Denoising High-field Multi-dimensional MRI with Local Complex PCA

Pierre-Louis Bazin[1,2], Anneke Alkemade[1], Wietske van der Zwaag[3], Matthan Caan[4], Martijn Mulder[1,5], Birte U. Forstmann[1]

[1] Integrative Model-based Cognitive Neuroscience research unit, Universiteit van Amsterdam, Amsterdam, The Netherlands; [2] Max-Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany; [3] Spinoza Centre for Neuroimaging, Amsterdam, The Netherlands; [4] Brain Imaging Centre, Amsterdam University Medical Center, Amsterdam, The Netherlands; [5] Psychology Department, Universiteit Utrecht, Utrecht, The Netherlands

Abstract

Modern high field and ultra high field magnetic resonance imaging (MRI) experiments routinely collect multi-dimensional data with high spatial resolution, whether multi-parametric structural, diffusion or functional MRI. While diffusion and functional imaging have benefited from recent advances in multi-dimensional signal analysis and denoising, structural MRI has remained untouched. In this work, we propose a denoising technique for multi-parametric quantitative MRI, combining a highly popular denoising method from diffusion imaging, over-complete local PCA, with a reconstruction of the complex-valued MR signal in order to define stable estimates of the noise in the decomposition. With this approach, we show signal to noise ratio (SNR) improvements in high resolution MRI without compromising the spatial accuracy or generating spurious perceptual boundaries.

1. Introduction

Ultra-high field magnetic resonance (MR) imaging at 7 Tesla and beyond has enabled neuroscientists to probe the human brain *in vivo* beyond the macroscopic scale (Weiskopf et al., 2015). In particular, quantitative MRI techniques (Cercignani et al., 2018) have become more readily available and offer the promise of quantitative information about the underlying microstructure. Unfortunately, as the size of the imaging voxel decreases well below the cubic millimeter, so does the signal to noise ratio (SNR), and the achievable resolution even when imaging a small portion of the brain remains limited, and requires multiple averages at the highest resolutions (Fedeeau et al., 2018; Fracasso et al., 2016; Stucht et al., 2015).

Multi-parametric quantitative methods acquire multiple images within a single sequence, in order to estimate the underlying quantity (Helms et al. 2008, Metere et al. 2017, Caan et al., 2018). This is comparable to diffusion-weighted MR imaging (DWI) where multiple images with different weighting are acquired. Recent work in diffusion analysis demonstrated that the intrinsic redundancy of the signal across these images can be employed to separate signal from noise with a principal component analysis (PCA) over small patches of the images (Manjon et al., 2013; Veraart et al., 2016). This principle can be transferred to multi-parametric quantitative MRI: in this work we present an extension of the PCA denoising to the recently described MP2RAGEME sequence, which provides estimates of R1 and R2* relaxation rates as well as quantitative susceptibility maps (QSM) in a very compact imaging sequence (Caan et al., 2018).

The MP2RAGEME data comes with additional challenges for the classical PCA approaches. First, the images are complex-valued, comprising of both magnitude and phase needed for the estimation of quantitative MRI parameters. This complex nature of the data needs to be taken into account. Second, they include only a few images compared to the high number of directions acquired in DWI. Interestingly, we can make fairly simple assumptions about the dimensionality of

the signal, as it contains primarily R1, R2*, susceptibility and proton density (PD) weighting. Taking these features into account, we were able to effectively denoise high resolution MP2RAGEME data without averaging, and maintain spatial precision. We studied the impact of denoising on the computation of quantitative maps, and its practical impact for delineating small, low contrast subcortical nuclei such as the habenula, a notoriously difficult small nucleus in the subcortex.

2. Materials and Methods

2.1. Data acquisition

Our denoising method focuses on the recently developed MP2RAGEME sequence (Caan et al., 2018) which combines multiple inversions and multiple echoes from a magnetisation-prepared gradient echo sequence (MPRAGE) to simultaneously obtain an estimate of quantitative T1, T2* and susceptibility. The specific protocol of interest is comprised of five different images: a first inversion with T1 weighting, followed by a second inversion with predominantly PD weighting and four echoes with increasing T2* weighting (Fig. 1).

Figure 1: The MP2RAGEME sequence: A: first inversion, B,D,C,E: second inversion, first to fourth echo (top: magnitude, bottom: phase).

For the experiments, we used a high-resolution imaging slab centered on the subcortex with the following parameters: resolution = 0.5mm isotropic, TR = 8.33s, TR_1 = 8ms, TR_2 = 32ms, TE_1 = 4.6ms, $TE_{2A\text{-}D}$ = 4.6/12.6/20.6/28.6ms, TI_1/TI_2 = 670/3738ms, $\alpha_1/\alpha_2 = 7/8^o$. Data was obtained for five healthy human subjects as part of an ongoing atlasing study approved by the Ethics Review Board of the Faculty of Social and Behavioral Sciences, University of Amsterdam, The Netherlands (approval number: 2016-DP-6897). All subjects provided written informed consent for the study.

2.2. Complex signal reconstruction

One of the main drawbacks of existing local PCA methods for denoising is that they work on magnitude images, which follow Rician distributions. A simple correction proposed in (Eichner et al., 2016) for removing the Rician bias in DWI averaging is to revert to the complex signal, taking into account the phase information. Phase contains additional variations due to non-local effects of air cavities around the brain, which bring severe ringing artifacts in the reconstructed data (Fig.2A). In (Eichner et al., 2016), the global phase information is removed with a total variation method, which generally respects the location of phase wraps. However, we found that this approach is not effective in regions where the phase wraps have high frequency, and there are residual phase artifacts in the local phase. While these have little consequences for their averaging application, they provide systematic variations for the PCA decomposition, which is undesirable here. We therefore start with a full phase unwrapping (Abdul-Rahman et al., 2005) followed by a total variation smoothing (Chambolle, 2004) of the unwrapped phase (Fig.2B). The residual phase variations are used as local phase, and combined with the magnitude to reconstruct real and imaginary parts of the complex signal (Fig.2C,D). Note that here, unlike (Eichner et al., 2016), we do not discard the imaginary part of the signal as it contains valuable information about the noise and residual anatomical information.

Figure 2: Phase preprocessing. A: original phase map, B: after unwrapping and global phase removal, C: reconstructed real signal, D: reconstructed imaginary signal.

2.3. Local PCA of complex signal

The complex signal, now comprising ten images all similar in nature, is then processed following the local overcomplete PCA approach of (Manjón et al., 2013). In short, the images are cut into small, overlapping patches of *NxNxN* voxels, and the *M* contrasts combined into a N^3xM matrix. The average patch value per contrast is subtracted, and the matrix decomposed via singular value decomposition (SVD) to yield the eigenvectors and associated singular values (square roots of the eigenvalues) of the covariance matrix across the patch. In this work, we used patches of size $N = 4$, and $M = 10$.

The overlapping patches are combined following the technique of (Manjón et al., 2013), weighting each patch by the number of kept eigenvectors $w_{patch} = \frac{1}{1+M_{kept}}$. Once recombined, the eigenvectors rapidly change from a rich information content, encoding boundaries, to pure noise (Fig. 3). However, the decision boundary between actual signal and noise is variable across the brain, due to the presence of different tissue types, multiple tissue boundaries, etc. In order to infer for each patch the number of eigenvectors to keep, we thus need to first quantify the expected noise distribution over the SVD.

Figure 3: Local PCA decomposition: first five eigenvectors (top) and singular values (bottom) from highest to lowest.

2.4. Estimating the noise

Noise estimation in advanced MRI is far challenging: variations in coil sensitivity, non-local susceptibility effects and dependencies to the B0 field as well as various acquisition techniques will affect the signal and the noise differently in different regions. The original local PCA denoising methods of (Manjón et al., 2013) used elaborate estimates of the magnitude image noise, taking into account its Rician properties. A recent extension of the work fit the expected distribution of eigenvalues from a random Rician noise matrix to determine its threshold (Veraart et al., 2016).

In this work, we propose a simpler approach based on two properties of the complex signal: first, that the noise is locally Gaussian, and second that the spread of singular values for Gaussian noise can be reasonably well approximated by a straight line (Fig.4A,B). In addition, this simple property is retained by interpolation (Fig.4C,D), which makes it possible to perform denoising after image registration, for instance when motion correction is needed as in DWI or fMRI, or after fat navigator-based motion correction in structural MRI (Gallichan and Marques, 2017).

Figure 4: Noise properties of the singular value decomposition (SVD). A: SVD of 15 simulated 4x4x4 voxel Gaussian noise patterns, B: Distribution of R^2 residuals over from fitting a line to the second half of 100,000 simulated noise patterns, C: SVD of 15 simulated 4x4x4 voxel Gaussian noise patterns after a random sub-voxel shift and trilinear interpolation, D: Distribution of R^2 residuals for 100,000 simulated noise patterns with the same random shift, E: Map of estimated number of kept eigenvectors from actual data, F: Map of the fitting R^2 residuals, G: Map of kept eigenvectors after co-registration and trilinear interpolation of the data, H: Map of the fitting R^2 residuals after co-registration.

Our algorithm to estimate the noise level proceeds as follows. First, a line is fitted to the $M/2 = 5$ lowest singular values of the patch decomposition using linear least squares. Then, every singular value above a factor of α above the expected noise level given by the fitted line is kept, while the others are removed. Thus, the main requirements of our method are that: 1) local signal variations across contrasts in each individual patch are Gaussian-distributed, and 2) the intrinsic dimension of the data is lower than half of the number of acquired images. The patches are then reconstructed and averaged across the image, the complex images separated into magnitude and phase, and finally the discarded global phase variations are reintroduced and wrapped, to obtain data as similar as possible to the original input. In addition, a map of the number of kept eigenvectors and of the noise fitting residuals weighted by patch are computed for quality control (Fig.4E,F,G,H). The complete denoising algorithm is available in Open Source as part of the IMCN Toolkit (https://www.github.com/imcn-uva/imcn-imaging/) and the Nighres library (Huntenburg et al., 2018; https://www.github.com/nighres/nighres/).

