

Joint Design of RF and Gradient Waveforms via Auto-differentiation for 3D Tailored Excitation in MRI

Tianrui Luo^{ID}, Douglas C. Noll^{ID}, *Senior Member, IEEE*,
 Jeffrey A. Fessler^{ID}, *Fellow, IEEE*, and Jon-Fredrik Nielsen

Abstract—This paper proposes a new method for joint design of radiofrequency (RF) and gradient waveforms in Magnetic Resonance Imaging (MRI), and applies it to the design of 3D spatially tailored saturation and inversion pulses. The joint design of both waveforms is characterized by the ODE Bloch equations, to which there is no known direct solution. Existing approaches therefore typically rely on simplified problem formulations based on, e.g., the small-tip approximation or constraining the gradient waveforms to particular shapes, and often apply only to specific objective functions for a narrow set of design goals (e.g., ignoring hardware constraints). This paper develops and exploits an auto-differentiable Bloch simulator to directly compute Jacobians of the (Bloch-simulated) excitation pattern with respect to RF and gradient waveforms. This approach is compatible with *arbitrary* sub-differentiable loss functions, and optimizes the RF and gradients directly without restricting the waveform shapes. For computational efficiency, we derive and implement explicit Bloch simulator Jacobians (approximately halving computation time and memory usage). To enforce hardware limits (peak RF, gradient, and slew rate), we use a change of variables that makes the 3D pulse design problem effectively unconstrained; we then optimize the resulting problem directly using the proposed auto-differentiation framework. We demonstrate our approach with two kinds of 3D excitation pulses that cannot be easily designed with conventional approaches: Outer-volume saturation (90° flip angle), and inner-volume inversion.

Index Terms—Auto-differentiable Bloch simulator, constrained joint pulse design, inner-volume inversion, large flip-angle pulse, outer-volume saturation, tailored RF pulse design.

I. INTRODUCTION

IN A magnetic resonance imaging (MRI) experiment, the dynamic system relationship between the applied

Manuscript received February 9, 2021; revised May 9, 2021; accepted May 20, 2021. Date of publication May 24, 2021; date of current version November 30, 2021. This work was supported by the National Institutes of Health under Grant R01EB023618 and Grant R21AG061839. (Corresponding author: Tianrui Luo.)

Tianrui Luo, Douglas C. Noll, and Jon-Fredrik Nielsen are with the Department of Biomedical Engineering, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: tianrui@umich.edu; dnoll@umich.edu; jfniese@umich.edu).

Jeffrey A. Fessler is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: fessler@umich.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TMI.2021.3083104>, provided by the authors.

Digital Object Identifier 10.1109/TMI.2021.3083104

radiofrequency (RF) and gradient magnetic fields, and the instantaneous spin magnetization change they induce, is concisely described by the Bloch equation. While it is straightforward to calculate the magnetization pattern resulting from a given set of RF and gradient waveforms and tissue parameters, *inverting* the Bloch equation to obtain the waveforms that produce a given desired excitation pattern can be challenging.

This Bloch inversion task is conventionally called an “RF pulse design” problem, reflecting the fact that the most common way to design excitation pulses in MRI is to predefine the gradients in some way, and then optimize only the (complex) RF waveform. Even with that simplification, the design problem remains non-linear and non-convex. Another common simplification is to apply the small-tip approximation [1] that can give reasonable excitation accuracy even for flip angles as high as 90°, at least for conventional 1D (slice-selective) excitations where the instantaneous flip angle during RF excitation remains relatively low. The small-tip approximation leads to a linear (Fourier) relationship between applied fields and the resulting magnetization pattern, and provides intuition about the excitation process by defining an “excitation k-space” trajectory and viewing RF transmission as depositing energy along that trajectory.

The more difficult problem of *jointly* optimizing both RF and gradient waveforms has been approached in various ways. Several methods are based on the small-tip approximation, and on optimizing the gradients over a restricted set of waveform shapes, such as “spoke” or “kt-point” locations in excitation k-space [2]–[6] or parameterized echo-planar or non-Cartesian trajectories [7]–[11]. A more general small-tip design approach for 3D tailored excitation used a B-spline parametrization of the gradient trajectory that is not restricted to particular fixed waveform shapes [12]. These approaches work well for small-tip excitations, but not for applications such as tailored saturation or inversion. In addition, even when the final desired flip angle is small, the *instantaneous* flip angle during RF excitation can be large enough to violate the small-tip assumption [13]. This model mismatch can cause noticeable differences between the Bloch-simulated excitation pattern and that predicted by the small-tip model used in the design.

Another limitation of previous approaches is that the design loss functions are typically limited to certain forms such as least squares (LS) based on the complex transverse excited magnetization, although adaptations to magnitude least squares (MLS) costs have been proposed [14]. Adding hardware constraints to the design formulation adds an additional

layer of complexity that is often either ignored during pulse design, or controlled indirectly via, e.g., Tikhonov regularization of the RF waveform [6].

This work¹ approaches the Bloch inversion task in a more direct and general way that is applicable to the joint design of RF and gradient waveforms for tailored multi-dimensional excitation in MRI. We temporally discretize the pulse, assuming piecewise constant gradient and RF within every time segment. Our method does not rely on the small-tip approximation, works for arbitrary sub-differentiable loss functions, and incorporates hardware constraints. Our approach contains three key elements: First, we derive analytic expressions for the Jacobian operations needed for the Bloch inversion for a unit (discrete) time step. Second, we incorporate these discrete-time Jacobian operations into an automatic differentiation framework [15], to obtain the Jacobian that relates the final magnetization pattern (at the end of the pulse) to the RF and gradient waveforms. Third, we enforce hardware limits by a change of variables that makes the optimization problem effectively unconstrained.

The paper is organized as follows. Section II gives a general form of the joint design problem, and derives the explicit Jacobians useful for accelerating the proposed auto-differentiation pulse design tools. Sections III and IV apply our pulse design tool to two large-tip excitation problems, and validate the results experimentally on a 3T MRI scanner. Sections V and VI discuss and conclude this work.

II. THEORY

A. Problem Formation

We discretize 3D space on a regular grid with a total n_M voxels (“spins”). These spins can have different parameters, e.g. T1, T2, and off-resonance. Let n_T denote the length (number of time points) of the pulse to be designed. For (single coil) joint design of complex RF waveform $b \in \mathbb{C}^{n_T}$ and gradient $g \in \mathbb{R}^{n_T \times 3}$ we are interested in tackling the following general problem:

$$\begin{aligned} \arg \min_{g \in \mathbb{R}^{n_T \times 3}, b \in \mathbb{C}^{n_T}} \quad & \mathcal{L} := f(M_T(g, b), M_D) + \lambda \mathcal{R} \\ \text{s. t.} \quad & \|b\|_\infty \leq b_{\max} \\ & \|g\|_{\infty, \infty} \leq g_{\max} \\ & \|Dg\|_{\infty, \infty} \leq s_{\max}, \end{aligned} \quad (1)$$

where \mathcal{L} is the loss function; $M_D \in \mathbb{R}^{n_M \times 3}$ is the target (Desired) magnetization pattern (a 3-dimensional magnetization vector at each spatial location); $M_T \in \mathbb{R}^{n_M \times 3}$ is the magnetization at the end of the pulse (time T) obtained by integrating the Bloch equation; f is the excitation error metric (e.g., a common choice is least-square error of transverse magnetization, i.e., $\|M_T[:, 1:2] - M_D[:, 1:2]\|_F^2$); and \mathcal{R} is an optional regularizer with weight λ (a common choice is $\mathcal{R} = \|b\|_2^2$ to control peak RF amplitudes and SAR indirectly). For the constraints, we have b_{\max} , g_{\max} , and s_{\max} for peak RF, gradient, and slew rate, respectively; $D \in \mathbb{R}^{n_T \times n_T}$ is the temporal difference matrix divided by δ_t , i.e., Dg takes the 1st order temporal derivative of g and yields the slew rate;

¹Open sourced, github.com/tianluo/AutoDiffPulses.

and $\|\cdot\|_\infty$, and $\|\cdot\|_{\infty, \infty}$ are entry-wise norm returning the largest absolute value of the operand elements.

Problem (1) is challenging for two main reasons: First, the objective is non-convex with respect to its arguments, and is constrained. Second, neither $M_T(g, b)$, nor its Jacobians $\partial M_T/\partial g$ and $\partial M_T/\partial b$ that would be needed to directly minimize (1), have an explicit expression in g and b . To the best of our knowledge, existing methods all deal with simplifications of problem (1) based on, e.g., the small-tip, or spin domain models. In this work, we minimize (1) directly, and assume only that the temporal integration of the Bloch equation is well-approximated by a discrete-time Bloch simulator.

B. Auto-Differentiation

We propose to compute the necessary derivatives² using auto-differentiation [17], such that problem (1) can be optimized for *arbitrary* error metric f and regularization \mathcal{R} . Auto-differentiation tools, e.g., PyTorch [15], decouples computations into stages, and constructs the Jacobian operations at each stage. These single-stage Jacobians are eventually combined using the chain rule. For instance, with a PyTorch based Bloch simulator that computes $M_T(g, b)$, one implicitly obtains $\partial M_T/\partial g$ and $\partial M_T/\partial b$. The loss derivatives with respect to the variables we wish to optimize, i.e., $\partial \mathcal{L}/\partial g$, and $\partial \mathcal{L}/\partial b$, can then be obtained by combining these expressions with $\partial \mathcal{L}/\partial M_T$. This approach allows us to directly optimize g and b with respect to arbitrary losses.

