Deep learning-assisted model-based off-resonance correction for non-Cartesian susceptibility weighted imaging

Guillaume Daval-Frérot, Aurélien Massire, Boris Mailhé, Mariappan Nadar, Blanche Bapst, Alain Luciani, Alexandre Vignaud, Philippe Ciuciu

FULL PAPER

Deep learning-assisted model-based off-resonance correction for non-Cartesian susceptility weighted imaging

Guillaume Daval-Frérot[1,2,3] | Aurélien Massire[1] | Boris Mailhe[4] | Mariappan Nadar[4] | Blanche Bapst[5,6] | Alain Luciani[7,8,9] | Alexandre Vignaud[2] | Philippe Ciuciu[2,3]

[1]Siemens Healthineers, Saint-Denis, 93210, France

[2]CEA, NeuroSpin, CNRS, Paris-Saclay University, Gif-sur-Yvette, 91191, France

[3]Inria, MIND, Palaiseau, 91120, France

[4]Siemens Healthineers, Digital Technology & Innovation, Princeton, 08540, NJ, USA

[5]Departments of Neuroradiology, AP-HP, Henri Mondor, University Hospital, Créteil, 94010, France

[6]EA 4391, Paris-Est-Créteil University, Créteil, 94010, France

[7]Departments of Medical Imaging, AP-HP, Henri Mondor, University Hospital, Créteil, 94010, France

[8]Faculty of Public Health, Paris-Est-Créteil University, Créteil, 94010, France

[9]INSERM IMRB, U955, Equipe 18, Créteil 94000, France

Correspondence
Philippe Ciuciu PhD, NeuroSpin, CEA, Gif-sur-Yvette, 91191, France
Email: philippe.ciuciu@cea.fr

Purpose: Patient-induced inhomogeneities in the static magnetic field cause distortions and blurring (off-resonance artifacts) during acquisitions with long readouts such as in susceptibility-weighted imaging (SWI). Conventional versatile correction methods based on extended Fourier models are too slow for clinical practice in computationally demanding cases such as 3D high-resolution non-Cartesian multi-coil acquisitions.
Theory: Most reconstruction methods can be accelerated when performing off-resonance correction, by reducing the number of iterations, compressed coils and correction components. Recent state-of-the-art deep learning architectures could help but are generally not adapted to corrupted measurements as they rely on the standard Fourier operator in the data consistency term. The combination of correction models and neural networks is therefore necessary to reduce reconstruction times.
Methods: Hybrid pipelines using UNets were trained stack-by-stack over 99 SWI 3D SPARKLING 20-fold accelerated acquisitions at 0.6mm isotropic resolution using different off-resonance correction methods. Target images were obtained using slow model-based corrections based on self-estimated ΔB_0 field maps. The proposed strategies, tested over 11 volumes, are compared to model-only and network-only pipelines.
Results: The proposed hybrid pipelines achieved scores competing with 2-3 times slower baseline methods, and neural networks were observed to contribute both as pre-conditioner and through inter-iteration memory by allowing more degrees of freedom over the model design.
Conclusion: A combination of model-based and network-based off-resonance correction was proposed to significantly accelerate conventional methods. Different promising synergies were observed between acceleration factors (iterations, coils, correction) and model/network that could be expanded in the future.

KEYWORDS
off-resonance correction, non-Cartesian imaging, deep learning, unrolled neural network, 3D SPARKLING

1 | INTRODUCTION

Many parallel imaging and compressed-sensing (CS) methods [1, 2, 3, 4, 5, 6, 7] have been proposed over the last two decades to accelerate magnetic resonance imaging (MRI) acquisitions. Non-Cartesian sampling patterns [8, 9] have recently gained popularity through their capability to better exploit longer but fewer readouts. In particular, the Spreading Projection Algorithm for Rapid K-space sampLING (SPARKLING), proposed for 2D [9] and 3D [10, 11] imaging, responds to all degrees of freedom offered by modern MR scanners [11] to fully explore k-space and match optimized target sampling densities. Susceptibility-weighted imaging (SWI) [12], commonly used in high resolution brain venography or traumatic brain injuries [13], has been recently studied with SPARKLING [11, 14] to reach acceleration factors (AF) superior to 15 in scan times compared to fully sampled Cartesian imaging in high resolution (0.6mm) isotropic brain imaging. However non-Cartesian sampling patterns tend to be more sensitive to off-resonance artifacts causing geometric distortions and image blurring [15], notably with long readouts (e.g. 20ms), thereby inducing k-space inconsistencies over the different gradient directions [16, 17]. These artifacts emerge mostly from patient-induced static B_0 field inhomogeneities, notably pronounced near air-tissue interfaces, for instance in the vicinity of nasal cavity and ear canals.

Diverse methods have been proposed in the literature to correct those artifacts during the acquisition or image reconstruction. The spherical harmonic shimming technique is the current standard for all systems [15, 18] but is generally limited to second or third-order harmonics, which already provide critical improvements. More advanced shim coil designs have been proposed recently [19, 20] but still face technical and theoretical limitations [21]. Post-processing methods can therefore be necessary as a complement for more demanding cases, such as Cartesian EPI [22, 23] where alternating gradient direction at every time frame can be used to deduce and revert off-resonance induced geometric distortions. However this technique is not applicable to non-Cartesian readouts (e.g. spirals [24, 25, 26], rosette [27], SPARKLING [9, 10, 11]) due to the multiple spatially-encoding gradients played simultaneously. Another less constraining and well-established method [28, 29, 30, 31, 32] consists in compensating the undesired ΔB_0 spatial variations by modifying the Fourier operator involved in image reconstruction in order to integrate prior knowledge on a ΔB_0 field map. This technique can be applied to any imaging setup but it considerably slows down (e.g. 15-fold) the image reconstruction process. The mandatory ΔB_0 field map is directly available for multi-echo acquisitions [15], but it necessitates to be either externally collected by extending the scan time, or estimated [33, 34, 35, 36, 37, 38, 14].

Non-Cartesian acquisition strategies enable shorter scan times at the cost of increased image reconstruction duration. However, taking into account off-resonance correction within extended forward and adjoint operators has a multiplicative effect that makes the processing excessively long. In the recent years, deep learning (DL) has emerged for MRI reconstruction as a means to allow for improved image quality and faster processing, by similarly pushing the computation cost to offline training sessions. However, state-of-the-art network architectures are mostly focused on undersampling artifacts [39, 40, 41, 42, 43] as they enforce data consistency with Fourier operators, which is inaccurate when dealing with off-resonance effects. More specific literature tends to invest most efforts into estimating the ΔB_0 field map [44, 45], already available in the context of SWI acquisitions (see [14] for details).

In this work, we study different approaches [29, 30, 31] to model compressed representations of the non-Fourier operator involved in the data consistency term, and compensate for them using neural networks. The proposed extended non-Cartesian Primal-Dual network (NC-PDNet) architectures [43] are trained to reproduce self-corrected [14] high resolution SWI volumes based on highly accelerated multi-coil 3D SPARKLING [11] trajectories (AF>17) at 3 Tesla and each obtained through 8-hour long reconstructions. This approximation allows us to reach a significant acceleration for model inversion with respect to three key el-

ements, namely the number of unrolled iterations, compressed coils and correction components (i.e. interpolators involved in the non-Fourier operator), all contributing multiplicatively to the reconstruction time. The results are then compared to both model-only (CS reconstruction with non-Fourier operator) and network-only (original NC-PDNet) pipelines over 11 dedicated volumes and further decomposed to analyze the contributions of both neural networks and partially correcting models with respect to the three sources of approximation using tailored off-resonance metrics. The various benefits of hybrid architectures are demonstrated, with observed synergies paving the way to more improvements on image quality.

2 | THEORY

2.1 | Image reconstruction

For convenience we define $M = N_c \times N_s$ the total number of samples (with N_c the number of spokes and N_s the number of samples per spoke) measured over the k-space Ω and $N = N_x \times N_y \times N_z$ the total number of voxels (with N_x, N_y and N_z the image dimension in voxels).