2.5. Quantitative Mapping

The main interest of quantitative MR mapping techniques such as the MP2RAGEME sequence is to obtain estimates of the MR parameters of T1, T2* relaxation times and susceptibility χ from the measured images. In order to do so, we used the look-up table method of (Marques et al., 2010) to generate T1 estimates, a simple regression in log domain to obtain T2* (Miller and Joseph, 1993) and the TGV-QSM reconstruction algorithm of (Langkammer et al., 2015) to create quantitative susceptibility maps (QSM), following the approach of (Caan et al., 2018) to estimate QSM for each of the three last echoes of the second inversion and take the median. T1 and T2* mapping methods are included in the Nighres library, and the TGV-QSM algorithm is also freely available (http://www.neuroimaging.at/pages/qsm.php).

2.6. Manual delineations

In order to obtain structure-specific measures of quality and to evaluate the practical benefits of the method for segmentation purposes, we performed a manual segmentation study. To evaluate the benefits of the denoising, we selected the habenula, a small and challenging structure ventral to the thalamus, which together with the pineal gland and stria medullaris form the epithalamus (Mai et al, 2016; Hikosaka, 2010; Strotmann et al., 2013). Since we cannot distinguish between the medial and lateral habenula on MRI, both were included in our delineations. Given its size, complex shape, and heterogeneous composition, the habenula requires high resolution images and detailed anatomical knowledge to be reliably delineated on MRI.

In this study, the habenula was delineated by two raters on the reconstructed quantitative T1 maps obtained with or without denoising. The consensus mask (regions labeled by all raters on all images) were used to measure noise properties, while the overlap of the masks was used to measure inter-rater reliability.

2.7. Co-registration to standard space

Previously proposed methods for local PCA denoising have shown decreased performance when handling interpolated data due to the changes in the noise distribution (Manjón et al., 2013; Veraart et al., 2016). To test the performance of our noise estimation approach, we co-registered our data to standard space by aligning it first with a whole brain image of the same subject and then to the MNI152 template at 0.5 mm resolution. Both registrations used the first inversion of the MP2RAGEME sequence, and were performed with linear registration in ANTS (Avants et al.,

2008). Phase images were unwrapped before transformation to avoid interpolating across phase wraps, and quantitative maps were computed before and after denoising in standard space.

2.8. Region labeling

Finally, to evaluate the signal improvement across a variety of structures, we used manual delineations of the striatum, globus pallidus internal and external segment, subthalamic nucleus, substantia nigra and red nucleus obtained from lower resolution MP2RAGEME images (0.64x0.64x0.7 mm resolution) with higher SNR. The delineations were performed independently by two raters and the conjunction between them was used as the structure mask. The masks were co-registered linearly to the high resolution slab based on their scanner coordinates followed by a rigid registration in ANTS. The intensities of the quantitative maps inside each region were used to compute signal-to-noise ratio statistics of the original images and reconstructed quantitative maps before and after denoising in original space.

3. Results

3.1. Impact on quantitative maps

As shown in Fig.5, the LCPCA denoising strongly impacts the estimation of quantitative signals, as the SNR gains of multiple images are combined. Visually, we found a strong gain in particular for R2* mapping, which is quite noise-sensitive with the low number of echo times of the MP2RAGEME sequence. While the improvements are subtle at the whole brain scale, they are very clear when focusing on small regions, particularly in the subcortex where the original SNR is low.

Figure 5: Comparison of quantitative MRI reconstruction: left, from the original MP2RAGEME images; right from the denoised result.

3.2. SNR Measures

The denoising systematically improved the SNR in all structures, although not identically across regions and contrasts (Fig.6). On individual MP2RAGEME images, the improvement was modest, while denoising had a stronger cumulative impact on quantitative map estimation, especially for R1. While R2* appears visually improved, the SNR measures were more similar partly due to the heterogeneity of R2* values in many anatomical structures. QSM is the least affected, both visually and quantitatively, which is expected as the quantitative susceptibility reconstruction algorithm includes its own spatial regularization method (in this case, total generalized variation (TGV)). Still, some more heterogeneous regions of the subcortex appear smoother after denoising. Interestingly, the highest gains were obtained for the smaller structures such as substantia nigra, sub-thalamic nucleus and red nucleus.

Figure 6: SNR improvement for the delineated structures and the different contrasts (from left to right: first inversion, second inversion echo 1 to 4, estimated R1 map, estimated R2 map, estimated QSM; mean and standard deviation across subjects).*

3.3. Impact on manual delineations

Manual delineations of the habenula are challenging, and inter-rater agreement is low. Our denoising improved the consensus (Fig. 7), but not to the point of an acceptable level of reproducibility for measuring anatomical quantities such as volume or shape.

Figure 7: Comparison of manual delineations of the lateral habenula with or without denoising. A: Inter-rater agreement between delineations, B: 3D rendering of the lateral habenula in one subject, as delineated by each rater (red and green outline, respectively), C: Inter-rater difference between delineations and average internal contrast variations.

3.4. Effects of interpolation

When performing the denoising step after image co-registration to a standard space and interpolation, we observe only small differences in the LCPCA algorithm behavior (Fig.4G,H). The interpolation procedure tends to slightly increase the dimension of the kept signal, probably due to the mixing of signals across voxels. However, the proposed dimensionality estimation procedure appears largely robust to interpolation.

4. Discussion

In this work we presented a new local PCA denoising technique, using the full complex signal acquired to correct acquisitions with multiple image contrasts such as DWI or quantitative MRI sequences. Even with a sequence like MP2RAGEME, where only five different images are used in order to recover T1, T2* and QSM, the proposed approach was able to effectively reduce the noise without introducing blurring artifacts. The homogeneity of structures compared to their boundary was improved over the basal ganglia structures, and delineation of small nuclei such as the habenula was shown to be more reproducible in high resolution images with high levels of imaging noise.

By using both magnitude and phase information, the proposed method captures the entire noise distribution, and stays in the statistically simpler domain of Gaussian perturbations. However, obtaining high quality phase images is not trivial for all scanners and imaging sequences and can be a limitation also in retrospective processing, as the phase is commonly discarded. Note also that certain acceleration techniques such as partial Fourier encoding may sacrifice the phase signal estimation. While there is no conceptual requirement to use both phase and magnitude in the denoising, lowering the number of dimensions reduces the applicability of the method: in the case of the MP2RAGEME using only five dimensions in order to separate four-dimensional signals from noise is challenging, and the linear fit of the noise to the last half of the local PCA eigenvalues becomes unreliable. Yet, preliminary experiments with other relaxometry methods such as multi-parameter mapping (Weiskopf et al., 2013), which acquire between 14 and 20 images, indicated that the separation of signal and noise based on magnitude alone was possible (unpublished data).

The main requirements of the noise estimation method are that: 1) local signal variations across contrasts in each individual patch are Gaussian-distributed, and 2) the intrinsic dimension of the data is generally lower than half of the number of acquired images. The first requirement is easily met in small patches, regardless of the type of data acquired. The second requirement depends on the type of MR sequence and tissue properties under consideration, but can be met simply by running twice the same sequence, as is commonly done for increasing SNR by classical averaging. Note also that in the case of a single contrast the PCA approach reduces to simple averaging with no added value.

The main interest of the PCA-based approaches for denoising high resolution images is that unlike most other methods, they do not impose spatial regularity but rather enforce regularity across contrasts, preserving small anatomical details without blurring or inducing artificial boundaries. In addition, while the impact of noise removal on conventional MR images remains subtle, it offers the option to push MR sequences beyond the usually accepted limits of thermal noise as long as enough signal remains to be reliably detected. As ultra-high field advances toward higher and higher resolutions, such denoising methods may become essential part of the imaging protocol for multi-contrast anatomical imaging in the same way they have become a standard tool for advanced DWI pre-processing (Veraart et al., 2016).

Finally, the proposed method is generally applicable to other relaxometry sequences such as multi-echo GRE, multi-echo MPRAGE or multi-parametric maps (MPM), as well as DWI or (multi-echo) fMRI. Other types of MR sequences that include multiple images with different contrast may benefit from this approach, as long as the total number of images is sufficient to properly estimate noise properties.

Conclusions

Here we presented a new local PCA method to denoise high resolution multi-parametric quantitative MRI data. Combining magnitude and phase data, we could differentiate between signal- and noise-induced variations with a simple model that is robust to interpolation. The resulting quantitative images are automatically regularized and additional anatomical detail is visible in low contrast regions. The denoising software is openly available as part of the IMCN Toolkit (https://www.github.com/imcn-uva/imcn-imaging/) and the Nighres library (https://www.github.com/nighres/nighres/). The proposed method can be extended to denoise other MR imaging sequences with similar properties, namely partially redundant contrasts and low intrinsic dimensionality. While the method is already efficient to denoise MR images acquired with cutting edge methods at the lower limits of SNR, we hope they may help further to push MR imaging toward acquiring even more challenging data where the noise may visually dominate but significant amounts of signal are still available.

References

Abdul-Rahman, H., Gdeisat, M., Burton, D., Lalor, M., 2005. Fast three-dimensional phase-unwrapping algorithm based on sorting by reliability following a non-continuous path, in: Osten, W., Gorecki, C., Novak, E.L. (Eds.), . Presented at the Optical Metrology, Munich, Germany, p. 32. https://doi.org/10.1117/12.611415

Caan, M.W.A., Bazin, P., Marques, J.P., Hollander, G., Dumoulin, S., Zwaag, W., 2018. MP2RAGEME: T 1 , T 2 * , and QSM mapping in one sequence at 7 tesla. Human Brain Mapping. https://doi.org/10.1002/hbm.24490

Cercignani, M., Dowell, N.G., Tofts, P.S. (eds), 2018. Quantitative MRI of the Brain: Principles of Physical Measurement, Second edition, 2018. . CRC Press. https://doi.org/10.1201/b21837

Chambolle, A., 2004. An algorithm for total variation minimization and applications. Journal of Mathematical imaging and vision 20, 89–97.

Eichner, C., Cauley, S.F., Cohen-Adad, J., Möller, H.E., Turner, R., Setsompop, K., Wald, L.L., 2015. Real diffusion-weighted MRI enabling true signal averaging and increased diffusion contrast. NeuroImage 122, 373–384. https://doi.org/10.1016/j.neuroimage.2015.07.074

Federau, C., Gallichan, D., 2016. Motion-Correction Enabled Ultra-High Resolution In-Vivo 7T-MRI of the Brain. PLOS ONE 11, e0154974. https://doi.org/10.1371/journal.pone.0154974

Fracasso, A., van Veluw, S.J., Visser, F., Luijten, P.R., Spliet, W., Zwanenburg, J.J.M., Dumoulin, S.O., Petridou, N., 2016. Lines of Baillarger in vivo and ex vivo: Myelin contrast across lamina at 7 T MRI and histology. NeuroImage 133, 163–175. https://doi.org/10.1016/j.neuroimage.2016.02.072

Gallichan, D., Marques, J.P., 2017. Optimizing the acceleration and resolution of three-dimensional fat image navigators for high-resolution motion correction at 7T. Magn Reson Med 77:547–558.

