

Application of Evolution Strategies to the Design of SAR Efficient Parallel Transmit Multi-Spoke Pulses for Ultra-High Field MRI

Boris Eberhardt, Benedikt A. Poser, N. Jon. Shah, and Jörg Felder, Member,

August 18, 2020

We present an evolution-strategy based approach to solve the magnitude least squares (MLS) design problem of low flip-angle slice-selective parallel transmit RF pulses for ultra-high field MRI using SAR and peak-RF-constraints. A combined transmit k-space trajectory and RF pulse weight optimization is proposed in two algorithmic steps. The first step is a coarse grid search to find an initial solution that fulfills all constraints for the subsequent multistage optimization. This avoids convergence to the next nearest local minimum. The second step attempts to refine the results using multiple evolution strategies. We compare the performance of our approach with the non-convex optimization methods described in the literature.

The proposed algorithm converges for phantom and in vivo data and only requires an initial estimate of the range of suitable regularization parameters. It demonstrates improved excitation homogeneity compared to published spoke-design methods and allows optimization for homogeneity with a subsequent reduction in the SAR burden. Moreover, excitation

*Manuscript received XXX XXX, XXX

[†]B. Eberhardt is with the Institute of Neuroscience and Medicine 4, Forschungszentrum Jülich, Jülich, and RWTH Aachen University, Germany (e-mail: b. eberhardt@fz-juelich.de).

[‡]J. Felder is with the Institute of Neuroscience and Medicine 4, Forschungszentrum Jülich, Jülich, Germany (e-mail: j.felder@fz-juelich.de).

[§]B. Poser is with the Department of Cognitive Neuroscience, Faculty of Psychology and Neuroscience, Maastricht University, Maastricht, Netherlands (e-mail: benedikt.poser@maastrichtuniversity.nl).

[¶]N. J. Shah is with the Institute of Neuroscience and Medicine 4 and 11, Forschungszentrum Jülich, Jülich, Germany, with the Department of Neurology, RWTH Aachen University, Aachen, Germany and with JARA-BRAIN, Translational Medicine, Aachen, Germany (e-mail: n.j.shah@fz-juelich.de)

homogeneity and the SAR burden can be balanced against each other, enabling a further reduction in SAR at the cost of minor relaxations in excitation homogeneity. This feature makes the algorithm a good candidate for SAR limited sequences in ultra-high field imaging. The algorithm is validated using phantom and in vivo measurements obtained with a 16-channel transmit array at 9.4 T.

1 Introduction

Magnetic resonance imaging (MRI) enables the examination of anatomical structure and metabolic function non-invasively [1]. Although standard clinical MRI systems currently operate between 1.5 and 3 T, ultra-high field (≥ 7 T) MRI is also evolving into a clinical modality due to its superior signal-to-noise ratio (SNR), increased BOLD contrast and higher spectral resolution [2]. However, despite these advantages, the use of ultra-high magnetic fields also leads to increased radio frequency (RF) field inhomogeneities and increased RF power deposition in the form of higher specific absorption rate (SAR) [3, 4].

One method to achieve homogeneous excitation in an UHF system is with the application of spoke RF-pulses [5]. In this method, RF pulses with individual complex weightings are played out in combination with concatenated slice-select gradients arranged at multiple locations in the transverse transmit k-space. The spoke pulses also allow SAR and power constraints to be explicitly incorporated [6, 7]. The design of the RF transmit pulses for application in spoke excitation methods is commonly based on the small flip-angle approximation (SFA) [8]. This is done for the sake of computational efficiency and commonly employs multiple transmit channels in parallel to decrease the duration of the excitation, as well as to increase the available degrees of freedom [9, 10]. In addition, spoke gradient trajectories are often used for slice-selective RF pulses to compensate for in-slice B_1^+ inhomogeneities, the feasibility of which has been demonstrated even at higher fields [11, 12, 13].

Optimizing for the magnitude of the excitation flip angle, instead of both the magnitude and the phase, improves excitation quality and is referred to as magnitude least squares (MLS) optimization [14]. The MLS pulse design proposed by Guérin et al. [7] explicitly includes local and global SAR as well as hardware constraints. Furthermore, they showed that the choice of an appropriate target phase can improve the quality of the resulting excitation profile and may lead to a reduction of local and global SAR. The study employed an MLS-designed target phase to solve a convex quadratically-constrained quadratic program (QCQP). The QCQP used Q-matrices [15], which were compressed with the virtual observation points algorithm [16] to impose local and global SAR constraints. In addition, peak- and average-power was limited according to the performance maxima of the hardware. The MLS optimization problem, with strict SAR and power constraints, was also implemented using an active-set (A-S) algorithm [6, 17].