In the absence of B_0 inhomogeneities, the reconstructed image $\widehat{x}$ can be obtained from the multi-channel k-space measurements $y = (y_q)_{q=1}^{Q}$ by solving:

$$\widehat{x} = \underset{x \in \mathbb{C}^N}{\operatorname{argmin}} \sum_{q}^{Q} \frac{1}{2} \| y_q - F_\Omega S_q x \|_2^2 + \mathcal{R}(x) \tag{1}$$

where Q is the number of channels, S_q is the sensitivity map of the q^{th} channel, and $\mathcal{R}$ a regularization function. The operator F_Ω is the non-uniform fast Fourier transform (NUFFT) defined through the ideal signal equation:

$$f(r) = \int_{T_{obs}} s(t)\, e^{ik(t)\cdot r}\, \mathrm{d}t \tag{2}$$

with T_{obs} the observation window in seconds, $s(t)$ the measured k-space sample at time t, $f(r)$ the object magnetization at position r, $k(t)$ the k-space position at time t. The sensitivity maps S_q can be externally acquired or estimated from the central θ% of the k-space [46], with the low frequency $F_{\Omega_{\theta\%}}$ operator:

$$S_q = \frac{F_{\Omega_{\theta\%}}^H y_q}{\sqrt{\sum_{p=1}^{Q} \| F_{\Omega_{\theta\%}}^H y_p \|_2^2}} \tag{3}$$

An efficient way to solve Eq. (1) is through proximal gradient descent:

$$w_{k+1} = x_k - \alpha_k \sum_{q=1}^{Q} S_q^H F_\Omega^H D \left(F_\Omega S_q x_k - y_q \right) \tag{4}$$

$$x_{k+1} = \text{prox}_{\mathcal{R}} \left(w_{k+1} \right) \tag{5}$$

where α_k is the step size at iteration k, and $\text{prox}_{\mathcal{R}}$ is a proximal operator associated with $\mathcal{R}$. The step described in Eq. (4) is often called data consistency. This basic approach can be extended to both CS and DL methods, using different algorithms [47, 48] and architectures [40, 42, 43] to reach improved image quality in fewer iterations. Hereafter, we replace $\text{prox}_{\mathcal{R}}$ with a neural network similarly to [43] and set $\alpha_k = \frac{1}{\beta}$ where β is the Lipschitz constant of the data consistency term, following [47].

2.2 | Signal correction

In order to apply ΔB_0 corrections, we need to extend the basic Fourier model from Eq. (2) to include the off-resonance effects [49]:

$$f(r) = \int_{T_{obs}} s(t)\, e^{i(k(t)\cdot r + \Delta\omega_0(r)t)}\, \mathrm{d}t \tag{6}$$

with $\Delta\omega_0(r) = \gamma \Delta B_0(r)$ the off-resonance frequency in radian at position r, γ the hydrogen gyromagnetic ratio and ΔB_0 the actual magnetic field deviation. This new signal formula is discretized as follows:

$$f(r_n) = \sum_{m=1}^{M} s(t_m)\, e^{i\Delta\omega_0(r_n)\, t_m}\, e^{ik(t_m)\cdot r_n}\,. \tag{7}$$

The term $\Delta\omega_0(r_n)\, t_m$ is dependent on both the k-space and image domain, which is not compatible with a regular Fourier transform. The approach initially proposed

by Noll et al. [28] and later extended in [29, 30, 31] amounts to splitting this exponential term into a sum of variables that are each dependent on a single domain:

$$e^{i\Delta\omega_0(r_n)\,t_m} = \sum_{\ell=1}^{L} b_{m,\ell}\, c_{\ell,n}\,. \tag{8}$$

This way, by combining Eqs. (7)– (8) we can factorize out the term dependent on the image domain and obtain a weighted sum of L regular Fourier transforms, with $L \ll M, N$:

$$f(r_n) = \sum_{\ell=1}^{L} c_{\ell,n} \sum_{m=1}^{M} s(t_m)\, b_{m,\ell}\, e^{ik(t_m)\cdot r_n}\,. \tag{9}$$

The coefficients $B = (b_{m,\ell}) \in \mathbb{C}^{M,L}$ and $C = (c_{\ell,n}) \in \mathbb{C}^{L,N}$ can be optimally estimated using the method proposed by Fessler et al. [31] considering the following matrix factorization problem:

$$\widehat{B}, \widehat{C} = \underset{B\in\mathbb{C}^{(M,L)}, C\in\mathbb{C}^{(L,N)}}{\operatorname{argmin}} \|E - BC\|_{Fro}^2 \tag{10}$$

with E the $M \times N$ matrix defined by $E_{m,n} = e^{i\Delta\omega_0(r_n)\,t_m}$. The optimal solution is obtained by decomposing E through Singular Value Decomposition (SVD) along either axis to obtain B or C, and then finding the other one as the least squares solution. We refer hereafter to this solution as Singular Vector Interpolation (SVI) coefficients.

Prior to this formalism, different solutions have been explored [49, 28, 50, 29, 1, 30] notably two that were respectively introduced by Man et al. [29] and Sutton et al. [30], and still currently used [51, 52]. Both methods were also tested in this work as optimality in regards to Eq. (10) does not imply optimality in regards to Eq. (1). Similarly, the solutions are given through least squares after fixing either B or C as follows:

$$e^{i\Delta\omega_0(r_n)\,t_m} = \sum_{\ell=1}^{L} e^{i\Delta\omega_{0,\ell}\,t_m}\, c_{\ell,n} \tag{11}$$

$$e^{i\Delta\omega_0(r_n)\,t_m} = \sum_{\ell=1}^{L} b_{m,\ell}\, e^{i\Delta\omega_0(r_n)\,t_\ell} \tag{12}$$

with $\Delta\omega_{0,\ell}$ and t_ℓ obtained by segmenting the off-resonance frequency range and the time window into L values, respectively. In what follows, the solution obtained from Eq. (11) named Multi-Frequency Interpolation [29] is referred to as MFI, and by analogy the solution from Eq. (12) with the method from [30] is referred to as MTI for Multi-Temporal Interpolation.

For all coefficients, the computational load can be considerably reduced by taking advantage of the spoke redundancy (i.e. using the same decomposition over the N_c spokes) and using histograms of the ΔB_0 field map to solve a weighted version of Eq. (10), typically decreasing the image dimensions $N = 384 \times 384 \times 208$ voxels to $N_b = 1000$ bins (see details in [31]). This way, the matrix E is reduced from $M \times N$ to $N_s \times N_b$ and therefore correction coefficients can be obtained in a few seconds for high resolution 3D volumes.

We obtain from Eq. (9) a pseudo-Fourier operator $F_{\Omega,\Sigma}$ (with Σ representing the interpolation processing) that can be implemented as a wrapper for any regular Fourier operator F_Ω and directly integrated into Eq. (1) and Eq. (3). The same remark also holds for the adjoint Fourier operator F_Ω^H.

2.3 | Accelerated reconstruction and correction

The above mentioned correction technique is convenient but still increases the computation cost by multiplying factors of an already time-consuming reconstruction in the case of 3D high-resolution non-Cartesian imaging. Starting from the algorithm presented in Eqs. (4)–(5) combined with the correction operator in Eq. (9), different ways to reduce the computational burden can be explored by decreasing: the number of proximal gradient iterations I, the number of channels Q, or the number of correction components L. All possibilities were considered and the last two are further explained in this subsection.

Different coil compression methods exist to decrease Q [53, 54, 55], but the most efficient ones, such as geometric coil compression [55], often exploit constraining k-space trajectory properties. More recent learning-

based techniques have not been explored but might be considered in the future. Meanwhile, we used the trajectory-independent method by Buehrer et al. [53] also based on SVD as it allows us to efficiently reduce Q while also ordering compressed channels by explained variance, as represented in Figure 1a and detailed afterwards. Some relations are observed between ΔB_0 field maps and compressed channels sensitivities in the supplementary materials (Figure S6) but without strong demarcations, and therefore only the first Q components are kept.

On the contrary, the SVI correction coefficients cover specific regions of the off-resonance spectrum (Figure S5). Simply using the first L components, also studied hereafter, would mostly shift the data consistency focus toward low off-resonance areas that cover a much broader part of the brain images. A possibility to extend the spectrum coverage is to change the components used for data consistency over consecutive iterations. The SVI coefficients are more convenient for this, as if we decide to integrate L_2 SVI components instead of L_1 with $L_2 > L_1$, the first L_1 are the same due to the orthogonality of the decomposition. This is not true for MFI and MTI methods as observed in Figure S3-S4. It ensures that while the first L_1 components carry the maximal amount of information, they are not redundant with the other $L_2 - L_1$ components. Diverse strategies have been considered, and the best performing one updates data consistency between the first L components and the $L + 1$ to $2L$ components over the iterations, called $\text{SVI}_{1/2}$ hereafter, to enforce fidelity toward either low or high off-resonance areas alternatively.

The goal is to provide improved reconstruction and correction while minimizing the processing time. $\mathcal{I}$, Q and L linearly multiply the time with exponentially decreasing quality contributions, as shown in Figure 1 and in supplementary materials with Table S4. Our approach therefore consists in decreasing time and allowing the neural networks and recover image quality based on largely complete information.

3 | METHODS

3.1 | Proposed pipelines

Following the recommendations from Ramzi et al. [42], the end-to-end NC-PDNet architecture [43] has been used as we consider 3D non-Cartesian SWI data. Particularly, we investigated the primal-only version where only the image domain processing is learned with an arbitrary neural network whereas data consistency is applied in k-space. However, in the context of 3D high resolution non-Cartesian and multi-coil imaging, the amount of GPU memory required for a complete training is excessive. All or part of that memory can be moved to CPU, but would result in a considerably longer training duration.

