mm²). Before acquisition, global static shimming was performed with the vendor-supplied

shimming routine, and the corresponding shim settings were recorded in the DICOM files. The obtained field maps were unwrapped using PRELUDE (Phase Region Expanding Labeller for Unwrapping Discrete Estimates).²⁵ Subsequently, the shim field of the global static shimming was removed to yield field maps as if acquired without any shim currents. Whole-brain masks were generated by applying BET (Brain Extraction Tool) to the magnitude images of the field map acquisition.²⁶

Simulated slice-wise shimming was performed on each subject in the database. To include information about through-slice field inhomogeneity for each slice, one neighboring slice on each side was included in the shim calculation.⁸ The calculation only included signal within the brain mask. All simulations were implemented in the MATLAB R2017a environment (MathWorks Inc, Natick, MA, USA) on a computer with 2.6 GHz Intel Core i7 CPU and 16 GB of RAM.

The unconstrained linear least-squares solution was calculated with the pseudoinverse of the shim matrix. A CLS solution was obtained with a Matlab-embedded solver using the interior-point algorithm. Hardware-constrained TSVD as described in Ref. [19] was implemented by iteratively increasing the truncation factor until the solution satisfied shim current constraints. The proposed Tikh-LC method was implemented with the boundaries set to improve the condition number by 20%, while keeping the increase in residual field SD below 1% of the residual field of the unconstrained solution. To quantify the performance of the different approaches, the shim current norm and the residual field SD after simulated shimming were calculated in each case. The unconstrained pseudoinverse solution was used as a reference point for all comparisons. In order to assess the trade-off between shim current and residual field SD with the different optimizations, we calculated the ratio of the average reduction in the current norm to the average degradation in the residual field SD, compared to the unconstrained solution. The average was taken over all slices in the database. We term this measure the “regularization efficiency”. Additionally, the norm of the shim steps between slices was quantified, including the step from the last to the first slice. As regularization reduces the shim current norm, the difference in shim currents between slices is also likely to be reduced, thereby yielding the acquisitions less sensitive to residual eddy current fields.

In order to determine the relative benefit of using Tikhonov regularization with a slice-wise variable regularization parameter, the Tikh-LC method was additionally compared to a constrained Tikhonov regularization with λ fixed at 0.1 (Tikh-F1, low weighting of solution norm), 1 (Tikh-F2, equal weighting of residual and solution norm),

1.7628 (Tikh-F3, median of all Tikh-LC regularization parameters), and 3.0168 (Tikh-F4, mean of all Tikh-LC regularization parameters). The fixed values of λ were chosen such as to span a range from overall less to overall more regularization than the slice-wise variable approach. In slices where the calculated shims exceeded hardware constraints, the solution that yielded lowest residual SD was adopted.

To assess the generalizability of the results to different higher-order shim systems, simulations were additionally performed for systems with up to third- and fourth-order shims, based on values from Ref. [19]. The results of these evaluations are given in the Supporting Information.

3.3 | In vivo validation

To evaluate the practical performance of the proposed Tikh-LC method in second-order slice-wise shimming of the brain, eight subjects who were not part of the database were scanned. In each subject, a B_0 field map (2 mm isotropic resolution, 75 axial slices, TE1 4.08 ms, Δ TE 1.02 ms, TR 1500 ms, flip angle 39°, FOV 220 × 220 mm²), and an EPI acquisition (2 mm isotropic resolution, 75 axial slices, TR 3000 ms, TE 28 ms, FOV 220 × 220 mm², GRAPPA factor 2, echo-spacing 0.72 ms) were obtained with a global static shim, and a slice-wise shim using the CLS, the TSVD and the Tikh-LC shim optimization. Shim optimization was performed based on a B_0 field map acquired with the system's baseline shim settings ("tune-up"). The global static shim settings were calculated by a pseudo-inverse of the shim matrix, ie, without regularization or current constraints, using the whole brain as ROI. The performance of the different methods was quantitatively evaluated by the residual field SD in the measured field maps, the shim current norm and the average shim step. To test for significant differences between the regularization methods, a Wilcoxon signed rank test was used, with significance thresholds at $P < .05$ (significant) and $P < .01$ (highly significant). A nonparametric test was chosen as the data was not normally distributed. No multiple comparison correction was performed. The EPI acquisitions were used to qualitatively evaluate geometric distortion and signal dropout.

4 | RESULTS

4.1 | Simulated shimming

Figure 3 displays results from a single subject, comparing simulated shimming using CLS, TSVD and Tikh-LC optimization across all slices of a field map. Regularization

improved the conditioning of the optimization problem in all slices for Tikh-LC, and in slices where the unconstrained solution exceeded current constraints for TSVD (Figure 3B). It can be noted that the current constraints were more likely to be reached at the edge of the FOV where the ROI sizes are small. In slices exceeding current constraints, the TSVD generally produced higher residual field SD than the CLS and Tikh-LC optimizations (Figure 3A). The difference is particularly apparent around slice 9, at the level of the orbitofrontal cortex, where the TSVD degraded field homogeneity by 15 Hz, as compared to 4 Hz and 5 Hz for CLS and Tikh-LC, respectively. The Tikh-LC shim optimization reduced the current norm in all slices, but without substantially degrading the residual field homogeneity (Figure 3A,C). The reduced current norm led to more stable shim settings across slices, as exemplified here with the ZY and Z2 shim channels (Figure 3D).