RF pulses in spoke excitation can be designed with a predetermined k-space trajectory or the trajectory can be designed jointly with the pulses [18]. Grissom et al.

[19] showed the benefits of the joint optimization of the target phase pattern and the transverse transmit k-space (k_x, k_y) -spoke locations; but also described the simplicity with which it is possible to converge to a local sub-optimal minimum of the MLS cost function. In [20], the (k_x, k_y) -spoke locations were not optimized jointly with the RF pulses; however, the authors tried to control local SAR with channel-dependent regularization parameters. To improve excitation quality, optimization of the (k_x, k_y) -position was performed in [21] and [22] without any SAR or hardware constraints. Gras et al. [23] recently confirmed the importance of the joint optimization and included explicit SAR and hardware constraints. Time-efficient methods to calculate individual spoke locations in the (k_x, k_y) -plane have been developed by designing the pulses in a time-reversed fashion [24] or by using "universal pulses" intended to cover a wide range of anatomies without per subject pulse optimization [25, 26]. However, due to the greedy selection of spoke locations and/or weights in a non-convex function search-space, the approaches described above are prone to converging to local optima.

In this work, we examine and compare global derivative-free optimization methods to find low excitation error, low SAR solutions to the MLS design problem for ultra-high fields. We propose the use of an interleaved sequence of separate global, population-based optimization methods (also termed "meta-heuristics," according to [27]) to find suitable spoke locations in transmit (k_x, k_y) -space, as well as optimal transmit channel weights. Compared to previous works, we closely examine the effect of finding appropriate initial values for the problem and investigate the use of evolution strategies that avoid local minima of the non-convex optimization landscape. With this approach, we are able to design pulses that significantly reduce RF inhomogeneities with respect to B_1^+ shimming approaches and outperform published spoke design algorithms while respecting local and global SAR safety constraints and peak RF power constraints. Moreover, the proposed algorithm allows a trade-off between excitation homogeneity and SAR burden, which is useful when a further reduction of SAR is desired.

2 Methods

2.1 General Problem Formulation

In the SFA, the transverse magnetization m at position \mathbf{x} can be expressed as

$$m(\mathbf{x}) = i\gamma m_0 \int_0^T b_1^+(\mathbf{x}, t) e^{i\mathbf{x}\mathbf{k}(t)} dt \quad (1)$$

with the equilibrium magnetization, m_0 , the local RF magnetic transmit field, $b_1^+(\mathbf{x}, t)$, and the excitation k-space trajectory, $\mathbf{k}(t)$ [8]. In the case of N_c transmit channels, (1) can be extended as a superposition of all transmit channels with the complex transmit sensitivity, $S_j(\mathbf{x})$, of the j^{th} transmit channel, following the spatial-domain approach

by Grissom et al. [28]

$$lm(\mathbf{x}) = i\gamma m_0 \sum_{j=1}^{N_c} S_j(\mathbf{x}) \cdot \int_0^T b_j(t) e^{i\gamma \Delta B_0(\mathbf{x})(t-T)} e^{i\mathbf{x}\mathbf{k}(t)} dt. \quad (2)$$

In (2), a main magnetic field map, $\Delta B_0(\mathbf{x})$, has been added to account for static field inhomogeneities and $b_j(t)$ is the complex excitation signal for each transmit element in the rotating frame of reference. Equation (2) can be discretized as a matrix equation of the form $\mathbf{m} = \mathbf{A}\mathbf{w}$, with a desired target excitation, \mathbf{m} , the system matrix, \mathbf{A} , and the complex channel weights, \mathbf{w} . The system matrix for N_r voxels for a single spoke has the dimensions $N_r \times N_c$. The elements of matrix \mathbf{A} for the s^{th} spoke subpulse are given as [29, 28]

$$a_{s,ij} = S_{ij} e^{i\mathbf{r}_i \mathbf{k}(t_s)} \sum_{l=1}^{N_{rf}} b_l e^{i\gamma \Delta B_{0i}(t_s + N_{rf}\Delta t - l\Delta t)}, \quad (3)$$