Helms, G., Dathe, H., Dechent, P., 2008. Quantitative FLASH MRI at 3T using a rational approximation of the Ernst equation. Magnetic Resonance in Medicine 59, 667–72.

Hikosaka, O., 2010. The habenula: from stress evasion to value-based decision-making. Nature Reviews Neuroscience 11, 503–513. https://doi.org/10.1038/nrn2866

Huntenburg, J.M., Steele, C.J., Bazin, P.-L., 2018. Nighres: processing tools for high-resolution neuroimaging. GigaScience 7. https://doi.org/10.1093/gigascience/giy082

Langkammer, C., Bredies, K., Poser, B.A., Barth, M., Reishofer, G., Fan, A.P., Bilgic, B., Fazekas, F., Mainero, C., Ropele, S., 2015. Fast quantitative susceptibility mapping using 3D EPI and total generalized variation. NeuroImage 111, 622–630. https://doi.org/10.1016/j.neuroimage.2015.02.041

Mai, J.K., Majtanik, M., Paxinos, G., 2016. Atlas of the human brain, 4. ed. ed. Elsevier AP, Amsterdam.

Manjón, J.V., Coupé, P., Concha, L., Buades, A., Collins, D.L., Robles, M., 2013. Diffusion Weighted Image Denoising Using Overcomplete Local PCA. PLoS ONE 8, e73021. https://doi.org/10.1371/journal.pone.0073021

Metere, R., Kober, T., Möller, H.E., Schäfer, A., 2017. Simultaneous Quantitative MRI Mapping of T1, T2* and Magnetic Susceptibility with Multi-Echo MP2RAGE. PLOS ONE 12, e0169265. https://doi.org/10.1371/journal.pone.0169265

Miller, A.J., Joseph, P.M., 1993. The use of power images to perform quantitative analysis on low SNR MR images. Magnetic resonance imaging 11, 1051–1056.

Strotmann, B., Kögler, C., Bazin, P.-L., Weiss, M., Villringer, A., Turner, R., 2013. Mapping of the internal structure of human habenula with ex vivo MRI at 7T. Frontiers in Human Neuroscience 7. https://doi.org/10.3389/fnhum.2013.00878

Stucht, D., Danishad, K.A., Schulze, P., Godenschweger, F., Zaitsev, M., Speck, O., 2015. Highest Resolution In Vivo Human Brain MRI Using Prospective Motion Correction. PLOS ONE 10, e0133921. https://doi.org/10.1371/journal.pone.0133921

Veraart, J., Novikov, D.S., Christiaens, D., Ades-aron, B., Sijbers, J., Fieremans, E., 2016. Denoising of diffusion MRI using random matrix theory. NeuroImage 142, 394–406. https://doi.org/10.1016/j.neuroimage.2016.08.016

Weiskopf N, Suckling J, Williams G, Correia MM, Inkster B, Tait R, Ooi C, Bullmore ET, Lutti A (2013) Quantitative multi-parameter mapping of R1, PD*, MT, and R2* at 3T: a multi-center validation. Frontiers in Neuroscience 7:95