To address this issue, a solution is to break down the learning process into stacks to avoid any restriction on the number of iterations or network size. An ideal case would be to have stacks consisting of at least one network pass followed by a data consistency block to still learn complementary features, however the memory requirement explodes at the transition between single channel image domain and multi-coil k-space data. The simplest proposition therefore, is to exclude the data consistency from the gradient computation by training the dataset image-to-image as represented in Figure 2 and then only updating the dataset in k-space at each stack. This modification requires adapting the feature originally named "memory" [40] and later "buffer" [42] as the multiple network output channels are not all included in the backpropagation graph anymore. Instead, the network output is a single complex-valued volume but that will still be stacked into a buffer of size N_B to be used as input for the next stack. That way, the overall pipeline keeps the expressiveness required to learn CS-like acceleration schemes as originally suggested [40]. The main drawback compared to end-to-end training is enforcing that each stack should yield the target results rather than letting the consecutive networks store intermediate states. However, this modification allows us to train architectures with an arbitrarily high number of coils and iterations.

The different partial correction strategies based on $F_{\Omega,\Sigma}$ discussed in Section 2.3 are studied within an unrolled architecture in combination with UNets in the image domain. Indeed, conventional data consistency would either enforce the off-resonance artifacts in spite of the networks, or at best would not contribute to their correction. For the main studies, the variables $\mathcal{I} = 5$, $Q = 5$ and $L = 5$ have been retained to allow for 3D 0.6mm isotropic corrections within 7 to 8 minutes of computing time at inference. The MFI, MTI and SVI coefficients are compared, along with the $SVI_{1/2}$ strategy, with $N_B = 3$ buffers and pre-computed density compensations [56] for the NUFFT operator from the `gpuNUFFT`[1] [57] and `pysap-mri`[2] [58] packages. The residual UNets are composed of three scales, each made of convolutional blocks of three layers with kernel size of $3 \times 3 \times 3$ and each followed by a ReLU activation (except for the final layer). The number of filters is doubled at each 3D $2 \times 2 \times 2$ downscale, with 16 filters at the first scale, and vice-versa for upscales. The complex nature of the input and output volumes is handled by considering real and imaginary parts as two separate real-valued channels. Each UNet is composed of 390,066 parameters with buffers and 388,338 without. They are trained for 300 epochs, resulting in 100-hour long trainings overall, with RAdam optimizer and learning rate of 5×10^{-4}. The minimized cost function is the sum of an L1 loss applied to complex-valued images and multiscale SSIM to the magnitude images. All training experiments were run on the Jean-Zay supercomputer over a single NVIDIA Tesla V100 GPU with 32GB of VRAM.

3.2 | Dataset

A total of 123 SWI volumes were acquired on patients with non-Cartesian 3D GRE sequences at 3T (Magnetom Prisma, Siemens Healthcare, Erlangen, Germany) with a 64-channel head/neck coil array. The protocol was approved by local and national ethical committees (IRB: CRM-2111-207). Patient demographics and a study flow diagram are provided in the supplementary materials (Section S1). The dataset covers a wide range of pathologies (aneurysm, sickle cell anemia, multiple sclerosis) and off-resonance related artifact levels.

Four different variations of the recently proposed full 3D SPARKLING [11] sampling pattern were used, but the vast majority (n=102) was acquired using the following parameters: a 0.6mm isotropic resolution, a field-of-view of 24 cm in-plane (N=384) over 12.5 cm (N_z=208), a readout of $T_{obs} = 20.48$ ms centered around an echo time $TE = 20$ ms, a repetition time $TR = 37$ ms. A dwell time of $\delta t = 2\ \mu$s was used to balance the small number of spokes N_c=4900, resulting in N_s=10,240 samples per spoke. The other closely related acquisition variations are similarly described in the supplementary materials (Section S1). An acquisition time of 3 min corresponds to an acceleration factor (AF) of 17, defined as follows:

$$AF = \frac{N \times N_z}{N_c}\,. \tag{13}$$

The ΔB_0 field maps were not acquired in order to avoid prolonging the exams. Instead self-estimated field maps were computed a posteriori using a recently published technique [14]. To generate the ground truths, the field maps were used for model-based correction using the method described in [31] with the SVI coefficients, resulting in approximately 8 hours long reconstructions/corrections with $\mathcal{I} = 20$, $Q = 20$ and $L = 20$ over a single NVIDIA Tesla V100 GPU.

The dataset was then split into training (n=99), validation (n=11) and testing (n=11) sets according to balanced age, gender, weight, off-resonance pre and post-correction visibility, and pathology type and visibility distributions. Two acquisitions were excluded due to strong motion (n=1) and insufficient off-resonance correction (n=1) caused by braces. In both cases, the self-estimated off-resonance correction still considerably improved the images. Minor quality concerns were raised regarding skin fat artifacts (n=6) and partially deactivated readout coils (n=2) without causing exclusion.

Although the trainings were carried out from complex-valued to complex-valued volumes, the SWI specific processing was applied afterwards for visualization and scoring as described in [12]. The low fre-

[1] `https://github.com/andyschwarzl/gpuNUFFT`

[2] `https://github.com/CEA-COSMIC/pysap-mri`

quencies were extracted by applying a Hanning window over the central third of k-space, before being removed from the phase image to obtain a high frequency map, subsequently normalized to produce a continuous mask.The magnitude image was multiplied five times by the mask, and a minimum-intensity projection (mIP) was computed using a thickness of 8mm. All post-processing steps were again run on the Jean-Zay supercomputer.

3.3 | Baselines

Diverse baseline models are proposed to assess the contributions of the different features involved in our pipeline. A first baseline detailed in the supplementary materials (Section S3) is given by replacing the neural networks by conventional CS reconstruction using sparsity promoting regularization in the wavelet domain (Symlet 8 basis decomposed over three scales). The L_1-norm was used for $\mathcal{R}$ function with $\lambda = 10^{-7}$ for thresholding the wavelet coefficients. To compensate the absence of buffers to learn acceleration schemes over iterations, the FISTA [47] algorithm available in the `pysap-mri` package was used. The data consistency is implemented with $F_{\Omega,\Sigma}$ with SVI coefficients reduced to $L = 5$. This baseline noted $\text{BASE}_{\mathcal{R}}$ is used to determine to what extent the contribution of neural networks is critical for improved image quality.

The second baseline BASE_C similarly replaces the correcting operator $F_{\Omega,\Sigma}$ with the regular NUFFT F_Ω while applying the same network-based regularization as the proposed pipeline.

Finally, the SVI and $\text{SVI}_{1/2}$ pipelines were tested without buffers, along with the BASE_C baseline. The goal was to assess that the buffer feature, modified to fit the stacked-training setup, was still relevant even without end-to-end training.

3.4 | Evaluation

The results are assessed according to several criteria: Image quality, off-resonance correction and speed. The first one is commonly assessed [59] using the structural similarity index (SSIM) and the peak signal-to-noise ratio (PSNR) measures. These metrics also cover off-resonance related image artifacts, and therefore to facilitate the interpretation of the results we also provided ΔB_0-weighted versions of SSIM and PSNR. In both cases, the voxel-wise score at position r was weighted by:

$$w_{\Delta B_0}(r) = \frac{|\Delta B_0(r)|}{\overline{|\Delta B_0|}} \tag{14}$$

with $\overline{|\Delta B_0|}$ the average value of the absolute ΔB_0 field map. In the case of PSNR, it corresponds to using a weighted mean-squared error (MSE) in the expression. Among those four metrics, the ΔB_0-weighted SSIM was chosen as the primary score for this study as it relates more to the visual impact of off-resonance artifacts, followed by classic SSIM for image quality in general. Since these criteria are not defined for complex-valued data, all metrics are applied after SWI processing to also account for magnetic susceptibility information. The individual scores were then compared between pipelines for both SSIM and ΔB_0-weighted SSIM in order to obtain p-values through two-sided Wilcoxon signed-rank tests, before applying a Benjamini-Hochberg correction to adjust the false discovery rate at a p-value of $p = 0.05$.

The final aspect we focus on in this study is the computation speed of the different pipelines. The main bottleneck is the NUFFT, and has therefore been the focus for more code optimizations[3,4]. Some other elements, such as the wavelet transform used for sparsity-promoting regularization or the computation of correcting coefficients $(\widehat{B}, \widehat{C})$, that are already fast in comparison but not negligible, did not benefit from similar attention. Consequently, these computations might irrelevantly change the duration. Additionally, as the proposed pipeline optimizations are solely related to the amount of individual calls to the NUFFT operator, we decided to track them as N_F to account for the numerical complexity rather than using misleading time mea-

[3] `https://github.com/andyschwarzl/gpuNUFFT`
[4] `https://github.com/chaithyagr/gpuNUFFT`

surements:

$$N_F = N_{F,S} + N_{F,R} \tag{15}$$

$$N_{F,S} = L_S \times Q \tag{16}$$

$$N_{F,R} = L \times Q \times (2 \times \mathcal{I} - 1) \tag{17}$$

with $N_{F,S}$ the number of calls related to the sensitivity map computation, and $N_{F,R}$ the number of calls related to image reconstruction. The sensitivity maps are systematically corrected with $L_S = 10$ components when $F_{\Omega,\Sigma}$ is used in Eq. (3), but not corrected at all when using F_Ω which is equivalent to $L_S = 1$. The reconstruction consists of one iteration for initialization and $\mathcal{I} - 1$ iterations using data consistency. Note that the forward and adjoint calls are similarly considered for simplicity reasons, although they differ over non-Cartesian grids.

