Received: 14 October 2020 | Revised: 8 January 2021 | Accepted: 11 January 2021

DOI: 10.1002/mrm.28706

FULL PAPER

Magnetic Resonance in Medicine

Region-optimized virtual (ROVir) coils: Localization and/or suppression of spatial regions using sensor-domain beamforming

Daeun Kim[1,2] | **Stephen F. Cauley**[3] | **Krishna S. Nayak**[1,2] | **Richard M. Leahy**[1,2] | **Justin P. Haldar**[1,2]

[1]Signal and Image Processing Institute, University of Southern California, Los Angeles, CA, USA

[2]Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA

[3]Deparment of Radiology, A. A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA, USA

Correspondence
Daeun Kim, University of Southern California, University Park Campus, 3710 McClintock Avenue, Ronald Tutor Hall (RTH) #317, Los Angeles, CA 90007, USA.
Email: daeunk@usc.edu

Funding information
This work was supported in part by NIH grants R01-MH116173, R01-NS074980, R01-NS089212, and R33-CA225400.

Purpose: In many MRI scenarios, magnetization is often excited from spatial regions that are not of immediate interest. Excitation of uninteresting magnetization can complicate the design of efficient imaging methods, leading to either artifacts or acquisitions that are longer than necessary. While there are many hardware- and sequence-based approaches for suppressing unwanted magnetization, this paper approaches this longstanding problem from a different and complementary angle, using beamforming to suppress signals from unwanted regions without modifying the acquisition hardware or pulse sequence. Unlike existing beamforming approaches, we use a spatially invariant sensor-domain approach that can be applied directly to raw data to facilitate image reconstruction.

Theory and Methods: We use beamforming to linearly mix a set of original coils into a set of region-optimized virtual (ROVir) coils. ROVir coils optimize a signal-to-interference ratio metric, are easily calculated using simple generalized eigenvalue decomposition methods, and provide coil compression.

Results: ROVir coils were compared against existing coil-compression methods, and were demonstrated to have substantially better signal suppression capabilities. In addition, examples were provided in a variety of different application contexts (including brain MRI, vocal tract MRI, and cardiac MRI; accelerated Cartesian and non-Cartesian imaging; and outer volume suppression) that demonstrate the strong potential of this kind of approach.

Conclusion: The beamforming-based ROVir framework is simple to implement, has promising capabilities to suppress unwanted MRI signal, and is compatible with and complementary to existing signal suppression methods. We believe that this general approach could prove useful across a wide range of different MRI applications.

KEYWORDS

accelerated acquisition, beamforming, multichannel MRI, reduced field-of-view imaging, signal suppression

1 | INTRODUCTION

MRI experiments are often designed to provide information about a specific spatial region of interest (ROI) inside the body. However, instead of only acquiring data and reconstructing images corresponding to the ROI, typical MRI acquisitions often reconstruct much larger spatial regions, including regions that can be uninteresting for the target application. Illustrations of this are shown in Figure 1.

These uninteresting spatial regions are often encoded and reconstructed purely out of the desire to avoid artifacts. A conventional MR receiver coil does not discriminate between interesting and uninteresting spatial regions, and will simply detect signal from all of the excited magnetization within the sensitive region of the coil. Uninteresting magnetization that has been excited cannot simply be ignored. If it is not adequately spatially encoded and reconstructed, signal from uninteresting regions can leak or alias into other spatial regions, potentially contaminating the reconstruction of the desired ROI.

Encoding and reconstructing uninteresting image regions can be inefficient and burdensome for data acquisition. For example, standard MRI sampling theory[1] dictates that reconstructing a larger field-of-view (FOV) necessitates increased k-space sampling density requirements. Conversely, it is also well-understood that spatial encoding requirements and image reconstruction complexity can be substantially reduced if the received signal only originates from a small spatially localized ROI.[2-4]

Because of these issues, substantial effort has been spent within the MR community to avoid signal from uninteresting spatial regions in order to simplify other aspects of the experiment. Examples include the development of methods that only produce a consistent homogeneous B_0 main magnetic field at a single spatial point within the object[5,6]; the use of localized surface coils that are designed to focus on a specific spatial ROI[7,8]; the development of spatially-selective radiofrequency excitation methods that only excite magnetization from a specific ROI, allowing for "zoomed," "inner volume," "single voxel," "line scan," "slice selective," or "reduced field-of-view" imaging[9-17]; and outer volume suppression methods designed to saturate the signal from uninteresting spatial regions.[9,18-21]

In this work, we introduce a novel approach to suppressing the signal from uninteresting spatial regions that is distinct from and complementary to previous methods. Rather than modifying the receiver hardware or the acquisition pulse sequence as in previous approaches, we instead identify simple linear combinations of the sensor-domain data obtained from an array of receiver coils that optimally suppress the unwanted signal. These linear combinations are designed to produce region-optimized virtual (ROVir) coils that simultaneously maximize the signal energy observed from interesting spatial regions and minimize the signal energy observed from undesired spatial regions. Importantly, these operations can be applied in the sensor domain (ie, to the raw multichannel k-space data prior to image reconstruction), which enables reduced k-space data sampling requirements and simplified image reconstruction. Moreover, our approach for generating ROVir coils is relatively simple, requiring only the specification of interesting and uninteresting spatial regions and the computation of a small generalized eigenvalue decomposition. The ROVir approach is also optimal with respect to an intuitive signal-to-interference ratio (SIR) metric. We will delay presenting a formal definition of SIR until the Theory section (after establishing the necessary mathematical notation), but as the name implies, it represents the ratio between signal energy from the spatial region of interest and the interference energy from unwanted signal regions.

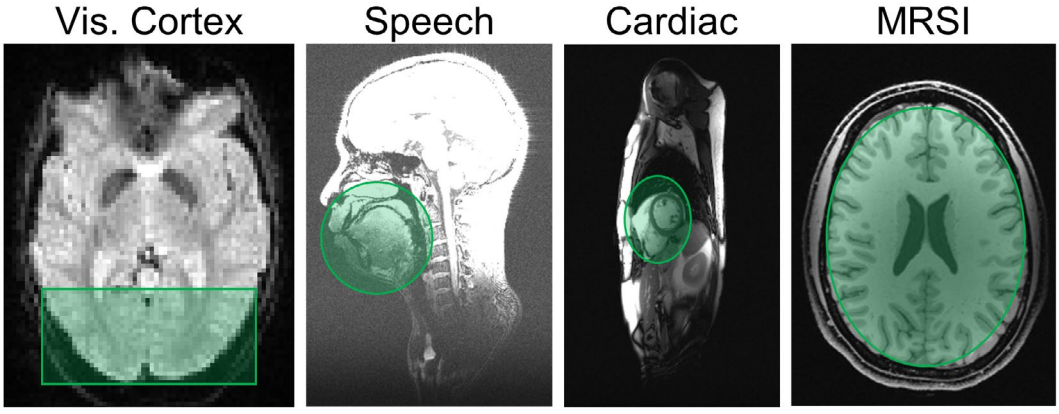

FIGURE 1 Illustration of applications where the ROI (marked in green) can be much smaller than the full FOV. In functional MRI of the visual cortex, we may only care about the signal from the visual cortex itself, but not the signal from other parts of the brain. In speech imaging, we may only care about the signal from the vocal tract, but not the signal from the brain, scalp, neck, or torso. In cardiac imaging, we may only care about the MRI signal from the heart, but not the signal from other parts of the body. In MR spectroscopic imaging (MRSI), we may only care about the signal from brain parenchyma, but not the signal from extracranial lipids

The concept of virtual coils was introduced many years ago for multichannel MRI. The basic principle is that, instead of working directly with the data from the original coils, it can be beneficial to synthesize virtual coils as simple linear combinations of the data from the original coils. This idea has been used successfully for a number of different purposes, including to eliminate inter-channel noise correlation,[22] to compress a large number of original real coils into a smaller number of virtual coils without loss of substantial information content,[23-27] and to simplify the use of k-space conjugate symmetry relationships in image reconstruction.[28,29] Although one of the previous approaches[23] tailors the design of virtual coils with respect to an ROI for coil-compression purposes, none of them attempt to suppress the signal from uninteresting spatial regions in the same way as ROVir.

The ROVir approach is based on classical spatial beamforming principles from sensor array signal processing,[30,31] which were originally introduced more than a half-century ago to maximize signals of interest and minimize interfering signals, with applications in radar, sonar, seismology, astronomy, ultrasonics, electroencephalography, magnetoencephalography, and other scenarios where data is acquired with a sensor array. Although MRI receiver arrays fit naturally within the broader class of sensor array acquisitions, to the best of our knowledge, these kinds of spatial beamforming principles have not been previously applied in the sensor domain (as in ROVir). However, it should be noted that spatial beamforming principles form the foundation of several popular existing methods for the voxelwise (spatially varying) image-domain combination of previously reconstructed multichannel images,[32-38] which in some cases where designed to suppress artifacts from unwanted spatial regions.[34-36,38] Note that, unlike ROVir's spatially invariant sensor-domain beamforming, spatially varying image-domain beamforming is generally an image postprocessing technique that cannot be applied directly to raw multichannel k-space data to facilitate the image reconstruction process.

Coil selection methods are an alternative to the use of virtual coils, and proceed by directly discarding channels from the original array that have too little desired signal or too much unwanted signal.[39,40] Coil selection is a sensor-domain approach similar to ROVir, although ROVir's use of optimal virtual coils/beamforming endows it with greater flexibility and better SIR than could generally be achieved by coil selection.

We believe that the ROVir concept may be useful across a wide range of different application domains, and later parts of this paper will show illustrations pertaining to several different scenarios. However, as this is intended to be a largely conceptual/theoretical paper, these illustrations are only meant to demonstrate the broad potential of ROVir, and are not intended to represent thorough evaluations of each case. We anticipate that each individual scenario will deserve deeper followup investigations to fully establish the performance differences and potential synergies between ROVir and other spatial localization/suppression techniques. Importantly, we do not envision ROVir as a full replacement for other methods, but rather as an additional useful tool that can be used in combination with other techniques to better achieve experimental imaging objectives.

During the preparation of this work, we became aware of a preliminary method reported by Cauley et al at a recent conference[41] that, like ROVir, aims to generate virtual coils that suppress information from undesired spatial regions. Although this alternative approach can also be effective, the virtual coils produced by this method will not achieve the SIR-optimality of ROVir and the optimization formulation is quite different. We believe that this alternative preliminary method is not competitive with ROVir, and the remainder of this paper presents ROVir without further discussion of the other approach.

