

Optimal control gradient precision trade-offs: Application to fast generation of DeepControl libraries for MRI

Mads Sloth Vinding^{a,*}, David L. Goodwin^{b,c,*}, Ilya Kuprov^d, Torben Ellegaard Lund^a

^a Center of Functionally Integrative Neuroscience (CFIN), Department of Clinical Medicine, Faculty of Health, Aarhus University, Denmark

^b Institute for Biological Interfaces 4 - Magnetic Resonance, Karlsruhe Institute for Technology (KIT), Karlsruhe, Germany

^c Chemistry Research Laboratory, University of Oxford, Mansfield Road, Oxford, UK

^d School of Chemistry, University of Southampton, Southampton SO17 1BJ, UK



ARTICLE INFO

Article history:

Received 9 July 2021

Revised 21 September 2021

Accepted 19 October 2021

Available online 27 October 2021

Keywords:

Optimal Control

Optimization Gradients

MRI

DeepControl

2D RF

ABSTRACT

We have recently demonstrated supervised deep learning methods for rapid generation of radiofrequency pulses in magnetic resonance imaging (<https://doi.org/10.1002/mrm.27740>, <https://doi.org/10.1002/mrm.28667>). Unlike the previous iterative optimization approaches, deep learning methods generate a pulse using a fixed number of floating-point operations - this is important in MRI, where patient-specific pulses preferably must be produced in real time. However, deep learning requires vast training libraries, which must be generated using the traditional methods, e.g., iterative quantum optimal control methods. Those methods are usually variations of gradient descent, and the calculation of the gradient of the performance metric with respect to the pulse waveform can be the most numerically intensive step. In this communication, we explore various ways in which the calculation of gradients in quantum optimal control theory may be accelerated. Four optimization avenues are explored: truncated commutator series expansions at zeroth and first order, a novel midpoint truncation scheme at first order, and the exact complex-step method. For the spin systems relevant to MRI, the first-order midpoint truncation is found to be sufficiently accurate, but also significantly faster than the machine precision gradient. This makes the generation of training databases for the machine learning methods considerably more realistic.

© 2021 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Optimal control theory (OCT) in magnetic resonance started at a similar time to hardware implementation of arbitrarily shaped pulses [1,2]. Numerical optimization was used to design highly selective excitation pulses [3] and population inversion pulses [4]. The early successes of using non-rectangular pulses led researchers to explore the potential of OCT. Selective excitation problems using numerically designed 90° and 180° pulses used OCT to solve Bloch equations with arbitrary pulse shapes [5].

One of the first successful applications of OCT in nuclear magnetic resonance spectroscopy (NMR) used an existing deuterium decoupling pulse sequence [6] as an initial guess to a gradient-based numerical optimization algorithm [7]. The unitary bounds on general NMR polarization transfer experiments were shown to

be much larger than the apparent limits of NMR polarization transfer using state-of-the-art experiments of that time [8,9]. Pulses with the shortest possible time were found for coherence transfer experiments in heteronuclear two and three spin systems [10,11]. The search for solutions nearer to these bounds became the work of the following decade.

In the context of magnetic resonance imaging (MRI), tailored radiofrequency (RF) pulses are used in advanced applications: reduced field-of-view imaging [12–14], spectral-spatial selectivity [12,15], as well as imaging with inhomogeneous RF fields (B_1^+) [16], off-resonance effects [17–20] or gradient imperfections [21]. Among several pulse design methods, e.g., small-tip-angle [22,23], Shinnar-LeRoux [24], magnitude least-squares [25], OCT is well suited for generating the necessary RF pulse shapes, including multi-channel pulses, several milliseconds long with microsecond time stepping and arbitrary flip angles [15,26–31].

One method, called gradient ascent pulse engineering (GRAPE), uses a piecewise-constant pulse approximation that is well suited for MRI [32]. Here, the term gradient is in relation to optimization and not to be confused with the magnetic field gradients. GRAPE

* Corresponding authors.

E-mail addresses: msv@cfm.au.dk (M.S. Vinding), david.goodwin@partner.kit.edu (D.L. Goodwin).

¹ These two authors contributed equally.

was initially designed using the gradient ascent method (linear convergence), and then extended to quasi-Newton (super-linear convergence) methods [33], and Newton–Raphson (quadratic convergence) methods [34]. The latter two require very accurate gradients [33–36]. Importantly, tailored RF pulses in MRI often need to be computed in real time – the patient is inside the scanner. However, GRAPE uses iterative optimization that has no wall clock time guarantees: for realistic systems, it can become impractically slow.

We have recently proposed a neural network based method for generating RF pulses to alleviate the run time problem [37–39]. This *DeepControl* framework requires training libraries with hundreds of thousands of RF pulses generated, e.g., with OCT, but a trained neural network predicts a pulse quickly, and with a hard guarantee on the wall clock time.

These vast libraries require an improvement in our ability to run OCT simulations and calls for optimization of the numerical efficiency of the GRAPE algorithm, in particular, in the part that deals with the gradient calculation.

In this paper, we report some efficiency improvements in the GRAPE gradient calculation process. We analyze the trade-offs between CPU time and convergence rate for four different levels of approximation to the GRAPE gradient: standard zero and first order approximation, a novel midpoint first order approximation proposed here, and the exact, complex-step gradient.

A note should be made here, in that this work is presented here in the context of MRI, but the core of this work can also be extended to the general area of OCT for pulsed control systems using the piecewise constant pulse approximation, particularly for areas using magnetic field gradients [40].

2. Optimal control theory

The equation of motion of a system of non-interacting spin-1/2 particles, ignoring relaxation, can be formulated as:

$$\dot{\mathbf{M}}(\mathbf{r}, t) = \mathbf{\Omega}(\mathbf{r}, t)\mathbf{M}(\mathbf{r}, t) \quad (1)$$

where $\dot{\mathbf{M}}(\mathbf{r}, t)$ is the time-derivative of $\mathbf{M}(\mathbf{r}, t)$, the instantaneous magnetization vector at a spatial location $\mathbf{r}$, and $\mathbf{\Omega}(\mathbf{r}, t)$ is the dynamics matrix containing external magnetic fields, defined below. By the nature of this study, we introduce discrete notation from this point. The spatial dimension is separated into P different locations, with $1 < p < P$, and for the p^{th} location $\mathbf{r}^{(p)} = [x^{(p)}, y^{(p)}, z^{(p)}]^T$ (T is the vector transpose operation).

The control of these dynamics is in the form of a piecewise-constant *control waveform* with a fixed total duration T , divided into N time steps of duration $\Delta t = T/N$, called a time slice. The magnetization vector we are attempting to control, $\mathbf{M}_n^{(p)} = [M_{x,n}^{(p)}, M_{y,n}^{(p)}, M_{z,n}^{(p)}]^T$, will have $N+1$ temporal instances for each spatial position, with an initial state $\mathbf{M}_0^{(p)}$ and a final state $\mathbf{M}_N^{(p)}$. Hence, for state vectors $0 < n < N$, the n^{th} control propagates the system from the $(n-1)^{\text{th}}$ time point to the n^{th} time point. The task for an OCT problem is to bring this final state as close as possible to a desired target state by maximizing the overlap of the target state, $\mathbf{M}_{\text{target}}^{(p)}$, and the final state, $\mathbf{M}_N^{(p)}$.