4 | RESULTS

The different pipelines are confronted one another hereafter, along with an ablation study of the modified buffer feature. Additional resources are provided in the supplementary materials, and notably more exhaustive CS baselines to quantify and understand the specific contributions of varying parameters $\mathcal{I}$, Q and L (Section S3). The number of coils Q in particular is shown to have little influence on off-resonance correction, and the coil sensitivity profiles normalized in Eq. (3) seem to have more impact on all metrics than information carried by additional coils. In contrast, increasing L contributes solely to off-resonance artifacts correction while augmenting $\mathcal{I}$ positively impacts both off-resonance and high frequency content.

4.1 | Correction pipelines

The different scores are reported in Table 1 and Figure 3, with corresponding pairwise p-values in Table 2, and illustrated on two acquisitions in Figures 4–5, respectively. Firstly, the baselines BASE_C and $\text{BASE}_\mathcal{R}$ in Table 1 demonstrate the expected positive contribution of a partially corrected Fourier operator $F_{\Omega,\Sigma}$ against a basic model (BASE_C), and that of stacked UNets against wavelet-based regularization ($\text{BASE}_\mathcal{R}$). The only statistically significant difference holds for ΔB_0-weighted SSIM, with a large improvement over $\text{BASE}_\mathcal{R}$. Therefore, the proposed UNet architecture, when combined with conventional operator F_Ω, is not capable of correcting off-resonance effects as much as a partially correcting operator $F_{\Omega,\Sigma}$ does.

Secondly, the different correction approaches (MFI, MTI, SVI) are combined with UNets through data consistency terms. All of them show significantly higher scores than both baselines. Their ΔB_0-weighted SSIM score follow the same ranking as those previously observed in Figure 1b prior to image reconstruction when solving Eq. (10): the SVI coefficients reach a better score as compared to MFI (second) and MTI (third) coefficients. This suggests that none of the diverse correction approaches carried by the different coefficients are better suited than another to help UNets compensate for the missing correction.

Finally, the $\text{SVI}_{1/2}$ pipeline improves the SVI approach with regard to ΔB_0-weighted SSIM, while not significantly deviating from it for classic SSIM. Exploring more components appears to guide the correction that networks alone could not achieve. This can be observed in Figure 4 near the bucco-nasal region visible with the sagittal views, where only $\text{SVI}_{1/2}$ recovers high-frequency details consistent with the target. The additional baselines from supplementary materials (Section S3) show that the $\text{SVI}_{1/2}$ pipeline competes with reconstruction 2-3 times slower depending on the evaluation criteria.

4.2 | Network and model contributions

The reconstruction steps are further decomposed between data consistency and regularization in Figure 6 and illustrated in Figure 7. For both metrics, the two baselines shown in Figure 6 with brown ($\text{BASE}_\mathcal{R}$) and gray (BASE_C) curves are again significantly different. The $\text{BASE}_\mathcal{R}$ curves show a slow start with oscillations between improving data consistency and degrading regularization but with a regular progression afterwards.

The oscillations are explained by the constant wavelet threshold, too large during early iterations when the overall magnitude was still too low. In contrast, the networks used in BASE_C considerably improve the initialization, then followed by oscillations caused by the data consistency. Those oscillations can be explained by the off-resonance artifacts enforced through the wrong Fourier model F_Ω but also by the sensitivity profiles altered by the normalization from Eq. (3) when reducing the number of channels Q. Both are illustrated in Figure 7 with green and blue arrows respectively.

The proposed pipelines SVI and $\text{SVI}_{1/2}$ are observed to combine a steeper initialization and better progression over iterations, but still with sensibility-related oscillations. The $\text{SVI}_{1/2}$ is however more robust to those as it alternates between large regions with heterogeneous sensitivities and smaller ΔB_0-specific regions. Overall, the UNets appear to contribute to all three studied factors, but the striking difference between BASE_C, SVI and $\text{SVI}_{1/2}$ pipelines suggests that networks only recover low image frequencies when correcting off-resonance. The partial correction model is therefore required anyway in those regions, but networks serve as an effective image pre-conditioner.

4.3 | Buffer feature ablation

The main pipelines have been evaluated with and without the proposed buffer feature modified to account for stacked training, namely SVI and $\text{SVI}_{1/2}$ along with BASE_C. The different results are shown in Table 3.

The baseline BASE_C that combines UNets with conventional Fourier data consistency shows ΔB_0 scores significantly worse when using the buffer feature. An interpretation would be that as consistency with corrupted brings back artifacts, as observed in Figure 7, adding memory of previous iterations only leads to a noisier learning process without additional information.

On the other hand, the $\text{SVI}_{1/2}$ approach shows significant and large improvements with buffers, more than for the SVI pipeline. However, it should be noted that $\text{SVI}_{1/2}$ classic scores without buffers are much lower than any other proposed pipeline, but still better than the baselines. This suggests that the exploratory correction strategy inherently degrades the overall reconstruction to advantage the off-resonance areas, but that buffer feature allows for a compensatory effect within neural networks.

5 | DISCUSSION

The diverse contributions of this article are discussed hereafter. Firstly, we compared conventional methods [29, 30, 31] to perform, to the best of our knowledge, the first partial off-resonance correction study. Secondly, we analyzed the deep learning contributions in a multi-parametric acceleration pipeline. Finally, we conducted a deep learning study over a fairly large in vivo dataset consisting of model-based self-estimated off-resonance corrected volumes [14].

5.1 | Network and model improvements

The network and model contributions when facing reduced number of iterations, compressed coils, and correction components have been studied in Section 4.

Overall, the number of coils Q was observed to impact metrics by changing the sensitivity distribution superficially rather than improving the signal. Other normalization [60] or coil compression methods [61] could be explored in the future to at least balance this scoring issue. The proposed UNets appeared to mostly improve the initialization overall and help progression through the buffer feature when applying the $\text{SVI}_{1/2}$ strategy otherwise less efficient. The latter shows how deep learning allows for more flexibility when designing hybrid algorithms. However, high-frequency details are still only recovered through model-based correction.

More elaborated strategies could be developed to distribute efficiently the correction components over iterations. Particularly, for in-out SPARKLING trajectories the MTI coefficients were shown to correlate with high/low image frequencies. A better balance between low frequencies covered by neural networks and high frequencies enforced by specifically selected correction

components might produce more reliable results.

5.2 | Dataset generation

The ΔB_0 map estimation technique developed in a previous publication [14] assumes the relationship between B_0 inhomogeneities and image phase to be dominant over other phase sources for large echo times (e.g. 20ms and higher) at 3T. It allows us to simply estimate the ΔB_0 field map retrospectively with minimal error and no motion-related mismatch for SWI acquisitions and alike. It also avoids any assumptions on the trajectory, and can therefore be used on any dataset matching the above mentioned acquisition setup and providing either the raw k-space data or complex-valued MR images.

The self-estimation method has been carefully assessed during the early stages of the study, and reached our expectations. However, two other competing methods were also considered, the first by Lee et al. [38] through simulation from a binary mask, and the second by Patzig et al. [35] through non-convex optimization. The simulation method in particular has been implemented in Python but the required mask obtained from the artifacted images was suboptimal. The main advantage of those three self-estimation methods for the purpose of our proposed acceleration pipeline is to require strictly the same data. This implies that the pipeline could be efficiently applied within the clinical context and additionally improve over sessions using the same acquisitions processed through longer and dedicated procedures.

6 | CONCLUSION

MR acquisitions based on non-Cartesian acquisitions tend to exploit longer but fewer complex readouts to reduce overall scan times, but suffer more from B_0 inhomogeneities. Diverse methods exist to compensate those artifacts, but the few non-constraining ones slow down reconstructions by an unacceptably large factor (10-20 at 3T). Deep learning techniques have been developed in the recent years, but mostly to estimate the ΔB_0 field map which can already be self-estimated using efficient models for SWI, or to compensate for undersampling artifacts in the CS setting. The proposed approach combines partial models and deep learning to ensure data fidelity at a low computation cost when addressing off-resonance artifacts and gain in image quality. It outlines both their individual and joint contributions. The MR volumes reconstructed in only 7-8 minutes are competing with baselines obtained in 30 minutes to approximate 8-hour long computations, and could be further accelerated by developing new hybrid strategies. Future work might also explore the parallel aspects of the correcting models to better fit memory constraints of the GPUs, or simply include undersampling compensation through self-supervised learning methods.