2 | THEORY

2.1 | Beamforming and virtual coils

For simplicity and without loss of generality, we will describe ROVir in the context of the standard sensitivity-encoded (SENSE) Fourier imaging model,[42] although the same principles also generalize naturally to non-Fourier models (including field inhomogeneity, relaxation, nonlinear gradient fields, etc.) of multichannel data acquisition. Under this model, the k-space data measured from the ℓth coil at k-space location $\mathbf{k}$ is given by

$$d_\ell(\mathbf{k}) = \int s_\ell(\mathbf{x}) f(\mathbf{x}) e^{-i2\pi \mathbf{k}\cdot\mathbf{x}} d\mathbf{x} + n_\ell(\mathbf{k}) \qquad (1)$$

for $\ell = 1, \ldots, N_c$, where $f(\mathbf{x})$ represents the complex-valued MR image as a function of the spatial position $\mathbf{x}$, $s_\ell(\mathbf{x})$ is the complex-valued spatial sensitivity profile of the ℓth receiver coil, N_c is the total number of coils in the receiver array, and $n_\ell(\mathbf{k})$ represents the thermal measurement noise associated with the ℓth receiver coil. We will further assume that the data has already been whitened,[22] such that the intercoil noise covariance matrix Ψ is simply the $N_c \times N_c$ identity matrix and all noise samples can be modeled as independent and identically distributed zero-mean complex Gaussian random variables.

For this paper, the main concern with Equation (1) is that the measured data may contain information from the sensitivity weighted image arising from spatial locations $\mathbf{x}$ that are outside the ROI. This is problematic because the signal from these spatial locations may interfere with data acquisition and image reconstruction for the spatial regions we actually care about. As a result, the measured data samples $d_\ell(\mathbf{k})$ from

each coil can be represented as a mixture of useful signal from the ROI together with unwanted signal from other spatial locations ("interference").

The main idea of beamforming is that, instead of working directly with the data from the original coils, we can instead work with data from virtual coils obtained from linear combinations of the original coils, where the linear combination weights are designed in an optimal way to maximize signal and minimize interference.[30,31] Consider the scenario in which the k-space data $d_\ell(\mathbf{k})$ from a set of N_c original coils are linearly transformed into data from a set of N_v virtual coils according to

$$v_j(\mathbf{k}) = \sum_{\ell=1}^{N_c} w_{\ell j} d_\ell(\mathbf{k}), \tag{2}$$

for $j = 1, \ldots, N_v$, where $v_j(\mathbf{k})$ is the k-space data for the jth virtual coil and the complex-valued scalars $\{w_{\ell j}\}$ are the linear combination weights. By inserting Equation (2) into Equation (1) and rearranging terms, it is straightforward to show that these new virtual coils obey a sensitivity-weighted Fourier encoding model that is identical to Equation (1), except with different sensitivity profiles and noise values. In particular,

$$v_j(\mathbf{k}) = \int \tilde{s}_j(\mathbf{x}) f(\mathbf{x}) e^{-i2\pi \mathbf{k}\cdot\mathbf{x}} d\mathbf{x} + \tilde{n}_j(\mathbf{k}), \tag{3}$$

with

$$\tilde{s}_j(\mathbf{x}) = \sum_{\ell=1}^{N_c} w_{\ell j} s_\ell(\mathbf{x}) \tag{4}$$

and

$$\tilde{n}_j(\mathbf{k}) = \sum_{\ell=1}^{N_c} w_{\ell j} n_\ell(\mathbf{k}) \tag{5}$$

for each virtual coil j. Importantly, an appropriate choice of the weighting coefficients $\{w_{\ell j}\}$ allows the sensitivity profiles $\tilde{s}_j(\mathbf{x})$ of the virtual array to be steered towards spatial locations within the ROI and away from unwanted spatial locations. An illustration of what can be achieved with this kind of steering is presented in Supporting Information Video S1, which shows ROVir results obtained from sweeping a circular-shaped ROI (ie, everything interior to the green circle) through the FOV for a 32-channel brain MRI dataset. As can be seen, the ROVir procedure effectively puts a spotlight on the ROI (marked in green), while suppressing unwanted signal from outside the ROI (ie, everything in the region marked in red).

It should be noted that while the sensitivity-based analysis above provides insight into the theoretical principles of ROVir, the ROVir framework described below does not actually require knowledge of the sensitivity profiles $s_\ell(\mathbf{x})$ or $\tilde{s}_j(\mathbf{x})$.

2.2 | Formulation of ROVir

For ROVir, we choose the $\{w_{\ell j}\}$ coefficients to maximize an SIR metric, which is defined as the ratio between the signal energy from the spatial region of interest and the interference energy from unwanted spatial regions. In particular, we use Ω to denote the spatial ROI. We can then define the total k-space signal energy in the jth virtual coil arising from the ROI using Parseval's theorem as

$$\begin{aligned} S(\mathbf{w}_j) &\triangleq \int_{\mathbf{x}\in\Omega} \left|\tilde{s}_j(\mathbf{x}) f(\mathbf{x})\right|^2 d\mathbf{x} \\ &= \int_{\mathbf{x}\in\Omega} \left|\sum_{\ell=1}^{N_c} w_{\ell j} g_\ell(\mathbf{x})\right|^2 d\mathbf{x} \\ &= \mathbf{w}_j^H \mathbf{A} \mathbf{w}_j, \end{aligned} \tag{6}$$

where $\mathbf{w}_j \in \mathbb{C}^{N_c}$ is the vector of weighting coefficients $\{w_{\ell j}\}$ for $\ell = 1, \ldots, N_c$, $g_\ell(\mathbf{x}) \triangleq s_\ell(\mathbf{x}) f(\mathbf{x})$ is the sensitivity weighted image for the ℓth coil, and the matrix $\mathbf{A} \in \mathbb{C}^{N_c \times N_c}$ is defined as

$$\mathbf{A} = \int_\Omega \mathbf{g}(\mathbf{x})^H \mathbf{g}(\mathbf{x}) d\mathbf{x}, \tag{7}$$

with

$$\mathbf{g}(\mathbf{x}) \in \mathbb{C}^{1\times N_c} \triangleq \left[g_1(\mathbf{x}) \cdots g_{N_c}(\mathbf{x}) \right]. \tag{8}$$

The $\mathbf{A}$ matrix can be interpreted as the inter-coil signal correlation matrix corresponding to the spatial region Ω.

Similarly, using Γ to denote the spatial region containing unwanted information, we define the interference energy in the jth virtual coil arising from these unwanted spatial regions as

$$\begin{aligned} I(\mathbf{w}_j) &\triangleq \int_{\mathbf{x}\in\Gamma} \left|\tilde{s}_j(\mathbf{x}) f(\mathbf{x})\right|^2 d\mathbf{x} \\ &= \mathbf{w}_j^H \mathbf{B} \mathbf{w}_j, \end{aligned} \tag{9}$$

where the matrix $\mathbf{B} \in \mathbb{C}^{N_c \times N_c}$ is defined as

$$\mathbf{B} = \int_\Gamma \mathbf{g}(\mathbf{x})^H \mathbf{g}(\mathbf{x}) d\mathbf{x}. \tag{10}$$

As above, the $\mathbf{B}$ matrix can be interpreted as the inter-coil signal correlation matrix corresponding to the spatial region Γ.

Having defined both signal and interference, we can then define the SIR for the jth virtual coil as the ratio between the two

$$SIR(\mathbf{w}_j) \triangleq \frac{\mathbf{w}_j^H \mathbf{A} \mathbf{w}_j}{\mathbf{w}_j^H \mathbf{B} \mathbf{w}_j}. \tag{11}$$

ROVir coils are obtained by finding weight vectors $\mathbf{w}_j$ that maximize this SIR criterion, thereby ensuring that the signal from the ROI is maximized while the interference is minimized.

Before describing the optimization of Equation (11), it is worth noting that an interesting relationship between this SIR criterion and optimality criteria have been previously used to construct virtual coils for the purpose of coil-compression.[23-27] In particular, many of the previous coil-compression methods have used singular value decomposition (SVD)/principal component analysis (PCA) to choose weight vectors that maximize the total amount of observed signal energy from the receiver array, without worrying about minimizing interference energy. In particular, it can be shown that the ROI-based version[23] of optimal SVD/PCA coil-compression technique is equivalent to finding weight vectors that maximize

$$SNR(\mathbf{w}_j) \triangleq \frac{\mathbf{w}_j^H \mathbf{A} \mathbf{w}_j}{\mathbf{w}_j^H \mathbf{w}_j}. \quad (12)$$

As can be seen, the ROVir approach reduces to the standard SVD/PCA coil-compression approach when the $\mathbf{B}$ matrix is chosen as the identity matrix. Note also that if the interference region Γ contains only whitened Gaussian noise, then the $\mathbf{B}$ matrix will be an empirical coil covariance matrix that is well-approximated by the identity matrix, and ROVir will behave similarly to the SVD/PCA approach.

2.3 | Solving for optimal ROVir weights

Simple methods for choosing coil-combination weights that maximize the SIR from Equation (11) are known from the beamforming literature.[30,31] In particular, Equation (11) has the form of a generalized Rayleigh quotient, and optimal weight vectors for this kind of optimization problem can be obtained using a simple generalized eigendecomposition.

Assume that $\mathbf{A}$ and $\mathbf{B}$ are both positive-semidefinite Hermitian-symmetric matrices, and that $\mathbf{B}$ has full rank. Then it can be shown[43] that there exists a set of N_c real-valued positive generalized eigenvalues λ_i and a corresponding linearly independent set of generalized eigenvectors $\mathbf{w}_i$, $i = 1, \ldots, N_c$, that satisfy

$$\mathbf{A}\mathbf{w}_i = \lambda_i \mathbf{B}\mathbf{w}_i. \quad (13)$$

These generalized eigenvalues and eigenvectors are easily computed using standard eigenvalue solvers,[44] and a generalized eigenvector for the pair of matrices $\mathbf{A}$ and $\mathbf{B}$ must also be a standard eigenvector for the single matrix $\mathbf{B}^{-1}\mathbf{A}$. Notably, these generalized eigenvectors are not usually orthogonal (as was the case when applying SVD/PCA), which can have the consequence of introducing noise correlation between the resulting virtual channels as described below. Instead, the generalized eigenvectors are "$\mathbf{B}$-orthogonal" in the sense that $\mathbf{w}_j^H \mathbf{B} \mathbf{w}_i = 0$ whenever $i \neq j$.