Figure 4 shows the results of simulated shimming for all slices in all subjects in the database, relating the reduced current norm to the increased residual field SD for the CLS, TSVD, and Tikh-LC methods, as compared to the unconstrained solution. In slices that do not exceed the current constraints (Figure 4A), CLS and TSVD are equal to the unconstrained optimization. In contrast, Tikh-LC provides a reduction in the current norm by 29% on average without degrading the SD of the residual field by more than 0.16 Hz on average (max 1 Hz) compared to the unconstrained method. In slices that do exceed the current constraints (Figure 4B), all three methods provide a reduction in the current norm at the cost of increased SD. The TSVD method delivered the largest average reduction in current norm (59%), as well as suffering the largest average increase in SD of the residual field (2.1 Hz). The Tikh-LC method yielded nearly as large reduction in current norm (53%) as TSVD, but at a much lower cost in increased residual field SD (0.9 Hz). The increased residual field SD of Tikh-LC was close to that of CLS (0.6 Hz), which however did not reduce the current norm by nearly as much (32%).

Figure 5 compares simulated shimming performance using the CLS, TSVD and Tikh-LC optimizations compared to Tikhonov regularization with fixed λ (Tikh-F1 – Tikh-F4). Most slices yielded unconstrained shim solutions below the current constraints, hence the median of the hardware-constrained methods is at zero reduction in condition number, current norm, and increase in residual field SD. However, as also noted in Figure 4, in slices where the unconstrained solution exceeds current constraints, the TSVD tends to regularize more than Tikh-LC, and thus both the average and the maximum increase in residual field SD is larger for TSVD than for Tikh-LC (average 0.6 Hz vs. 0.4 Hz, max 27.2 Hz vs. 14.2 Hz). Despite larger residual field SD,

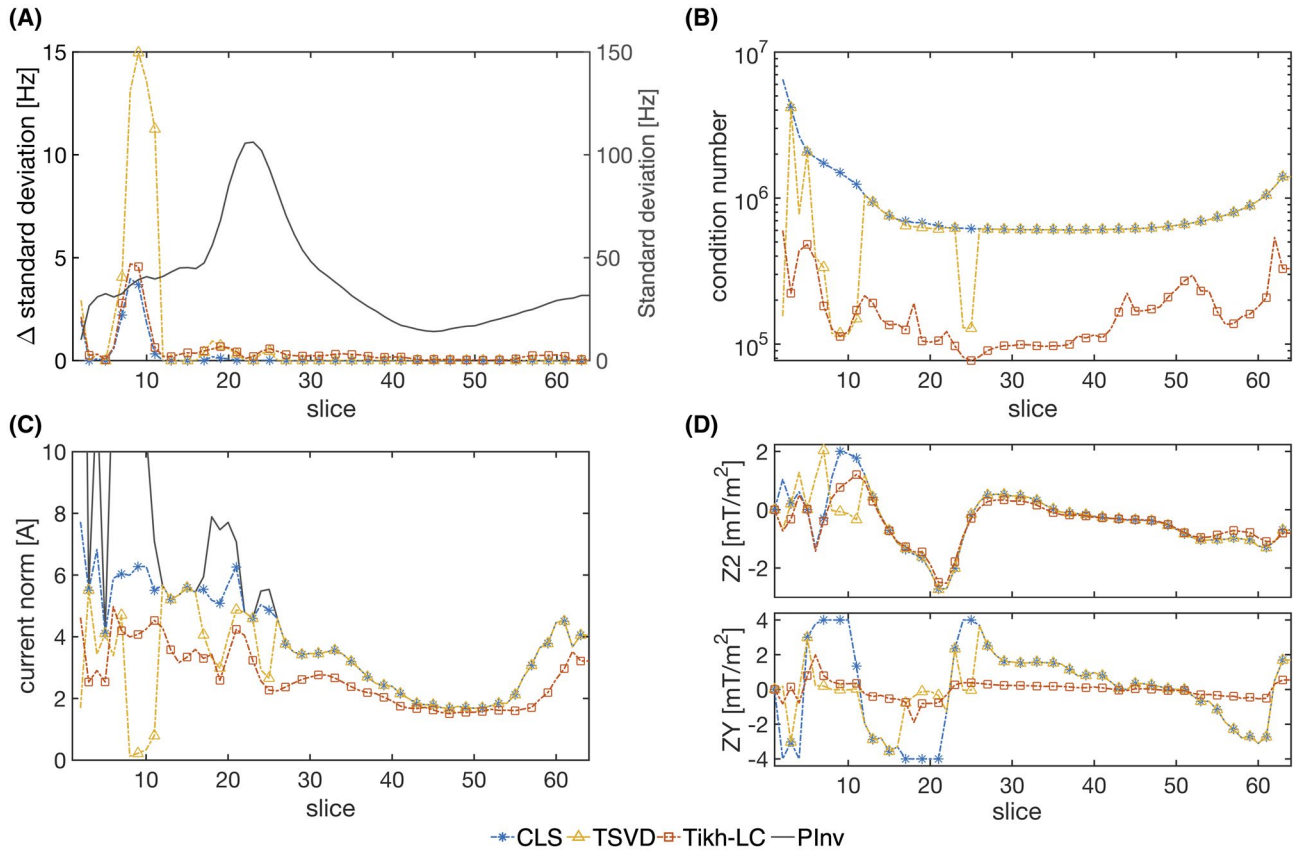


FIGURE 3 Comparison of simulated shimming performance of CLS (blue), TSVD (yellow), and Tikh-LC (red) shim optimization across all slices of a field map acquired from one subject. Subplot (A) shows the residual field SD of the unconstrained solution (black line, y-axis to the right), and the increase in residual field SD with the different regularization methods (blue/yellow/red lines, y-axis to the left). Subplot (B) indicates the condition number, (C) shows the current norm, and (D) plots the amplitude of the Z2 and ZY shims for each slice

the average reduction in current norm is lower for TSVD compared to Tikh-LC (17% vs. 36%), as the latter effectively reduces the current norm in all slices. Moreover, the shim step norm was larger on average for TSVD as compared to Tikh-LC (1.1 A vs. 0.7 A).