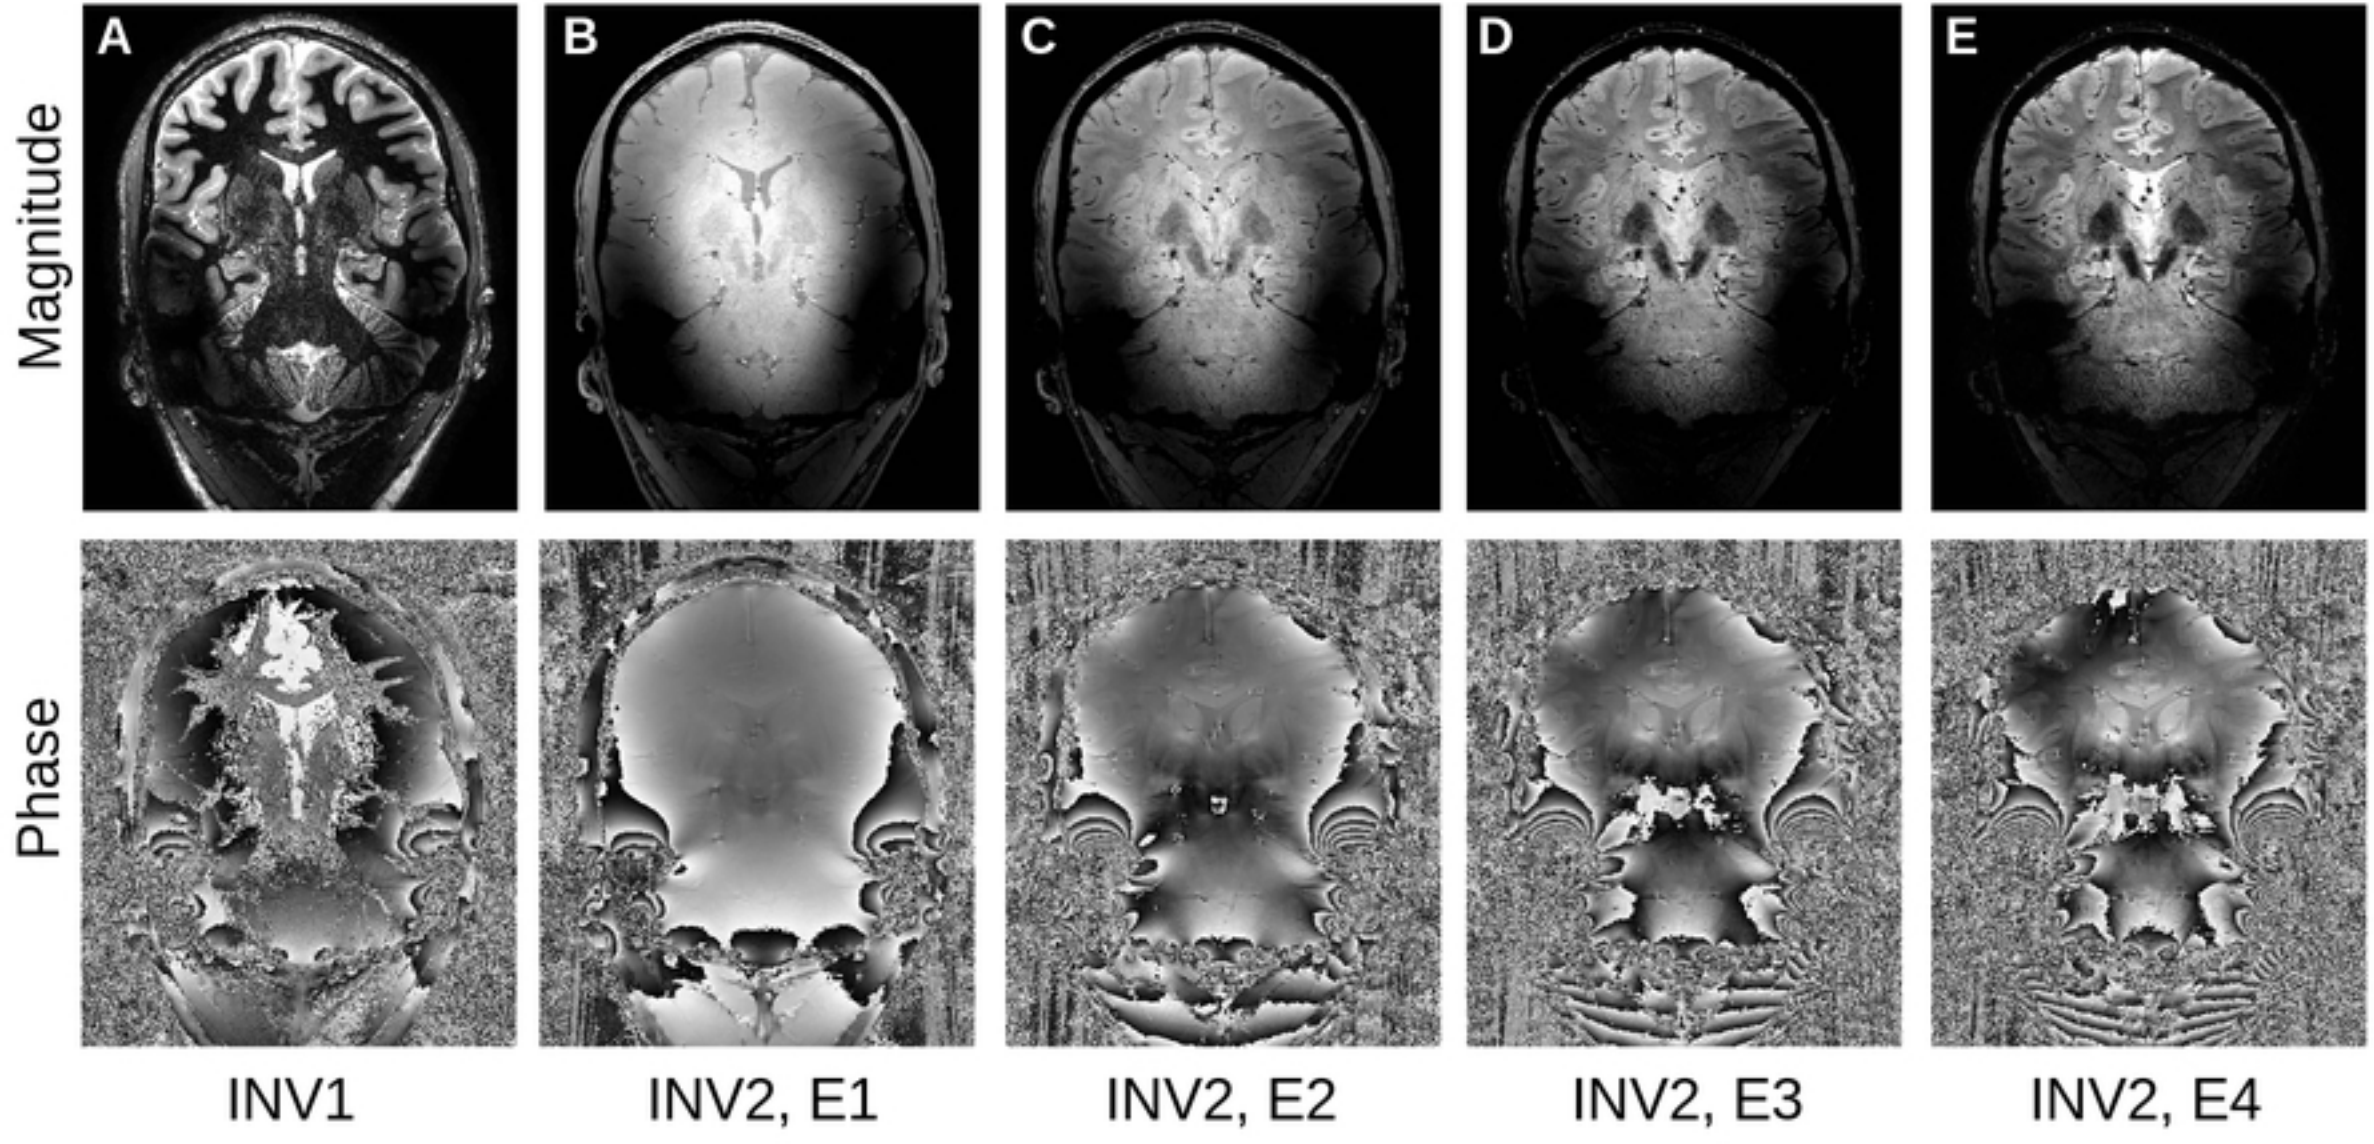

Figure 1

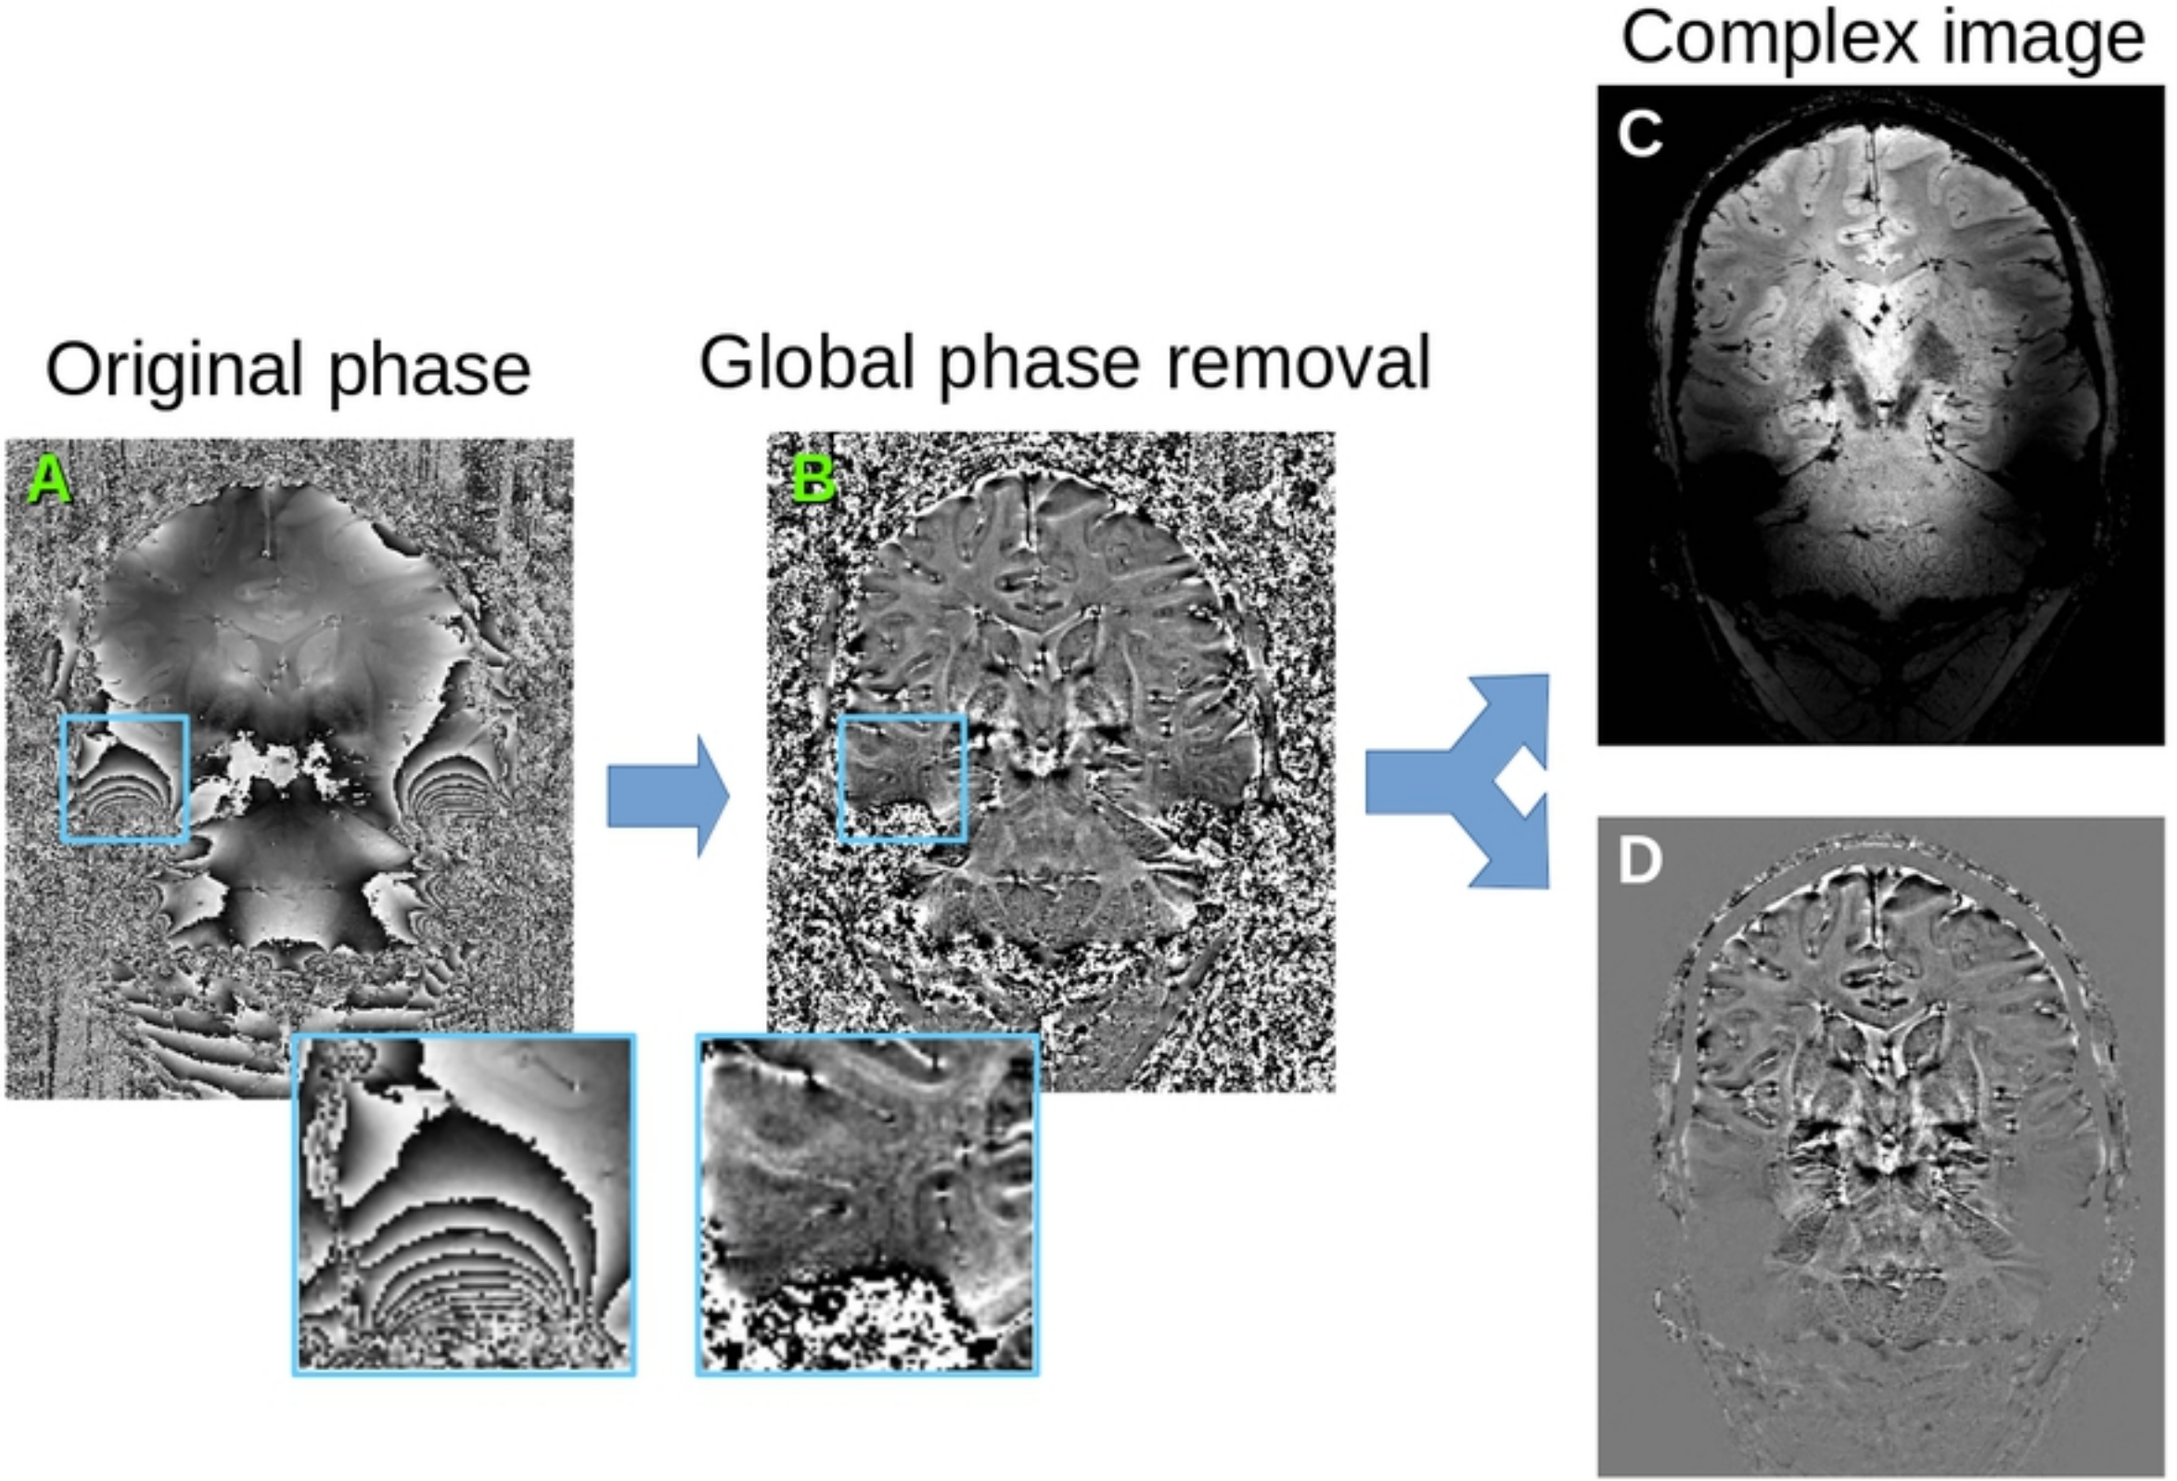

Figure 2