Furthermore, if the eigenvalues are ordered such that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{N_c} > 0$, then it is known that the first eigenvector $\mathbf{w}_1$ maximizes the SIR criterion from Equation (11) while the bottom eigenvector $\mathbf{w}_{N_c}$ minimizes the SIR.[30,31] Further, it can be shown that projecting the data onto the N_v-dimensional subspace spanned by the top-N_v generalized eigenvectors $\mathbf{w}_1, \ldots, \mathbf{w}_{N_v}$ is optimal in the sense that it maximizes the ratio between the retained signal energy and the retained interference energy among all possible N_v-dimensional subspace projections.[45]

This leads to a simple recipe for choosing the $\{ w_{\ell j} \}$ coefficients needed to form optimal ROVir coils:

1. Form matrices $\mathbf{A}$ and $\mathbf{B}$ corresponding to Equations (7) and (10).
2. Compute the generalized eigenvalue decomposition for matrices $\mathbf{A}$ and $\mathbf{B}$, and order the generalized eigenvalues and eigenvectors such that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{N_c} > 0$. The eigenvectors associated with distinct eigenvalues are unique up to scaling ambiguities, and it is convenient to work with unit-normalized versions of the eigenvectors such that $\| \mathbf{w}_j \|_2 = 1$ for each $j = 1, \ldots, N_c$.
3. (Optional) The weight vectors $\mathbf{w}_i$ and $\mathbf{w}_j$ obtained from the generalized eigenvalue decomposition are unlikely to be orthogonal unless $\mathbf{B}$ happens to be a diagonal matrix. This means that there may be nonzero noise correlation between the different virtual coils, since the noise correlation between virtual channel i and virtual channel j will be equal to $\mathbf{w}_j^H \Psi \mathbf{w}_i = \mathbf{w}_j^H \mathbf{w}_i$. If noise-whitened virtual channels are desired, such whitening can be achieved by applying a simple Gram-Schmidt orthonormalization procedure[44] to the top-N_v generalized eigenvectors. Note that this orthonormalization procedure does not change the subspace spanned by the top-N_v generalized eigenvectors, and therefore does not change the optimality characteristics of the remaining steps. If this orthogonalization step is not utilized, then achieving optimal image SNR will generally require this new source of noise correlation to be properly accounted for if coil combination and/or multichannel image reconstruction steps are subsequently applied to ROVir data.[22,32,46]
4. Use the top-N_v generalized eigenvectors $\mathbf{w}_j$ for $j = 1, \ldots, N_v$ (or their orthonormalized versions), to define the linear combination weights according to $w_{\ell j} = [\mathbf{w}_j]_\ell$.
5. Form virtual coil data following Equation (2).

A representative implementation of the basic ROVir framework is illustrated in Figure 2.

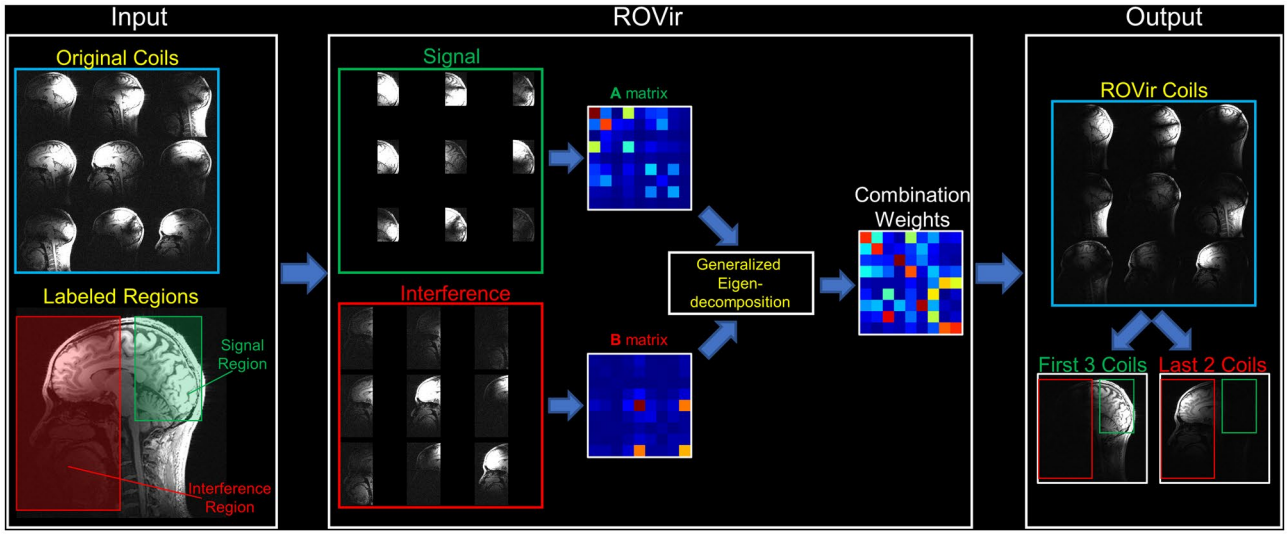

FIGURE 2 Representative illustration of ROVir. Starting from an original set of coils and labeled signal and interference regions, the ROVir approach forms region-specific intercoil correlation matrices **A** and **B**, which are used to compute coil combination weights $w_{\ell j}$ that optimize the signal-to-interference ratio through a generalized eigendecomposition procedure. The ROVir coils that result from using these combination weights have excellent separation between the signal and interference regions, and are also highly effective at compressing the high-dimensional coil array into a much smaller number of relevant coils

All of the ROVir steps are straightforward, except perhaps for forming the **A** and **B** matrices in the first step. The potential challenge in forming these matrices is that (1) it appears that prior knowledge of each of the individual sensitivity-weighted coil images $g_\ell(\mathbf{x})$ must be available and (2) it is necessary to integrate continuous-valued images. However, both of these potential challenges are easily overcome. The first issue can be solved by forming images from prescan/calibration data, and computing the optimal coil combination weights with respect to those images. The acquisition of prescan/calibration data is a standard part of most MRI protocols and should not be burdensome, and the ROVir coils obtained from such data should still have excellent SIR characteristics for any subsequent datasets as long as the relative spatial distribution of image energy is not substantially changed. The second potential challenge is also easily solved, noting that the integrals in Equations (7) and (10) are easily approximated with discrete sums over reconstructed voxels that fall within Ω (for Equation 7) or Γ (for Equation 10).

Although the ROVir procedure is robust in all of the practical situations that we have tried, we should note that the generalized eigendecomposition procedure can be unstable if the matrix **B** is either rank-deficient or approximately rank deficient. In principle, this kind of problem could happen if the spatial region Γ is too small (ie, fewer than N_c voxels were used in the construction of the matrix **B**) or if the sensitivity maps do not have significant spatial variations within the region Γ. A standard mitigation approach is to use simple regularization of **B** to ensure it has full numerical rank, that is, replacing **B** with $\mathbf{B}+\xi\mathbf{I}$, where **I** is the $N_c \times N_c$ identity matrix and ξ is a small positive constant.

2.4 | Coil compression and automatic choice of N_v

In order to suppress the unwanted signal from the spatial region Γ, it is necessary to choose $N_v < N_c$, which has the practical effect of discarding those virtual coils with the lowest SIR while retaining the virtual coils with the highest SIR. As a result, ROVir should not only be viewed as a signal-suppression method, but it can also be viewed as a coil compression method. Coil compression methods map a high-dimensional coil array into a much lower dimensional coil array, resulting in substantial reductions in memory requirements and computational complexity for downstream image reconstruction.

Most coil-compression approaches attempt to maximize the signal energy that is present in the top-N_v virtual coils without considering interference, which is distinct from ROVir's approach of maximizing the SIR of the top-N_v virtual coils. As a result, ROVir is theoretically guaranteed to be no better (and likely slightly worse) at preserving signal energy than traditional coil-compression methods, though is also theoretically guaranteed to have no worse (and likely much better) interference suppression characteristics.

One of the nice features of the ROVir approach is that the number of virtual channels N_v does not need to be a user-selected parameter, and instead can be chosen automatically based on the quantitative signal and interference characteristics of each channel. In particular, given a weight vector $\mathbf{w}_j$, it is straightforward to evaluate the signal and interference components for the corresponding ROVir coil using Equations (7) and (10), and it is thus also straightforward to simply choose

N_v by only retaining the virtual coils that have sufficiently high signal and sufficiently low interference components. The examples shown later in this paper use different methods for automatically selecting N_v based on several different application-dependent considerations (eg, depending on whether it is more important to maximize signal or to minimize interference in the given application).

3 | METHODS

To evaluate the characteristics of ROVir, we applied the approach to several different datasets representing several different imaging applications, including reduced-FOV brain imaging, reduced-FOV speech imaging, reduced-FOV cardiac imaging, and outer-volume suppression in the brain. All in vivo datasets were acquired under IRB-approved written informed consent. Because the scenarios and datasets we use are all different and largely independent of each other, to improve the readability and flow of our article, we will describe the details of each scenario sequentially within the Results section itself.

In all cases, the performance of ROVir was evaluated both qualitatively and quantitatively. Qualitatively, we show visualizations of the images obtained using ROVir. When displaying coil-combined images, we properly account for noise correlation by utilizing the previously described Gram-Schmidt orthogonalization procedure. Quantitatively, we used Equation (11) to measure the SIR of each virtual coil, and also used Equations (7) and (10) to separately quantify the signal energy and interference energy present in each virtual coil. We also quantified the percentage of retained ROI signal and the percentage of retained interference obtained after projecting the original data into the subspace corresponding to the top-N_v ROVir coils. Specifically, assuming that the collection of N_v weight vectors $\mathbf{w}_j$ has been orthonormalized and that a matrix $\mathbf{W} \in \mathbb{C}^{N_c \times N_v}$ is formed from the eigenvectors such that the jth column of $\mathbf{W}$ is equal to $\mathbf{w}_j$, then $\mathbf{W}\mathbf{W}^H \in \mathbb{C}^{N_c \times N_c}$ is the unique orthogonal projection matrix for the span of the top-N_v virtual coil combination weights.[47] As a result, the percentage of retained ROI signal can be computed as

$$\frac{\|\mathbf{W}\mathbf{W}^H\mathbf{A}\mathbf{W}\mathbf{W}^H\|_F}{\|\mathbf{A}\|_F} \times 100\% \tag{14}$$

and the percentage of retained interference can be computed as

$$\frac{\|\mathbf{W}\mathbf{W}^H\mathbf{B}\mathbf{W}\mathbf{W}^H\|_F}{\|\mathbf{B}\|_F} \times 100\%. \tag{15}$$

We have used $\|\cdot\|_F$ to denote the Frobenius norm in these expressions, which is the standard/natural choice for measuring matrix energy.