For the Tikhonov regularization with fixed λ , increasing the regularization parameter yielded successively lower condition number, reduced current norm, increased residual field SD and decreased shim step norm, as expected (Figure 5). In comparison with Tikh-LC, the average and maximum increase in SD of Tikh-LC corresponded to the lowest fixed regularization levels (Tikh-F2 and Tikh-F1, respectively), whereas the average reduction in current norm was more similar to higher regularization levels (between Tikh-F2 and Tikh-F3). The Tikh-LC regularization had the highest regularization efficiency of the methods compared (Figure 5F).


4.2 | In vivo validation

Figure 6 compares field maps and EPI images acquired with global shimming and slice-wise shimming based on

CLS, TSVD, and Tikh-LC shim optimization in one subject. To visualize image distortion artifacts, the EPI images are overlaid with green contour outlines generated from the magnitude image of the field map acquisition with global static shimming. Slice-wise shimming generally yielded better field homogeneity and reduced distortion artifacts compared to global static shimming. EPI image quality was comparable between CLS and Tikh-LC shim optimization, while the latter yielded a substantial reduction in the current norm. TSVD yielded similar image quality in most slices, but increased image distortion and signal loss in some slices at the level of the brainstem and cerebellum.

A quantitative comparison of shimming results in the eight subjects measured with the different shim optimization methods is shown in Figure 7. In all subjects, the slice-wise shimming substantially improved field homogeneity as compared to global static shimming. The average of the residual field SD within the brain was 44.4 Hz for static shimming, and for slice-wise shimming it was 34.3 Hz with CLS, 34.8 Hz with TSVD, and 34.0 Hz with Tikh-LC shim optimization. A Wilcoxon signed rank test showed a significant difference in the residual field SD

(A) slices not exceeding current constraints



(B) slices exceeding current constraints




FIGURE 4 Comparison of simulated shimming performance with the CLS (yellow), the TSVD (red), and the Tikh-LC (blue) method on all slices for all subjects in the database. The scatter plots (each dot representing one slice) show the relative current norm versus the increase in residual field SD, as compared to the unconstrained method. The two subplots show slices where the unconstrained solution (A) does not exceed current constraints, and (B) does exceed the current constraints. Along each axis, the corresponding box plots are shown, with the middle line/cross indicating the median/average, the box representing the 25th and 75th percentile, and the dashed lines showing the full range. Note the different scaling of the x-axis in (A) and (B)

between TSVD and Tikh-LC, but not between CLS and TSVD or CLS and Tikh-LC. The current norm was significantly lower for both the TSVD (3.0 A) and the Tikh-LC (2.8 A) shim optimization as compared to CLS (3.9 A). In addition to the reduced current norm, the regularized shim optimizations yielded significantly lower shim steps as compared to the CLS, with the Tikh-LC method delivering the lowest shim step norm, indicating smoother shim changes across slices.

5 | DISCUSSION

We have here presented a novel approach to regularized shim optimization based on a ROI-wise variable Tikhonov regularization. The proposed method allows to select the degree of regularization for each ROI separately in order to ensure shim robustness and hardware compliance, while avoiding excessive residual field variance due to over-regularization. The selection of the regularization parameter happens in a fully automated procedure based on a combination of target boundaries and an L-curve search method. The feasibility of the proposed approach was demonstrated for slice-wise shimming of the brain at 7T, in simulation and in direct measurements. The variable Tikhonov regularization effectively reduced the current norm, while preserving good field homogeneity. Interestingly, the proposed regularization even yielded slightly lower average residual field variance than the

theoretically optimal constrained optimization in measured field maps, while at the same time delivering almost 30% reduction in the current norm. These results suggest that the variable regularization can increase the robustness of the shim optimization to noise and measurement errors, without incurring a penalty in increased field variance in practical applications.

A key feature of the proposed Tikh-LC regularization is the variable selection of the regularization parameter per ROI. Similarly to previously proposed methods, the regularization parameter is adjusted such as to ensure hardware compliance. Additionally, we include boundaries to ensure at least a minimal degree of regularization in each ROI to increase the robustness of the shim optimization problem. The boundaries define a reasonable domain for the regularization, in terms of quantities that are more intuitive to work with than the regularization parameter itself. The exact values for the boundaries may be chosen depending on the specific application. As an upper bound, we set the increase in residual field SD to be no more than 1% of the residual field of the unconstrained solution, in most cases corresponding to <1 Hz. A field offset of 1 Hz corresponds to a shift of about 0.04 pixels in the EPI acquisitions obtained for evaluation of the shimming in this work, and was thus deemed to be largely negligible. Within the range defined by the boundaries, we then used an automatic L-curve search algorithm to select the optimal regularization parameter. Alternatively, other methods for automatic parameter selection may be employed, such as