C. Explicit Jacobian Operations

Auto-differentiation tools provide implicit Jacobian operations (also known as the default backward operations in auto-differentiation context) formed from tracking all elementary computations (e.g., addition, multiplication, etc). Such tools also allow users to substitute default Jacobian operations with their own implementations. In practice, such *explicitly* implemented Jacobian operations can be more efficient both computationally and memory-wise. Bloch simulation is typically the most computationally expensive stage in relating pulse waveforms to objective costs. Having explicit Jacobians of the Bloch simulator can therefore accelerate the computation.

To derive discrete time (δ_t) explicit Jacobians in the rotating frame, for all magnetic spins, we assume equilibrium spin magnitudes of 1, relaxation constants $e_1 := \exp(-\delta_t/T_1)$, $e_2 := \exp(-\delta_t/T_2)$, and gyromagnetic ratio γ . At time t , the rotating frame effective magnetic field (B-effective), $B_t \in \mathbb{R}^3$, causes the magnetic spin state $m_t \in \mathbb{R}^3$, to precess (rotate) about an axis $u_t := B_t/\|B_t\|_2$ by angle $\phi_t := -\gamma \delta_t \|B_t\|_2$. One iteration of discrete time Bloch simulation can be expressed as:

$$m_{t+1} = E R_t m_t + e, \quad (2)$$

where $E := \text{diag}([e_2, e_2, e_1])$, $e := [0, 0, 1 - e_1]^T$ model the relaxations; $R_t = \cos(\phi_t)I + (1 - \cos(\phi_t))u_t u_t^T + \sin(\phi_t)[u_t]_\times$ models the rotation; I is the 3D identity matrix, $\text{diag}([1, 1, 1])$; and $[u_t]_\times$ denotes the cross product matrix of u_t , i.e., $[u_t]_\times m_t = u_t \times m_t$. The rotation matrix R_t is spatially dependent, as it accounts for B-effective which incorporates

²Or Clarke generalized subdifferentials for non-smooth objectives [16].

applied gradients, off-resonance, etc. Relaxation terms, as they depend on the underlying tissue property, are generally also spatially dependent. To avoid notation clutter, we have not indicated those spatial dependencies in Eq. (2); rather, Eq. (2) can be considered to hold for a single spin isochromat, with the appropriate R_t and relaxation terms for that isochromat.

One can verify the following recursive expressions for partial derivatives of the loss with respect to m_t and B_t :

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial m_t} &= R_t^\top E \frac{\partial \mathcal{L}}{\partial m_{t+1}} =: h_t, \\ \frac{\partial \mathcal{L}}{\partial B_t} &= \gamma \delta_t / \phi_t (u_t u_t^\top - I) \frac{\partial \mathcal{L}}{\partial u_t} - \gamma \delta_t \frac{\partial \mathcal{L}}{\partial \phi_t} u_t, \\ \frac{\partial \mathcal{L}}{\partial u_t} &= \phi_t \left(c_t (m_t u_t^\top + m_t^\top u_t I) + s_t [m_t]_\times \right) E h_{t+1}, \\ \frac{\partial \mathcal{L}}{\partial \phi_t} &= ([u_t]_\times R_t m_t)^\top E h_{t+1}, \end{aligned} \quad (3)$$

where $c_t := (1 - \cos(\phi_t)) / \phi_t$, $s_t := \sin(\phi_t) / \phi_t$. Given $\partial B_t / \partial g_t$ and $\partial B_t / \partial b_t$ (which are easy to compute), we obtain the necessary derivatives for the joint optimization by the chain rule:

$$\frac{\partial \mathcal{L}}{\partial g_t} = \frac{\partial \mathcal{L}}{\partial B_t} \cdot \frac{\partial B_t}{\partial g_t}, \quad \frac{\partial \mathcal{L}}{\partial b_t} = \frac{\partial \mathcal{L}}{\partial B_t} \cdot \frac{\partial B_t}{\partial b_t}.$$

Using the explicit Jacobians in (3) for the Bloch simulator operations halved both the computation time and memory use compared to the default implicit Jacobian operations provided by PyTorch (v1.3).

The remaining Jacobians, such as $\partial \mathcal{L} / \partial M_T$, $\partial B_t / \partial b_t$, and $\partial B_t / \partial g_t$, typically do not involve complicated computations. Also, they can vary with different objectives, e.g., switching from LS to MLS; or with different excitation settings, e.g., uniform vs non-uniform transmit sensitivities. For program generality, we left these remaining Jacobians to be obtained implicitly by the auto-differentiation framework.

D. Constraints

Constrained optimization often requires extra effort to ensure solution feasibility, such as feasible set projection and constraint substitution with penalizations. This would involve crafting projection algorithms, and tuning penalty parameters. For problem (1), in the absence of convexity, we use a change of variables [18], [19] that converts the problem into an effectively unconstrained one and avoids such extra effort during optimization.

Let $s \in \mathbb{R}^{n_T \times 3}$ denote the slew rate, i.e., $s = Dg$. Define $\tilde{s} = \tan(\pi/2 \cdot s/s_{\max})$; $\tilde{\rho} := \tan(\pi/2 \cdot |b|/b_{\max})$, $\theta := \angle b$. We automatically have $\|b\|_\infty \leq b_{\max}$ and $\|s\|_{\infty, \infty} \leq s_{\max}$ always satisfied (Fig. 1). Thus, we reformulate problem (1) as:

$$\begin{aligned} \arg \min_{\tilde{s} \in \mathbb{R}^{n_T \times 3}; \tilde{\rho}, \theta \in \mathbb{R}^{n_T}} \quad & \mathcal{L} := f(M_T(g, b), M_D) + \lambda \mathcal{R} \\ \text{s. t.} \quad & \|g\|_{\infty, \infty} \leq g_{\max} \\ & Dg = 2s_{\max} / \pi \cdot \tan^{-1}(\tilde{s}) \\ & b = 2b_{\max} / \pi \cdot \exp(i\theta) \tan^{-1}(\tilde{\rho}). \end{aligned} \quad (4)$$

In practice, for change of variable, \tan^{-1} can be replaced with any other *strictly monotone* function, e.g., sigmoid, that maps an unconstrained domain to an interval.

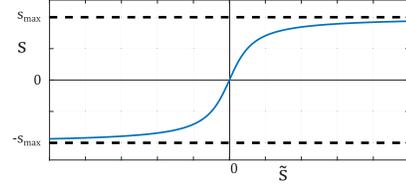


Fig. 1. Turning constrained slew rate s into unconstrained \tilde{s} , by change of variable \tan^{-1} .

Empirically, for 3D tailored pulse design, we observe that, with extended kt-points initializations [12], gradient amplitudes are well below typical max gradient constraints (55 Gcm^{-1}) prior to and throughout the optimization procedure. Hence, while problem (4) is still constrained formally, its max gradient is practically inactive. We thus treated it as an unconstrained problem for the results shown in this paper.

E. Optimization Algorithm

We select initial waveforms g and b that satisfy the constraints. To minimize (4), we alternately update $\tilde{\rho}$, θ , and \tilde{s} , as shown in Algorithm 1. This alternating strategy is commonly used in existing joint design approaches [7], [12], and helps reduce the problem size for the L-BFGS algorithm used in updating the pulse. With auto-differentiation, the optimization algorithm can be formulated without reference to the specific loss function, as demonstrated with very different design problems in section III. We use the L-BFGS optimizer provided by PyTorch for updating the variables within an iteration. The number of iterations may depend on pulse initializations. We empirically choose $N = 10$ for experiments in this work. This choice can vary with applications.

Algorithm 1 Alternating Minimization

- 1: **Inputs:** Variables: g, b ; Number of iterations: N
 - 2: Compute $\tilde{\rho}, \theta$ and \tilde{s} from g, b
 - 3: **for** $n = 1$ to N **do**
 - 4: Fix \tilde{s} ; Optimize $\tilde{\rho}, \theta$, using L-BFGS
 - 5: Fix $\tilde{\rho}, \theta$; Optimize \tilde{s} , using L-BFGS
 - 6: **end for**
 - 7: Compute g and b , from $\tilde{\rho}, \theta$ and \tilde{s}
 - 8: **return** g, b
-

III. METHODS

To demonstrate the utility and generality of our approach, we designed two different kinds of 3D tailored pulses: outer-volume (OV) saturation, and inner-volume (IV) inversion.

A. 3D Outer-Volume Saturation Pulse Design

Outer-volume saturation pulses can be used to limit the imaging field of view (FOV), and hence has the potential to reduce both the time needed for data acquisition as well as motion artifacts from, e.g., the chest wall or abdomen in body imaging applications [20]–[22]. OV saturation pulses should ideally have a high flip angle in the OV region (e.g., 90 degrees), while leaving the IV unperturbed. These pulses are typically followed immediately by a gradient crusher. Since the phase of OV magnetization prior to the

crusher is unimportant, we use MLS loss in design, and include a regularization term on RF power to indirectly control SAR as well as to demonstrate the generality of our approach:

$$\mathcal{L}_{90} = |||M_T[:, 1:2]| - |M_D[:, 1:2]|||_2^2 + \lambda ||b||_2^2, \quad (5)$$

where, $|M[:, 1:2]| := \text{abs}(M[:, 1] + iM[:, 2])$, is a vector function computing magnitudes of spin transverse magnetizations. For the target excitation profile, we set rows in M_D to $[1, 0, 0]$ for OV spins, and $[0, 0, 1]$ for IV spins. We implemented this loss in PyTorch to obtain the Jacobian $\partial \mathcal{L}_{90} / \partial M_T$ as described in the Theory section.