where S_{ij} is the discretized complex transmit sensitivity of the j^{th} transmit channel at position \mathbf{r}_i relative to the isocenter of voxel number i . Here the slice-selective spoke subpulse is discretized into N_{rf} time-steps with Δt being the dwell-time and b_l the complex waveform amplitude at discrete time step l . $\mathbf{k}(t_s)$ is the product of the gyromagnetic ratio with the total gradient moment from the end of the s^{th} subpulse to the end of the pulse train at time point t_s . For N_s spokes, the single spoke system matrices and weights are concatenated

$$\mathbf{A}\mathbf{w} = \begin{bmatrix} \mathbf{A}_1 & \dots & \mathbf{A}_{N_s} \end{bmatrix} \begin{bmatrix} \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_{N_s} \end{bmatrix} = \mathbf{m}. \quad (4)$$

The optimization problem can now be expressed as a minimization of the difference between $\mathbf{A}\mathbf{w}$ and \mathbf{m} .

This can be solved in a least squares (LS) way by

$$\mathbf{w} = \arg \min_{\mathbf{w}} \{ \|\mathbf{A}\mathbf{w} - \mathbf{m}\|_2^2 + \lambda \|\mathbf{w}\|_2^2 \} \quad (5)$$

with a scalar Tikhonov regularization parameter, λ , that penalizes large values of \mathbf{w} and thus the RF power is traded off against excitation homogeneity. However, this usually results in higher flip angle errors than using the magnitude least squares (MLS) formulation [14, 30]

$$\mathbf{w} = \arg \min_{\mathbf{w}} \{ \|\mathbf{A}\mathbf{w} - \mathbf{m}\|_2^2 + \lambda \|\mathbf{w}\|_2^2 \} \quad (6)$$

as the MLS problem does not put constraints on the excitation phase and thus exhibits more degrees of freedom. The optimum regularization value depends on the properties of the transmit coil array and slice orientation and can be derived a priori for each coil

array using the L-curve methodology [31]. The goal of the MLS optimization is to find a set of complex transmit-channel weights, $\mathbf{w} \in \mathbb{C}^{N_c \times N_s}$, with a minimum root mean square error (RMSE)

$$RMSE = \sqrt{\frac{\sum_{i=1}^{N_r} (|\mathbf{A}\mathbf{w}|_i - |\mathbf{m}|_i)^2}{N_r}}, \quad (7)$$

while still fulfilling the local SAR, global SAR and peak RF constraints. Commonly, the local SAR matrices, \mathbf{Q}_i , which are compressed to N_{VOPs} virtual observation points (VOPs) [16], are used to calculate the local SAR burden as described in [29]. Local SAR is constrained to be below the IEC limit lsAR_{\max} of 10 W kg^{-1} . In neuroimaging, the global SAR value should be below the limit gSAR_{\max} of 3.2 W kg^{-1} , which is the IEC limit for the head. The global SAR can be calculated using the global SAR matrix $\langle \mathbf{Q} \rangle$ or – if this is not available – it may be derived from averaging over the local SAR matrices as described in [15]. The individual channel weights are constrained in magnitude so as to not exceed a maximum value V_{\max} . V_{\max} is given by the maximum voltage the power amplifier can supply or by the maximum voltage the RF coil can support per channel, whichever is less. The constrained optimization problem can thus be formulated as

$$\begin{aligned} \text{ll} \quad & \arg \min_{\mathbf{w}, k_x, k_y} f = \sqrt{\frac{\sum_{i=1}^{N_r} (|\mathbf{A}\mathbf{w}|_i - |\mathbf{m}|_i)^2}{N_r}} \\ \text{s.t.} \quad & \sum_{s=1}^{N_s} \mathbf{w}_s^H \cdot \mathbf{Q}_i(\mathbf{r}) \cdot \mathbf{w}_s \leq \text{lsAR}_{\max} \quad \forall i = 1 \dots N_{VOPs} \\ & \sum_{s=1}^{N_s} \mathbf{w}_s^H \cdot \langle \mathbf{Q} \rangle \cdot \mathbf{w}_s \leq \text{gSAR}_{\max} \\ & |\mathbf{w}_{s,j}| \leq V_{\max} \quad \forall s = 1 \dots N_s, \quad \forall j = 1 \dots N_c \end{aligned} \quad (8)$$

where the superscript H denotes the conjugate transpose.