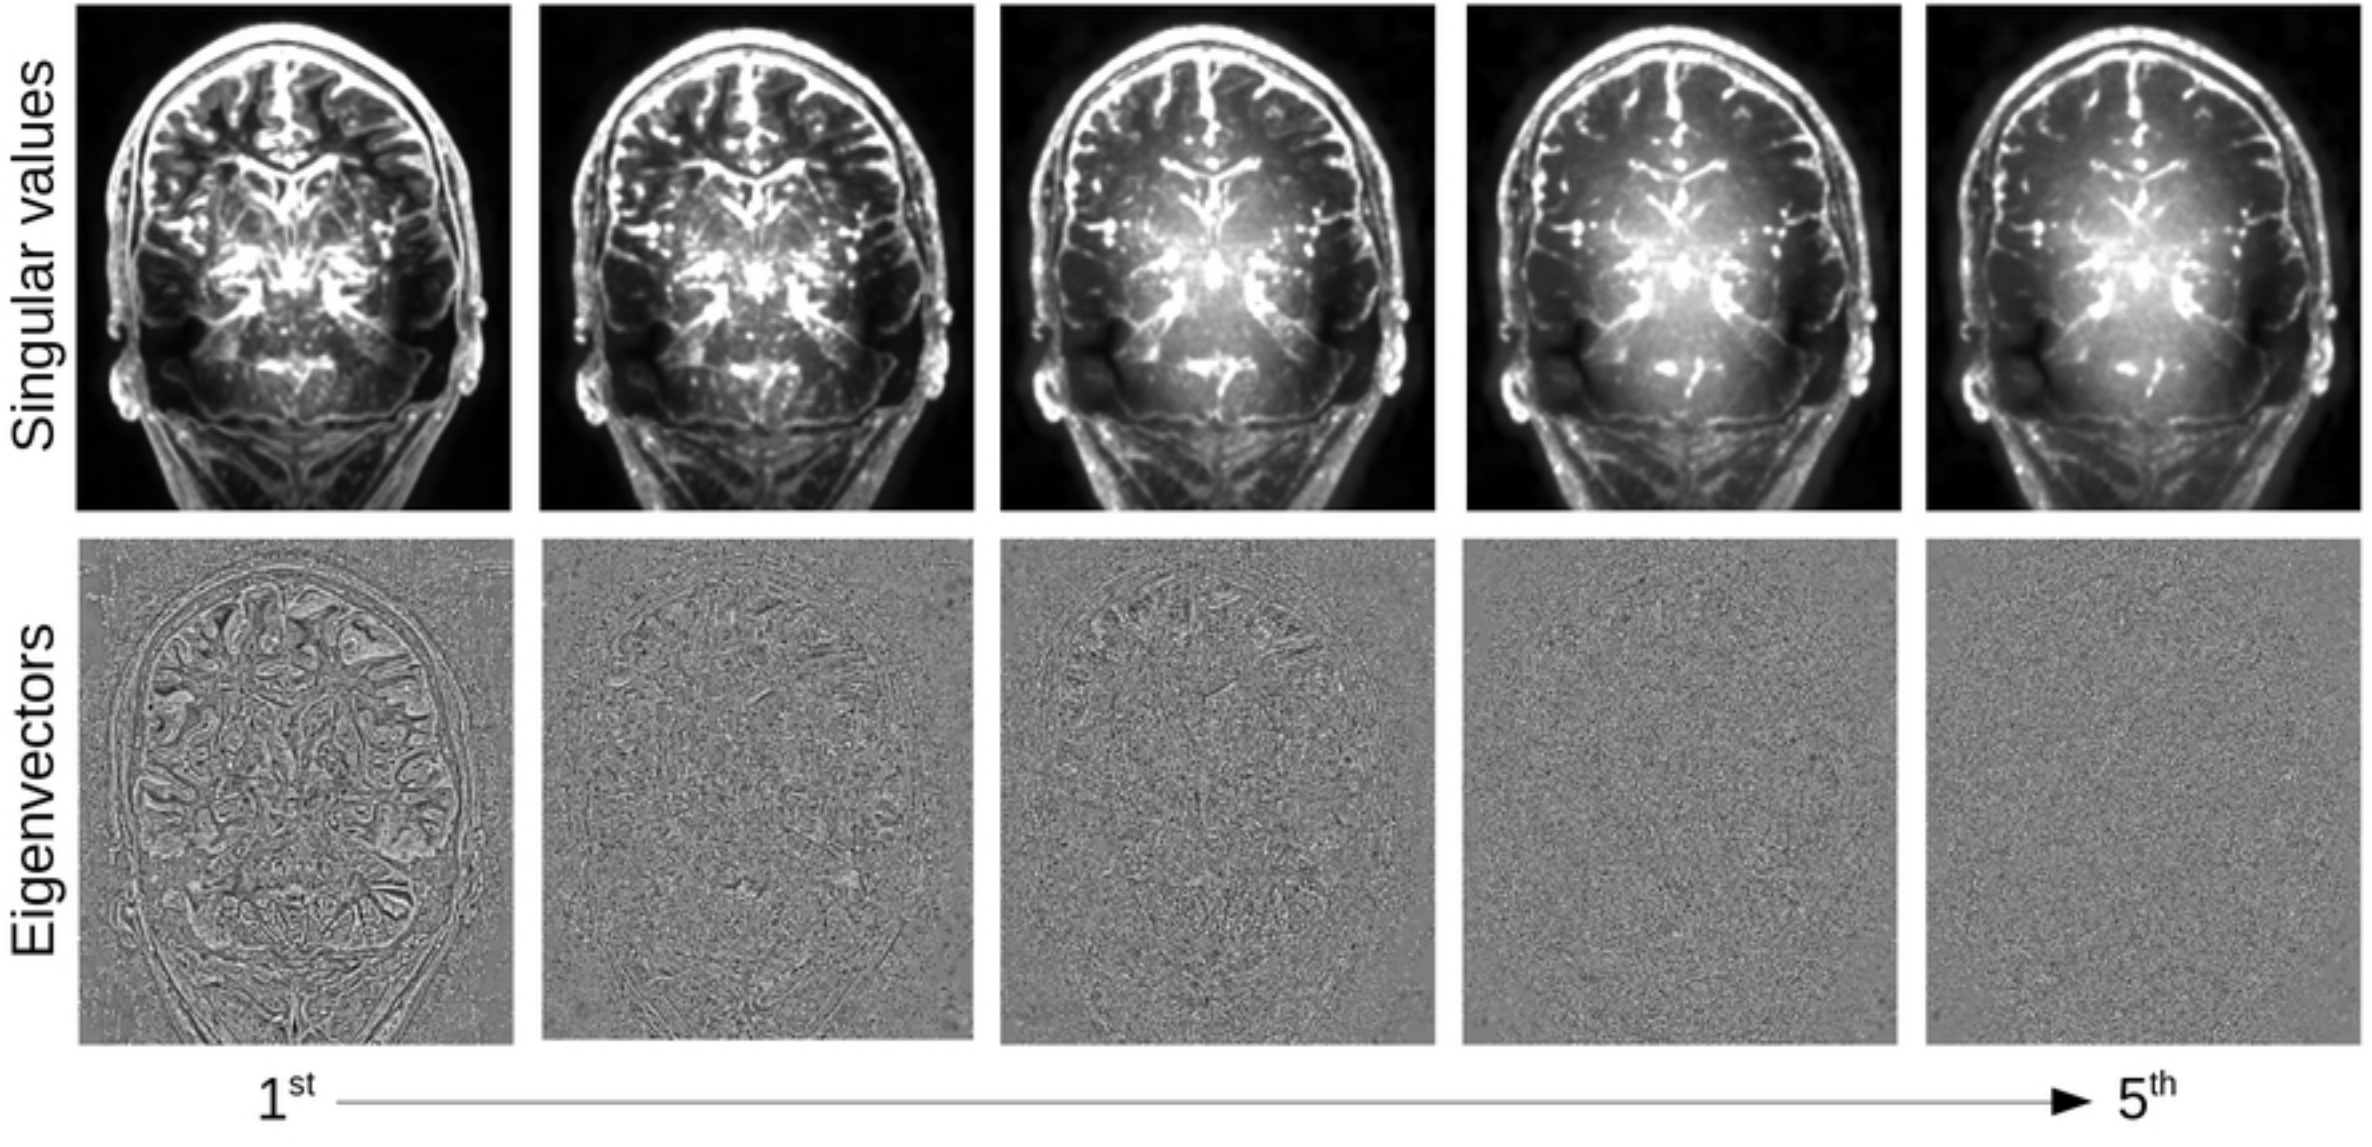

Figure 3

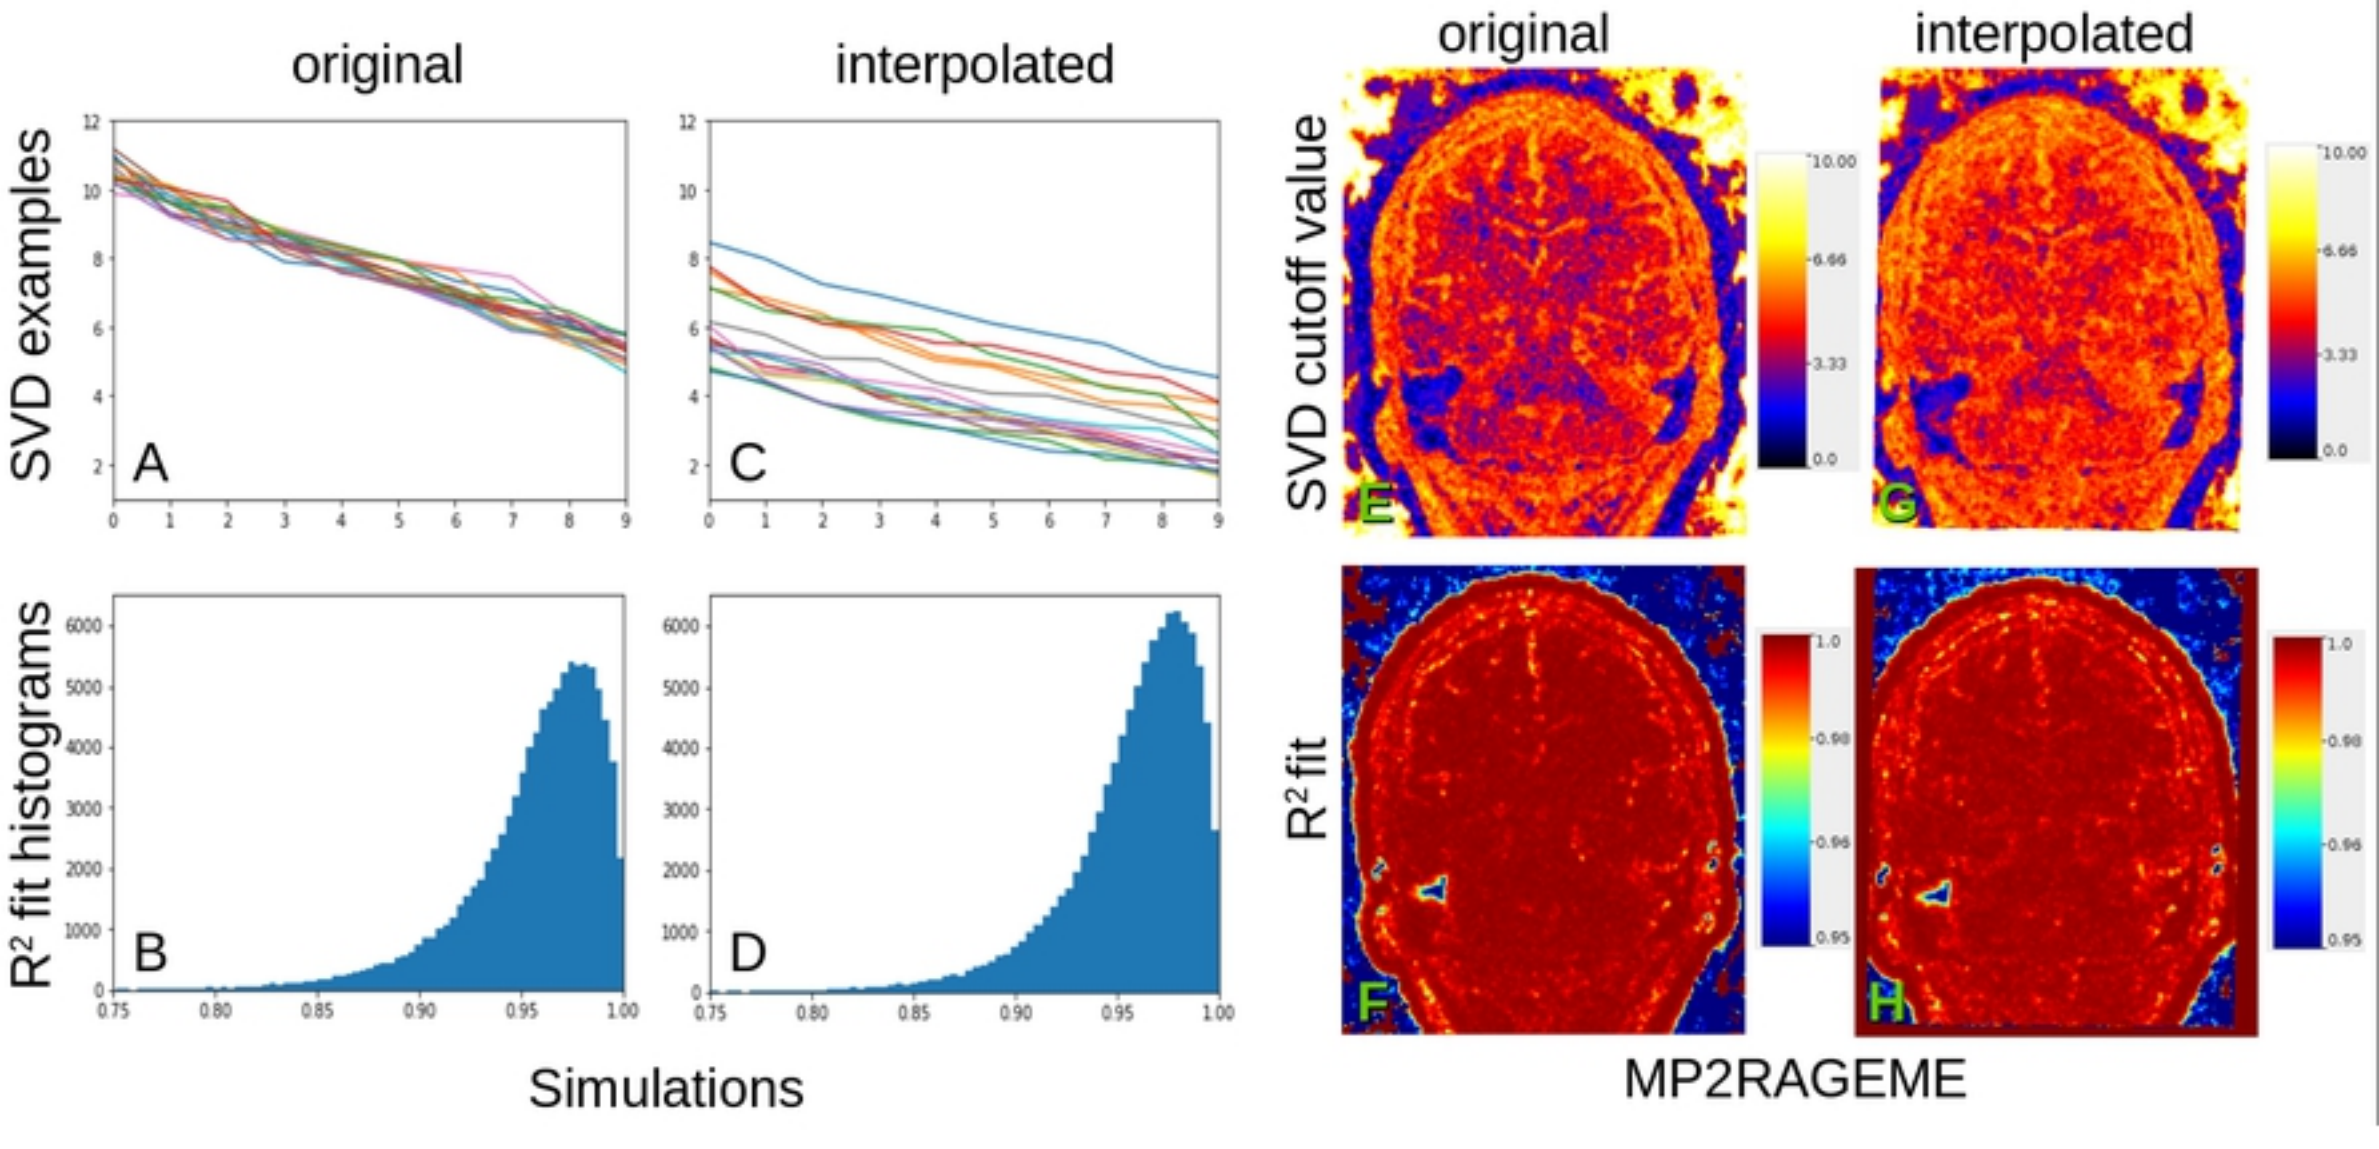
original
interpolated
SVD examples
A
C
R² fit histograms
B
D
Simulations
original
interpolated
SVD cutoff value
E
G
R² fit
F
H
MP2RAGEME