Dynamics of Eq. (1) are governed by $\mathbf{\Omega}$, and the piecewise-constant formulation of this is

$$\mathbf{\Omega}_n^{(p)} = \gamma \begin{bmatrix} 0 & +\mathbf{G}_n^T \mathbf{r}^{(p)} + B_z^{(p)} & -B_{y,n}^{(p)} \\ -\mathbf{G}_n^T \mathbf{r}^{(p)} - B_z^{(p)} & 0 & +B_{x,n}^{(p)} \\ +B_{y,n}^{(p)} & -B_{x,n}^{(p)} & 0 \end{bmatrix}, \quad (2)$$

which includes the RF fields in the Zeeman rotating frame, $B_{x,n}^{(p)}$ and $B_{y,n}^{(p)}$, the magnetic field gradient, $\mathbf{G}_n = [G_{x,n}, G_{y,n}, G_{z,n}]^T$, the magnetic field inhomogeneity, $B_z^{(p)}$, and the gyromagnetic ratio, γ .

Xu et al. [30] have shown how to extend these equations to multiple RF channels, $1 < l < N_{\text{Tx}}$, i.e., parallel transmit (pTx) with complex coil sensitivity patterns, $s_l^{(p)}$:

$$B_{x,n}^{(p)} = \sum_{l=1}^{N_{\text{Tx}}} [\text{Re}(s_l^{(p)})c_{u,l,n} - \text{Im}(s_l^{(p)})c_{v,l,n}] \quad (3)$$

$$B_{y,n}^{(p)} = \sum_{l=1}^{N_{\text{Tx}}} [\text{Im}(s_l^{(p)})c_{u,l,n} + \text{Re}(s_l^{(p)})c_{v,l,n}] \quad (4)$$

Here, $c_{j,l,n} \in \{c_{u,l,n}, c_{v,l,n}\}$ are the RF control sequences to be optimized to achieve the aim, and all other parameters are considered fixed. For generality, we introduce the control $c_{j,l,n}$, with the subscript j as reference to either u or v , when it is possible to imply to either.

Discrete solutions to Eq. (1), in the piecewise-constant approximation, can be written as

$$\mathbf{M}_n^{(p)} = \mathbf{U}_n^{(p)} \mathbf{M}_{n-1}^{(p)} \quad (5)$$

where $\mathbf{M}_0^{(p)}$ is a given initial state of the system and $\mathbf{U}_n^{(p)}$ is a time-propagator describing a rotation on the Bloch sphere, which can be calculated with the matrix exponential:

$$\mathbf{U}_n^{(p)} = e^{\mathbf{\Omega}_n^{(p)} \Delta t} \quad (6)$$

Note, the magnetic field gradient waveform is specified for the entire pulse duration and kept fixed, although it could also be included as a control to be optimized.

Our performance functional, describing our aim, is the projection of $\mathbf{M}_N^{(p)}$, onto the target state $\mathbf{M}_{\text{target}}^{(p)}$, summed and normalized over all spatial points:

$$J = \frac{1}{P} \sum_{p=1}^P \mathbf{M}_N^{(p)\top} \cdot \mathbf{M}_{\text{target}}^{(p)} \quad (7)$$

Thus, J is the quantity we want to maximize and is termed the *fidelity*. Further to this fidelity measure, which is dependent on only the final state of the system and termed the *terminal cost*, an additional term can be included, which is termed a *running cost* and depends on the state of the system over all of the pulse duration. The running cost can be regularization terms, to penalize untargeted control behavior, e.g., excess power or jagged waveforms [34], or boundary constraints. In this work we use simple boundary constraints, which are detailed in [29].

3. Fidelity gradient calculations

The formal theory of OCT introduces an *adjoint state* of the system, $\mathbf{L}_n^{(p)}$, with the detailed derivation published elsewhere [41]. This adjoint state of the system can be interpreted as the propagation of the target, backwards in time from $t = t_N$ to $t = t_0$, and for a particular time slice this is

$$\mathbf{L}_{n-1}^{(p)} = \mathbf{U}_n^{(p)\top} \mathbf{L}_n^{(p)} \quad (8)$$

where $\mathbf{L}_N^{(p)} = \mathbf{M}_{\text{target}}^{(p)}$. The adjoint state is used in calculation of the fidelity gradients, $\nabla J_{u,l,n}$ and $\nabla J_{v,l,n}$, with respect to $c_{u,l,n}$ and $c_{v,l,n}$, enabling the use of gradient-following numerical optimization methods to maximize J .

The scope of this study is to cast light on a number of different strategies for calculating $\nabla J_{u,l,n}$ and $\nabla J_{v,l,n}$. Common prerequisites for the different gradient strategies are a forward time-

propagation of the magnetization from the initial to the final state for all positions by Eq. (5) and a backward time-propagation of the adjoint state for all positions by Eq. (8).

3.1. Standard approximate gradients

Derivatives of functions of matrices, $f(\mathbf{\Omega})$, are called *directional derivatives*, or *Gâteaux derivatives*, which are defined in Ref. [42] as

$$D_{\Theta}(f(\mathbf{\Omega})) \triangleq \lim_{h \rightarrow 0} \frac{f(\mathbf{\Omega} + h\mathbf{\Theta}) - f(\mathbf{\Omega})}{h} = \left. \frac{d}{dh} \right|_{h=0} f(\mathbf{\Omega} + h\mathbf{\Theta}) \quad (9)$$

where $\mathbf{\Theta}$ is the operator matrix for the control in question and elaborated below.

Eq. (9) is a similar form to the finite difference equation and the Taylor series truncated to first order, and should be considered the formal definition of the derivative of the function of a matrix.