In principle, small-tip based 3D tailored design approaches can also be applied to this loss by scaling the designed RF pulse to attain the desired 90° flip (although the resulting pulse may exceed peak RF limits). We therefore compare our approach with the small-tip method in [12], starting with the same initial b and g waveforms in both cases (initialized as described in III-C).

B. Inner-Volume Inversion Pulse Design

Next we designed another type of excitation pulse that is difficult to design using conventional approaches: an IV inversion pulse. Such a pulse may be useful for, e.g., selective inversion of arterial blood for flow territory mapping in perfusion imaging. For this pulse we propose a very different excitation loss based on the *longitudinal* magnetization:

$$\mathcal{L}_{180} = ||M_T[:, 3] - M_D[:, 3]||_2^2 + \lambda ||b||_2^2. \quad (6)$$

We set rows in M_D to $[0, 0, 1]$ for OV spins, and $[0, 0, -1]$ for IV spins. We also implement this loss in PyTorch.

C. Pulse Initializations

The loss in problem (4) is non-convex, and the choice of initial g and b waveforms influences the final excitation result. How best to initialize these waveforms is an open problem. In [12], the initialization problem (in the context of small-tip 3D tailored excitation) was addressed by evaluating two popular choices for the excitation k-space trajectory, stack-of-spirals and SPINS [11], along with a novel alternative approach, “extended kt-points”, that chooses gradients based on the desired (target) excitation pattern. Sun *et al.* showed that the extended kt-points approach produces comparable or better excitation accuracy than stack-of-spirals or SPINS [12], so we chose it for the experiments in this work.

Once the gradients were initialized in this way, we initialized b using the approach in [23]. These initial RF waveforms were scaled down when necessary to satisfy the b_{\max} constraint.

D. B-Effective Computation

The particular form of B-effective depends on the excitation objective and other application-specific components. Besides the RF and gradient waveforms, it often also contains an off-resonance map (that may vary with time) and transmit sensitivity maps. Other factors such as gradient non-linearity can also be included in B-effective. For the single transmit coil phantom studies in this work, B-effective accounts for RF, gradient, and a static off-resonance map. Specifically,

TABLE I
PULSE DURATION AND TR/TE

Parameters	OV90	IV180	IV180M
Pulse Duration	6.5 ms	4.8 ms	4.5 ms
TR / TE	2 s / 15 ms	3 s / minimum TE	

at time t , let $b_t \in \mathbb{C}$ and $g_t \in \mathbb{R}^3$ denote the instantaneous RF and gradients, respectively. For position r relative to the scanner iso-center, with off-resonance $\omega(r)$, the instantaneous B-effective is:

$$B_t = [\Re(b_t), \Im(b_t), \langle g_t, r \rangle + \omega(r)/\gamma], \quad (7)$$

where \Re and \Im extract the real and imaginary component, respectively.

E. Phantom Experiments

We performed validation experiments in an Agar phantom on a GE MR750 3T scanner. Fig. 2 illustrates the experimental setup, including the prescribed IV and OV regions. All experiments used the same observed off-resonance map in the pulse design (Fig. 2). We used $T_1 = 1.47$ s, $T_2 = 70$ ms in the Bloch simulation during designs. For all studies, we conducted the 3D design on a $32 \times 32 \times 20$ voxel grid of FOV $24 \times 24 \times 24$ cm; with RF power weighting coefficient $\lambda = 4$, and constraints: $b_{\max} = 0.25$ G, $g_{\max} = 5$ G cm⁻¹, $s_{\max} = 12$ G cm⁻¹ ms⁻¹. We quantified excitation performance in simulations with normalized root mean squared error (NRMSE). Spins in “don’t care” regions (Fig. 2) were excluded when calculating the NRMSE. We ran our design programs on an NVidia 2080 Ti graphics card. With the settings above, our method uses around 1.1 GB GPU RAM for all 3 designs. This includes the intrinsic GPU RAM usage of the PyTorch environment.

We performed three different experiments: 1. OV 90° excitation using the cuboid target pattern shown in Fig. 2 (OV90); 2. IV inversion using that same cuboid target pattern (IV180); and 3. IV inversion with a block-M target pattern (IV180M). The experiments were implemented using a vendor-agnostic platform for rapid prototyping of MR pulse sequences [24], [25]. IV dimensions were $9 \times 9 \times 4.8$ cm³ and $9 \times 12.8 \times 4.8$ cm³ for the cuboid and block-M target patterns, respectively. We used a single channel transmit/receive bird-cage coil for all experiments and assumed uniform RF transmit sensitivity during pulse design. To mitigate Gibbs ringing artifacts, we acquired the phantom images at a matrix size of $120 \times 120 \times 48$ and then downsized in image space to match the design grid size $32 \times 32 \times 20$.

We used long TR to wait for spin full recovery from saturation and inversion. For the OV90 experiment, as a substantial volume of the phantom is excited with large angle, we used TE = 15 ms to intentionally decay signal intensity and avoid saturating amplifiers in signal receiver during acquisition. We use minimum TE for the inversion experiments.

For inversion performance validation, we use the sequence in Fig. 3 to obtain both phase and magnitude phantom images, with tip-down angle set to 10°. We expect a π phase difference between inverted (IV) non-inverted (OV) regions,

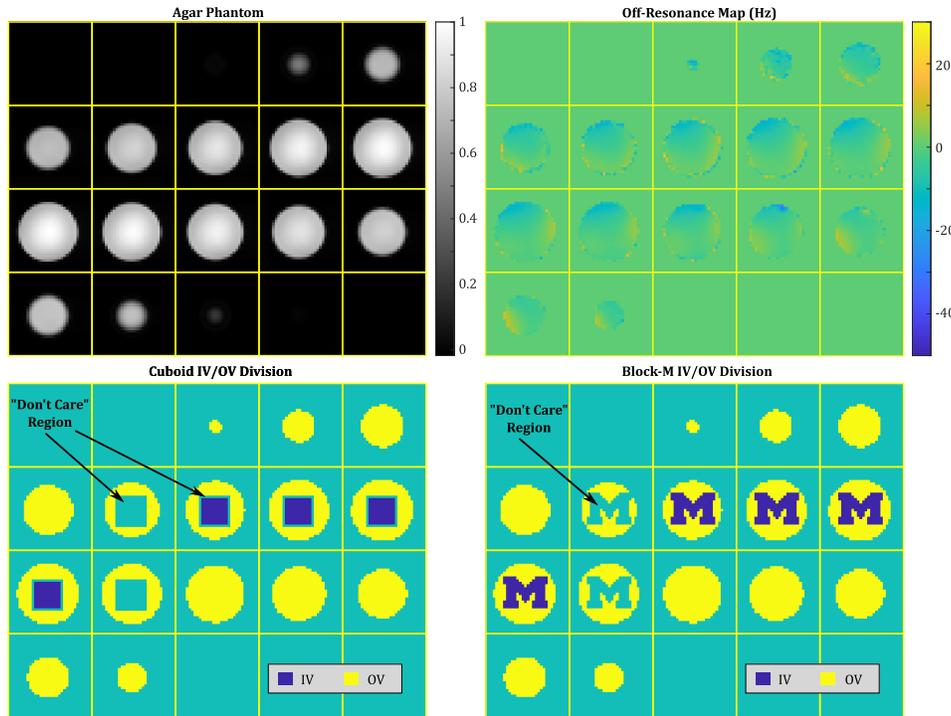


Fig. 2. Experimental phantom, and the two target patterns (IV/OV divisions) used in our experiments. Top left: Magnitude image of the uniform Agar phantom. Top right: Observed field map, used in the pulse design to account for B0 inhomogeneity. A conservative mask that is 1-voxel-wide larger than the phantom support was used to ensure that the phantom boundary was included in the design. This expanded mask is the likely cause for the relatively large B0 values in some pixels at the edge of the mask (that are likely just outside the phantom). Bottom: Cuboid (left) and “block-M” (right) target patterns. We prescribed a “don’t care” (region with arrows) at the boundary between the IV and OV regions that is excluded when calculating the design loss. For the cuboid pattern, the don’t care region included the entire 3D IV/OV boundary, whereas for the block-M pattern, only the top and bottom slices (slices 6 and 11; slice numbers increase left-to-right and top-to-bottom) were included due to the low in-plane spatial resolution of the design grid.



Fig. 3. Schematic diagram of the imaging sequence used to characterize (validate) the 3D tailored inversion pulses.

as the excitation pulse should tip inverted and non-inverted spins in opposite directions. In addition to the IV180 and IV180M excitation pulses, we imaged the phantom using the same sequence settings (TE/TR, flip angle, matrix size) using a conventional slab-selective Shinnar-Le Roux (SLR) pulse. We normalized the inversion images using this “non-inversion” image to eliminate receive coil sensitivity weighting (both magnitude and phase) in the inversion images. For completeness, the unnormalized images are shown in supplemental materials (Fig. S5).

IV. RESULTS

A. OV90

Figure 4 shows the OV saturation pulses obtained with the proposed method and the small-tip approach in [12], and Fig. 5 shows the corresponding phantom imaging results. To keep the small-tip RF pulse within peak amplitude limits, we applied VERSE [26] near the end of the pulse. Our approach required 10 min for design, longer than the small-tip approach (2 min). While the RF waveforms differ markedly, the (excitation)

k-space trajectories are more similar, though differences are clearly observed in the 3D trajectory plot (Fig. 4).

We observe excellent agreement between simulated and acquired excitation patterns (Fig. 5). Also, the proposed method produces much lower excitation error than the small-tip design (46% lower NRMSE error overall); this is expected as the small-tip assumption is violated after scaling the RF to attain the desired 90° flip angle in the OV, which reduces the error in the OV at the expense of increased error in the IV.