Acknowledgments

The concepts and information presented in this article are based on research results that are not commercially available. Future availability cannot be guaranteed. This work was granted access to the HPC resources of IDRIS under the allocation 2021-AD011011153 made by GENCI. We thank Dr. Zaccharie Ramzi and Chaithya G R for their previous works and open-source contributions. Additionally we thank Cecilia Garrec for the scientifc editing in English.

References

1 Pruessmann KP, Weiger M, Börnert P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. Magnetic Resonance in Medicine 2001;46(4):638–651.

2 Griswold MA, et al. Generalized autocalibrating partially parallel acquisitions (GRAPPA). Magnetic Resonance in Medicine 2002;47(6):1202–1210.

3 Lustig M, et al. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magnetic Resonance in Medicine 2007;58(6):1182–1195.

4 Seiberlich N, Breuer FA, Blaimer M, Barkauskas K, Jakob PM, Griswold MA. Non-Cartesian data reconstruction

using GRAPPA operator gridding (GROG). Magnetic Resonance in Medicine 2007;58(6):1257–1265.

5 Vasanawala SS, Alley MT, Hargreaves BA, Barth RA, Pauly JM, Lustig M. Improved pediatric MR imaging with compressed sensing. Radiology 2010;256(2):607–616.

6 Boyer C, et al. On the generation of sampling schemes for MRI. SIAM Imaging Sciences 2016;9(4):2039–2072.

7 Luo T, Noll DC, Fessler JA, Nielsen JF. A GRAPPA algorithm for arbitrary 2D/3D non-Cartesian sampling trajectories with rapid calibration. Magnetic resonance in Medicine 2019;82(3):1101–1112.

8 Polak D, Cauley S, Huang SY, Longo MG, Conklin J, Bilgic B, et al. Highly-accelerated volumetric brain examination using optimized wave-CAIPI encoding. Journal of Magnetic Resonance Imaging 2019;50(3):961–974.

9 Lazarus C, Weiss P, Chauffert N, Mauconduit F, El Gueddari L, Destrieux C, et al. SPARKLING: variable-density k-space filling curves for accelerated T2*-weighted MRI. Magnetic Resonance in Medicine 2019;81(6):3643–3661.

10 Lazarus C, Weiss P, El Gueddari L, Mauconduit F, Massire A, Ripart M, et al. 3D variable-density SPARKLING trajectories for high-resolution T2*-weighted magnetic resonance imaging. NMR in Biomedicine 2020;33(9):e4349.

11 Chaithya GR, Weiss P, Daval-Frérot G, Massire A, Vignaud A, Ciuciu P. Optimizing full 3D SPARKLING trajectories for high-resolution Magnetic Resonance Imaging. IEEE Transactions on Medical Imaging 2022;.

12 Haacke EM, Xu Y, Cheng YCN, Reichenbach JR. Susceptibility weighted imaging (SWI). Magnetic Resonance in Medicine 2004;52(3):612–618.

13 Wu Z, Mittal S, Kish K, Yu Y, Hu J, Haacke EM. Identification of calcification with MRI using susceptibility-weighted imaging: a case study. Journal of Magnetic Resonance Imaging 2009;29(1):177–182.

14 Daval-Frérot G, Massire A, Mailhé B, Nadar MS, Vignaud A, Ciuciu P. Iterative static ΔB_0 field map estimation for off-resonance correction in non-Cartesian susceptibility weighted imaging. Magnetic Resonance in Medicine 2022;.

15 de Graaf RA, Juchem C. B0 Shimming Technology. In: Magnetic Resonance Technology; 2016.p. 166–207.

16 Yudilevich E, Stark H. Spiral sampling in magnetic resonance imaging-the effect of inhomogeneities. IEEE transactions on medical imaging 1987;6(4):337–345.

17 Chaithya GR, Daval-Frérot G, Massire A, Mailhe B, Nadar M, Vignaud A, et al. MORE-SPARKLING: Non-Cartesian trajectories with Minimized Off-Resonance Effects. In: Proceedings of the 30th Annual Meeting of ISMRM London; 2022. p. 1611.

18 Gruetter R. Automatic, localized in vivo adjustment of all first-and second-order shim coils. Magnetic Resonance in Medicine 1993;29(6):804–811.

19 Aghaeifar A, Zhou J, Heule R, Tabibian B, Schölkopf B, Jia F, et al. A 32-channel multi-coil setup optimized for human brain shimming at 9.4 T. Magnetic Resonance in Medicine 2020;83(2):749–764.

20 Pinho-Meneses B, Amadon A. A fieldmap-driven few-channel shim coil design for MRI of the human brain. Physics in Medicine & Biology 2021;66(1):015001.

21 Pinho-Meneses B, Amadon A. Physical limits to human brain B0 shimming with spherical harmonics, engineering implications thereof. Magnetic Resonance Materials in Physics, Biology and Medicine 2022;p. 1–19.

22 Andersson JL, Skare S, Ashburner J. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. Neuroimage 2003;20(2):870–888.

23 Fritz L, Mulders J, Breman H, Peters J, Bastiani M, Roebroeck A, et al. Comparison of EPI distortion correction methods at 3T and 7T. In: Proceedings of the Annual Meeting of OHBM Hamburg, Germany; 2014. .

24 Weiger M, Pruessmann KP, Österbauer R, Börnert P, Boesiger P, Jezzard P. Sensitivity-encoded single-shot spiral imaging for reduced susceptibility artifacts in BOLD fMRI. Magnetic Resonance in Medicine 2002;48(5):860–866.

25 Wilm BJ, Barmet C, Gross S, Kasper L, Vannesjo SJ, Haeberlin M, et al. Single-shot spiral imaging enabled by an expanded encoding model: Demonstration in diffusion MRI. Magnetic resonance in medicine 2017;77(1):83–91.

26 Kasper L, Engel M, Barmet C, Haeberlin M, Wilm BJ, Dietrich BE, et al. Rapid anatomical brain imaging using spiral acquisition and an expanded signal model. NeuroImage 2018;168:88–100.

27 Noll DC. Multishot rosette trajectories for spectrally selective MR imaging. IEEE transactions on medical imaging 1997;16(4):372–377.

28 Noll DC, Meyer CH, Pauly JM, Nishimura DG, Macovski A. A homogeneity correction method for magnetic resonance imaging with time-varying gradients. IEEE transactions on medical imaging 1991;10(4):629–637.

29 Man LC, Pauly JM, Macovski A. Multifrequency interpolation for fast off-resonance correction. Magnetic Resonance in Medicine 1997;37(5):785–792.

30 Sutton BP, Noll DC, Fessler JA. Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities. IEEE transactions on medical imaging 2003;22(2):178–188.

31 Fessler JA, Lee S, Olafsson VT, Shi HR, Noll DC. Toeplitz-based iterative image reconstruction for MRI with correction for magnetic field inhomogeneity. IEEE Transactions on Signal Processing 2005;53(9):3393–3402.

32 Ostenson J, Robison RK, Zwart NR, Welch EB. Multifrequency interpolation in spiral magnetic resonance fingerprinting for correction of off-resonance blurring. Magnetic resonance imaging 2017;41:63–72.

33 Man LC, Pauly JM, Macovski A. Improved automatic off-resonance correction without a field map in spiral imaging. Magnetic Resonance in Medicine 1997;37(6):906–913.

34 Sutton BP, Noll DC, Fessler JA. Dynamic field map estimation using a spiral-in/spiral-out acquisition. Magnetic Resonance in Medicine 2004;51(6):1194–1204.

35 Patzig F, Wilm B, Gross S, Brunner D, Pruessmann KP. Off-Resonance Self-Correction for Single-Shot Imaging. In: Proceedings of the 28th Annual Meeting of ISMRM Virtual; 2020. p. 3400.

36 Patzig F, Wilm B, Pruessmann KP. Off-Resonance Self-Correction by Implicit B0-Encoding. In: Proceedings of the 29th Annual Meeting of ISMRM Virtual; 2021. p. 666.

37 Salomir R, de Senneville BD, Moonen CT. A fast calculation method for magnetic field inhomogeneity due to an arbitrary distribution of bulk susceptibility. Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering: An Educational Journal 2003;19(1):26–34.

38 Lee SK, Hwang SH, Barg JS, Yeo SJ. Rapid, theoretically artifact-free calculation of static magnetic field induced by voxelated susceptibility distribution in an arbitrary volume of interest. Magnetic Resonance in Medicine 2018;80(5):2109–2121.

39 Muckley MJ, Riemenschneider B, Radmanesh A, Kim S, Jeong G, Ko J, et al. Results of the 2020 fastmri challenge for machine learning mr image reconstruction. IEEE transactions on medical imaging 2021;40(9):2306–2317.