Note, that by replacing the objective function, f , in (8) with the maximum local SAR burden, the optimization target can be recast to a problem that reduces SAR instead of RMSE. The SAR optimization is expressed as

$$\begin{aligned} \text{ll} \quad & \arg \min_{\mathbf{w}, k_x, k_y} f = \max(\sum_{s=1}^{N_s} \mathbf{w}_s^H \cdot \mathbf{Q}_i(\mathbf{r}) \cdot \mathbf{w}_s) \\ & \quad \forall i = 1 \dots N_{VOPs} \\ \text{s.t.} \quad & \text{RMSE} \leq \text{RMSE}_{\max} \\ & \sum_{s=1}^{N_s} \mathbf{w}_s^H \cdot \langle \mathbf{Q} \rangle \cdot \mathbf{w}_s \leq \text{gSAR}_{\max} \\ & |\mathbf{w}_{s,j}| \leq V_{\max} \quad \forall s = 1 \dots N_s, \quad \forall j = 1 \dots N_c. \end{aligned} \quad (9)$$

No explicit local SAR constraint is needed in 9, as we start from a feasible local SAR solution which can only be minimized.

2.2 Evolution Strategies

Evolutionary computing (EC) aims to capitalize on collective phenomena in adaptive populations underlying birth and death, variation and selection, in a generational loop

[32]. The main variants of EC are evolutionary programming, genetic algorithms and evolution strategies (ES). While ES originally used a set of rules to optimize real world experiments, e.g. [33], they were quickly applied to find extreme values in real valued functions. In contrast to classical numerical optimization, the convergence of ES cannot always be mathematically proven (although some progress has been made, e.g. [34]) and an intuitive understanding is that ES are prone to premature saturation when the mutation strength becomes too small [32]. The effect is comparable to being stuck in a local minimum when using gradient descent numerical optimization techniques in non-convex optimization problems. However, ES have proven successful in solving non-convex optimization problems, e.g. [35], as well as in the global optimization of constraint problems [36]. This motivates the use of ES in the optimization problem presented here. The objective function, f , is to be minimized with respect to the real-valued input variables, \mathbf{x} . \mathbf{x} can be the three-dimensional vector (k_x, k_y, λ) , the target phase vector $\boldsymbol{\varphi}_{\text{target}}$, the weights \mathbf{w} , or a combination of these vectors. As \mathbf{w} is complex-valued, its real and imaginary parts are concatenated into a real-valued vector.

ES are population-based approaches that start with a population of N_{parent} parent solutions and generate $N_{\text{offspring}}$ offspring solutions [34]. ES usually apply two kinds of variation operators, recombination and mutation, for this purpose. Individual offspring solutions are created by recombination of suitable parent solutions. The recombination operator is often carried out as a weighted sum of the individual solution vectors. Its goal is to strengthen the beneficial information from all selected solutions and to diminish sub-optimal information content. The mutation operator is applied to the recombined solutions by adding values that are sampled from probabilistic search distributions, often from multivariate normal distributions scaled by the standard deviation σ . $\sigma^{(g)}$ can be considered a step size for the mutation in generation g . A large $\sigma^{(g)}$ can explore the search space and may escape local minima by making the perturbations large, whereas a small $\sigma^{(g)}$ can fine-tune a solution. Finally, the selection scheme determines which subset of solutions are selected from the population to become parents for the next generation. Two kinds of selection schemes are commonly used: An $(N_{\text{parent}} + N_{\text{offspring}})$ selection scheme chooses from both the N_{parent} and $N_{\text{offspring}}$ solutions and keeps the all-time best N_{parent} solutions over all generations. In a $(N_{\text{parent}}, N_{\text{offspring}})$ selection scheme all N_{parent} solutions die after each single generation and only a subset of the $N_{\text{offspring}}$ solutions survive into the next generation to become the new parents [34].

Modern implementations of ES automatically adapt the $\sigma^{(g)}$ to the problem domain. The covariance matrix adaptation ES (CMA-ES) [34, 37] for example updates the whole covariance matrix $\mathbf{C}^{(g)}$: $\mathbf{x}^{(g+1)} \sim \mathbf{m}^{(g)} + \sigma^{(g)} \cdot \mathcal{N}(\mathbf{0}, \mathbf{C}^{(g)})$ – with g being the generation index, $\mathbf{m}^{(g)}$ the mean value, and sampling points $\mathbf{x}^{(g+1)}$ – of the search distribution during the generational loop [34]. Therefore, they also consider correlations between the problem dimensions. The CMA-ES uses a $(N_{\text{parent}, W}, N_{\text{offspring}})$ selection scheme which employs a weighted average (denoted by W) of the $N_{\text{parent}} \leq N_{\text{offspring}}$ parents with the highest fitness value to generate the next generation mean value for generation $(g+1)$. A record of the step-size σ and the covariance matrix updates is kept over all time steps of the evolution. This evolution path is exponentially smoothed

to filter out noisy updates and improves the adaption process by keeping interrelated information from the search between the generations [34]. With each generation the multivariate normal search distribution from which new candidate solutions are sampled from is updated. If the fitness function does not improve within a defined tolerance, the algorithm has converged.