B. IV180

Fig. 6 shows the results of the cuboid IV inversion experiment. Pulse waveforms and images from simulation and phantom experiments are shown. Pulse design took 6 min. Simulations and acquired images are in excellent agreement, and indicate successful inversion within the IV with errors mainly located at the IV/OV boundary as expected. In particular, we observe dark bands along the IV/OV boundary in the magnitude image. Spins in this region are not fully inverted, resulting in low signal intensities in the magnitude image. The phase image shows an abrupt π transition at the IV/OV boundary, indicating successful IV inversion.

C. IV180M

Fig. 7 shows the results of the block-M IV inversion experiment. Pulse design took 6 min. Slew rates are near the limit, similar to the IV180 experiment. Simulation and acquired

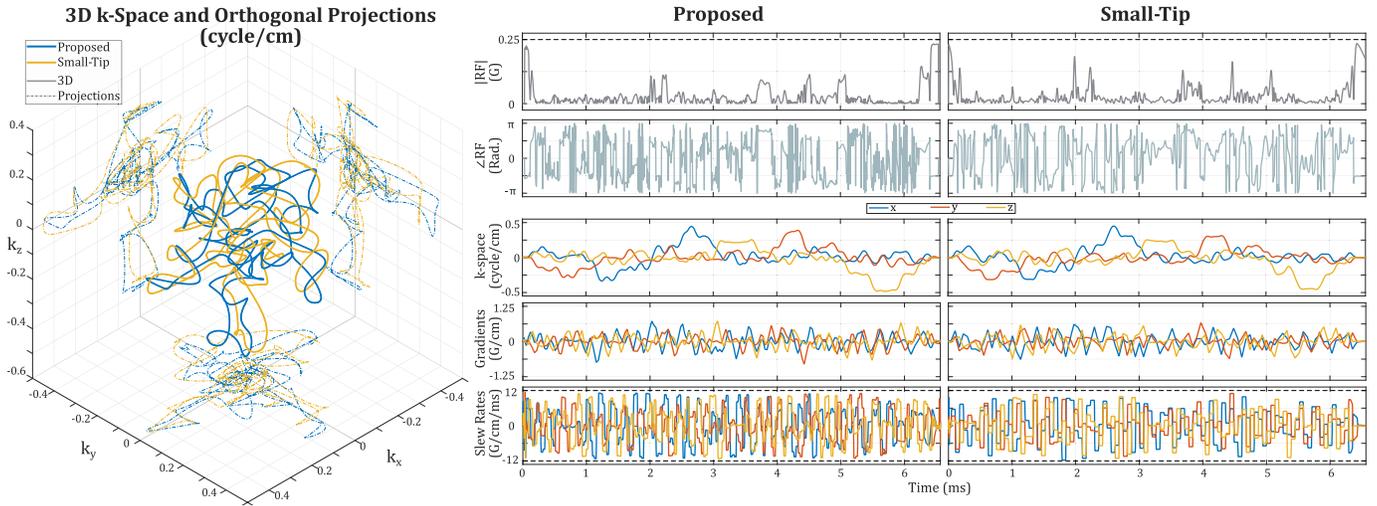


Fig. 4. OV saturation pulses for the cuboid IV (Fig. 2), designed with our approach (Proposed) and the small-tip method in [12] (experiment OV90). The left panel shows the 3D k-space trajectories and their orthogonal projections: The two trajectories explore largely overlapping regions in excitation k-space. The right two panels show RF, gradient, and slew rate waveforms. Both designs satisfy the constraints, but for the small-tip design it was necessary to apply the VERSE [26] algorithm near the end of the pulse (see Discussion). Gradient peak amplitudes remain quite small ($\ll 5 \text{ G cm}^{-1}$), whereas the gradient slew rates are frequently near their limit.

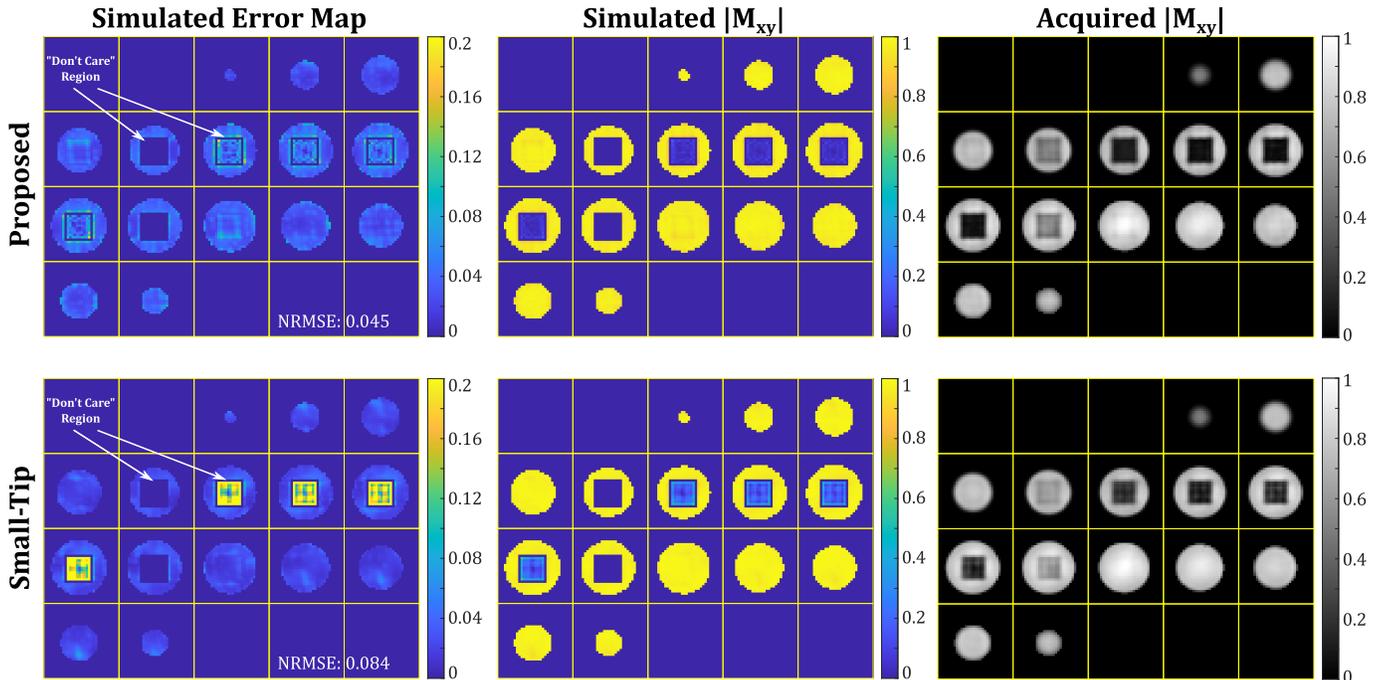


Fig. 5. Experimental validation of the pulses shown in Fig. 4. The left panel shows the error map from simulation. Our approach has much smaller ($\sim 46\%$) NRMSE in simulation compared to Small-Tip. Acquired results (right) agree with the simulations (middle). Small-Tip approach has larger error inside the IV: This is expected, as the method produces only small-tip pulses, that we then scaled to meet the large-tip objective. The scaling increases excitation error inside IV while reducing error in the OV. Our approach directly designs large-tip pulses without this type of ‘scaling’ error.

images again indicate successful inversion. The NRMSE is larger than in the IV180 experiment, suggesting a trade-off between geometry complexity and excitation accuracy. Excitation error is largest near the in-plane edge of the block-M, where target M_z changes sharply from 1 to -1 . We again observe dark bands along the IV/OV boundary, and an abrupt phase change across that boundary, as expected.

V. DISCUSSION

We have demonstrated a new approach to joint multi-dimensional excitation pulse design that directly optimizes

both RF and gradient waveforms. Our approach is not limited to small-tip design problems, and is compatible with quite general loss/design functions such as those that involve longitudinal and/or magnitude magnetization. We validated our approach with 3D tailored large-tip objectives. For this type of application, the “extended kt-points” small-tip initialization [12] led to excellent large-tip results.

We chose to implement our auto-differentiable Bloch simulator with B-effective as its input for its generality: one can possibly prepend to it arbitrary functions that compute B-effective from various parameters, such as multi-coil parallel