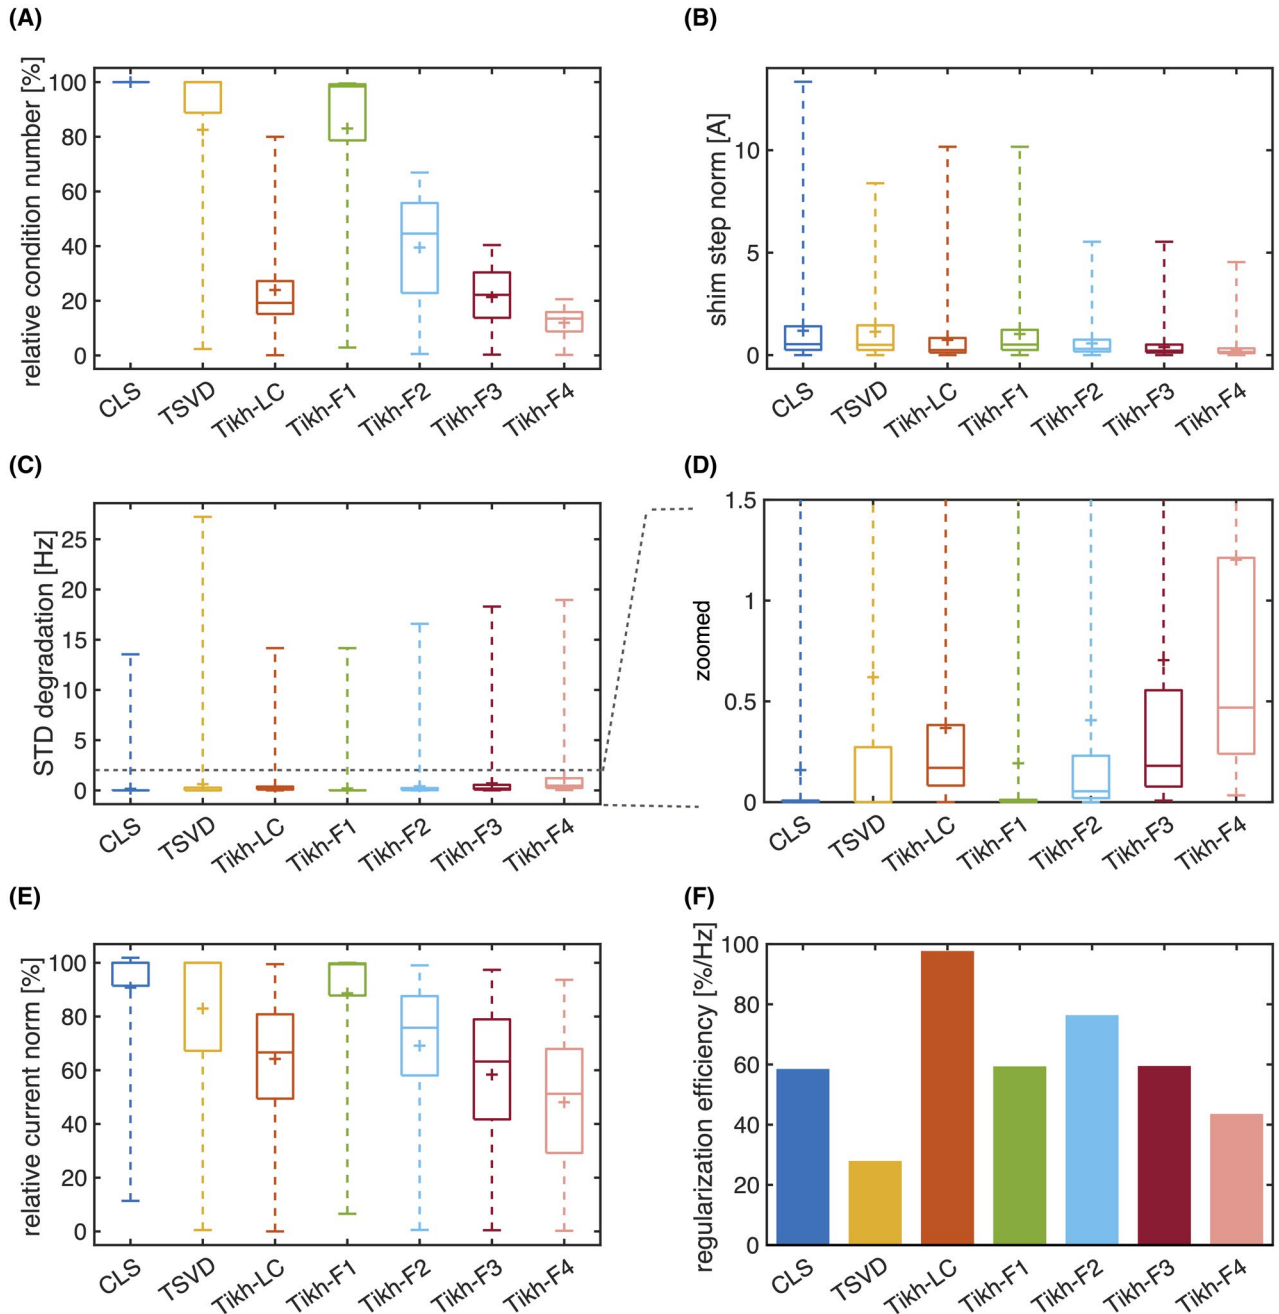


FIGURE 5 Comparison of simulated shimming performance of the CLS, TSVD, Tikh-LC, and Tikh-F1-F4 regularization for all subjects in the database. Plots A-E indicate relative condition number, shim step norm, SD degradation and relative current norm of the different methods, compared to the PINV method, respectively. The middle line/cross indicates the median/average, the box represents the 25th and 75th percentile, and the dashed line shows the full range. Plot F shows the regularization efficiency of the shim optimizations, calculated as the ratio of the average reduction in current norm to the average increase in residual field SD over all slices in the database