40 Adler J, Öktem O. Solving ill-posed inverse problems using iterative deep neural networks. Inverse Problems 2017;33(12):124007.

41 Schlemper J, Caballero J, Hajnal JV, Price A, Rueckert D. A deep cascade of convolutional neural networks for MR image reconstruction. In: International conference on information processing in medical imaging Springer; 2017. p. 647–658.

42 Ramzi Z, Ciuciu P, Starck JL. Benchmarking MRI reconstruction neural networks on large public datasets. Applied Sciences 2020;10(5):1816.

43 Ramzi Z, Chaithya G, Starck JL, Ciuciu P. NC-PDNet: a Density-Compensated Unrolled Network for 2D and 3D non-Cartesian MRI Reconstruction. IEEE Transactions on Medical Imaging 2022;.

44 Zeng DY. Acceleration of 3D Magnetic Resonance Angiography. Stanford University; 2020.

45 Haskell M. Deep Learning Field Map Estimation with Model-Based Image Reconstruction for Off-Resonance Correction of Brain Images Using a Spiral Acquisition. In: ISMRM workshop on Data Sampling and Image Reconstruction Sedona, AZ, United States; 2020. p. 43.

46 El Gueddari L, Lazarus C, Carrié H, Vignaud A, Ciuciu P. Self-calibrating nonlinear reconstruction algorithms for variable density sampling and parallel reception MRI. In: 2018 IEEE 10th Sensor Array and Multichannel Signal Processing Workshop (SAM) Sheffield, UK: IEEE; 2018. p. 415–419.

47 Beck A, Teboulle M. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. IEEE transactions on image processing 2009;18(11):2419–2434.

48 Kim D, Fessler JA. Adaptive restart of the optimized gradient method for convex optimization. Journal of Optimization Theory and Applications 2018;178(1):240–263.

49 Maeda A, Sano K, Yokoyama T. Reconstruction by weighted correlation for MRI with time-varying gradients. IEEE transactions on medical imaging 1988;7(1):26–31.

50 Noll DC. Reconstruction techniques for magnetic resonance imaging. PhD thesis, Stanford University; 1991.

51 Gerhalter T, Gast LV, Marty B, Uder M, Carlier PG, Nagel AM. Assessing the variability of 23Na MRI in skeletal muscle tissue: Reproducibility and repeatability of tissue sodium concentration measurements in the lower leg at 3 T. NMR in Biomedicine 2020;33(5):e4279.

52 Hardee JE, Phaneuf C, Cope L, Zucker R, Gearhardt A, Heitzeg M. Neural correlates of inhibitory control in youth with symptoms of food addiction. Appetite 2020;148:104578.

53 Buehrer M, Pruessmann KP, Boesiger P, Kozerke S. Array compression for MRI with large coil arrays. Magnetic Resonance in Medicine 2007;57(6):1131–1139.

54 Doneva M, Börnert P. Automatic coil selection for channel reduction in SENSE-based parallel imaging. Magnetic Resonance Materials in Physics, Biology and Medicine 2008;21(3):187–196.

55 Zhang T, Pauly JM, Vasanawala SS, Lustig M. Coil compression for accelerated imaging with Cartesian sampling. Magnetic Resonance in Medicine 2013;69(2):571–582.

56 Pipe JG, Menon P. Sampling density compensation in MRI: rationale and an iterative numerical solution. Magnetic Resonance in Medicine 1999;41(1):179–186.

57 Knoll F, Schwarzl A, Diwoky C, Sodickson DK. gpuNUFFT-an open source GPU library for 3D regridding with direct Matlab interface. In: Proceedings of the 22nd Annual Meeting of ISMRM Milan, Italy; 2014. p. 4297.

58 El Gueddari L, Chaithya G, Ramzi Z, Farrens S, Starck S, Grigis A, et al. PySAP-MRI: a Python Package for MR Image Reconstruction. In: ISMRM workshop on Data Sampling and Image Reconstruction Sedona, AZ, United States; 2020. .

59 Hore A, Ziou D. Image quality metrics: PSNR vs. SSIM. In: 2010 20th international conference on pattern recognition IEEE; 2010. p. 2366–2369.

60 Herterich R, Sumarokova A, Coil Sensitivity Estimation and Intensity Normalisation for Magnetic Resonance Imaging (Dissertation). KTH Royal Institute of Technology; 2019. Retrieved from http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-263149.

61 Cole E, Meng Q, Pauly J, Vasanawala S. Learned Compression of High Dimensional Image Datasets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2022. p. 1748–1752.

(a) Compressed channels total variance proportions.

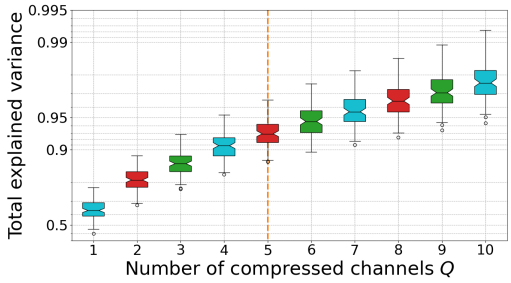

(b) Correction matrix factorization NRMSE.

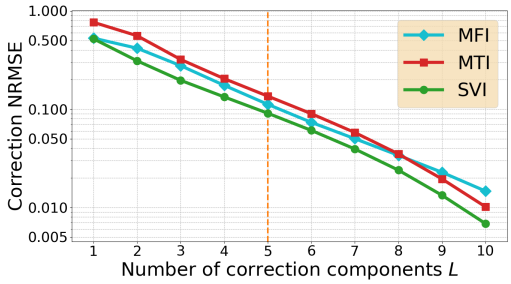

FIGURE 1 Expected model-based contributions from coil compression and partial correction. The contributions from coil compression and partial correction are each considered prior to reconstruction on reversed logarithmic scales. For coil compression, the total explained variance known using SVD is given for varying Q. For off-resonance correction, the normalized root-mean-square error from Eq. (10) is given for MFI, MTI and SVI coefficients with varying L. In both cases, the values are obtained from the 99 training acquisitions.

Figures and Tables

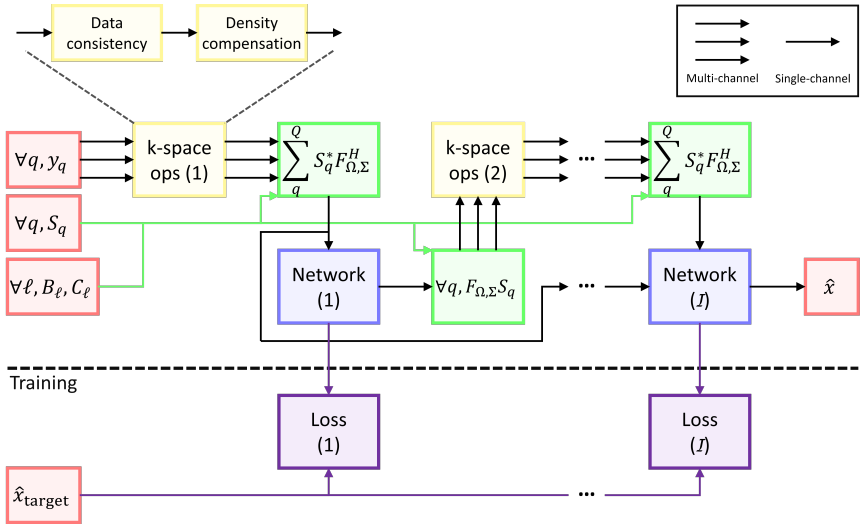

FIGURE 2 **Illustration of the proposed cross-domain pipeline with stacked-training.**
The proposed pipeline is shown for Q compressed coils y_q, the corresponding sensitivity maps S_q, and L correction coefficients B_ℓ and C_ℓ used in the partially-correcting operator $F_{\Omega,\Sigma}$, over $\mathcal{I}$ unrolled iterations. Each stack is trained independently from the first to last iteration, by checkpointing the entire dataset when reaching single-channel image domain.

TABLE 1 **Scores for different correction coefficients and pipelines over the test dataset.**
Various pipelines with $\mathcal{I} = 5$, $Q = 5$ and $L = 5$ are summarized with their data consistency and regularization terms and evaluated using classic and ΔB_0-weighted SSIM and PSNR scores.