Since ROVir can also be viewed as a coil-compression technique in addition to an SIR-maximization technique, we also performed a quantitative comparison between ROVir and 2 popular SVD-based coil-compression techniques.[23,24] The method we label as "ROI-SVD" is based on maximizing signal energy in the ROI using Equation (12), using the same ROI Ω that is used for ROVir. The method we label simply as "SVD" is also based on maximizing signal using Equation (12), but focuses on maximizing signal energy across the entire FOV without focusing narrowly on the ROI. This can be achieved by modifying the ROI Ω so that it equals the entire FOV.

4 | RESULTS

4.1 | Quantitative evaluations and performance comparisons

In this subsection, we perform a detailed/systematic quantitative evaluation of ROVir with comparisons against the SVD and ROI-SVD coil-compression approaches for a 2D sagittal brain imaging scenario, with real in vivo T1-weighted data acquired using a 32-channel array on a 3 T scanner. We consider a situation where we are interested in the signal from the brain, but not in the signal originating from other parts of the head, neck, or torso. As shown in the left side of Figure 3, the original coils have substantial sensitivity beyond just the brain.

Suppressing the unwanted signal from the head, neck, and torso could be beneficial in certain application scenarios. Not only could this enable the use of a smaller FOV without incurring aliasing artifacts, but it could also be useful for reducing gradient nonlinearity artifacts that could arise if one of the uninteresting parts of the image (eg, the torso) is far away from isocenter. To accomplish this goal, we defined a signal region Ω and an interference region Γ as also shown in the left side of Figure 3. While we could have designed the Ω and Γ regions to be immediately adjacent to one another, we instead decided to place a small gap between them. This gap serves the same purpose as a "transition band" in digital filter design and radiofrequency pulse design,[48,49] and allows for better signal retention in the "passband" (ie, Ω) and better interference suppression in the "stopband" (ie, Γ). In particular, since Maxwell's equations generally result in coil sensitivity maps that are spatially smooth, it is difficult to sharply switch between retaining one region and suppressing an immediately adjacent region, and overall performance can be substantially better if the transition is allowed to be more gradual.

Figure 4 shows the original set of 32 coils, as well as the full sets of 32 virtual coils obtained using SVD, ROI-SVD, and ROVir. As should be expected from previous literature, the SVD approach effectively compacts as much energy as

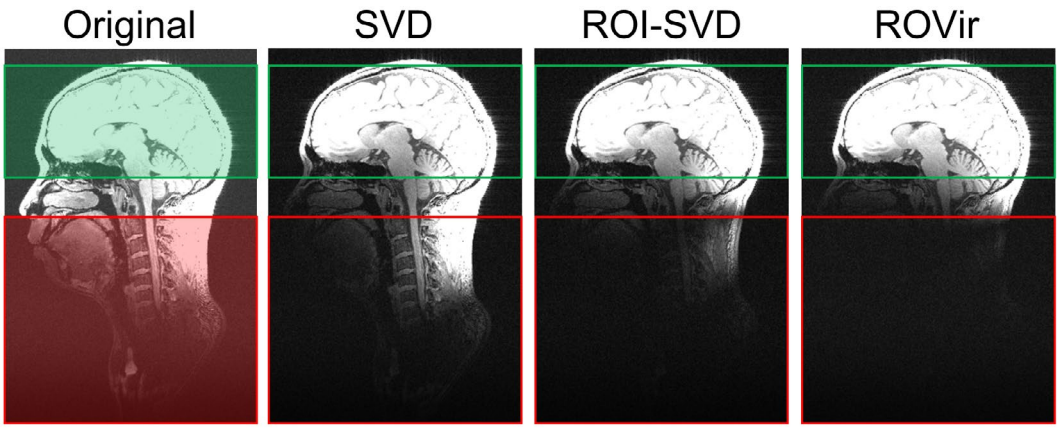

FIGURE 3 The leftmost image shows a sagittal T1-weighted brain image obtained by applying root-sum-of-squares coil combination to an original 32-channel dataset. The ROI signal that was used in this case (corresponding to the brain) is marked in green, while the "interference" region is marked in red. The remaining 3 images show the results of combining a much smaller number of virtual coils obtained by SVD, ROI-SVD, and ROVir, respectively. The number of virtual channels N_v was determined automatically, based on the desire to retain at least 95% of the signal from the ROI

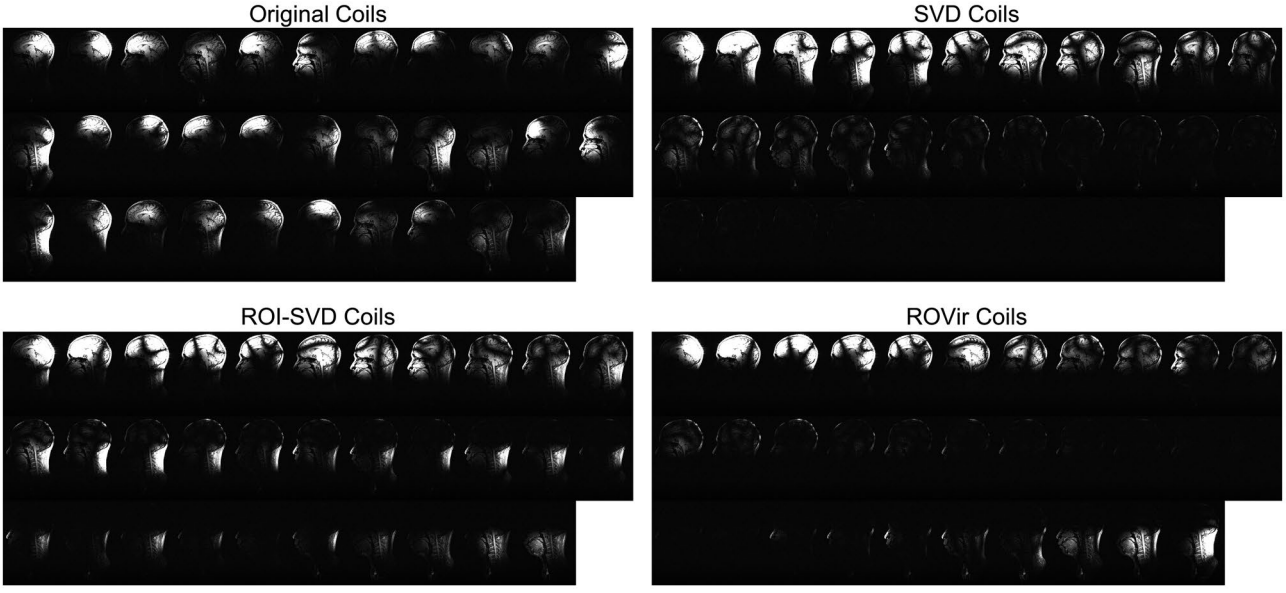

FIGURE 4 The individual coil images corresponding to the 32-channel brain imaging scenario from Figure 3. We show the original coils as well as the virtual coils obtained using SVD, ROI-SVD, and ROVir. For each set of virtual coils, the coils are shown in order from left to right and top to bottom

possible from the full FOV into the first few virtual coils, causing the last few virtual coils to have very little information content. This allows for easy coil-compression by keeping the top-N_v virtual coils and discarding the remainder, although is not effective in suppressing interference—the first few SVD virtual coils still contain a substantial amount of energy from the neck and torso. As should also be expected from previous literature, the ROI-SVD virtual coils have slightly better characteristics than the SVD virtual coils within the ROI. Instead of compacting energy from the full FOV into the leading virtual coils, the ROI-SVD approach instead only compacts the energy from the ROI Ω into the first few virtual coils, without caring about what happens outside the ROI. As a result, the first few ROI-SVD virtual coils retain more ROI signal than the SVD virtual coils. However, there is still no attempt to suppress interference or separate signal from interference, and the first few ROI-SVD virtual channels can still contain a substantial amount of interference. In contrast to both of these approaches, ROVir effectively compacts as much energy as possible from the ROI signal into the first few virtual coils while also compacting as much energy as possible from the interference region into the last few virtual coils. The middle virtual coils in between these 2 extremes contain relatively little energy from either the ROI signal or the interference.

As shown in Figure 5 (and as should be expected from theory), the first few ROI-SVD virtual coils maximize the amount of signal from the ROI, although we also observe that the difference between the amount of ROI signal between the first few ROI-SVD virtual coils and the first few ROVir coils is relatively small. Importantly (and as also should be expected from theory), the first few ROVir coils have substantially better SIR values and substantially lower interference levels than the other 2 approaches.

In order to take advantage of the interference-suppression capabilities of ROVir, it is necessary to only keep the first few virtual coils, while discarding the rest. There are different strategies that could be used for this. One approach could be to set the number of virtual channels N_v as the smallest number such that the percentage of retained signal is no smaller than a given tolerance. For example, if we decide that we want to make sure that at least 95% of the ROI signal is retained, we can use Figure 5 to automatically determine that we should use $N_v = 6$ for ROVir, $N_v = 5$ for SVD, and $N_v = 4$ for ROI-SVD. Using these N_v parameters leads to the coil-combined root-sum-of-squares images shown on the right side of Figure 3. It is both visually and quantitatively obvious that ROVir retains substantially less interference (1.9%) than either SVD (98.0%) or ROI-SVD (4.0%), while by design, all 3 approaches have a similar amount of retained ROI signal (SVD: 98.1%; ROI-SVD: 98.3%; ROVir: 96.2%).

The previous approach to selecting N_v prioritizes signal retention, but other approaches could also be applied if SIR or interference suppression were higher priorities. For example, a rule that prioritizes SIR might keep those virtual coils that have an SIR value that is larger than a given tolerance, for example, keep only those virtual coils with SIR > 2. This choice would lead to $N_v = 13$ for ROVir (retained ROI signal = 98.1%; retained interference = 3.2%), $N_v = 10$ for SVD (retained ROI signal = 94.9%; retained interference = 45.6%), and $N_v = 9$ for ROI-SVD (retained ROI signal = 99.98%; retained interference = 8.8%). As another example, a rule that prioritizes interference suppression might choose N_v as the largest number such that the percentage of retained interference is no larger than a given tolerance, eg, choose N_v under the constraint that the retained interference is <3%, which would lead to $N_v = 11$ for ROVir (retained ROI signal = 97.5%; retained interference = 2.9%), $N_v = 1$ for SVD (retained ROI signal = 70.8%; retained interference = 2.6%), and $N_v = 2$ for ROI-SVD (retained ROI signal = 82.2%; retained interference = 2.4%). In all of these cases, ROVir demonstrates either substantially less interference for a similar amount of retained ROI signal or substantially more ROI signal for a similar amount of interference.