For the dynamic system of Eq. (6) we are considering here, the *directional derivative* can be written as [43]

$$D_{\Theta}(e^{\mathbf{\Omega}\Delta t}) = \left[\int_0^{\Delta t} e^{\mathbf{\Omega}s} \mathbf{\Theta} e^{-\mathbf{\Omega}s} ds \right] e^{\mathbf{\Omega}\Delta t} \quad (10)$$

$$= \left\{ \mathbf{\Theta}, \int_0^{\Delta t} e^{\mathbf{\Omega}s} ds \right\} e^{\mathbf{\Omega}\Delta t} = \left\{ \mathbf{\Theta}, \frac{e^{\mathbf{\Omega}\Delta t} - \mathbb{I}}{\mathbf{\Omega}} \right\} e^{\mathbf{\Omega}\Delta t} = \left[\sum_{r=0}^{\infty} \frac{(\Delta t)^{r+1}}{(r+1)!} \left\{ \mathbf{\Theta}, \mathbf{\Omega}^r \right\} \right] e^{\mathbf{\Omega}\Delta t} \quad (11)$$

where $\{\mathbf{A}, \mathbf{B}\}$ denote the anti-commutator of the matrices $\mathbf{A}$ and $\mathbf{B}$. The first order approximation, $r = 1$, corresponds to Eq. (12) in the original GRAPE paper [32]. For the controls $c_{u,l,n}$ and $c_{v,l,n}$ [30], the corresponding control operators are

$$\mathbf{\Theta}_{u,l}^{(p)} = \begin{bmatrix} 0 & 0 & -\text{Im}(s_l^{(p)}) \\ 0 & 0 & +\text{Re}(s_l^{(p)}) \\ +\text{Im}(s_l^{(p)}) & -\text{Re}(s_l^{(p)}) & 0 \end{bmatrix} \quad (12)$$

$$\mathbf{\Theta}_{v,l}^{(p)} = \begin{bmatrix} 0 & 0 & -\text{Re}(s_l^{(p)}) \\ 0 & 0 & -\text{Im}(s_l^{(p)}) \\ +\text{Re}(s_l^{(p)}) & +\text{Im}(s_l^{(p)}) & 0 \end{bmatrix} \quad (13)$$

In addition, if we optimized the magnetic field gradient waveform, e.g., $G_{g,n}$ with $g \in \{x, y, z\}$, its control operator is

$$\mathbf{\Theta}_{G_g}^{(p)} = \begin{bmatrix} 0 & +g^{(p)} & 0 \\ -g^{(p)} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, g \in \{x, y, z\} \quad (14)$$

In our vectorized model, Eqs. (2), (4)–(8), we consider the first two standard gradients approximations denoted STD-0 and STD-1 precise to zeroth and first order with $r = 0$ and $r = 1$ in Eq. (11), respectively:

$$\nabla J_{j,l,n}^{(\text{STD-0})} = \frac{\Delta t}{P} \sum_{p=1}^P \mathbf{L}_n^{(p)\top} [\mathbf{\Theta}_{j,l}^{(p)}] \mathbf{M}_{n-1}^{(p)} + \mathcal{O}(\Delta t) \quad (15)$$

$$\nabla J_{j,l,n}^{(\text{STD-1})} = \frac{\Delta t}{P} \sum_{p=1}^P \mathbf{L}_n^{(p)\top} [\mathbf{\Theta}_{j,l}^{(p)} \mathbf{U}_n^{(p)}] \mathbf{M}_{n-1}^{(p)} + \mathcal{O}(\Delta t^2) \quad (16)$$

where the $\mathcal{O}$ -notation quantifies the error, which depends mainly on the size of the time step [33].

Had we included regularization in Eq. (7) this would also be present in the gradient terms above [34].

The zeroth and first-order standard gradient approximations are widely used and fast to compute [44], but for a quasi-Newton method they would quickly corrupt the Hessian estimation and slow the super-linear convergence to linear, especially for long time increments [33]. De Fouquieres et al. [33] show how gradients computed to increasing precision improve the performance of quasi-Newton methods.

3.2. Midpoint approximate gradients

In this paper, we propose a *midpoint* variant of the standard approximate gradient to partly overcome the convergence problem of the standard approximate gradients. The midpoint variant was inspired by the central finite-difference gradient, although it has a deep theoretical basis in the *exponential midpoint method* [45–47,43,48].

The midpoint method for this manuscript is

$$\nabla J_{j,l,n}^{(\text{MID})} = \frac{\Delta t}{2P} \sum_{p=1}^P \mathbf{L}_n^{(p)\top} [\mathbf{U}_{n-1}^{(p)} \mathbf{\Theta}_{j,l}^{(p)} + \mathbf{\Theta}_{j,l}^{(p)} \mathbf{U}_n^{(p)}] \mathbf{M}_{n-1}^{(p)} + \mathcal{O}(\Delta t^2) \quad (17)$$

which is an average of the $(n-1)^{\text{th}}$ and n^{th} gradient elements; centering the derivative to the midpoint of the $(n-1)^{\text{th}}$ and n^{th} time point. The propagator at t_0 is the identity matrix so, when expanded, $\mathbf{L}_1^{(p)\top} \mathbf{U}_0^{(p)} \mathbf{\Theta}_{j,l}^{(p)} \mathbf{M}_0^{(p)} = \mathbf{L}_1^{(p)\top} \mathbf{\Theta}_{j,l}^{(p)} \mathbf{M}_0^{(p)}$.

Although this may seem a simple extension to Eq. (16) without obvious benefit, since both have an error $\mathcal{O}(\Delta t^2)$, a deeper read to the literature shows the $\mathcal{O}(\Delta t^2)$ error term in Eq. (16) has a constant hidden in the $\mathcal{O}$ -notation. The *hidden constant* depends on the time derivatives of $\mathbf{M}$ and the bounds of the $\mathbf{\Omega}$ [46–48]. The midpoint method in Eq. (17) does not suffer from this error dependency [45] and can be important, when space discretization is used [43]. It should be emphasized that the exponential midpoint method [45–47,43,48] is designed for approximation of time propagators, which is not the case in this study; time propagators in this study are exact, and the derivatives of time propagators, $\mathbf{D}_{j,l,n}^{(p)}$ defined by Eq. (11), are approximated using a midpoint method.

3.3. Exact Gradient

Floether et al. [49] presented a method to compute exact control gradients using an auxiliary matrix approach [50], and Goodwin et al. [35] extended this to the second order derivatives needed for a Hessian.

We recently adopted the exact gradient calculation in our OCT framework, and described it in Ref. [29]. Briefly,

$$\nabla J_{j,l,n}^{(\text{EXACT})} = \frac{1}{P} \sum_{p=1}^P \mathbf{L}_n^{(p)\top} [\mathbf{D}_{j,l,n}^{(p)}] \mathbf{M}_{n-1}^{(p)} \quad (18)$$

The gradient in Eq. (18) is exact to machine precision [35]. The essence of this exact approach is the derivation of $\mathbf{D}_{j,l,n}^{(p)}$ in Eq. (18). This can be obtained from the upper right corner of an auxiliary matrix [50,49,35] but for real matrices $\mathbf{\Omega}_n^{(p)}$ and $\mathbf{\Theta}_{j,l}^{(p)}$, the compact method of *complex-step approximation* [51] can be used to calculate the directional-derivative, $\mathbf{D}_{j,l,n}^{(p)}$. This method gives the propagator and derivative as the real and imaginary parts of a complex matrix, respectively:

$$\mathbf{U}_n^{(p)} + i\mathbf{D}_{j,l,n}^{(p)} = \exp \left\{ \left[\mathbf{\Omega}_n^{(p)} + i\mathbf{\Theta}_{j,l}^{(p)} \right] \Delta t \right\}, \quad (19)$$

where i is the imaginary unit, and $\Omega_n^{(p)}$ and $\Theta_{j,l}^{(p)}$ must be real matrices, which can be the case in the single-spin model.