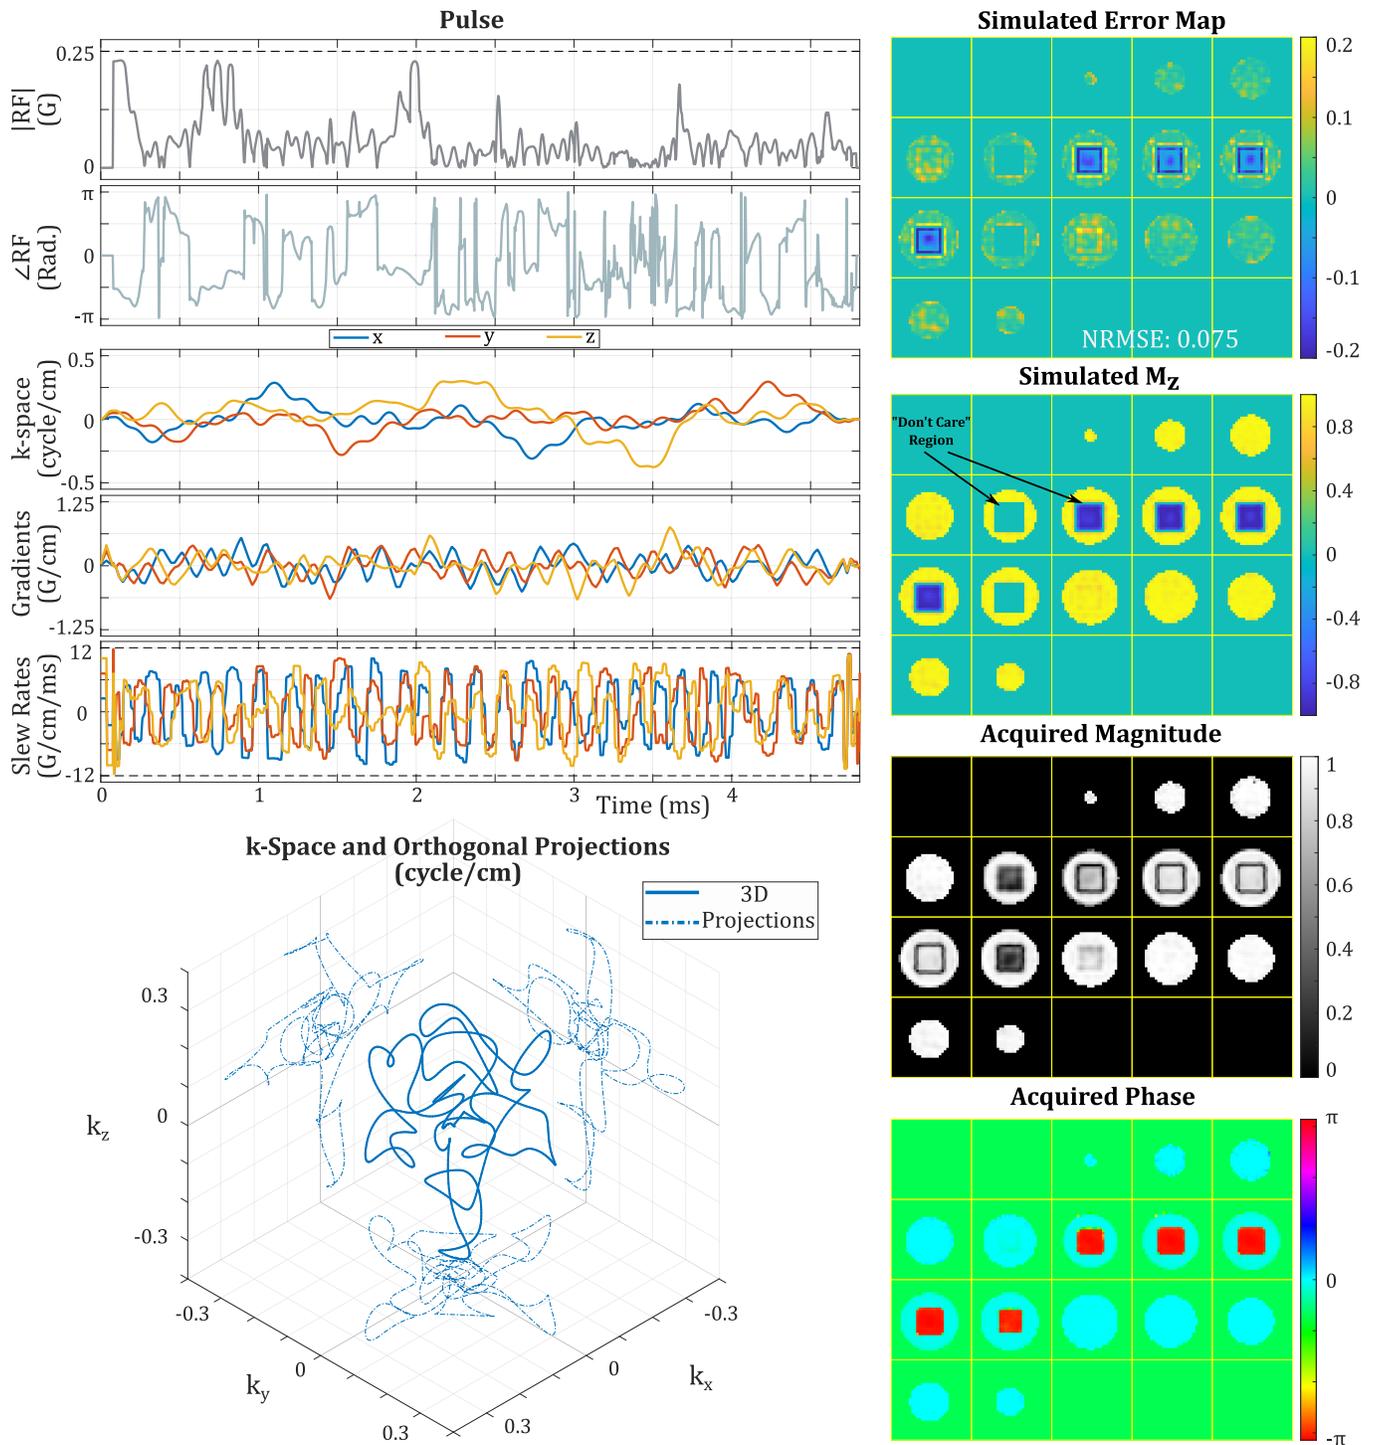


Fig. 6. IV inversion results for the cuboid IV pattern (Experiment IV180). As desired, the 4.2 ms pulse satisfies all constraints. The gradient waveform is again far from its peak constraint of 5 G cm^{-1} . Compared to the OV90 experiment, the pulse has more extreme slew rate waveforms. The acquired magnitude and phase (i.e., the “observed inversion”) were obtained with the sequence in Fig. 3. We observe good agreement between simulated and acquired inversion patterns. The designed pulse successfully inverts the IV, as indicated by similar magnitude image intensity in the IV and OV regions (apart from transmit/receive coil shading) and a π phase shift across the IV/OV boundary. The dark bands in the acquired images at the IV/OV boundary are due to spin saturation from incomplete inversion (and overlap substantially with the prescribed “don’t care” region).

transmit (pTx) sensitivities, spin movements, and even non-linear response of gradient amplifiers, gradient delays, etc. This choice that favors generality may require more memory than software designs that take RF and gradients as inputs directly, and may require more expensive hardware with

adequate memory for high-dimensional design problems. In particular, an interface that uses RF, gradient and spin location inputs requires a memory size proportional to $(N_T + N_T \times 3 + N_M \times 3)$, whereas our interface requires memory proportional to $(N_M \times 3 \times N_T + N_M \times 3)$. Our implementation can find use