It should be noted that the plots shown in Figure 5 only describe a spatially averaged summary of the retained ROI signal and interference characteristics, but do not provide

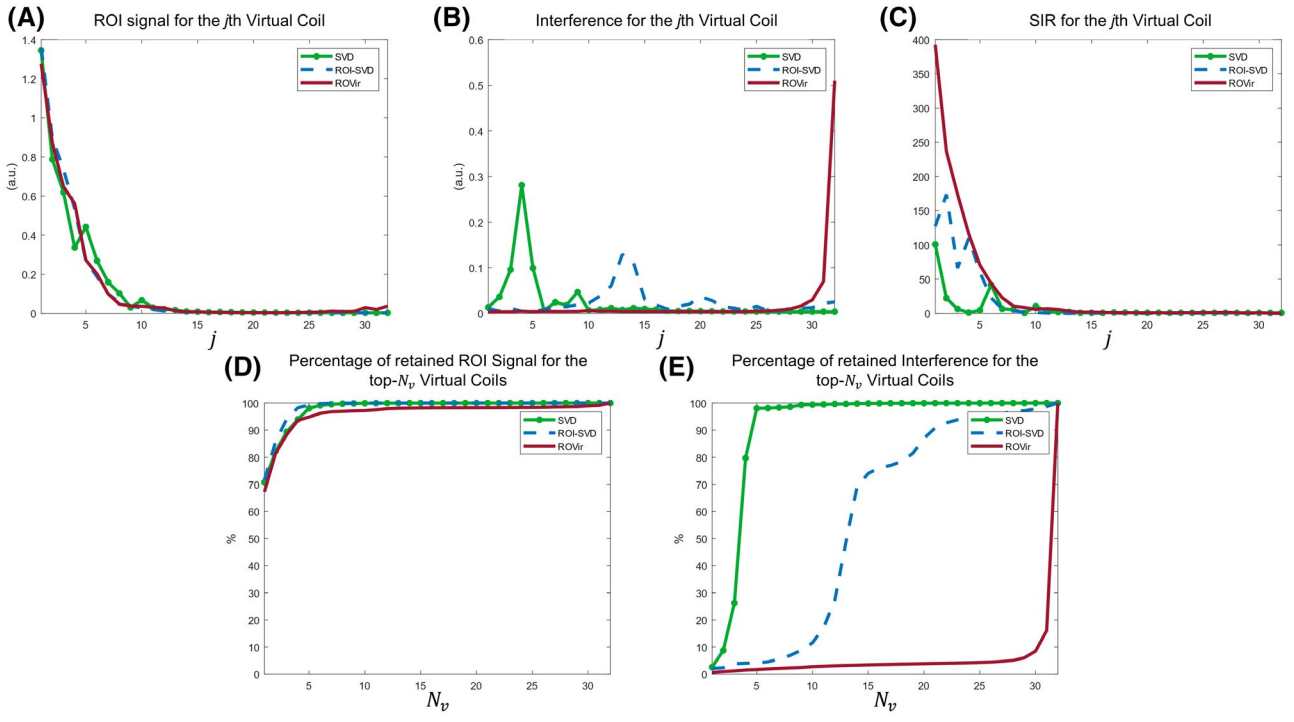

FIGURE 5 Plots of the signal and interference characteristics of SVD, ROI-SVD, and ROVir. A, Plots of the ROI signal (calculated using Equation 7). B, Plots of the interference (calculated using Equation 10). C, Plots of the SIR (calculated using Equation 11). D, Plots of the percentage of retained ROI signal (calculated using Equation 14). E, Plots of the percentage of retained interference (calculated using Equation 15). The relative signal and interference to the combination of the original 32-channel array as a function of the total number of coils to be combined. For the calculation, images in SNR units were used

information about their spatial distributions. To gain better insights into spatially varying and N_v-dependent characteristics, Figure 6 shows images of the relative image intensity (computed voxelwise, and shown as a percentage of the intensity that would be obtained using the full set of 32 coils) in the root-sum-of-squares images after projecting the data onto the subspace spanned by the top-N_v virtual coils. These results show that the ROVir virtual coils can retain nearly 100% of the original ROI signal for most voxels within the ROI, while simultaneously very-substantially suppressing interference from all parts of the interference region. One notable limitation of ROVir relative to ROI-SVD is that ROVir appears to have attenuated the ROI signal near the transition band slightly more than ROI-SVD. However, this may be seen as a reasonable compromise given ROVir's greatly superior capabilities for suppressing interference.

The results shown above were based on computing coil combination weights from high-quality calibration data. For practical applications, it may be worthwhile to understand how sensitive ROVir is to the quality of the calibration data. Supporting Information Figure S1 shows results obtained when computing coil combination weights from calibration images obtained with different resolutions, using the same data and ROIs as in Figure 3. In all cases, the spatial resolution of the original calibration data was reduced by zeroing-out high-frequency k-space information and applying a Hamming window to minimize Gibbs ringing. The results show that while there are relatively major changes in the visual appearance of the calibration images, the characteristics of the ROVir are relatively consistent, with the biggest differences appearing near the transition region (ie, with progressive reduction in the retained signal energy from the brain stem and cerebellum as the spatial resolution of the calibration data worsens).

While the previous results explored the basic interference-suppression characteristics of ROVir with some amount of depth, the following sections are designed to more-quickly illustrate some of the potential uses of ROVir for a variety of different imaging applications.

4.2 | Applications to reduced FOV brain and vocal tract imaging

One potential application of ROVir could be to reduce the size of the image support, enabling the use of a smaller FOV. For example, Figure 7 shows a non-Cartesian imaging scenario where the same 32-channel sagittal brain dataset from the previous example is used to simulate spiral k-space data with a trajectory designed for a small FOV. Although the Nyquist criterion might be satisfied for an image with small spatial support, the data is effectively undersampled for the actual image in this case.

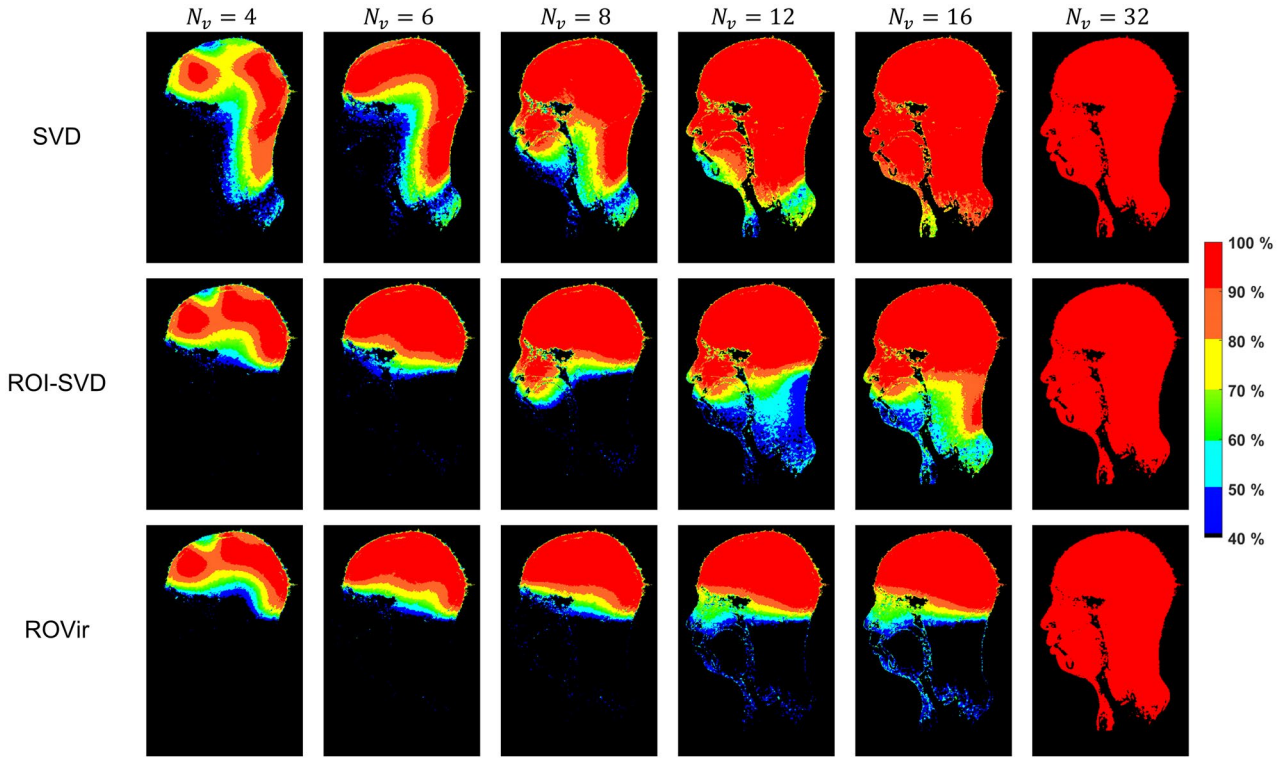

FIGURE 6 Images showing the spatial distribution of the image intensity (computed voxelwise, and shown as a percentage of the intensity that would be obtained using the full set of 32 coils) that is retained in root-sum-of-squares coil-combined images after projecting the original data onto the subspace spanned by the top-N_v virtual coils for SVD, ROI-SVD, and ROVir