Considering any one given spatial point, p , the complex-step differentiation involves 3×3 complex matrices instead of 6×6 real matrices. With that, there is potential to improve memory allocation. Complex-step differentiation is also known to have better roundoff error tolerance in finite precision arithmetic [51].

3.4. Accuracy tests

The relative accuracy of the gradients calculated with the methods in Eq. (15), Eq. (16), and Eq. (17), is shown in Fig. 1. The error in the accuracy is calculated with

$$\xi = \frac{\|\nabla J - \nabla J^{(\text{EXACT})}\|_2}{\|\nabla J^{(\text{EXACT})}\|_2} \quad (20)$$

and averaged over 20 instances of random pulses. Each pulse has a duration of $T = 6.39$ ms, with the random distribution between ± 1 kHz for RF controls, c_u and c_v , and ± 160 kHz for magnetic field gradient controls, c_G . The tests were done firstly (top, Fig. 1) with time slice, Δt , ranging logarithmically in 20 steps from 63.9 ns to 0.639 ms for an offset of $B_z^{(p)} = 0$ Hz. These limits were chosen to extend the analysis to time scales also relevant to NMR [40]. Secondly (bottom, Fig. 1), the offset was picked randomly in a bandwidth of 160 kHz for a fixed $\Delta t = 10$ μ s. The number of offset points is approximately one point per 2.1 kHz of bandwidth, and they are spread randomly within the given bandwidth. The initial state of the system is all spins having positive z -magnetization, and the target state is all spins having positive x -magnetization i.e. a 90° rotation around the y -axis of the Bloch sphere.




Fig. 1. Relative accuracy of approximate gradients compared to the exact gradient in Eq. (18). Both plots show the relative accuracy of the standard zeroth order (STD-0), Eq. (15), first order (STD-1), Eq. (16), and first order midpoint (MID), Eq. (17), gradients. (Top) Errors as a function of time slice, Δt , with the dashed vertical line being the Δt used in the study. (Bottom) Errors as a function of offset for $\Delta t = 10$ μ s, corresponding to the dashed vertical line in the top plot. Both plots show the error, when field gradient controls (c_G) are used together with RF controls (c_u, c_v) with dotted lines, and when it is just RF controls (as in this study) with full lines.

The accuracy comparison in Fig. 1 shows that the MID gradient gives an increase in accuracy, even though the method of calculation in Eq. (17) requires only the insignificant computational cost of a matrix addition per gradient element. The dashed vertical line in the top panel of Fig. 1 indicates the time slice, $\Delta t = 10$ μ s, used in the results that follow. This time slice is chosen as it is physically relevant to the MRI hardware, but also from approximately this time slice the MID gradient becomes increasingly more effective than the STD-0 and STD-1 gradients for decreasing time slices. The pulse duration, bandwidths, and control amplitudes were based on the settings described in the next section.

4. Methods

An OCT framework was implemented in MATLAB (Mathworks, Natick, MA, USA), with hard constraints on the RF pulse amplitudes using the `fmincon` optimization function supplemented with a performance gradient function and the necessary housekeeping procedures. We focus here on the gradients required by the quasi-Newton option of `fmincon`, although we are not limited to GRAPE.

The four gradients types have been implemented [52] in the Bloch framework used in Refs. [13,41,15,28,29,37,38].

A common feature for these calculations is that all magnetization and adjoint states must be known at a given time slice, n , and/or at the adjacent time slice, $n - 1$. Conducting full magnetization forward propagation, and adjoint state backward propagation solves this issue. In principle, however, the magnetization vector can be calculated as gradient element calculations proceed, i.e., $\mathbf{U}_n^{(p)}$ can also be extracted from Eq. (19), in the same way as propagator recycling implemented for Hessian calculations [53], possibly with efficient matrix caching [34].

The most efficient code for parallelizing the propagation and gradient elements is specific to computing architecture. The two most obvious parameters considered are the amount of CPU cores available, and the storage size needed for all matrices from forward and backward propagation. Our computations were performed on a 28-CPU core, Intel Xeon Gold 5120, 2.2 GHz workstation with 384 GB of RAM.

The EXACT gradient of Eq. (19) was assessed regarding parallelization efficiency. An L-curve analysis is shown in Figure S1 in the Supplementary Material. The EXACT gradient was on the present workstation most efficient with parallelization on around 23 out of the 28 available workers. Accordingly, the reported computation times for the EXACT gradient are with 23-fold parallelization. The STD-0, STD-1 and MID gradients are efficient without parallelization, and we report computation times for the fastest (non-parallelized) implementations we could develop.

We evaluated optimization performance with the four different gradients by computing single-channel 2DRF pulses with a library of total 150 different flip-angle (FA) maps of three different target categories:

1. 50 binary images, where ten randomly placed points were dilated and merged to a single binary shape with no holes. These are denoted BW (black-white), see Figure S2.
2. 50 normalized gray-scale images made binary with a 0.5-threshold, denoted GrBW, see Figure S3.
3. 50 normalized gray-scale images, denoted Gr (Gray), see Figure S4.

Images were taken from the *ImageNet* database [54], and the nominal target FA was 30° . The target magnetization x - and z -component levels were dictated by the image intensity levels multiplied with the nominal FA. The target y -component was 0. The

initial magnetization was longitudinal. The Gr category represents our ability to induce spatially variant FAs, while the binary categories are two random ad hoc ways to obtain shaped excitation patterns, mimicking anatomic regions of interest.

The (fixed) gradient waveform accommodating the spatial selectivity of the 2DRF pulses formed an inward 16-turn spiral in the excitation k-space (full Nyquist sampling) with a duration of 6.39 ms, a field of excitation of 25 cm, and limited to 40 mT/m amplitude and 180 T/m/s slew rate [55,26]. To support works like that of Zhang et al. [56], where the gradient waveforms are also predicted along with the RF waveforms, our choice of spiral could be used as an initial starting point in the field gradient optimization, if we activated the corresponding optimization gradients by Eq. (14) etc. Until then, we are limited to a single pulse duration, spatial resolution and field of excitation. The common time slice for RF and field gradients was 10 μ s. RF pulses were constrained below 1 kHz (nutration rate) amplitude using the boundary option implemented in *fmincon*. The spatial grid was 64×64 with a mask formed by a centrally placed super ellipse, which reduced the number of spatial points to 2919. The quasi-Newton method with STD-0, STD-1, MID and EXACT gradients was limited with the algo-

rithm termination conditions of reaching 100 iterations or a minimum functional/step-size convergence tolerance of 1×10^{-6} .

5. Results

For all three categories (BW, GrBW and Gr), all optimizations with approximate gradients converged. The optimizations with EXACT gradients were stopped in 32% of all cases by the iteration limit.

Fig. 2 shows the NRMSE and FA statistics from all Bloch simulations. The cases are sorted by the NRMSEs of the EXACT gradient) within each target category. All FA maps are shown in Figure S2-4 in the Supplementary Material.