Figure 4

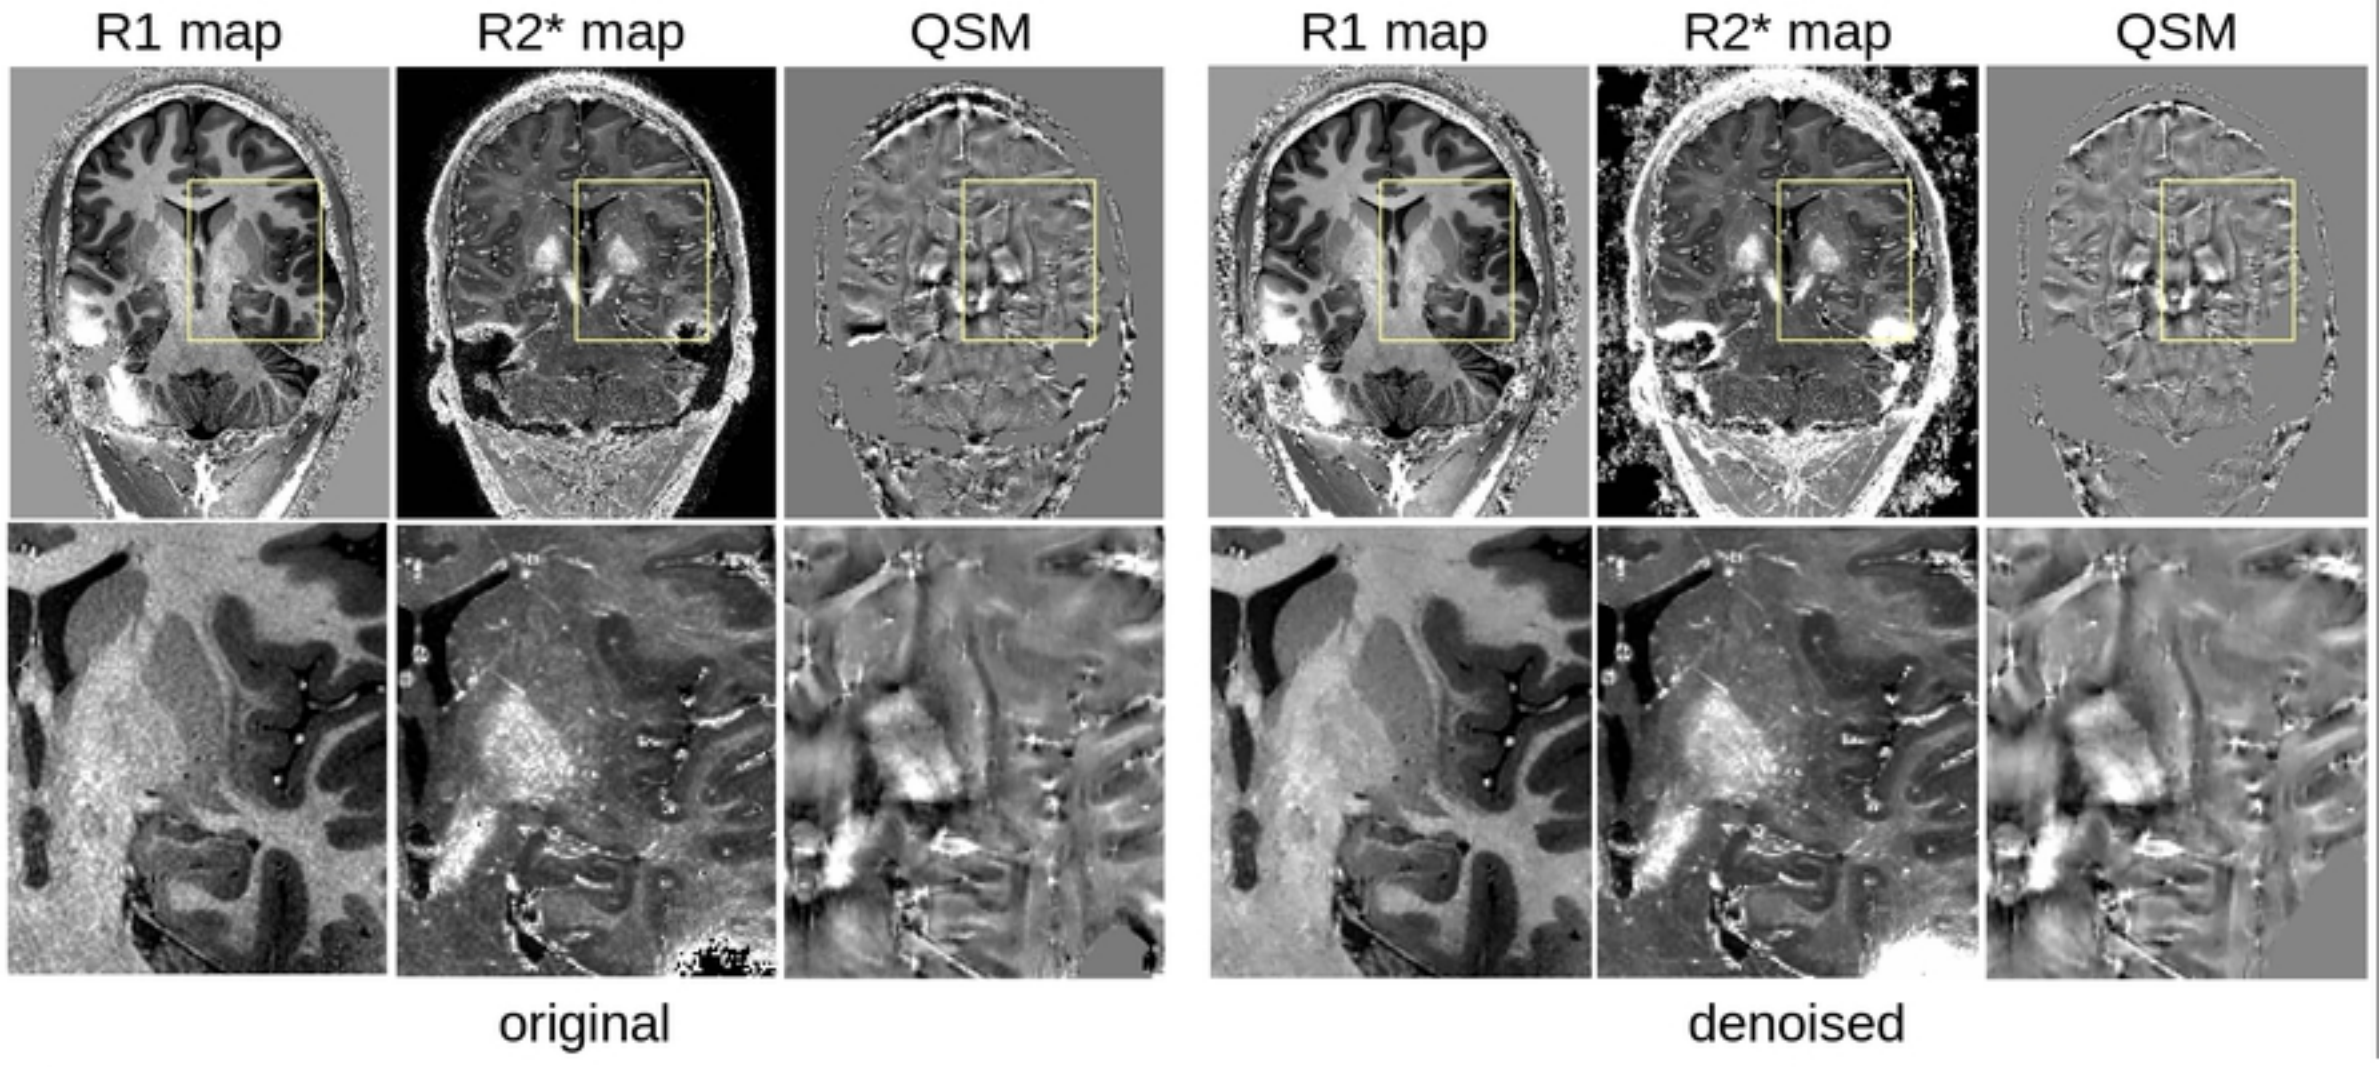

Figure 5

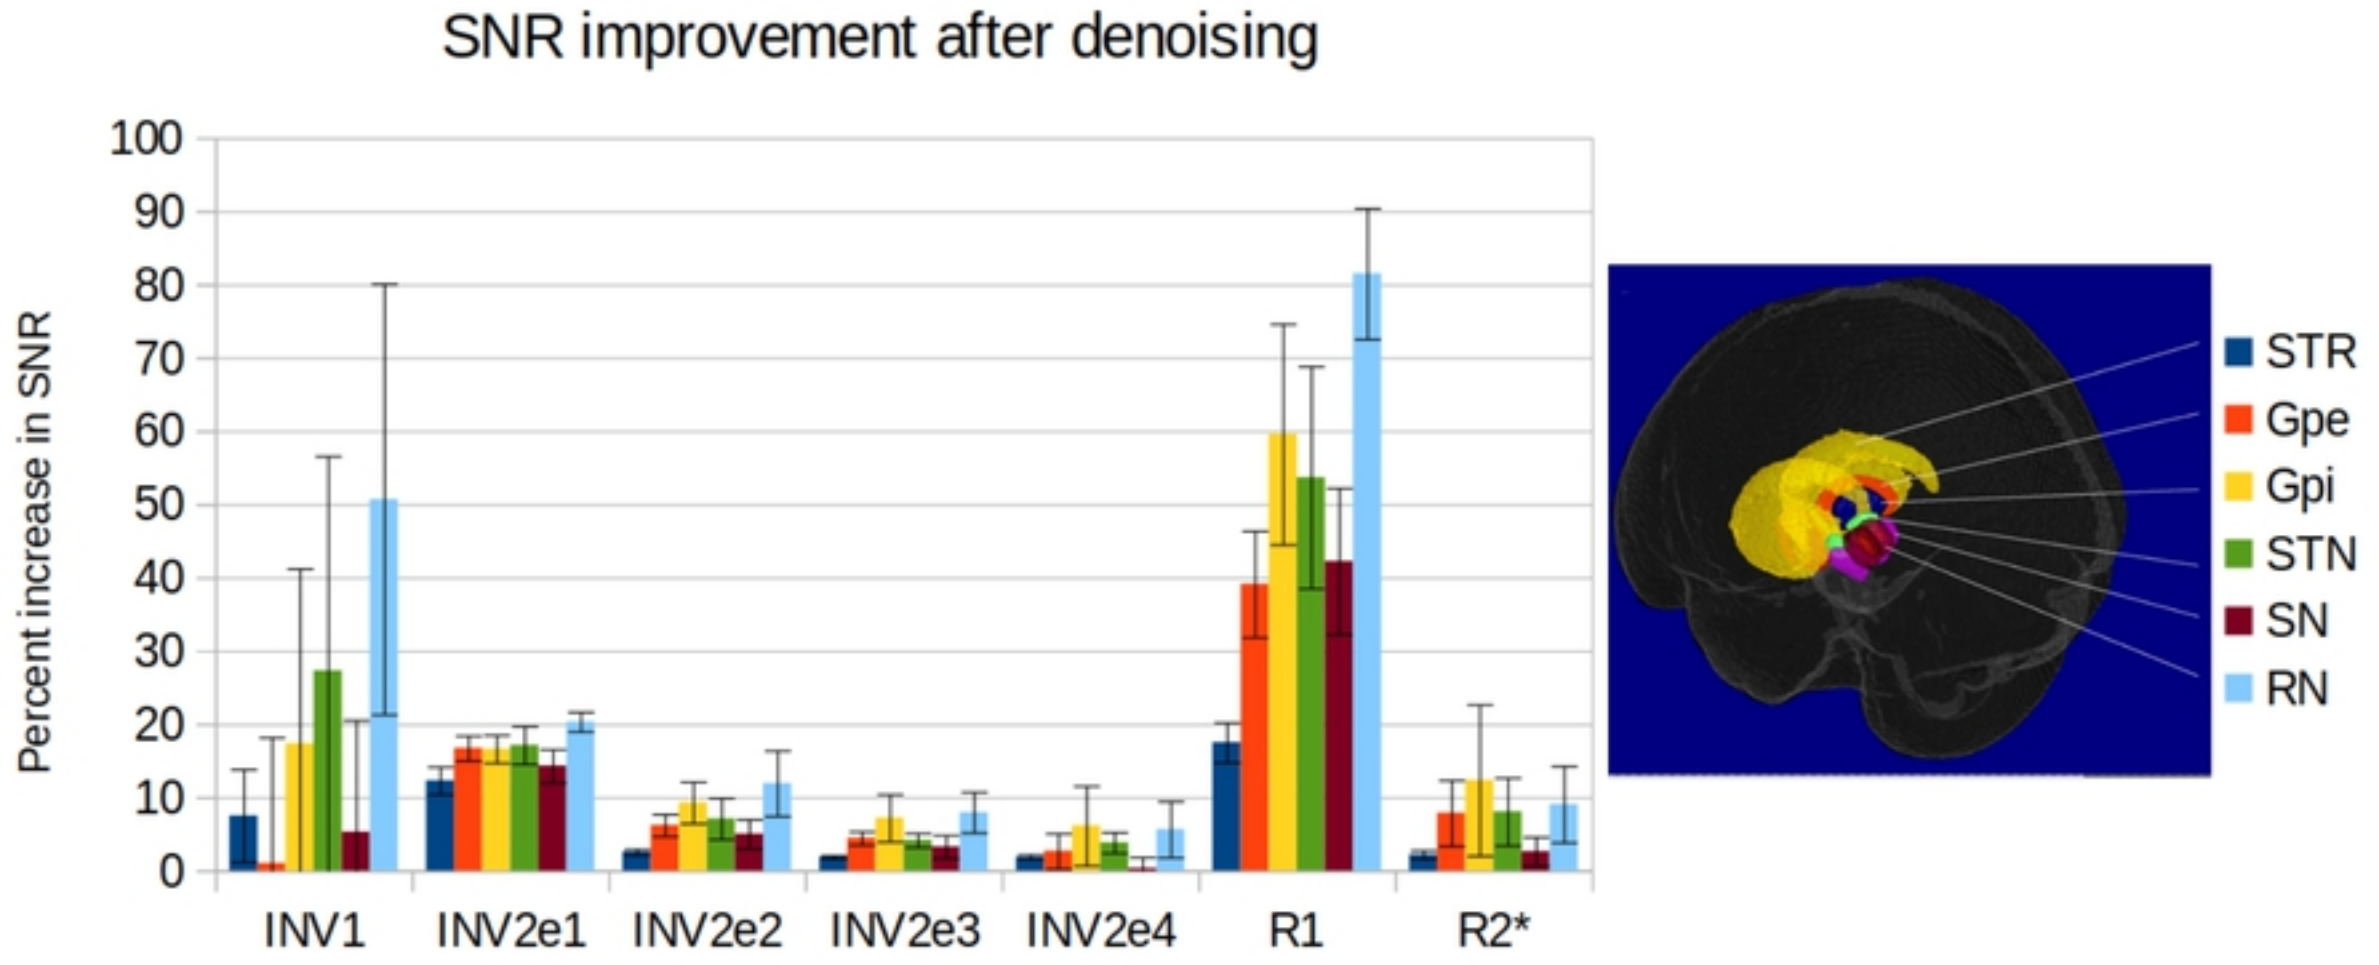
SNR improvement after denoising
Percent increase in SNR
100
90
80
70
60
50
40
30
20
10
0
INV1
INV2e1
INV2e2
INV2e3
INV2e4
R1
R2*
STR
Gpe
Gpi
STN
SN
RN