Pipeline	Data consistency	Regularization	Classic		ΔB_0-weighted	
			SSIM	PSNR	SSIM	PSNR
BASE_C	F_Ω	UNet	0.9421	33.32	0.8947	23.50
$\text{BASE}_\mathcal{R}$	$F_{\Omega,\Sigma}$ (SVI)	Wavelet	0.9392	31.32	0.9249	23.33
MFI	$F_{\Omega,\Sigma}$ (MFI)	UNet	0.9596	34.96	0.9475	26.42
MTI	$F_{\Omega,\Sigma}$ (MTI)	UNet	0.9601	35.13	0.9452	26.11
SVI	$F_{\Omega,\Sigma}$ (SVI)	UNet	0.9613	**35.24**	0.9498	26.57
$\text{SVI}_{1/2}$	$F_{\Omega,\Sigma}$ (SVI)	UNet	**0.9616**	35.19	**0.9541**	**27.31**

TABLE 2 Pairwise comparison p-values for all coefficients using two-sided Wilcoxon signed-rank tests on classic and ΔB_0-weighted SSIM scores with Benjamini-Hochberg correction.
The pipeline SSIM and ΔB_0-weighted SSIM scores averaged in Table 1 are tested for statistical significance. A Friedman test was applied over each metric results and obtained $p < 10^{-7}$, before applying pairwise two-sided Wilcoxon signed-rank tests with a global Benjamini-Hochberg correction to adjust the false discovery rate.

(a) Classic SSIM

	$BASE_C$	$BASE_{\mathcal{R}}$	MFI	MTI	SVI	$SVI_{1/2}$
$BASE_C$		-	**	**	**	**
$BASE_{\mathcal{R}}$	-		**	**	**	**
MFI	**	**		-	**	**
MTI	**	**	-		**	**
SVI	**	**	**	**		-
$SVI_{1/2}$	**	**	**	**	-	

(b) ΔB_0-weighted SSIM

	$BASE_C$	$BASE_{\mathcal{R}}$	MFI	MTI	SVI	$SVI_{1/2}$
$BASE_C$		**	**	**	**	**
$BASE_{\mathcal{R}}$	**		**	**	**	**
MFI	**	**		*	**	**
MTI	**	**	*		**	**
SVI	**	**	**	**		**
$SVI_{1/2}$	**	**	**	**	**	

$p > 0.05$ -
$p < 0.05$ *
$p < 0.005$ **

TABLE 3 Ablation study of the buffer feature over different architectures on testing dataset.
The best performing pipelines (SVI and $SVI_{1/2}$) are evaluated with and without buffers, along with the baseline $BASE_C$ based on the non-correcting Fourier operator F_Ω. The classic and ΔB_0-weighted scores are averaged over the testing dataset, with statistical significance assessed through pairwise two-sided Wilcoxon signed-rank tests over the SSIM scores with a global Benjamini-Hochberg correction.

Pipeline	Buffers	Classic			ΔB_0-weighted		
		p	SSIM	PSNR	p	SSIM	PSNR
$BASE_C$	-	-	**0.9427**	**33.40**	**	**0.8979**	**23.69**
	3		0.9421	33.32		0.8947	23.50
SVI	-	*	0.9603	35.11	*	0.9490	**26.57**
	3		**0.9613**	**35.24**		**0.9498**	**26.57**
$SVI_{1/2}$	-	**	0.9575	34.71	**	0.9493	26.98
	3		**0.9616**	**35.19**		**0.9541**	**27.31**

$p > 0.05$ -
$p < 0.05$ *
$p < 0.005$ **

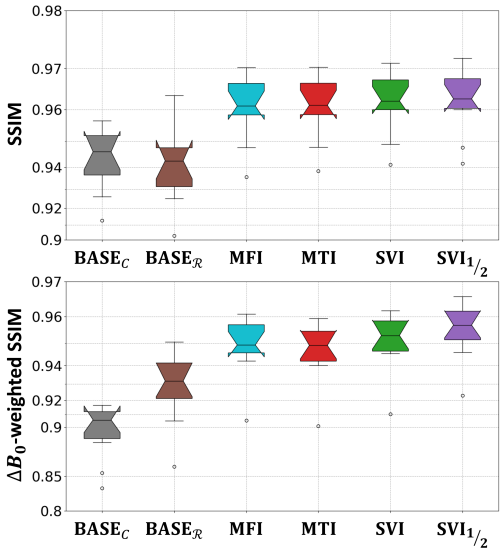

FIGURE 3 Detailed SSIM scores distributions for different correction coefficients and pipelines over testing dataset.
The SSIM and ΔB_0-weighted SSIM distributions averaged over testing dataset in Table 1 are detailed for all pipelines over reversed logarithmic scales. The first quartile Q_1, median and third quartile Q_3 are shown as notched boxes. Maximum and minimum scores are delimited by the whiskers, except for outliers defined as values farther from nearest quartile than $1.5 \times (Q_3 - Q_1)$ and shown as points.

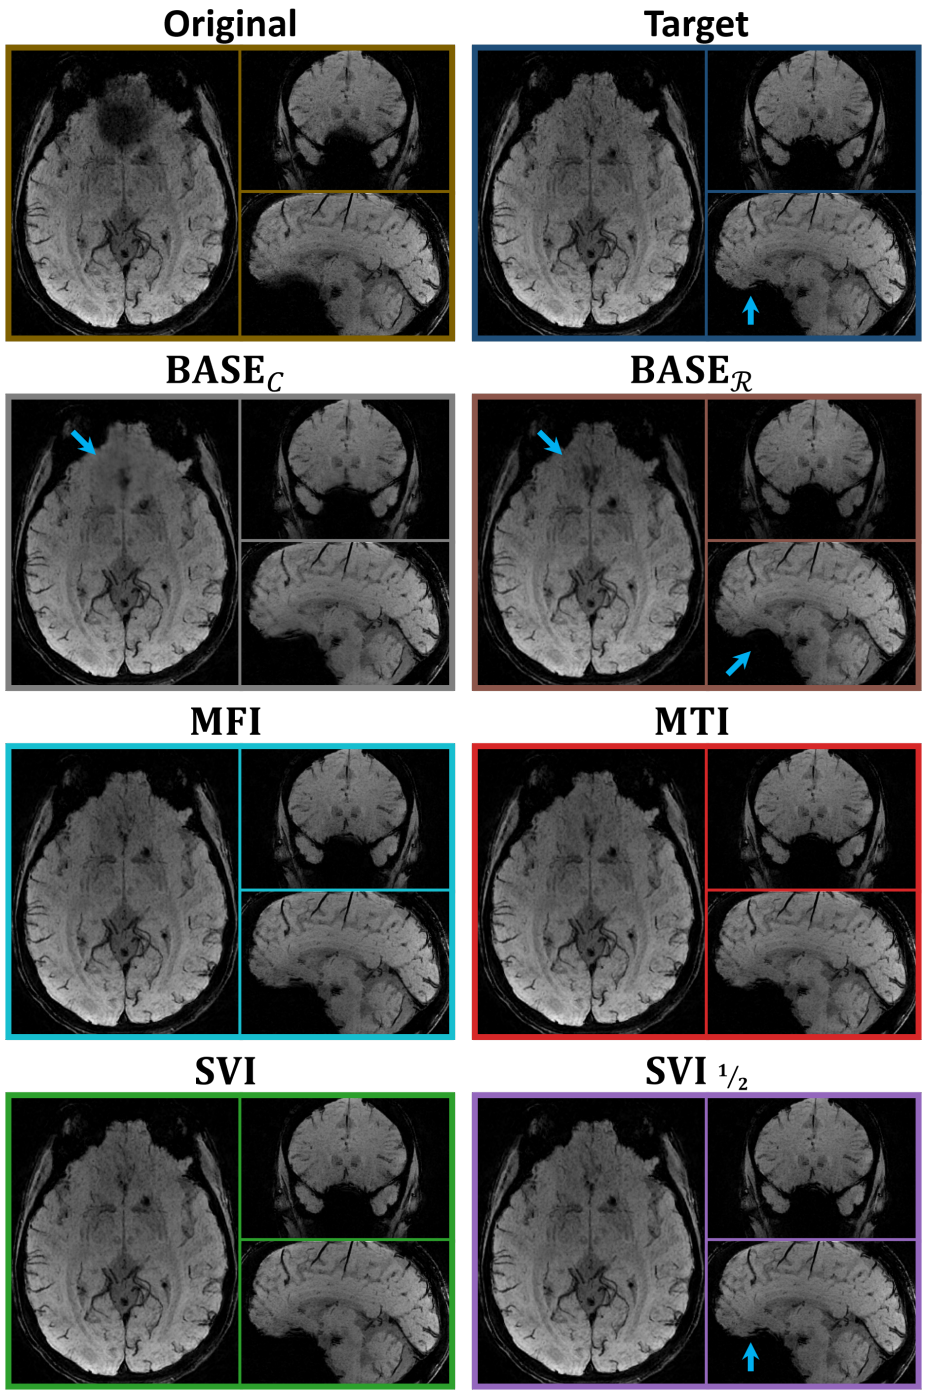

FIGURE 4 Reconstructed SWI images over all pipelines (first example).
The reconstructed SWI images with three 4 mm in-plane (axial, sagittal, coronal) minimum intensity projections are provided for all color-coded pipelines (following the convention adopted in Figure 3 in regard to the color coding) for acquisition #68 from testing dataset. Various details are pointed out with blue arrows.