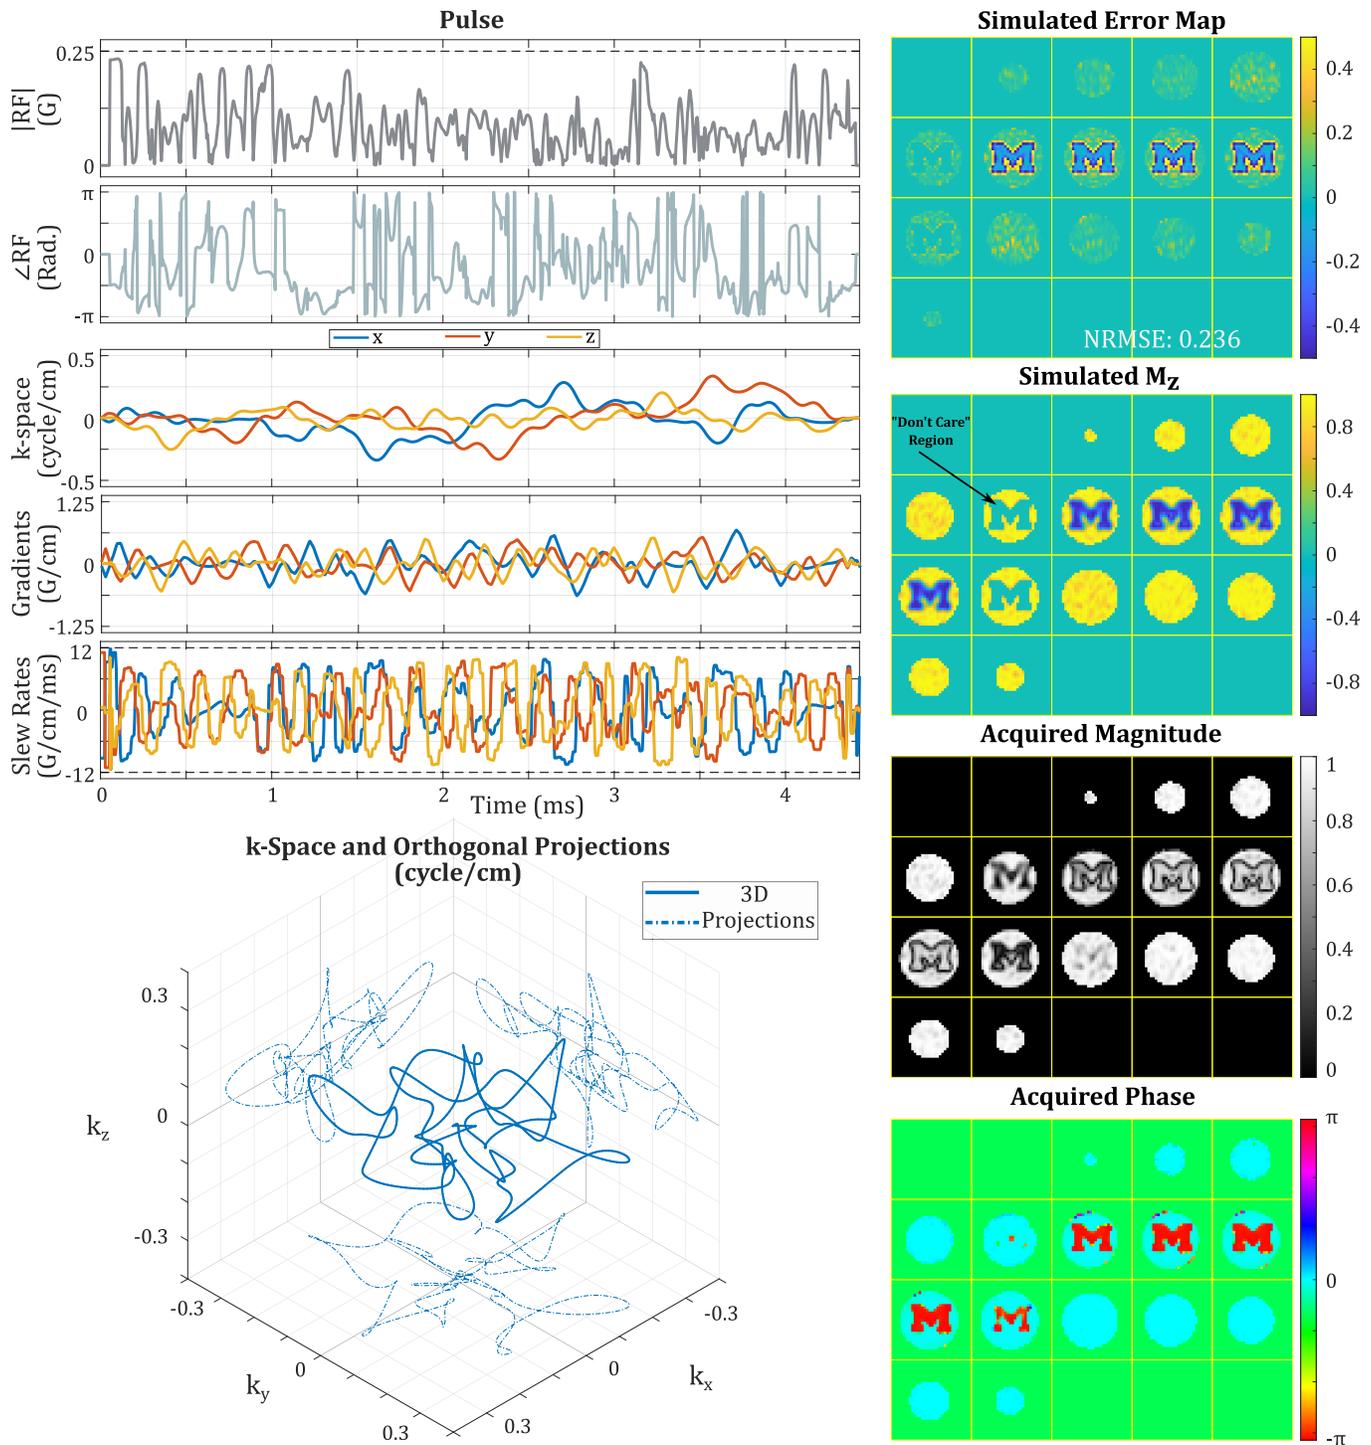


Fig. 7. IV inversion results for the block-M target pattern (Experiment IV180M). The 4.5 ms pulse satisfies all constraints. The gradient waveforms are again well below the peak amplitude constraint of 5 G cm^{-1} , and slew rates are near the constraint for significant portions of the waveform duration. As in Fig. 6, the pulse successfully inverts the IV. The dark bands in the acquired magnitude image at the IV/OV boundary are due to saturation effects arising from the finite resolution (excitation k-space extent) of the pulse – even though only slices 6 and 11 were included in the “don’t care” region in the design due to the low in-plane spatial resolution of the design grid (see Fig. 2).

in different scenarios for proof-of-principle designs that one could then follow by customized simulators that meet specific computational requirements.

For the Bloch simulator, one may alternatively consider using the hard pulse approximation, which splits the instantaneous rotation matrix R_t into two rotations: RF rotation, and transversal rotation due to the applied gradients and

off-resonance. The hard pulse approximation is the basis for the SLR pulse design algorithm, and is crucial for the development of that algorithm. In our case, however, such splitting actually increases the number of elementary computations: when multiplying a vector, RF and transversal rotations require 9 and 4 multiplications, respectively, while direct multiplication by R_t requires only 9. We therefore believe that

the hard pulse approximation does not confer any particular advantages on our approach.

Apart from the explicit Jacobians introduced here, additional steps may be taken to reduce computation time. Computation time is primarily determined by pulse length, and not on the grid size (number of voxels) since computations are done voxel-wise and can be easily parallelized to within GPU RAM limits. Apart from increasing the simulation expense, longer pulses may also slow down Algorithm 1, since we used L-BFGS for updating RF and gradients. In the future, to shorten the optimization time for online pulse design tasks, it may be helpful to use coarser δ_t in the Bloch simulation [27] (here we used $4\ \mu\text{s}$ to match our scanner's hardware dwell time), or parameterize the gradient waveforms to reduce the optimization problem size (e.g., using B-splines as in [12]).

For the experiments presented, we used voxel resolution $7.5 \times 7.5 \times 12\ \text{mm}^3$ and grid size $32 \times 32 \times 20$ for the pulse design. For more complex target excitation patterns and/or a larger FOV (e.g., as in the ISMRM parallel transmit pulse design challenge [28]), it may be desirable to increase the spatial resolution (maximum extent in excitation k-space) and/or grid size for finer excitation accuracy control. For instance, with a larger grid size, we would have space for in-plane “don't care” region for the IV180M experiment, which may help reduce excitation error. A larger grid size will increase the memory usage in simulation, for which the use of multiple graphics cards may be needed to parallelize simulations across voxels.

In the OV90 experiment (Figs. 4–5), we were able to apply VERSE [26], [29] to the pulse designed with the small-tip approach [12] to avoid violating the RF amplitude limit (after scaling to attain 90° flip angle in the OV). However, this adjustment was possible only because the 6.5 ms pulse happened to exceed peak RF only near the end of the pulse, allowing us to apply the VERSE strategy in a relatively straightforward way. In the more general case, where peak RF is exceeded during the middle of the pulse, it is more difficult to apply the VERSE technique to 3D RF pulses such as those designed here. We found empirically that in the OV90 experiment, shorter pulses designed with the small-tip approach tended to exceed peak RF during one or more intermediate intervals (after scaling), and that we were therefore unable to carry out an effective experimental evaluation for the purposes of the comparison presented here (Figs. 4–5). The proposed approach avoids this difficulty because peak RF is constrained as described in II-D; our approach was in fact able to design a shorter (4 ms) OV saturation pulse with the same excitation error as in Fig. 5 (not shown). In future work, a design approach that integrates VERSE into our method may be useful to further shorten a pulse for a given excitation objective.

Pulse design problems are in general non-convex in terms of b and g . Due to a lack of theoretical tools for non-convex problem convergence analysis, it is unclear how to best design an optimization algorithm for such problems *a priori*. In Algorithm 1, instead of simultaneously updating both RF and gradient waveforms, we chose to update them alternately as often done in existing small-tip joint designs [2], [7], [12], [30]. In supplemental Figs. S6-S7, we compared the alternating scheme with the simultaneous scheme, and found

empirically that the alternating scheme optimizes faster than the simultaneous scheme for the specific problem settings we have in this work. Unfortunately, the non-convexity prevents us from fully comprehending this behavior, and we make no claims that the alternating approach used here is optimal over the many possible alternatives. Iteration stopping criteria for updating b and g are also commonly chosen *ad hoc*. Besides limiting the maximum number of iterations as we have done in Algorithm 1, another option can be setting a threshold to assert large-enough loss decreases and/or updates of the variables at each iteration. However, due to the \tan^{-1} change of variables, a minuscule update of b and s near their limits will be mapped to a vast difference in the optimization variables $\tilde{\rho}$ and \tilde{s} , respectively. As an alternative to the updates of variables, one can threshold the norms of variable derivatives, or the change of b and s as the iteration stopping criteria.

Our approach may remind readers of optimal control (OC) based pulse design methods [31], [32]. Comparing the OC formulation with our approach, we can make the following observations (ignoring the penalization and relaxation terms): Eq. (2) is the forward propagation of spin states in OC (state equation); the first identity in Eq. (3) is the backward propagation of OC Lagrange multiplier (costate equation); the second identity in Eq. (3) is the derivative for iteratively optimizing B-effective as the control. Being one step in the computation of excitation losses, our auto-differentiable Bloch simulator enables reusing the forward and backward iterations regardless of the actual design loss function chosen, and propagating the derivatives to the actual controls, i.e., the RF and gradient waveforms. In future works, it may be beneficial to employ tools from control research. For instance, optimization algorithms from OC may accelerate or replace algorithm 1.

Like many other pulse design works [2], [7], [12], [30], the experiments we presented in this work assumes that the off-resonance map is known. We have not attempted to enforce robustness to unknown off-resonance patterns, however the auto-differentiation nature of this work allows incorporating such robustness into the loss as an error metric or regularization for the joint design. In particular, noting the relation of our tool to the control framework, we anticipate that incorporation of robust control methods can improve robustness to off-resonance errors.

A major advantage of our approach is that it enables designs involving arbitrary loss functions, enabling novel design formulations that have so far not been tractable. For example, we demonstrated in (6) a loss involving only longitudinal magnetization. Other possibilities may include the addition of constraints or regularization terms involving specific absorption rate (SAR) or peripheral nerve stimulation (PNS). Another important feature is that the method back-propagates derivatives throughout the Bloch simulator, which may facilitate development of neural network based pulse design approaches.