the Generalized Cross-Validation method. The variable selection of λ allowed the optimization to tune the level of regularization for each slice, such as to take advantage of both lower and higher regularization levels. The advantage of this approach was reflected in the regularization efficiency, which showed that the Tikh-LC regularization on average yielded the largest relative reduction in current norm per increase in residual field SD compared to any

of the methods compared, including Tikhonov regularization with λ fixed at different levels.

We chose Tikhonov regularization as the base regularization approach for two main reasons: (i) Tikhonov regularization yields the lowest attainable current norm for any given increase in residual field SD, because it directly minimizes a weighted combination of the solution residual and the solution norm, and (ii) it has a smoothly

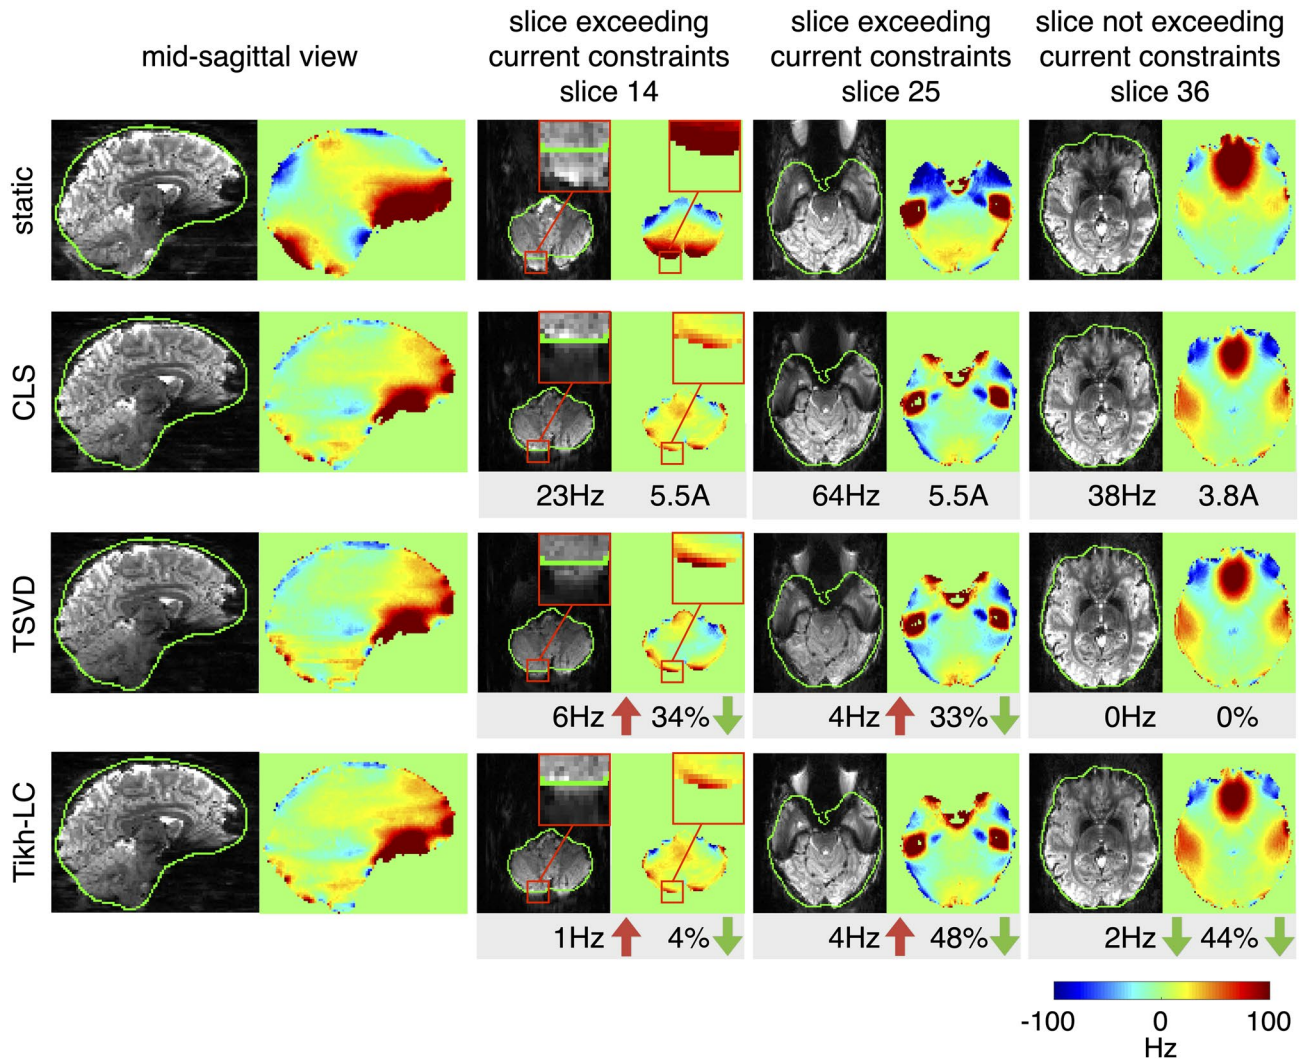


FIGURE 6 EPI images and field maps acquired using the different shim optimization methods in one subject. The columns from left to right display a sagittal image and transversal images at the brain stem and the mid-brain regions. The rows from top to bottom show images acquired with global static shimming, and slice-wise shimming based on CLS, TSVD, and Tikh-LC shim optimization, respectively. The EPI images are overlaid with green contour outlines generated from the field map acquisition with global static shimming. Current norm and residual field SDs are reported below each transversal slice for the CLS optimization. For TSVD and Tikh-LC the difference to CLS is shown

adjustable regularization parameter, allowing for flexible tuning of the regularization level of each ROI. A variable regularization approach could, however, also be performed based on other regularization methods, such as TSVD. In TSVD, the level of regularization corresponds to the number of singular values truncated. For low-order shim systems, with few shim channels in total, the shim matrix has a limited number of singular values. The discretization of regularization levels is therefore rather crude, which may lead to abrupt changes in the shim settings when truncating singular values. This effect may explain the tendency of the TSVD to over-regularize in slices exceeding current constraints in our data. For shim systems with a larger number of shim channels, the discretization effect is likely to be lower, and the two regularization approaches may be closer to interchangeable.

The proposed regularization was shown to effectively reduce shim steps between slices, which can make it less sensitive to residual eddy currents in dynamic shimming applications. It has previously been proposed to use Tikhonov regularization with an added differential term to reduce shim steps.¹⁶ Potentially, the two approaches could be combined, although it may make the automated selection of regularization parameters more computationally demanding.

The Tikh-LC method depends on determining the solution to the shim optimization problem for a large range of λ before selecting one, in a procedure repeated separately for each ROI. One concern may thus be related to computational costs of the method. However, given that there is an analytic solution to the linear least-squares optimization with Tikhonov regularization, and