Figure 6

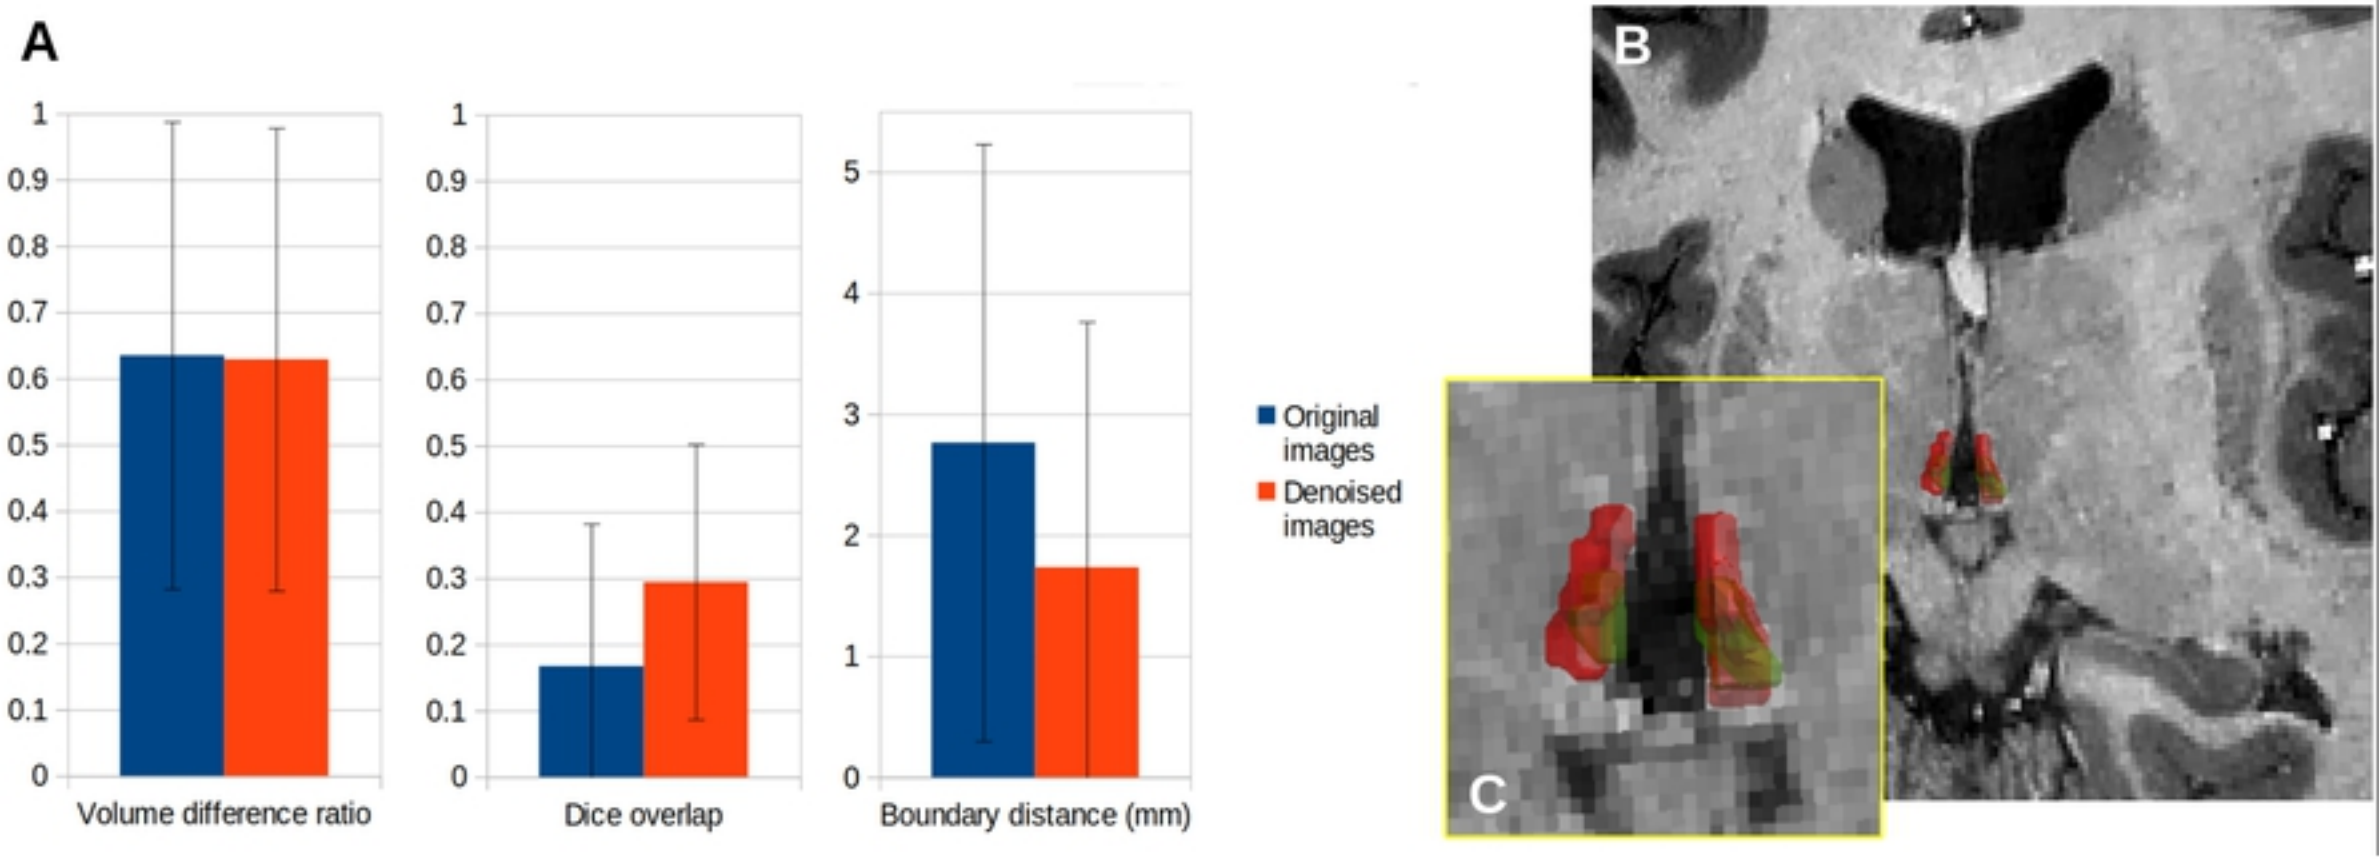
A
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Volume difference ratio
Dice overlap
5
4
3
2
1
Boundary distance (mm)
Original images
Denoised images
B
C

Figure 7