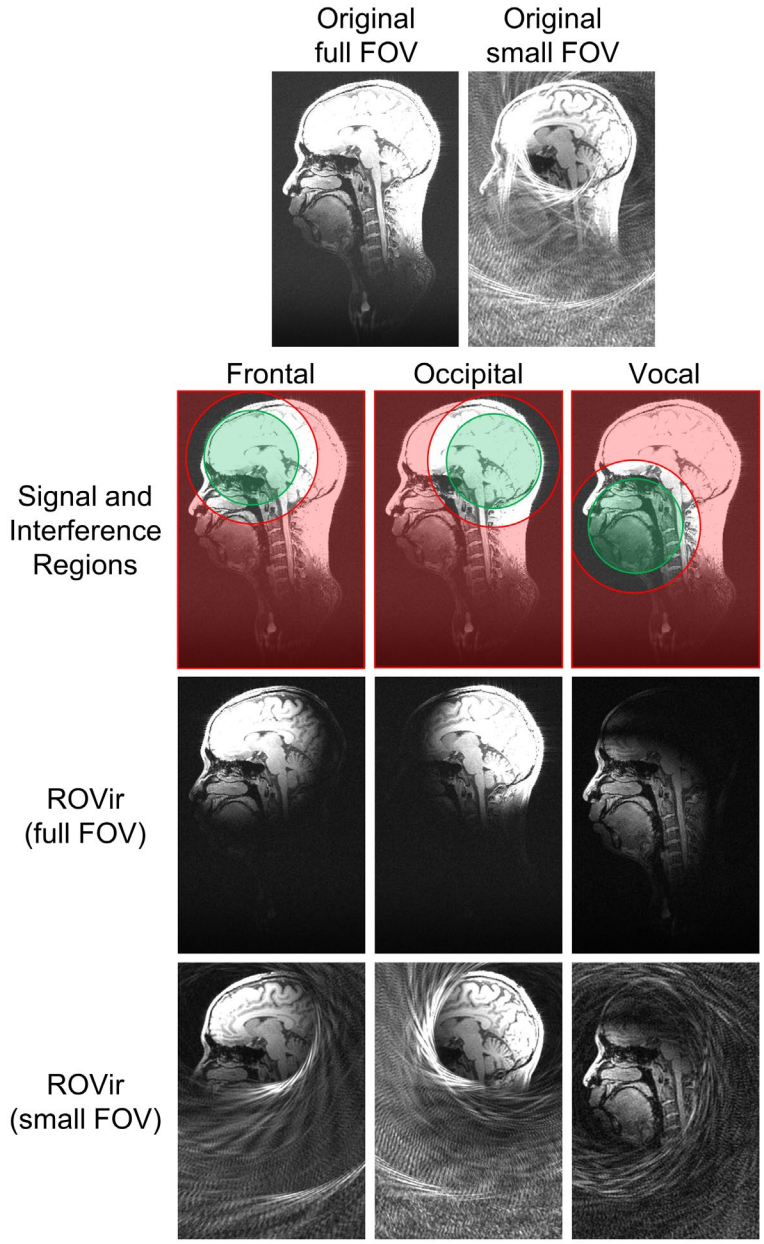

FIGURE 7 Application of ROVir to reduce the size of the FOV in non-Cartesian sagittal brain imaging, which enables highly effective mitigation of aliasing artifacts when using a spiral k-space trajectory designed for a smaller FOV. The signal ROIs are marked in green, while the interference regions are marked in red

As can be seen, if standard gridding-based reconstruction[50] is applied to the original set of 32 coils and coil-combined using root-sum-of-squares, severe aliasing artifacts occur that contaminate most of the image. However, if prescan data were used to learn appropriate ROVir combination weights that suppress all signal except for that within a much smaller ROI, then these same combination weights can be directly applied to the undersampled spiral k-space data (or, equivalently, to the multichannel aliased images obtained after gridding reconstruction) to reduce the effective support of the image post facto, thereby avoiding the problematic aliasing artifacts.

This is illustrated for 3 different ROIs in Figure 7: a frontal-lobe brain ROI, an occipital-lobe brain ROI, and a vocal tract ROI. In each case, we chose the N_v parameter of ROVir automatically based on an SIR criterion, keeping all channels with $\mathrm{SIR} > 1$. This led to $N_v = 6$ for the frontal lobe (with 95.3% retained signal but only 1.4% retained interference), $N_v = 7$ for the occipital lobe (with 96.7% retained signal but only 2.8% retained interference), and $N_v = 2$ for the vocal tract (with 56.0% retained signal but only 0.8% retained interference). In each of the different cases, the use of ROVir allows for a clear, aliasing-free depiction of the ROI. Importantly, all of these different results were obtained from the same initial set of undersampled spiral k-space data, and we could have easily used the beam-steering capabilities of ROVir (as was illustrated in Supporting Information Video S1) to target arbitrary other ROIs if so desired.

In effect, using ROVir in this way can be viewed as a virtual-coil based version of the parallel imaging with localized sensitivities (PILS) technique.[4] In PILS, a specially designed receiver array is utilized where the coil elements have localized and

non-overlapping sensitivities. In contrast, instead of relying on specialized hardware design, we instead use the beam-steering capabilities of ROVir to construct virtual localized sensitivities from an original set of potentially overlapping sensitivities. And unlike PILS, the ROVir coils do not have to be fixed prior to the experiment, they can instead be steered on-the-fly in real-time to meet the needs of the application, or even modified in post-processing long after the subject has left the scanner.

Similar to PILS, the ROVir approach can greatly simplify image reconstruction compared to traditional parallel imaging reconstruction methods, since image reconstruction can be performed using simple Fourier reconstruction techniques. This is quite different from more sophisticated reconstruction methods that must attempt to correct aliasing artifacts, potentially resulting in spatially varying noise amplification. In addition, the fact that ROVir frequently uses N_v values that are much smaller than N_c can reduce data processing and memory requirements even further. Note also that, since the ROVir approach still often keeps more than 1 virtual coil, ROVir does not have to be used exclusively with simple Fourier reconstruction techniques, and can also be potentially used in combination with parallel imaging and/or other advanced reconstruction techniques to enable even further experimental accelerations.

One of the other important observations we can take from this example is that the performance of ROVir varied considerably from the brain ROIs (where the retained signal was >95%) compared to the vocal tract ROI, where the retained signal was only slightly above 50% for a similar level of interference suppression. This difference in performance is highly related to the array geometry, since the soccerball-style headcoil that was used in this experiment was designed primarily for brain imaging and had fewer coil elements in close proximity to the mouth.[51] This highlights the fact that the performance of ROVir will be highly dependent on both the characteristics of the signal ROI and interference region, as well as the geometry of the receiver array.

While the previous case showed sagittal non-Cartesian brain imaging with a reduced FOV, the results shown in Figure 8 show similar reduced FOV results for axial Cartesian brain imaging, with data acquired at 3 T using the same 32-channel head coil as in the previous case. In this case, retrospective uniform 3× undersampling of k-space lines along the anterior-posterior axis leads to classical Cartesian aliasing artifacts for the original 32-channel data. ROVir was designed to retain signal from 4 different rectangular-shaped ROIs, which we name as anterior, central-anterior, central-posterior, and posterior. In each case, N_v was selected automatically based on an SIR criterion, keeping all channels with SIR > 2. This lead to the following parameters: $N_v = 4$ for the anterior ROI (50.4% retained signal but only 0.1% retained interference); $N_v = 3$ for the central-anterior ROI (19.7% retained signal but only 0.3% retained interference); $N_v = 4$ for the central-posterior ROI (20.0% retained signal but only 0.4% retained interference); and $N_v = 6$ for the posterior ROI (46.4% retained signal but only 0.5% retained interference). As can be seen from Figure 8, the use of ROVir is quite successful at reducing the effective size of the FOV, thereby mitigating aliasing artifacts. However, the quantitative numbers presented above also demonstrate that ROVir had better performance (ie, better ROI signal retention) for the anterior and posterior ROIs than it did for the more centralized ROIs, which is again a reflection of the geometric characteristics of the receiver array, since the soccerball-style coil elements generally have more sensitivity to the periphery than to deeper parts of the brain.

4.3 | Applications to reduced FOV cardiac imaging

ROVir concepts can also be potentially powerful for non-brain applications, and Figure 9 shows an example application to reduced FOV cardiac imaging. Specifically, we show data from 2 fully sampled 32-channel CINE acquisitions from the 2013 ISMRM Reconstruction challenge (https://www.ismrm.org/challenge), including 1 central 4-chamber view and 1 mid-papillary short-axis view. In both cases, the original FOV is substantially larger than the cardiac region that is of primary interest to the application. However, by using ROVir to suppress regions that would otherwise cause aliasing, it is possible in both cases to achieve uniform undersampling (equivalent to assuming a smaller FOV) along the phase encoding dimension without incurring the major aliasing artifacts that would otherwise occur. We show results for the 4-chamber with retrospective 2× undersampling and results for the short-axis view with retrospective 3× undersampling. In both cases, N_v was selected automatically based on an SIR criterion, keeping all channels with SIR > 2. This resulted in $N_v = 4$ for the 4-chamber view (with 66.9% retained signal but only 0.2% retained interference) and $N_v = 5$ for the short-axis view (with 90.6% retained signal but only 0.2% retained interference). Although both approaches are successful, it should be noted that substantially more signal was retained for the short-axis view compared to the 4-chamber view, which is again related to the geometric characteristics of the receiver array in relation to the ROIs.

4.4 | Application to brain outer volume suppression

As a last example of ROVir, we consider a brain imaging scenario that is typical of MR spectroscopic imaging (MRSI),