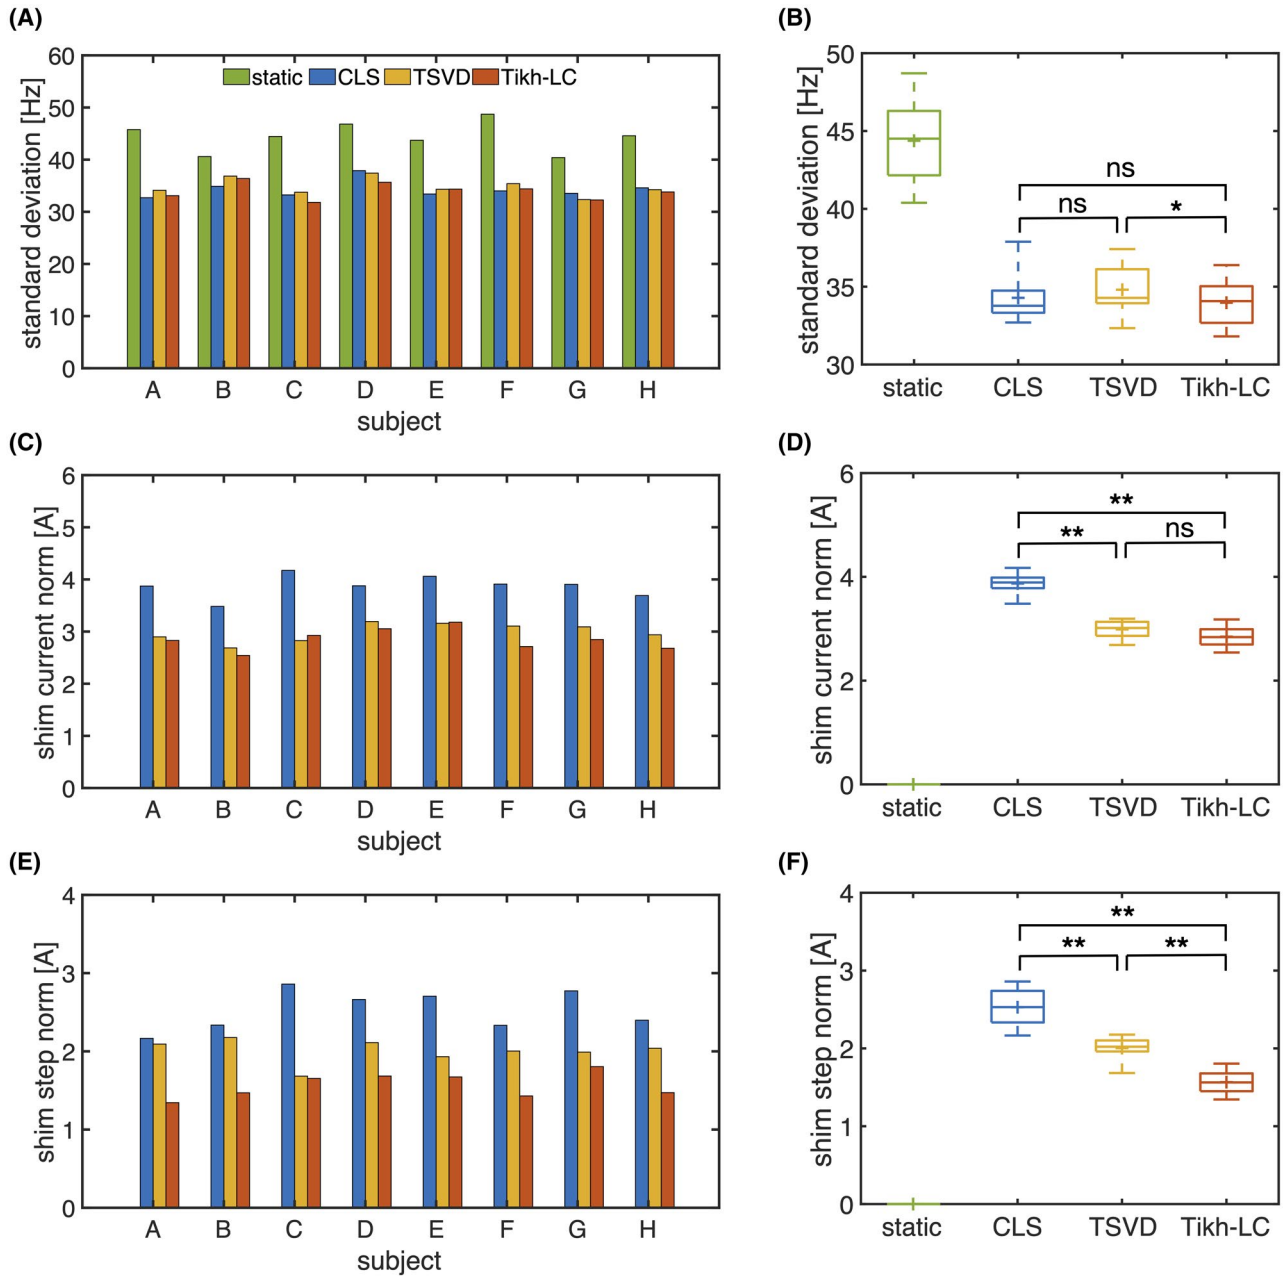


FIGURE 7 Shimming results in all eight subjects with direct measurements of field maps shimmed with a static shim, and a slice-wise shim calculated based on the CLS, the TSVD, and the Tikh-LC optimization. The rows show the SD of the residual field within the brain (top row), the shim current norm (middle row), and the shim step norm (bottom row), and the columns show results from each individual subject (left column), and the distribution over all subjects (right column). Significant differences between the slice-wise shimming optimizations was evaluated with a Wilcoxon signed rank test. ns: not significant, * P -value < .05, ** P -value < .01

that the number of shim channels is generally low, it is computationally feasible to obtain each solution with direct calculation of the pseudoinverse of the regularized problem. Constrained least-squares optimization, on the other hand, requires iterative optimization, as there is generally no analytic solution to the problem with inequality constraints. While Tikh-LC was slowest of the methods compared, the computation times for all methods were on the same order of magnitude and

typically required <1 s for whole-brain shimming with the second-order shim system evaluated in this work. The computation time is thus well feasible for practical shimming applications.

In this work, the variable regularization was applied to second-order slice-wise shimming of the brain. The proposed method, however, generalizes to arbitrary ROIs and shim systems. Slice-wise ROIs are practical for multi-slice 2D imaging, but are not optimal for field