Table 1 lists iteration and computation time statistics (mean and standard deviation) for the three target categories (BW, GrBW and Gr) pooled together. Of interest (see the diagonal entries), the MID and STD-1 gradients spend on average roughly the same amount of iterations and time to finish, around 14 iterations in 15 s, respectively. The EXACT method spends on average 74 iterations and nearly 300 s to finish.




Fig. 2. Plots of NRMSE and FA statistics for the different target categories: BW (left column), GrBW (center column), Gr (right column). The case numbers are sorted based on the NRMSE values from the optimization, i.e., the error of the Bloch simulated magnetization with respect to the target magnetization, based on the EXACT gradient results. For the Gr target category, the entire pattern is considered as inside the ROI, thus there are no plots for outside areas. In the same category, the expected mean and standard deviation FA inside the ROI is not expected or supposed to be 30° and 0°, respectively, but may vary according to the gray levels.

Table 1

Iteration and computation times (mean $\pm$ standard deviation). Diagonal entries with two mean $\pm$ standard deviation pairs correspond to total use of a given method, i.e., computation times include gradient and approximate Hessian estimation, step size search, hard-constraints, and all other internal routines of the interior-point algorithm used within MATLAB's `fmincon` function. The upper triangular entries include three mean $\pm$ standard deviation pairs (from top to bottom): 1) the number of iterations it takes for the method in the column to outperform the method in the row considering the NRMSE value; 2) the approximate time it takes the column method to overtake the row method; 3) the approximate time the row method was overtaken by the column method. The approximate times are estimated by the ratio of iterations needed for crossing and the total number of iterations multiplied by the total computations times for all iterations. Target categories (BW, GrBW and Gr) are pooled for the measures in this Table.

Gradient Method	STD-0	STD-1	MID	EXACT	
STD-0	6.5 $\pm$ 0.8	5.0 $\pm$ 0.4	4.9 $\pm$ 0.4	4.8 $\pm$ 0.4	Iterations
	11.9 $\pm$ 0.6	5.6 $\pm$ 1.0	5.7 $\pm$ 0.7	19.8 $\pm$ 1.7	Seconds
		9.3 $\pm$ 1.0	9.1 $\pm$ 0.7	9.0 $\pm$ 1.7	Seconds
STD-1		14.0 $\pm$ 6.3	7.0 $\pm$ 1.1	6.6 $\pm$ 1.1	Iterations
		14.7 $\pm$ 2.9	7.9 $\pm$ 0.9	27.0 $\pm$ 4.5	Seconds
			7.7 $\pm$ 1.2	7.2 $\pm$ 1.1	Seconds
MID			13.1 $\pm$ 2.5	8.1 $\pm$ 1.5	Iterations
			14.9 $\pm$ 1.6	33.1 $\pm$ 5.9	Seconds
				9.2 $\pm$ 1.1	Seconds
EXACT				73.5 $\pm$ 25.4	Iterations
				303.8 $\pm$ 109.7	Seconds




Fig. 3. FA maps of the best (left), closest-to-mean (middle) and worst (right) cases, when contrasting the NRMSE of the FA maps of the MID gradient with respect to the EXACT gradient. The corresponding STD-0 and STD-1 results are shown as well. Differences are shown in the lower part. Only pulses from converged solutions were considered for this Figure.




Fig. 4. NRMSE histograms contrasting the MID gradient against the EXACT gradient, for the BW (top), GrBW (middle) and Gr (bottom) target categories. The best, closest-to-mean and worst case examples shown in Fig. 3 are herein signified by the dashed lines. The histograms contain all cases, i.e., also pulses terminated by reaching the maximum number of iterations, why the histogram edges extend below and beyond the best and worst case levels, respectively.

In comparison (see the off diagonal entries), the MID gradient is not outperformed by the EXACT gradient until it reaches 8 iterations or 33 s, which is equivalent to around 9 s for the MID gradient. However, as shown in Fig. 2 the MID and EXACT gradients have near identical performance. This is supported FA maps displayed in

Fig. 3. In Figures S5-S13 in the Supplementary Material, we show the pulses corresponding to the FA maps of Fig. 3.

Judged from the NRMSE of FA maps (MID against EXACT), we compared the best, closest-to-mean, and worst cases for converged pulses. As reference, we also show the results from the corresponding STD-0 and STD-1 gradients, as well as the difference maps, highlighting the increasing errors using less precise gradients.

Fig. 4 shows the NRMSE histograms of all pulses, where the right skew reveals that the mean NRMSE is close to the best cases, and the worst cases are rare.

6. Discussion

We have presented a midpoint first order gradient approximation, denoted MID, with a similar computation time to a standard first order gradient approximation computation, and the accuracy similar to the exact, complex-step gradient for typical MRI settings.

The average time per iteration for the present settings increase by a factor of around 1.01 going from the standard first order to the midpoint gradient method, however by a factor of five going from the standard first order to the exact, complex-step gradient bearing in mind the exact, complex-step gradient exploit 23-fold parallelization. With four CPU cores, typically available on a laptop, we estimate the complex-step gradient will further require a factor of two of computation time.

The standard and the midpoint gradients are efficient without parallelization, and easily vectorized. Several designs for the approximate gradients have been tested for the RF pulse design. However, we could not find any implementation, where either gradient becomes faster from parallelization for the present application: the bottleneck of transferring data to/from the workers dominates the workload. For single-channel RF optimization and the spatial grid size (P) and number of time-steps (N) as reported above, the standard gradients are slightly faster being run in a for-loop over time steps, rather than as a more vectorized version. There was no significant difference between the for-loop and vectorized versions of the midpoint gradient. As a side-note, when running multi-channel RF optimization (not the case herein), both the standard and midpoint gradients benefit from the vectorized version and the time difference between them diminishes. However, there remains a for-loop over the number of RF channels currently. Accordingly, approximate gradient computation times were reported herein for the fastest, non-parallelized implementations we have found.

Optimization speed is inherently important for in vivo experiments, but also for generating AI training libraries. Considering the vast training libraries required for *DeepControl* [37,38], the non-parallelized, vectorized computation of the standard and midpoint approximate gradients is not a problem, when we exploit all available CPU cores with an *outer parallelization* over many independent pulses.

It is beyond the scope of this work to investigate how the qualities of the different gradients influence the *DeepControl* framework, but we do expect infidelities of a given library to propagate through to the final trained neural network, and that predicted pulses never perform better than the representative training library. This study leaves the question of what is more important: fast computation time or pulse perfection, as an additional option for the user. We also propose the midpoint gradient method for rapidly producing a good initial guess that may be handed over to a slower, exact gradient method for finalizing.

While it is possible to further improve the accuracy of the standard or midpoint approximate gradient by shortening the time slice, this may not be feasible on realistic hardware. The accuracy of the midpoint gradient corresponds roughly to that of the stan-

standard 0th order gradient, STD-0, with a halved time slice. This was found by running the same experiments (data not shown) with the STD-0 for a time slice of 5 μ s instead of 10 μ s.