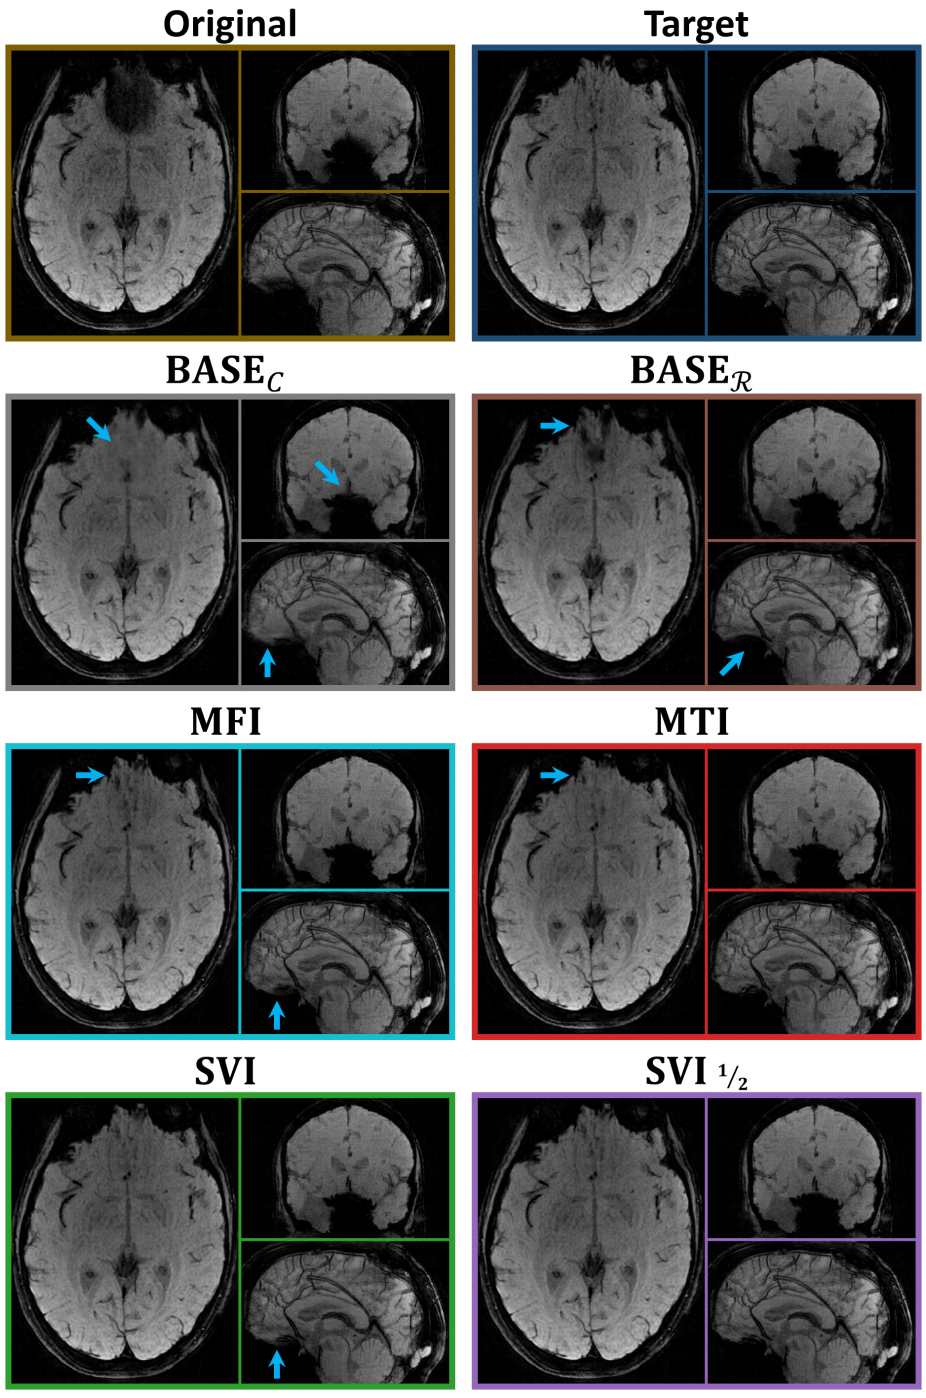

FIGURE 5 **Reconstructed SWI images over all pipelines (second example).**
The reconstructed SWI images with three 4 mm in-plane (axial, sagittal, coronal) minimum intensity projections are provided for all color-coded pipelines (following the convention adopted in Figure 3 in regard to the color coding) for acquisition #106 from testing dataset. Various details are pointed out with blue arrows.

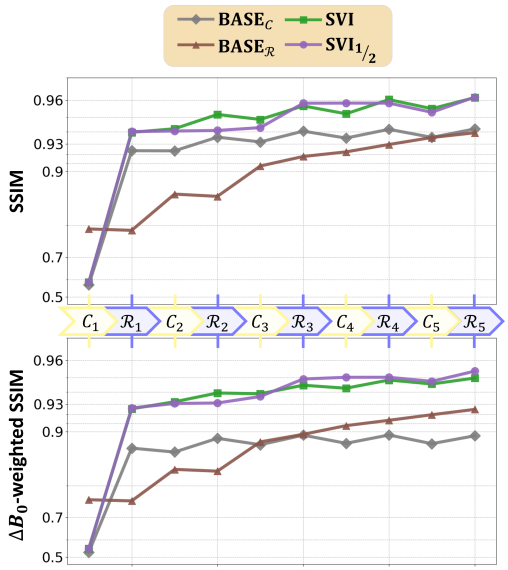

FIGURE 6 Testing scores evolution over reconstruction.
The main pipelines are decomposed between data consistency (C_i) and regularization ($\mathcal{R}_i$) steps to provide the classic and ΔB_0-weighted SSIM scores over the testing dataset over reversed logarithmic scales. The MFI and MTI pipelines are omitted for readability as both are redundant with SVI pipeline.

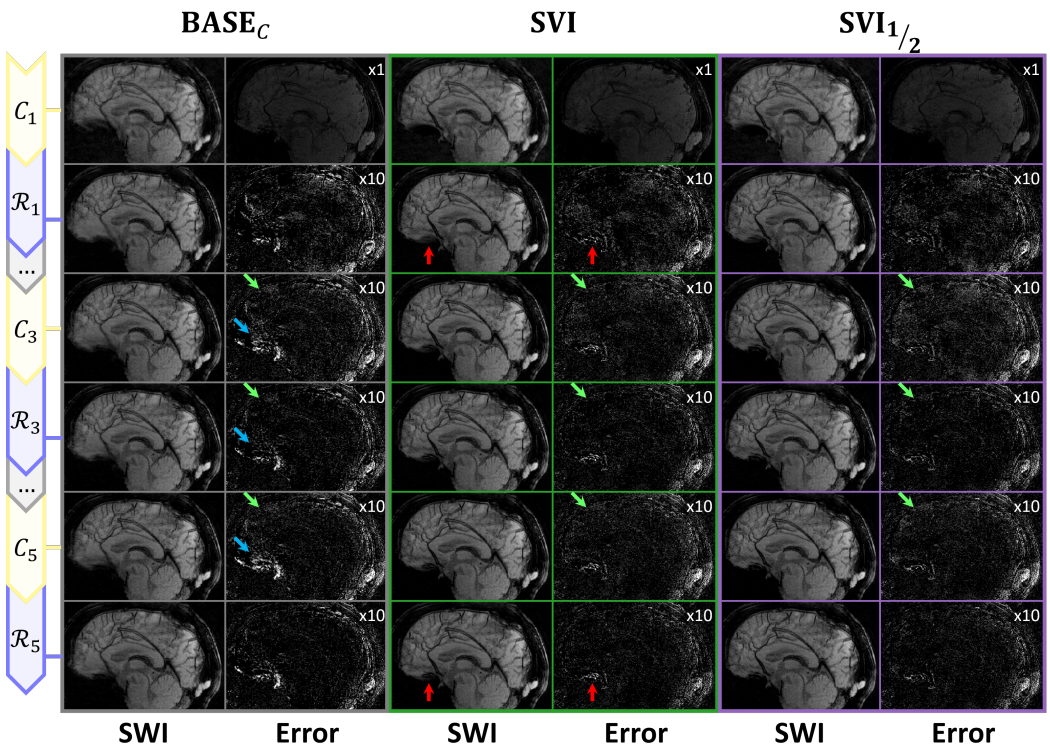

FIGURE 7 SWI images evolution over reconstruction.
The intermediate SWI images with 4 mm minimum intensity projection are provided over reconstruction decomposed into data consistency (C_i) and regularization ($\mathcal{R}_i$) steps for acquisition #106 from testing dataset. Various details are pointed out with arrows: coil sensitivity degraded by data consistency (green), off-resonance degraded by data consistency (blue) and off-resonance degraded by regularization (red).