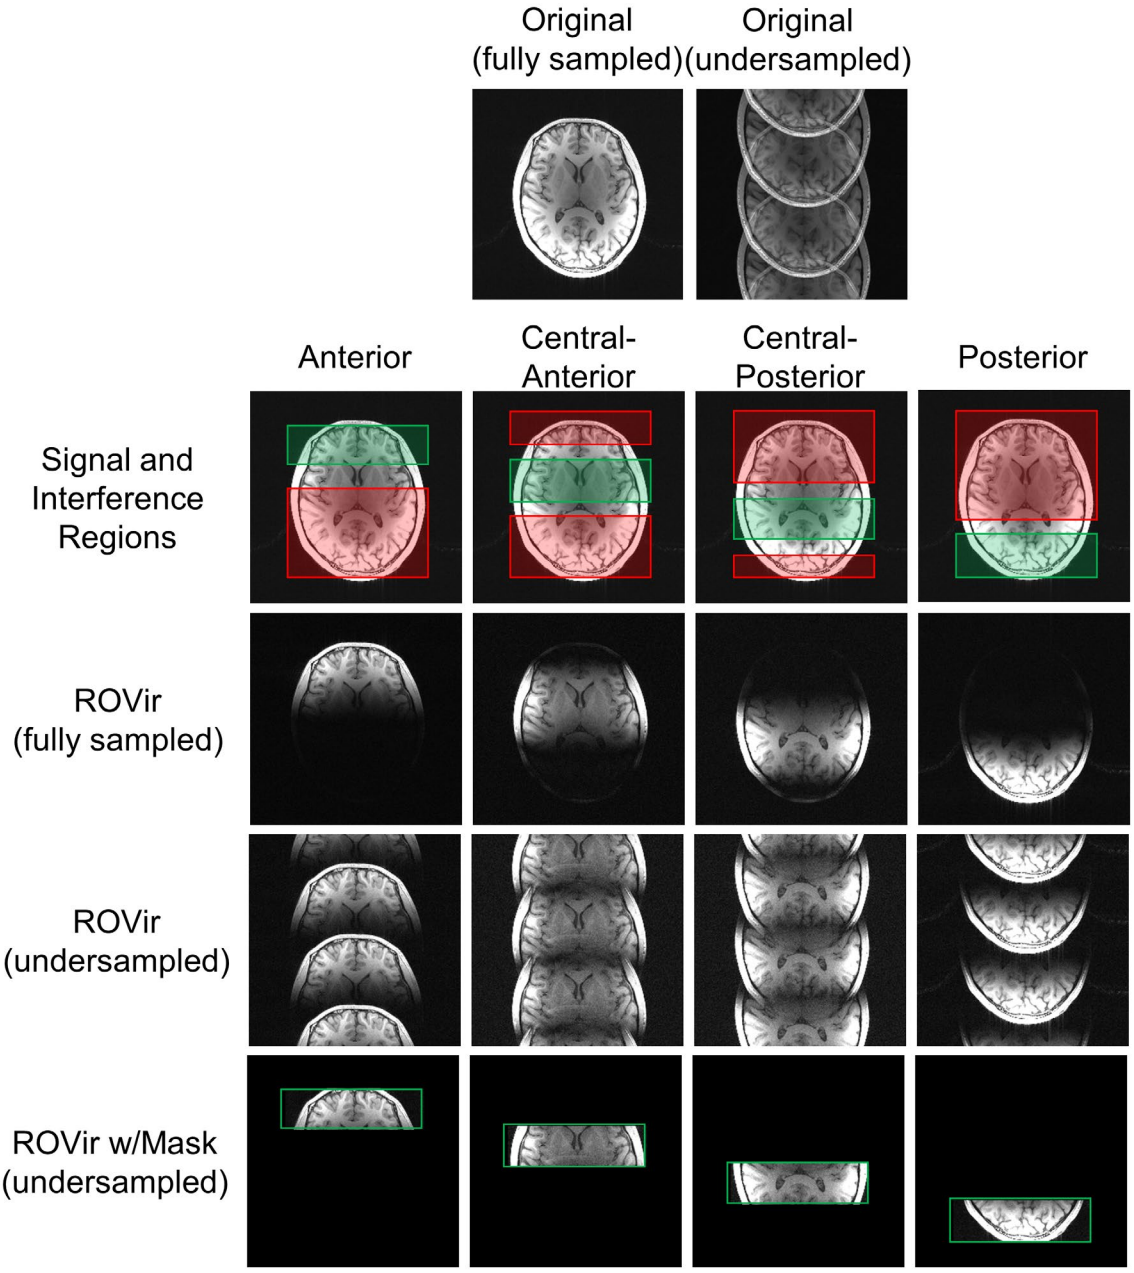

FIGURE 8 Application of ROVir to reduce the size of the FOV in Cartesian axial brain imaging, which enables highly effective mitigation of aliasing artifacts when using uniform 3× Cartesian undersampling. The signal ROIs are marked in green, while the interference regions are marked in red

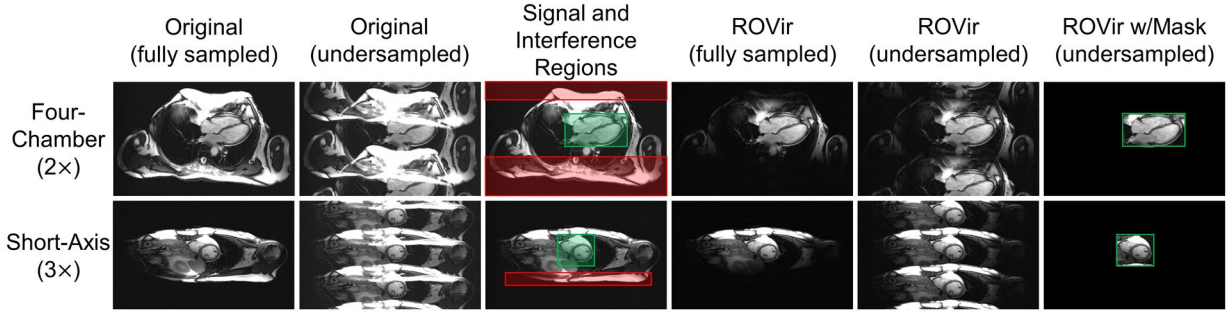

FIGURE 9 Application of ROVir to reduce the size of the FOV and enable undersampling in Cartesian cardiac MRI. The signal ROIs are marked in green, while the interference regions are marked in red

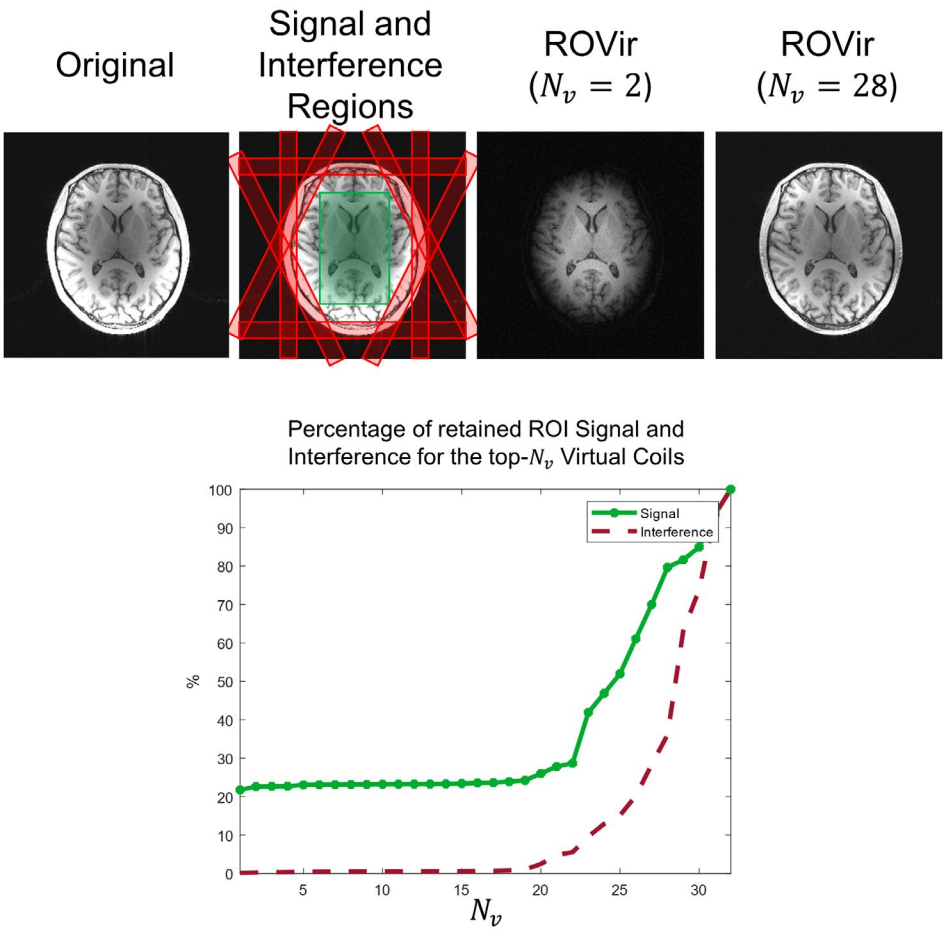

FIGURE 10 Application of ROVir for outer volume suppression. The signal ROIs are marked in green, while the interference regions are marked in red. Plots of the percentage of retained ROI signal (calculated using Equation 14) and the percentage of retained interference (calculated using Equation 15) are also shown

in which we desire to detect signal from brain parenchyma but want to avoid potential contamination from strong signal originating from the scalp (eg, extracranial lipid signals). Traditionally, extracranial signal is often avoided by combining spatially selective excitation methods with outer-volume suppression methods. Unfortunately, this process can be imperfect, and unsuppressed extracranial signal remains a long-standing problem.

Figure 10 illustrates the application of ROVir for the suppression of extracranial signal, where the interference regions have been designed to mimic outer-volume saturation bands that are currently used in MRSI. This result is based on the same 32-channel axial brain dataset from Figure 8. As can be seen from Figure 10, ROVir does a moderately good job of separating the brain parenchyma signal of interest from the extracranial interference, despite the fact that the soccerball-style array geometry (with every coil element having strong sensitivity to the periphery and weaker sensitivity to the deeper brain) may not be very well-suited to this task. Various trade-offs are possible here between interference suppression and signal retention, and we show root-sum-of-squares images for 2 different extremes: $N_v = 2$ (with 22.6% retained signal but only 0.2% retained interference) and $N_v = 28$ (with 79.7% retained signal but only 36.1% retained interference). Although neither of these extremes may be good enough to allow ROVir to be the sole outer volume suppression method in practical MRSI applications, we believe that the use of ROVir in combination with traditional outer volume suppression could potentially represent a substantial improvement over just using the traditional methods by themselves. We also anticipate that even better ROVir outer volume suppression performance might be achievable using an alternative array geometry.

5 | DISCUSSION AND CONCLUSIONS

This paper described the ROVir framework, which applies theoretically optimal beamforming principles to enable the suppression of uninteresting spatial regions and retention of interesting spatial regions. This is achieved by leveraging

the spatial sensitivity variations inherent to multichannel MRI data. The strong potential of ROVir was demonstrated in a number of different application scenarios, including brain, vocal tract, and cardiac imaging, with proof-of-principle results that we believe are quite promising. The scope of our presentation was intentionally broad, and we anticipate that further more-focused prospective studies will be needed to more thoroughly investigate the benefits of ROVir in various imaging scenarios. Importantly, although the results we have presented have exclusively focused on the use of ROVir by itself to demonstrate its capabilities and characteristics, it should be noted that ROVir is completely compatible with existing hardware-, sequence-, and reconstruction-based approaches. We believe that the best practical performance will be achieved through synergistic combinations of different approaches rather than through the exclusive use of ROVir.

Although we believe that ROVir by itself was oftentimes surprisingly effective, we also observed that practical performance had a strong dependence on the geometric properties of the receiver array. As such, this kind of beamforming approach may have major implications for the optimal design of new receiver arrays in the future.

Finally, while this work focused on the application of beamforming principles to linearly mix the data from multiple coils, it is notable that similar ideas could also be applied to optimally linearly mix other kinds of MRI data. For example, it may be possible to apply beamforming principles to linearly mix measured k-space data in order to perform optimal linear spatial localization (instead of relying on classical Fourier reconstruction techniques or other specialized linear estimation designs[52,53]), or to apply beamforming principles to linearly mix different contrast-encoded images to achieve optimal linear spectral-localization of different spectral tissue compartments in applications like multicomponent diffusometry and relaxometry.[54,55] We expect that these will be promising directions for future research.