The choice for target categories reflect our previous [37–39] and future *DeepControl* experiments, which will be described in a subsequent publication. Of the three target categories, the GrBW category stood out with higher NRMSEs than the other two. Considering just binary targets (BW and GrBW), as shown by Figures S2–3 of FA maps sorted after NRMSE, we observe that smaller targets yield the lowest NRMSEs, and the larger targets with more jagged edges result in higher NRMSEs. This trend is consistent for all gradient methods, see Fig. 2.

On field-gradient optimization, we have included examples of multiband pulse optimizations with both the midpoint and exact gradients in the Supplementary Material [57]. We contrasted the gradient computation times of RF-only optimizations, i.e. the standard slice selective gradient was fixed, against RF plus field-gradient optimizations. We found that the midpoint method including both RF and a field gradient in the optimization, was 21 to 28 times faster than the exact method without and with the field gradient also being optimized, respectively. That was with 23-fold CPU parallelization in the exact method and single thread computation in the midpoint method. Hence, on field-gradient optimization, e.g., to support works like that of Zhang et al. [56], the midpoint gradient method should still be attractive.

We acknowledge that routines operating in the so-called small-tip-angle regime pose a robust alternative to OCT in terms of speed and accuracy due to the linear Fourier relation existing between the pulse waveform and the excitation pattern. As shown in Ref. [37], we trained networks for both small and large flip angles, with Fourier and OCT based algorithms generating the libraries. However, we note that the OCT framework we use enables arbitrary flip angles and phases for each individual spatial and spectral position. It is due to *DeepControl* experiments that we chose in this study to target a nominal FA of 30° (in the small-flip-angle regime) together with OCT, which was not strictly necessary for this FA, but chosen for several other reasons, one being because our *DeepControl* experiments benefit from hard pulse constraints, which we have so far only implemented in the OCT framework. Yet hard constraints do exist in various other pulse designs [21,58]. It is important to mention that the individual gradient computation times do not change for other flip angles.

Further to the benefits outlined for the midpoint method, Eqs. (15) and (16) are good approximations only when solutions to the Bloch equations are mildly oscillating [59]. For MRI, mild oscillations are not to be expected, as spatial discretization and in particular field gradient controls introduce highly oscillating parts to the Bloch equations.

7. Conclusion

Pulse waveform optimization gradient calculation using a midpoint first order approximation was evaluated and found to provide a significant efficiency improvement. This illustrates clearly that the current trend of using exact gradients in all situations should be revised – tailored approximations are much cheaper and can be just as accurate. Approximate gradients have been observed to make quasi-Newton pseudo-Hessian less efficient at accelerating convergence, but we have not encountered this phenomenon here, likely because the terminal fidelity requirements were less stringent than they are in quantum technology applications.

The midpoint gradient method has been tested herein for 2DRF pulses, with the intent to create a large 2DRF neural network training library. Nevertheless, the midpoint and the other gradient

methods discussed herein are general for the OCT setting presented, and therefore also applicable to other target shapes and pulse types, e.g. multidimensional parallel transmit pulses of any target flip angle [29], and other control types, e.g. field gradients and novel multi-channel shim waveforms [60]. These applications have a large number of channels, and would benefit significantly from the fast gradient computation and the relatively high accuracy of the midpoint approximate gradient method outlined in this paper.

Data Availability

The data that supports the findings of this study are available within the article and its Supplementary Material. Our code repository at <https://github.com/madsakre/GradientPrecision> will contain the proposed gradient implementations including capabilities for parallel transmit.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

MSV thanks Villum Fonden, Eva og Henry Fraenkels Mindefond, Harboefonden, and Kong Christian Den Tiendes Fond. DLG thanks Burkhard Luy for his generous support and advice.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.jmr.2021.107094>.

References

- [1] M.S. Silver, R.I. Joseph, D.I. Hoult, Highly selective $\pi/2$ and π pulse generation, *Journal of Magnetic Resonance* (1969) 59 (2) (1984) 347–351, [https://doi.org/10.1016/0022-2364\(84\)90181-1](https://doi.org/10.1016/0022-2364(84)90181-1).
- [2] A.J. Temps, C.F. Brewer, Synthesis of arbitrary frequency domain transmitting pulses applicable to pulsed NMR instruments, *Journal of Magnetic Resonance* (1969) 56 (3) (1984) 355–372, [https://doi.org/10.1016/0022-2364\(84\)90300-7](https://doi.org/10.1016/0022-2364(84)90300-7).
- [3] D.J. Lurie, A systematic design procedure for selective pulses in NMR imaging, *Magn. Reson. Imaging* 3 (3) (1985) 235–243, [https://doi.org/10.1016/0730-725x\(85\)90352-2](https://doi.org/10.1016/0730-725x(85)90352-2).
- [4] W.S. Warren, Effects of arbitrary laser or NMR pulse shapes on population inversion and coherence, *J. Chem. Phys.* 81 (12) (1984) 5437–5448, <https://doi.org/10.1063/1.447644>.
- [5] S. Conolly, D. Nishimura, A. Macovski, Optimal control solutions to the magnetic resonance selective excitation problem, *IEEE Trans. Med. Imaging* 5 (2) (1986) 106–115, <https://doi.org/10.1109/tmi.1986.4307754>.
- [6] K.V. Schenker, D. Suter, A. Pines, Broadband heteronuclear decoupling in the presence of homonuclear dipolar and quadrupolar interactions, *Journal of Magnetic Resonance* (1969) 73 (1) (1987) 99–113, [https://doi.org/10.1016/0022-2364\(87\)90228-9](https://doi.org/10.1016/0022-2364(87)90228-9).
- [7] T.O. Levante, T. Bremi, R.R. Ernst, Pulse-sequence optimization with analytical derivatives. Application to deuterium decoupling in oriented phases, *Journal of Magnetic Resonance, Series A* 121 (2) (1996) 167–177, <https://doi.org/10.1006/jmra.1996.0157>.
- [8] S.J. Glaser, T. Schulte-Herbrüggen, M. Sieveking, O. Schedletsky, N.C. Nielsen, O. W. Sørensen, C. Griesinger, Unitary control in quantum ensembles: Maximizing signal intensity in coherent spectroscopy, *Science* 280 (5362) (1998) 421–424, <https://doi.org/10.1126/science.280.5362.421>.
- [9] T.S. Untidt, S.J. Glaser, C. Griesinger, N.C. Nielsen, Unitary bounds and controllability of quantum evolution in NMR spectroscopy, *Mol. Phys.* 96 (12) (1999) 1739–1744, <https://doi.org/10.1080/00268979909483117>.
- [10] N. Khaneja, R. Brockett, S.J. Glaser, Time optimal control in spin systems, *Phys. Rev. A* 63 (3) (2001) 032308, <https://doi.org/10.1103/physreva.63.032308>.
- [11] N. Khaneja, S.J. Glaser, R. Brockett, Sub-Riemannian geometry and time optimal control of three spin systems: Quantum gates and coherence transfer, *Phys. Rev. A* 65 (3) (2002) 032301, <https://doi.org/10.1103/physreva.65.032301>.