homogeneity due to their large FOV in two dimensions.³ Other types of ROIs may be used for example in imaging with multiple FOVs or in multi-voxel spectroscopy. Recent trends in shim hardware development has been towards utilizing a high number of non-spherical harmonic shim coils that are brought closer to the anatomy,²⁷ sometimes even integrated with the RF coil itself.^{28,29} The field profiles of such local shimming systems are likely to be highly non-orthogonal and will vary substantially between slices. Shimming with local coils is, thus, also likely to benefit from robust shim calculation that automatically adjusts to the conditions within each ROI.

6 | CONCLUSIONS

We have here presented a method for higher-order shim current optimization, based on Tikhonov regularization with a variable regularization parameter, λ . For each ROI (eg, slice), λ is adjusted to the local conditions of the shim optimization problem in a fully automated procedure. The benefit of the proposed method was demonstrated for second-order slice-wise shimming of the brain at 7T, yielding substantially reduced shim current norm without degrading shim quality. Increased robustness of the shim optimization may increase the utility and applicability of higher-order shim updating.

ACKNOWLEDGEMENTS


We are much grateful for technical support from Andrew Dewdney (Siemens Healthineers, Erlangen, Germany), who enabled the higher-order dynamic shimming setup on our 7T system.

DATA AVAILABILITY STATEMENT

The Tikh-LC shim optimization code will be made available via <https://git.fmrib.ox.ac.uk/shim>.