ORCID

Daeun Kim http://orcid.org/0000-0002-9887-5598

REFERENCES

1. Liang ZP, Lauterbur PC. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. New York: IEEE Press; 2000.
2. Hu X, Parrish T. Reduction of field of view for dynamic imaging. *Magn Reson Med*. 1994;31:691-694.
3. Nagle SK, Levin DN. Multiple region MRI. *Magn Reson Med*. 1999;41:774-786.
4. Griswold MA, Jakob PM, Nittka M, Goldfarb JW, Haase A. Partially parallel imaging with localized sensitivities (PILS). *Magn Reson Med*. 2000;44:602-609.
5. Hinshaw WS. Image formation by nuclear magnetic resonance: the sensitive-point method. *J Appl Phys*. 1976;47:3709-3721.
6. Gordon RE, Hanley PE, Shaw D, et al. Localization of metabolites in animals using 31P topical magnetic resonance. *Nature*. 1980;287:736-738.
7. Ackerman JJH, Grove TH, Wong GG, Gadian DG, Radda GK. Mapping of metabolites in whole animals by ^{31}P NMR using surface coils. *Nature*. 1980;283:167-170.
8. Ackerman JJH. Surface (local) coils as NMR receivers. *Concepts Magn Reson*. 1990;2:33-42.
9. Garroway AN, Grannell PK, Mansfield P. Image formation in NMR by a selective irradiative process. *J Phys C*. 1974;7:L457-L462.
10. Lauterbur PC, Kramer DM, House WV, Chen CN. Zeugmatographic high resolution nuclear magnetic resonance spectroscopy. Images of chemical inhomogeneity within macroscopic objects. *J Am Chem Soc*. 1975;97:6866-6868.
11. Mansfield P, Maudsley AA. Line scan proton spin imaging in biological structures by NMR. *Phys Med Biol*. 1976;21:847-852.
12. Aue WP, Muller S, Cross TA, Seelig J. Volume-selective excitation. A novel approach to topical NMR. *J Magn Reson*. 1984;56:350-354.
13. Feinberg DA, Hoenninger JC, Crooks LE, Kaufman L, Watts JC, Arakawa M. Inner volume MR imaging: technical concepts and their applications. *Radiology*. 1985;156:743-747.
14. Ordidge RJ, Connelly A, Lohman JAB. Image-selected in vivo spectroscopy (ISIS). A new technique for spatially selective NMR spectroscopy. *J Magn Reson*. 1986;66:283-294.
15. Bottomley PA. Spatial localization in NMR spectroscopy in vivo. *Ann NY Acad Sci*. 1987;508:333-348.
16. Mansfield P, Ordidge RJ, Coxon R. Zonally magnified EPI in real time by NMR. *J Phys E*. 1988;21:275-280.
17. Frahm J, Bruhn H, Gyngell ML, Merboldt KD, Hanicke W, Sauter R. Localized high-resolution proton NMR spectroscopy using stimulated echoes: initial applications to human brain in vivo. *Magn Reson Med*. 1989;9:79-93.
18. Haase A. Localization of unaffected spins in NMR imaging and spectroscopy (LOCUS spectroscopy). *Magn Reson Med*. 1986;3:963-969.
19. Singh S, Rutt BK, Henkelman RM. Projection presaturation: a fast and accurate technique for multidimensional spatial localization. *J Magn Reson*. 1990;87:567-583.
20. leRoux P, Gilles RJ, McKinnon GC, Carlier PG. Optimized outer volume suppression for single-shot fast spin-echo cardiac imaging. *J Magn Reson Imaging*. 1998;8:1022-1032.
21. Luo Y, de Graaf RA, DelaBarre L, Tannus A, Garwood M. BISTRO: an outer-volume suppression method that tolerates RF field inhomogeneity. *Magn Reson Med*. 2001;45:1095-1102.
22. Pruessmann KP, Weiger M, Bornert P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med*. 2001;46:638-651.
23. Buehrer M, Pruessmann KP, Boesiger P, Kozerke S. Array compression for MRI with large coil arrays. *Magn Reson Med*. 2007;57:1131-1139.
24. Huang F, Vijayakumar S, Li Y, Hertel S, Duensing GR. A software channel compression technique for faster reconstruction with many channels. *Magn Reson Imaging*. 2008;26:133-141.
25. Doneva M, Bornert P. Automatic coil selection for channel reduction in SENSE-based parallel imaging. *Mag Reson Mater Phys Bio Med*. 2008;21:187-196.

26. King SB, Varosi SM, Duensing GR. Optimum SNR data compression in hardware using an Eigencoil array. *Magn Reson Med*. 2010;63:1346-1356.
27. Zhang T, Pauly JM, Vasanawala SS, Lustig M. Coil compression for accelerated imaging with Cartesian sampling. *Magn Reson Med*. 2013;69:571-582.
28. Blaimer M, Gurberlet M, Kellman P, Breuer FA, Kostler H, Griswold MA. Virtual coil concept for improved parallel MRI employing conjugate symmetric signals. *Magn Reson Med*. 2009;61:93-102.
29. Haldar JP, Setsompop K. Linear predictability in magnetic resonance imaging reconstruction: leveraging shift-invariant Fourier structure for faster and better imaging. *IEEE Signal Process Mag*. 2020;37:69-82.
30. vanVeen BD, Buckley KM. Beamforming: a versatile approach to spatial filtering. *IEEE ASSP Mag*. 1988;5:4-24.
31. Monzingo RA, Haupt RL, Miller TW. *Introduction to Adaptive Arrays*. Raleigh: SciTech Publishing; 2011.
32. Roemer PB, Edelstein WA, Hayes CE, Souza SP, Mueller OM. The NMR phased array. *Magn Reson Med*. 1990;16:192-225.
33. Walsh DO, Gmitro AF, Marcellin MW. Adaptive reconstruction of phased array MR imagery. *Magn Reson Med*. 2000;43:682-690.
34. Kellman P, McVeigh ER. Ghost artifact cancellation using phased array processing. *Magn Reson Med*. 2001;46:335-343.
35. Eggers H, Börner P, Boernert P. Real-time partial parallel spiral imaging with localized sensitivities. *Proc Int Soc Magn Reson Med*. 2001;1772.
36. Lin FH, Witzel T, Zeffiro TA, Belliveau JW. Linear constraint minimum variance beamformer functional magnetic resonance inverse imaging. *NeuroImage*. 2008;43:297-311.
37. Ebbini E, DelaBarre L, Vaughan JT, Gopinath A. Digital beam forming in MRI. *Proc IEEE BenMAS*. 2014.
38. Mandava S, Keerthivasan MB, Martin DR, Altbach MI, Bilgin A. Radial streak artifact reduction using phased array beamforming. *Magn Reson Med*. 2019;81:3915-3923.
39. Xue Y, Yu J, Kang HS, Englander S, Rosen MA, Song HK. Automatic coil selection for streak artifact reduction in radial MRI. *Magn Reson Med*. 2012;67:470-476.
40. Holme HCM, Frahm J. Sinogram-based coil selection for streak artifact reduction in undersampled radial real-time magnetic resonance imaging. *Quant Imaging Med Surg*. 2016;6:552-556.
41. Cauley S, Polak D, Liu W, et al. Geometric coil mixing (GCM) to dampen confounding signals in MRI reconstruction. *Proc Int Soc Magn Reson Med*. 2019;449.
42. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med*. 1999;42:952-962.
43. Fukunaga K. *Introduction to Statistical Pattern Recognition*. San Diego, CA: Academic Press; 1990.
44. Golub GH, Van Loan CF. *Matrix Computations*. Baltimore, MD: Johns Hopkins University Press; 2013.
45. Koles ZJ, Lind JC, Flor-Henry P. Spatial patterns in the background EEG underlying mental disease in man. *Electroencephalogr Clin Neurophysiol*. 1994;91:319-328.
46. Kim D, Polimeni J, Setsompop K, Haldar JP. On coil combination with optimal SNR for linear multichannel k-space reconstruction methods. *In Proceedings of the ISMRM & SMRT Virtual Conference & Exhibition*. 2020;3430.
47. Luenberger DG. *Optimization by Vector Space Methods*. New York: John Wiley & Sons; 1969.
48. Oppenheim AV, Schafer RW, Buck JR. *Discrete-Time Signal Processing*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall, Inc.; 1999.
49. Pauly J, le Roux P, Nishimura D, Macovski A. Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm. *IEEE Trans Med Imaging*. 1991;10:53-65.
50. Jackson JI, Meyer CH, Nishimura DG, Macovski A. Selection of a convolution function for Fourier inversion using gridding. *IEEE Trans Med Imaging*. 1991;10:473-478.
51. Wiggins GC, Triantafyllou C, Potthast A, Reykowski A, Nittka M, Wald LL. 32-channel 3 Tesla receive-only phase-array head coil with soccer-ball element geometry. *Magn Reson Med*. 2006;56:216-223.
52. Bakir T, Reeves SJ. A filter design method for minimizing ringing in a region of interest in MR spectroscopic images. *IEEE Trans Med Imaging*. 2000;19:585-600.
53. Haldar JP, Hernando D, Liang ZP. Shaping spatial response functions for optimal estimation of compartmental signals from limited Fourier data. *Proc IEEE Int Symp Biomed Imaging*. 2007;1364-1367.
54. Varadarajan D, Haldar JP. A theoretical signal processing framework for linear diffusion MRI: implications for parameter estimation and experiment design. *NeuroImage*. 2017;161:206-218.
55. Kim D, Wisnowski JL, Nguyen CT, Haldar JP. Multidimensional correlation spectroscopic imaging of exponential decays: from theoretical principles to in vivo human applications. *NMR Biomed*. 2020;e4244.

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the Supporting Information section.

FIGURE S1 Illustration of the sensitivity of ROVir to the quality of the calibration data used to compute the coil combination weights. Each column corresponds to calibration data with different spatial resolution. The full-resolution case matches the results shown in Figure 3

VIDEO S1 Illustration of using beamforming to steer the sensitivity of the virtual array to different spatial positions. The video shows ROVir results obtained from seeping a circular-shaped ROI (ie, everything interior to the green circle) through the FOV for a 32-channel brain MRI dataset

How to cite this article: Kim D, Cauley SF, Nayak KS, Leahy RM, Haldar JP. Region-optimized virtual (ROVir) coils: Localization and/or suppression of spatial regions using sensor-domain beamforming. *Magn Reson Med*. 2021;86:197–212. https://doi.org/10.1002/mrm.28706