- [12] C. Ma, D. Xu, K.F. King, Z.-P. Liang, Reduced field-of-view excitation using second-order gradients and spatial-spectral radiofrequency pulses, *Magn. Reson. Med.* 69 (2) (2013) 503–508, <https://doi.org/10.1002/mrm.24259>.
- [13] I.I. Maximov, M.S. Vinding, D.H. Tse, N.C. Nielsen, N.J. Shah, Real-time 2D spatially selective MRI experiments: Comparative analysis of optimal control design methods, *J. Magn. Reson.* 254 (2015) 110–120, <https://doi.org/10.1016/j.jmr.2015.03.003>.
- [14] E.U. Saritas, C.H. Cunningham, J.H. Lee, E.T. Han, D.G. Nishimura, DWI of the spinal cord with reduced FOV single-shot EPI, *Magn. Reson. Med.* 60 (2) (2008) 468–473, <https://doi.org/10.1002/mrm.21640>.
- [15] M.S. Vinding, C. Laustsen, I.I. Maximov, L.V. Søgaard, J.H. Ardenkjær-Larsen, N. C. Nielsen, Dynamic nuclear polarization and optimal control spatial-selective 13C MRI and MRS, *J. Magn. Reson.* 227 (2013) 57–61, <https://doi.org/10.1016/j.jmr.2012.12.002>.
- [16] T.E. Skinner, N.I. Gershenson, M. Nimbalkar, S.J. Glaser, Optimal control design of band-selective excitation pulses that accommodate relaxation and RF inhomogeneity, *J. Magn. Reson.* 217 (2012) 53–60, <https://doi.org/10.1016/j.jmr.2012.02.007>.
- [17] J.P. Stockmann, L.L. Wald, In vivo B 0 field shimming methods for MRI at 7 T, *NeuroImage* 168 (2018) 71–87, <https://doi.org/10.1016/j.neuroimage.2017.06.013>.
- [18] J.P. Stockmann, T. Witzel, B. Keil, J.R. Polimeni, A. Mareyam, C. LaPierre, K. Setsompop, L.L. Wald, A 32-channel combined RF and B0 shim array for 3T brain imaging, *Magn. Reson. Med.* 75 (1) (2016) 441–451, <https://doi.org/10.1002/mrm.25587>.
- [19] C. Juchem, T.W. Nixon, S. McIntyre, V.O. Boer, D.L. Rothman, R.A. de Graaf, Dynamic multi-coil shimming of the human brain at 7T, *J. Magn. Reson.* 212 (2) (2011) 280–288, <https://doi.org/10.1016/j.jmr.2011.07.005>.
- [20] C. Juchem, T. W. Nixon, S. McIntyre, D. L. Rothman, R. A. de Graaf, Magnetic field homogenization of the human prefrontal cortex with a set of localized electrical coils, *Magnetic Resonance in Medicine* doi:10.1002/mrm.22164.
- [21] C. S. Aigner, A. Rund, S. Abo Seada, A. N. Price, J. V. Hajnal, S. J. Malik, K. Kunisch, R. Stollberger, Time optimal control-based RF pulse design under gradient imperfections, *Magnetic Resonance in Medicine* doi:10.1002/mrm.27955.
- [22] J. Pauly, D. Nishimura, A. Macovski, A k-space analysis of small-tip-angle excitation, *Journal of magnetic resonance* 81 (1989) 43–56.
- [23] C.-y. Yip, J.A. Fessler, D.C. Noll, Iterative RF pulse design for multidimensional, small-tip-angle selective excitation, *Magn. Reson. Med.* 54 (4) (2005) 908–917, <https://doi.org/10.1002/mrm.20631>, ISSN 0740-3194, 1522-2594, URL: <http://doi.wiley.com/10.1002/mrm.20631>.
- [24] J. Pauly, P. Le Roux, D. Nishimura, A. Macovski, Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm (NMR imaging), *IEEE Trans. Med. Imaging* 10 (1) (1991) 53–65, <https://doi.org/10.1109/42.75611>.
- [25] K. Setsompop, L. Wald, V. Alagappan, B. Gagoski, E. Adalsteinsson, Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 Tesla with eight channels, *Magn. Reson. Med.* 59 (4) (2008) 908–915, <https://doi.org/10.1002/mrm.21513>, ISSN 07403194, 15222594, URL: <http://doi.wiley.com/10.1002/mrm.21513>.
- [26] D. Lee, M. Lustig, W.A. Grissom, J.M. Pauly, Time-optimal design for multidimensional and parallel transmit variable-rate selective excitation, *Magn. Reson. Med.* 61 (6) (2009) 1471–1479, <https://doi.org/10.1002/mrm.21950>.
- [27] J. Ulloa, M. Guarini, A. Guesalaga, P. Irarrazaval, Chebyshev Series for Designing RF Pulses Employing an Optimal Control Approach, *IEEE Trans. Med. Imaging* 23 (11) (2004) 1445–1452, <https://doi.org/10.1109/TMI.2004.835602>.
- [28] M.S. Vinding, D. Brenner, D.H.Y. Tse, S. Vellmer, T. Vosegaard, D. Suter, T. Stöcker, I.I. Maximov, Application of the limited-memory quasi-Newton algorithm for multi-dimensional, large flip-angle RF pulses at 7T, *Magn. Reson. Mater. Phys., Biol. Med.* 30 (1) (2017) 29–39, <https://doi.org/10.1007/s10334-016-0580-1>.
- [29] M.S. Vinding, B. Guérin, T. Vosegaard, N.C. Nielsen, Local SAR, global SAR, and power-constrained large-flip-angle pulses with optimal control and virtual observation points, *Magn. Reson. Med.* 77 (1) (2017) 374–384, <https://doi.org/10.1002/mrm.26086>.
- [30] D. Xu, K.F. King, Y. Zhu, G.C. McKinnon, Z.-P. Liang, Designing multichannel, multidimensional, arbitrary flip angle RF pulses using an optimal control approach, *Magn. Reson. Med.* 59 (3) (2008) 547–560, <https://doi.org/10.1002/mrm.21485>.
- [31] W. Grissom, C.-Y. Yip, Z. Zhang, V.A. Stenger, J.A. Fessler, D.C. Noll, Spatial domain method for the design of RF pulses in multicore parallel excitation, *Magn. Reson. Med.* 56 (3) (2006) 620–629, <https://doi.org/10.1002/mrm.20978>.
- [32] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, S.J. Glaser, Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* 172 (2) (2005) 296–305, <https://doi.org/10.1016/j.jmr.2004.11.004>.
- [33] P. de Fouquieres, S. Schirmer, S. Glaser, I. Kuprov, Second order gradient ascent pulse engineering, *J. Magn. Reson.* 212 (2) (2011) 412–417, <https://doi.org/10.1016/j.jmr.2011.07.023>.
- [34] D.L. Goodwin, I. Kuprov, Modified Newton-Raphson GRAPE methods for optimal control of spin systems, *J. Chem. Phys.* 144 (20) (2016) 204107, <https://doi.org/10.1063/1.4949534>.