ORCID

Yuhang Shi  <https://orcid.org/0000-0001-8852-6100>

Signe Johanna Vannesjo  <https://orcid.org/0000-0003-2432-4192>

REFERENCES

- Chmurny GN, Hoult DI. The ancient and honourable art of shimming. *Concepts Magn Reson*. 1990;2:131-149.
- Schneider E, Glover G. Rapid in vivo proton shimming. *Magn Reson Med*. 1991;18:335-347.
- Poole M, Bowtell R. Volume parcellation for improved dynamic shimming. *Magn Reson Mater Phys*. 2008;21:31-40.
- Blamire A, Rothman D, Nixon T. Dynamic shim updating: a new approach towards optimized whole brain shimming. *Magn Reson Med*. 1996;36:159-165.
- Morrell G, Spielman D. Dynamic shimming for multi-slice magnetic resonance imaging. *Magn Reson Med*. 1997;38:477-483.
- Koch K, McIntyre S, Nixon T, Rothman D, De Graaf R. Dynamic shim updating on the human brain. *J Magn Reson*. 2006;180:286-296.
- Juchem C, Nixon TW, Diduch P, Rothman DL, Starewicz P, De Graaf RA. Dynamic shimming of the human brain at 7 T. *Concept Magn Reson B*. 2010;37B:116-128.
- Fillmer A, Vannesjo SJ, Pavan M, Scheidegger M, Pruessmann KP, Henning A. Fast iterative pre-emphasis calibration method enabling third-order dynamic shim updated fMRI. *Magn Reson Med*. 2016;75:1119-1131.
- Webb P, Macovski A. Rapid, fully automatic, arbitrary-volume in vivo shimming. *Magn Reson Med*. 1991;20:113-122.
- Hoult DI. "Shimming" on spatially localized signals. *J Magn Reson*. 1987;73:174-177.
- Wen H, Jaffer FA. An in vivo automated shimming method taking into account shim current constraints. *Magn Reson Med*. 1995;34:898-904.
- Schär M, Kozerke S, Fischer SE, Boesiger P. Cardiac SSFP imaging at 3 Tesla. *Magn Reson Med*. 2004;51:799-806.
- Hetherington HP, Chu W-J, Gonen O, Pan JW. Robust fully automated shimming of the human brain for high-field 1H spectroscopic imaging. *Magn Reson Med*. 2006;56:26-33.
- Pan JW, Lo K-M, Hetherington HP. Role of very high order and degree B0 shimming for spectroscopic imaging of the human brain at 7 tesla. *Magn Reson Med*. 2012;68:1007-1017.
- Fillmer A, Kirchner T, Cameron D, Henning A. Constrained image-based B0 shimming accounting for "local minimum traps" in the optimization and field inhomogeneities outside the region of interest. *Magn Reson Med*. 2015;73:1370-1380.
- Schwerter M, Hetherington H, Moon CH, et al. Interslice current change constrained B0 shim optimization for accurate high-order dynamic shim updating with strongly reduced eddy currents. *Magn Reson Med*. 2019;82:263-275.
- Klassen LM, Menon RS. Robust automated shimming technique using arbitrary mapping acquisition parameters (RASTAMAP). *Magn Reson Med*. 2004;51:881-887.
- Kim D-H, Adalsteinsson E, Glover GH, Spielman DM. Regularized higher-order in vivo shimming. *Magn Reson Med*. 2002;48:715-722.
- Nassirpour S, Chang P, Fillmer A, Henning A. A comparison of optimization algorithms for localized in vivo B0 shimming. *Magn Reson Med*. 2018;79:1145-1156.
- Hansen PC. REGULARIZATION TOOLS: a Matlab package for analysis and solution of discrete ill-posed problems. *Numer Algor*. 1994;6:1-35.
- Hansen PC, O'Leary DP. The use of the L-curve in the regularization of discrete ill-posed problems. *SIAM J Sci Comput*. 1993;14:1487-1503.
- Jehenson P, Westphal M, Schuff N. Analytical method for the compensation of eddy-current effects induced by pulsed magnetic field gradients in NMR systems. *J Magn Reson*. 1990;90:264-278.
- van Vaals JJ, Bergman AH. Optimization of eddy-current compensation. *J Magn Reson*. 1990;90:52-70.
- Bhogal A, Versluis M, Koonen J, et al. Image-based method to measure and characterize shim-induced eddy current fields. *Concepts Magn Reson*. 2013;42:245-260.

25. Jenkinson M. Fast, automated, N-dimensional phase-unwrapping algorithm. *Magn Reson Med*. 2003;49:193-197.
26. Smith SM. Fast robust automated brain extraction. *Hum Brain Mapp*. 2002;17:143-155.
27. Juchem C, Nixon TW, McIntyre S, Boer VO, Rothman DL, de Graaf RA. Dynamic multi-coil shimming of the human brain at 7 T. *J Magn Reson*. 2011;212:280-288.
28. Han H, Song AW, Truong T-K. Integrated parallel reception, excitation, and shimming (iPRES). *Magn Reson Med*. 2013;70:241-247.
29. Stockmann JP, Witzel T, Keil B, et al. A 32-channel combined RF and B0 shim array for 3T brain imaging. *Magn Reson Med*. 2016;75:441-451.

SUPPORTING INFORMATION

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

FIGURE S1 Simulated performance of CLS, TSVD, and Tikh-LC shim optimization on 2nd-, 3rd- and 4th order shim

systems. The plots show the increase in residual field standard deviation (A), the relative current norm (B), the regularization efficiency (C) and the relative condition number (D)

TABLE S1 Pre-emphasis parameters for self- and cross-term eddy current compensation (ECC) for second-order shim channels. Amplitudes are defined in percentage of the input shim step for the 2nd order self- and cross-terms, and in Hz for the cross-terms to B_0

TABLE S2 Measured shim coil sensitivities for the second-order shim channels

How to cite this article: Shi Y, Clare S, Vannesjo SJ. Shim optimization with region of interest-specific Tikhonov regularization: Application to second-order slice-wise shimming of the brain. *Magn Reson Med*. 2021;00:1–13. <https://doi.org/10.1002/mrm.28951>