A limitation of our method is that it only works for fixed pulse length, as determined by the initial waveforms. As shown in the pulse plots in Figs. 4, 6, and 7, there are temporal intervals where neither the RF, gradients, nor slew rates are hitting their constraints. This may suggest that the pulses can be shortened without sacrificing excitation accuracy. Pulse shortening can be formed as a minimum-time pulse design

problem [31] in the OC context. Noting the relation of our method to the OC approach, for future work, we expect employing existing OC tools to be helpful in overcoming the fixed length limitation.

VI. CONCLUSION

In this work, we have proposed a novel approach based on auto-differentiation tools for the joint design of RF and gradient waveforms, and validated it with multi-dimensional spatially tailored excitation tasks in MRI. Using short (<5 ms) excitation pulses and single (body) coil RF transmission, we demonstrated experimentally that even a fairly complex 3D spatial pattern (block-M) can be selectively inverted. Our method is not limited to specific design objectives. To reduce computation time and memory requirements, we derived explicit Jacobians for the Bloch simulator, as the simulation steps are typically the most computationally demanding. We used a change of variables to enforce hardware limits, enabling use of simpler unconstrained optimization. We anticipate that the proposed method will be useful for a broad range of excitation pulse design problems in MRI.

REFERENCES

- [1] J. Pauly, D. Nishimura, and A. Macovski, "A K-space analysis of small-tip-angle excitation," *J. Magn. Reson.*, vol. 81, no. 1, pp. 43–56, Jan. 1989.
- [2] C. Ma, D. Xu, K. F. King, and Z.-P. Liang, "Joint design of spoke trajectories and RF pulses for parallel excitation," *Magn. Reson. Med.*, vol. 65, no. 4, pp. 973–985, Apr. 2011, doi: [10.1002/mrm.22676](https://doi.org/10.1002/mrm.22676).
- [3] A. C. Zelniski, L. L. Wald, K. Setsompop, V. K. Goyal, and E. Adalsteinsson, "Sparsity-enforced slice-selective MRI RF excitation pulse design," *IEEE Trans. Med. Imag.*, vol. 27, no. 9, pp. 1213–1229, Sep. 2008.
- [4] M. A. Cloos *et al.*, "KT-points: Short three-dimensional tailored RF pulses for flip-angle homogenization over an extended volume," *Magn. Reson. Med.*, vol. 67, no. 1, pp. 72–80, Jan. 2012, doi: [10.1002/mrm.22978](https://doi.org/10.1002/mrm.22978).
- [5] D. Yoon, J. A. Fessler, A. C. Gilbert, and D. C. Noll, "Fast joint design method for parallel excitation radiofrequency pulse and gradient waveforms considering off-resonance," *Magn. Reson. Med.*, vol. 68, no. 1, pp. 278–285, Jul. 2012, doi: [10.1002/mrm.24311](https://doi.org/10.1002/mrm.24311).
- [6] W. A. Grissom, M.-M. Khalighi, L. I. Sacolick, B. K. Rutt, and M. W. Vogel, "Small-tip-angle spokes pulse design using interleaved greedy and local optimization methods," *Magn. Reson. Med.*, vol. 68, no. 5, pp. 1553–1562, Nov. 2012, doi: [10.1002/mrm.24165](https://doi.org/10.1002/mrm.24165).
- [7] C.-Y. Yip, W. A. Grissom, J. A. Fessler, and D. C. Noll, "Joint design of trajectory and RF pulses for parallel excitation," *Magn. Reson. Med.*, vol. 58, no. 3, pp. 598–604, Sep. 2007.
- [8] C. J. Hardy, P. A. Bottomley, M. O'Donnell, and P. Roemer, "Optimization of two-dimensional spatially selective NMR pulses by simulated annealing," *J. Magn. Reson.*, vol. 77, no. 2, pp. 233–250, Apr. 1988.
- [9] M. Davids, L. R. Schad, L. L. Wald, and B. Guérin, "Fast three-dimensional inner volume excitations using parallel transmission and optimized K-space trajectories," *Magn. Reson. Med.*, vol. 76, no. 4, pp. 1170–1182, Oct. 2016, doi: [10.1002/mrm.26021](https://doi.org/10.1002/mrm.26021).
- [10] S. Tingting, X. Ling, T. Guisheng, C. Jieru, L. Feng, and S. Crozier, "Advanced three-dimensional tailored RF pulse design in volume selective parallel excitation," *IEEE Trans. Med. Imag.*, vol. 31, no. 5, pp. 997–1007, May 2012. [Online]. Available: <http://ieeexplore.ieee.org/document/6094221/>
- [11] S. J. Malik, S. Keihaninejad, A. Hammers, and J. V. Hajnal, "Tailored excitation in 3D with spiral nonselective (SPINS) RF pulses," *Magn. Reson. Med.*, vol. 67, no. 5, pp. 1303–1315, May 2012, doi: [10.1002/mrm.23118](https://doi.org/10.1002/mrm.23118).
- [12] S. Hao, J. A. Fessler, D. C. Noll, and J.-F. Nielsen, "Joint design of excitation K-space trajectory and RF pulse for small-tip 3D tailored excitation in MRI," *IEEE Trans. Med. Imag.*, vol. 35, no. 2, pp. 468–479, Feb. 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7268765/>
- [13] H. Sun, "Topics in steady-state MRI sequences and RF pulse optimization," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., Univ. Michigan, Ann Arbor, MI, USA, 2015.
- [14] K. Setsompop, L. L. Wald, V. Alagappan, B. A. Gagoski, and E. Adalsteinsson, "Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 tesla with eight channels," *Magn. Reson. Med.*, vol. 59, no. 4, pp. 908–915, Apr. 2008, doi: [10.1002/mrm.21513](https://doi.org/10.1002/mrm.21513).
- [15] A. Paszke *et al.*, "Automatic differentiation in pytorch," in *Proc. NIPS Autodiff Workshop*, 2017. [Online]. Available: <https://openreview.net/forum?id=BJJsrnmcZ>
- [16] F. H. Clarke, "Nonsmooth analysis and optimization," in *Proc. Int. Congr. Mathematicians*, vol. 5, 1983, pp. 847–853.
- [17] A. Griewank and A. Walther, *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. Philadelphia, PA, USA: SIAM, Jan. 2008. [Online]. Available: <https://doi.org/10.1137%2F1.9780898717761>
- [18] F. S. Sisser, "Elimination of bounds in optimization problems by transforming variables," *Math. Program.*, vol. 20, no. 1, pp. 110–121, Dec. 1981.
- [19] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [20] T. Luo, D. C. Noll, and J.-F. Nielsen, "3D inner volume imaging, with 3D tailored outer, volume suppression RF pulses," in *Proc. ISMRM*, 2019, p. 4631. [Online]. Available: <http://archive.ismrm.org/2019/4631.html>
- [21] D. Mitsouras, R. V. Mulkern, and F. J. Rybicki, "Strategies for inner volume 3D fast spin echo magnetic resonance imaging using nonselective refocusing radio frequency pulses a," *Med. Phys.*, vol. 33, no. 1, pp. 173–186, 2005.
- [22] B. Wilm, J. Svensson, A. Henning, K. Pruessmann, P. Boesiger, and S. Kollias, "Reduced field-of-view MRI using outer volume suppression for spinal cord diffusion imaging," *Magn. Reson. Med.*, vol. 57, no. 3, pp. 625–630, 2007.
- [23] C.-Y. Yip, J. A. Fessler, and D. C. Noll, "Iterative RF pulse design for multidimensional, small-tip-angle selective excitation," *Magn. Reson. Med.*, vol. 54, no. 4, pp. 908–917, Oct. 2005, doi: [10.1002/mrm.20631](https://doi.org/10.1002/mrm.20631).
- [24] J.-F. Nielsen and D. C. Noll, "TOPPE: A framework for rapid prototyping of MR pulse sequences," *Magn. Reson. Med.*, vol. 79, no. 6, pp. 3128–3134, Jun. 2018. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26990>
- [25] K. J. Layton *et al.*, "Pulseq: A rapid and hardware-independent pulse sequence prototyping framework," *Magn. Reson. Med.*, vol. 77, no. 4, pp. 1544–1552, Apr. 2017.
- [26] S. Conolly, D. Nishimura, A. Macovski, and G. Glover, "Variable-rate selective excitation," *J. Magn. Reson.*, vol. 78, no. 3, pp. 440–458, Jul. 1988.
- [27] T. Luo, D. C. Noll, J. A. Fessler, and J.-F. Nielsen, "Multi-scale accelerated auto-differentiable bloch-simulation based joint design of excitation RF and gradient waveforms," in *ISMRM*, 2021, p. 3958.
- [28] W. A. Grissom, K. Setsompop, S. A. Hurley, J. Tsao, J. V. Velikina, and A. A. Samsonov, "Advancing RF pulse design using an open-competition format: Report from the 2015 ISMRM challenge," *Magn. Reson. Med.*, vol. 78, no. 4, pp. 1352–1361, Oct. 2017. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26512>
- [29] B. A. Hargreaves, C. H. Cunningham, D. G. Nishimura, and S. M. Conolly, "Variable-rate selective excitation for rapid MRI sequences," *Magn. Reson. Med.*, vol. 52, no. 3, pp. 590–597, 2004.
- [30] Z. Cao, M. J. Donahue, J. Ma, and W. A. Grissom, "Joint design of large-tip-angle parallel RF pulses and blipped gradient trajectories," *Magn. Reson. Med.*, vol. 75, no. 3, pp. 1198–1208, 2016.
- [31] S. Conolly, D. Nishimura, and A. Macovski, "Optimal control solutions to the magnetic resonance selective excitation problem," *IEEE Trans. Med. Imag.*, vol. I-5, no. 2, pp. 106–115, Jun. 1986. [Online]. Available: <http://ieeexplore.ieee.org/document/4307754/>
- [32] W. A. Grissom, D. Xu, A. B. Kerr, J. A. Fessler, and D. C. Noll, "Fast large-tip-angle multidimensional and parallel RF pulse design in MRI," *IEEE Trans. Med. Imag.*, vol. 28, no. 10, pp. 1548–1559, Oct. 2009. [Online]. Available: <http://ieeexplore.ieee.org/document/4915785/>

Supporting Information: Joint Design of RF and gradient waveforms via auto-differentiation for 3D tailored excitation in MRI

Tianrui Luo, Douglas C. Noll, Jeffrey A. Fessler, Jon-Fredrik Nielsen

I. ADDITIONAL SIMULATION RESULTS

For completeness, we present here both the initial pulses designed as described in III-C in the main text, and the corresponding optimized pulses obtained with the proposed approach. In the case of OV90, we also include the small-tip pulse. In each case, we show the simulated excitation pattern for each pulse.