- [35] D.L. Goodwin, I. Kuprov, Auxiliary matrix formalism for interaction representation transformations, optimal control, and spin relaxation theories, *J. Chem. Phys.* 143 (8) (2015) 084113, <https://doi.org/10.1063/1.4928978>.
- [36] I. Kuprov, C.T. Rodgers, Derivatives of spin dynamics simulations, *J. Chem. Phys.* 131 (23) (2009) 234108, <https://doi.org/10.1063/1.3267086>.
- [37] M.S. Vinding, B. Skyum, R. Sangill, T.E. Lund, Ultrafast (milliseconds), multidimensional RF pulse design with deep learning, *Magn. Reson. Med.* 82 (2) (2019) 586–599, <https://doi.org/10.1002/mrm.27740>.
- [38] M.S. Vinding, C.S. Aigner, S. Schmitter, T.E. Lund, DeepControl: 2DRF pulses facilitating B_1 inhomogeneity and B_0 off-resonance compensation in vivo at 7T, *Magn. Reson. Med.* 85 (6) (2021) 3308–3317.
- [39] M. S. Vinding, C. S. Aigner, J. P. Stockmann, B. Guérin, S. Schmitter, T. E. Lund, DeepControl: AI-powered slice flip-angle homogenization by 2D RF pulses, in: 29th Annual Meeting of International Society of Magnetic Resonance in Medicine, Virtual, 0785, 2021b.
- [40] J.-N. Dumez, Spatial encoding and spatial selection methods in high-resolution NMR spectroscopy, *Progress in nuclear magnetic resonance spectroscopy* 109 (2018) 101–134.
- [41] M.S. Vinding, I.I. Maximov, Z. Tošner, N.C. Nielsen, Fast numerical design of spatial-selective rf pulses in MRI using Krotov and quasi-Newton based optimal control methods, *J. Chem. Phys.* 137 (5) (2012) 054203, <https://doi.org/10.1063/1.4739755>.
- [42] N. J. Higham, Functions of Matrices: Theory and Computation, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, ISBN 978-0-89871-646-7, 2008.
- [43] S. Blanes, F. Casas, J.-A. Oteo, J. Ros, The Magnus expansion and some of its applications, *Physics reports* 470 (5–6) (2009) 151–238, <https://doi.org/10.1016/j.physrep.2008.11.001>.
- [44] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquieres, A. Gruslys, S. Schirmer, T. Schulte-Herbrüggen, Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework, *Physical Review A* 84 (2), doi:10.1103/PhysRevA.84.022305.
- [45] M. Hochbruck, C. Lubich, On Krylov subspace approximations to the matrix exponential operator, *SIAM Journal on Numerical Analysis* 34 (5) (1997) 1911–1925, <https://doi.org/10.1137/S0036142995280572>.
- [46] C. Lubich, Integrators for quantum dynamics: A numerical analyst's brief review, in: J. Grotendorst, D. Marx, A. Muramatsu (Eds.), Quantum Simulation of Complex Many-Body Systems: From Theory to Algorithms, vol. 10 of NIC Series, Jülich, 459–466, 2002.
- [47] M. Hochbruck, C. Lubich, On Magnus integrators for time-dependent Schrödinger equations, *SIAM journal on numerical analysis* 41 (3) (2003) 945–963, <https://doi.org/10.1137/S0036142902403875>.
- [48] M. Hochbruck, A. Ostermann, Exponential integrators, *Acta Numerica* 19 (2010) 209–286, <https://doi.org/10.1017/S0962492910000048>.
- [49] F.F. Floether, P.d. Fouquieres, S.G. Schirmer, Robust quantum gates for open systems via optimal control: Markovian versus non-Markovian dynamics, *New J. Phys.* 14 (7) (2012) 073023, <https://doi.org/10.1088/1367-2630/14/7/073023>.
- [50] C. Loan, Computing integrals involving the matrix exponential, *IEEE Trans. Autom. Control* 23 (3) (1978) 395–404.
- [51] A.H. Al-Mohy, N.J. Higham, The complex step approximation to the Fréchet derivative of a matrix function, *Numerical Algorithms* 53 (1) (2009) 133, <https://doi.org/10.1007/s11075-009-9323-y>.
- [52] M. S. Vinding, code repository: <https://github.com/madsakre>, GitHub, 2021.
- [53] D. L. Goodwin, Advanced Optimal Control Methods for Spin Systems, Ph.D. thesis, University of Southampton, URL: <https://doi.org/10.5258/soton/t0003>, 2017.
- [54] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, *Commun. ACM* 60 (6) (2017) 84–90, <https://doi.org/10.1145/3065386>.
- [55] M. Lustig, S.-J. Kim, J. M. Pauly, A fast method for designing time-optimal gradient waveforms for arbitrary k-space trajectories, *IEEE Transactions on Medical Imaging* 27 (6) (2008) 866–873, ISSN 0278-0062, 1558-254X, doi:10.1109/TMI.2008.922699, URL: <http://ieeexplore.ieee.org/document/4483775/>.
- [56] Y. Zhang, K. Jiang, W. Jiang, N. Wang, A. J. Wright, A. Liu, J. Wang, Multi-task convolutional neural network-based design of radio frequency pulse and the accompanying gradients for magnetic resonance imaging, *NMR in Biomedicine* 34 (2), ISSN 0952-3480, 1099-1492, doi:10.1002/nbm.4443, URL: <https://onlinelibrary.wiley.com/doi/10.1002/nbm.4443>.
- [57] D.J. Larkman, J.V. Hajnal, A.H. Herlihy, G.A. Coutts, I.R. Young, Use of multicore arrays for separation of signal from multiple slices simultaneously excited, *J. Magn. Reson. Imaging* 13 (2001) 313–317.
- [58] A. Hoyos-Idrobo, P. Weiss, A. Massire, A. Amadon, N. Boulant, On Variant Strategies to Solve the Magnitude Least Squares Optimization Problem in Parallel Transmission Pulse Design and Under Strict SAR and Power Constraints, *IEEE Trans. Med. Imaging* 33 (3) (2014) 739–748, <https://doi.org/10.1109/TMI.2013.2295465>.
- [59] A. Borzi, G. Ciarrella, M. Sprengel, Formulation and Numerical Solution of Quantum Control Problems, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, ISBN 978-1-61197-483-6, 2017.
- [60] M. S. Vinding, T. E. Lund, J. P. Stockmann, B. Guérin, INSTANT (INtegrated Shimming and Tip-Angle NormalizaTion): 3D flip-angle mitigation using joint optimization of RF and shim array currents, in: 28th Annual Meeting of International Society of Magnetic Resonance in Medicine, Virtual, 0612, 2020.