In addition to IV90, IV180, IV180M, we include here a cuboid IV inversion pulse based on the B0 field map acquired in the brain of a healthy volunteer. This is done to demonstrate the feasibility of designing an IV inversion pulse with the proposed approach using a more realistic B0 map than that shown in Fig. 3. In addition, for that simulation experiment we compare the optimized pulse with a pulse obtained by only optimizing the RF waveform, i.e., keeping the gradient waveforms fixed at their initial shapes. This is done to assess the relative importance of also optimizing the gradient waveforms.

The key takeaways from these figures are: (1) The excitation patterns produced by the initial pulses are substantially inferior to the optimized patterns. (2) The initial and optimized excitation k-space trajectories tend to be similar, suggesting that a local minimum is obtained. (3) The initial and optimized RF waveforms (amplitude and phase), on the other hand, differ markedly from each other, suggesting relatively weak dependence on initial RF waveform. This may be due to the fact that each RF sample can be optimized independently of the other samples, unlike gradient waveforms that are subject to slew rate constraints. (4) Despite the similarity between the initial and optimized gradient waveforms, the optimized waveforms produce a more accurate excitation than the pulse obtained by optimizing only the RF waveform (Fig. S4).

A. OV90

B. IV180

C. IV180M

D. Brain Simulation

II. UNNORMALIZED INVERSION IMAGES

In Fig. S5, we show the unnormalized images for the 3D spatially tailored inversion experiments (IV180, IV180M). In the main text, we normalized the ‘‘Cuboidal’’ and ‘‘Block-M’’ images by element-wise division by the ‘‘No Inversion’’ image, which eliminates the image intensity and phase variations due to receiver coil sensitivity.

III. ALTERNATING VS SIMULTANEOUS MINIMIZATION

Here we compare the alternating optimization used in the main text with a simultaneous update scheme that optimizes b (RF waveform) and g (the three gradient waveforms) together at each iteration rather than fixing one and optimizing the other. We observe empirically that for the \mathcal{L}_{90} and \mathcal{L}_{180} losses defined in the main text, with extended kt-points initializations, the alternating update decreases the design losses faster than the simultaneous update. However, the two objectives are both non-convex in terms of b and g , which makes this behavior difficult to analyze. We therefore cannot claim that this alternating scheme will outperform simultaneous updates in the general case (i.e., for all other design problems).

IV. IMPACT OF A SMALL GRADIENT DELAY

On modern MRI scanners, the physically realized gradient waveforms are typically slightly misaligned in time relative to the RF waveform, even after the vendor’s built-in gradient delay correction is applied. This delay is on the order of the gradient sampling (dwell, or raster) time, which on our scanner is 4 μ s. To assess robustness against such delays, we simulated the excitation produced by the OV90 and IV180 pulses for delays of 4 μ s and -4μ s. As shown in Fig. S8, such delays led to excitation patterns that are visually nearly indistinguishable from the original patterns with no delay, and degraded the performance of our designed pulses by less than 1 percentage point in NRMSE.

V. IMPACT OF INCORRECT OFF-RESONANCE

Motion or respiratory effects can cause mismatch between the acquired and actual off-resonance patterns. To assess robustness against such mismatch, we simulated the excitation produced by the same OV90 and IV180 pulses under different off-resonance maps: (i) the acquired off-resonance map, and (ii) the acquired off-resonance map scaled by a multiplicative factor of 3. As shown in Fig. S9, such mismatch led to a very similar excited patterns with about 1 percentage point increases in NRMSE.

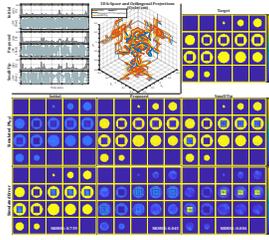


Fig. S1. OV saturation pulses (and results) for the cuboid IV, designed with our approach (Proposed) and Sun's small-tip method (experiment OV90 in the main manuscript). On the 3D k-space plot, the three trajectories explore largely overlapping regions in excitation k-space. Our approach has much smaller (-46%) NRMSE in simulation compared to Small-Tip. Small-Tip approach has larger error inside the IV: This is expected, as the method produces only small-tip pulses, that we then scaled to meet the large-tip objective. The scaling increases excitation error inside IV while reducing error in the OV. Our approach directly designs large-tip pulses without this type of 'scaling' error.

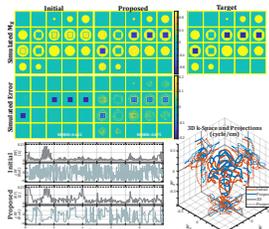


Fig. S2. Cuboid IV inversion (experiment IV180 in the main manuscript). The lower left panel compares the RF waveforms of the initial and optimized (Proposed) pulses. The lower right panel compares the k-space trajectories of the two pulses.

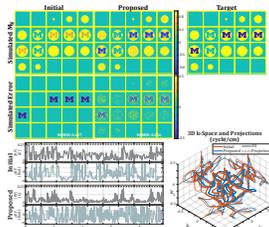


Fig. S3. Block-M IV inversion pulse (experiment IV180M in the main manuscript). The lower left panel compares the RF waveforms of the initial and optimized (Proposed) pulses. The lower right panel compares the k-space trajectories of the two pulses.

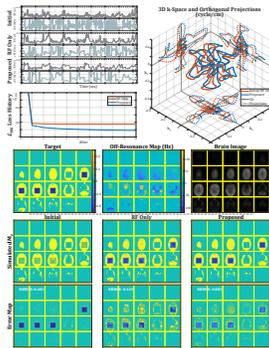


Fig. S4. Cuboid IV inversion based on a B0 (off-resonance) map obtained in a volunteer. Three pulses are compared: (1) Initial, (2) 'RF Only', obtained by keeping the gradients fixed at their initial shapes and optimizing only the RF waveform with the proposed auto-differentiation approach, and (3) the proposed jointly optimized pulse. While our optimized k-space trajectory is similar to the initial k-space trajectory, the jointly optimized pulse (Proposed) attains an excitation accuracy (NRMSE: 8.1%) that is 37% better than the 'RF Only' pulse (NRMSE: 12.9%). This improvement is also reflected in the convergence (\mathcal{L}_{180} loss history) plot.

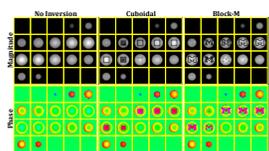


Fig. S5. Images from the IV180 and IV180M experiments. "Cuboidal" and "Block-M" are the raw images of the IV180 and IV180M experiments, respectively. "No Inversion" was acquired using the same sequence (TR/TE, flip-angle, readout trajectory), except the inversion pulse in the sequence has RF amplitude set to 0. The 3 sets of magnitude and phase images shown here share the same image intensity and phase variations due to receiver coil sensitivity.

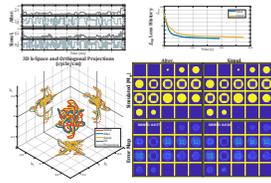


Fig. S6. Comparison of alternating (Alter.) and Simultaneous (Simul.) minimization. The two approaches find similar but different local minima according to the RF and k-space plots (Left Panel). On the Top Right, the loss for the first 40 iterations is plotted. The computation time for each iteration is slightly longer for the simultaneous L-BFGS updates. The simultaneous updating scheme converges slightly slower, while the eventual excitation performance of the two schemes is comparable (Lower Right).

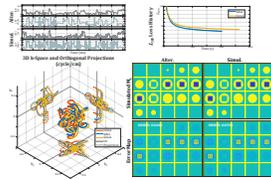


Fig. S7. Same comparison as in Fig. S6, for the IV180 design. The simultaneous approach attains a slightly better inversion, but ends with a higher loss value (likely due to the RF power penalization term).

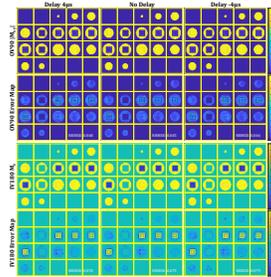


Fig. S8. Simulated excitation performances of our designed OV90 and IV180 pulses under different delays of the applied gradient fields.

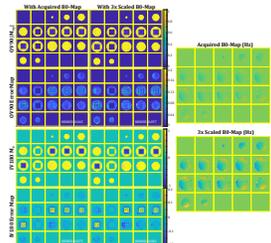


Fig. S9. Simulated excitation performances of our designed OV90 and IV180 pulses under different off-resonance maps.