2.3 Proposed Algorithm

The final goal of the algorithm is to find the optimal complex channel weights, \mathbf{w} , for the spoke RF-pulses and the optimal (k_x, k_y) -spoke locations in transmit k -space. Due to the highly non-convex optimization landscape, which exhibits a large number of local minima (compare Fig. 1), we propose a two-step algorithm as described in Fig. 2 to jointly optimize spoke locations and channel weights. In an initial coarse search, starting values are determined for the second optimization step which uses ES to further improve the excitation pulses. As the optimum Tikhonov regularization, λ , in (5) or (6) is not usually predetermined and is dependent on the transmit coil array configuration and the orientation of the slice to be examined [38], the proposed algorithm uses an implementation that computes channel weights for multiple values of the regularization parameter λ . This allows λ to be fine-tuned during both steps of the optimization.

In the first step, a low resolution (k_x, k_y) -grid is initialized. We use $N_k = 10$ positions in each direction, equally spaced between $-1/10$ cm and $1/10$ cm – which is approximately equal to $-2/\text{FoV}$ and $2/\text{FoV}$, with FoV being the desired acquisition field of view. The regularization vector is initialized with values covering a range of magnitudes, e.g. $\boldsymbol{\lambda} = [0 \ 0.1 \ 1 \ 10]$. The target excitation, \mathbf{m} , is used to mask the desired excitation region. The masks were created by thresholding the summed B_1^+ magnitude maps acquired with the DREAM sequence as detailed in [29]. In essence, the coarse grid search of the first optimization step evaluates each grid position using a concatenated LS and MLS optimization. The LS step – implemented according to [39] – is used to compute starting values, \mathbf{w}_{init} , for the pulse weights, which are then refined by the following MLS optimization. The MLS problem is solved using the variable exchange (VE) approach described previously [14] and is solved for all values of the regularization parameter vector. Note that, the VE method itself is an iterative process that updates for channel weights and a target phase, $\boldsymbol{\varphi}_{\text{target}}$, interchangeably and is referred to as the VE MLS loop in the following. The target phase is introduced as a factor $\mathbf{z} = \exp i \cdot \boldsymbol{\varphi}_{\text{target}}$ multiplied with the target excitation, \mathbf{m} , as described in [14].

For the grid search, $\boldsymbol{\varphi}_{\text{target}}$ is initialized as a random vector with values uniformly distributed in the interval $(-\pi, \pi)$ as no prior knowledge of the target phase is available. Possible (k_x, k_y) positions and regularization parameters, $\lambda \in \boldsymbol{\lambda}$, are evaluated from the predefined grid-locations, as described above. Each single point (k_x, k_y, λ) corresponds to a single full evaluation of the VE MLS loop, which each returns locally optimized pulse weights, \mathbf{w}_{MLS} . After finishing the evaluation of all initial grid positions, all solutions that violate peak power or SAR constraints are discarded. From the remaining solutions, the one with the lowest remaining RMSE is maintained as

the initial input parameters $s_{init} = (\lambda, \mathbf{w}, \boldsymbol{\varphi}_{\text{target}}, k_x, k_y)$. Note that the initial target phase $\boldsymbol{\varphi}_{\text{target}}$ is different from the phase pattern $\boldsymbol{\varphi}_{\text{MLS}}$, which is the phase of the returned complex excitation pattern $\mathbf{A}\mathbf{w}_{\text{MLS}}$.

If no feasible solutions remain, the first step grid-search optimization is repeated with a vector λ of increased (we simply double the initial values) regularization parameters until at least one solution is found that fulfills all constraints. The solution which fulfills all constraints and has the smallest RMSE is used as the initial solution for optimization in step two. The single set of feasible input parameters, s_{init} , for the VE MLS loop with the lowest resulting RMSE is kept during all stages of any subsequent VE based optimization. Whenever a better combination of values for s_{init} is found, it gets updated. Therefore, the most current best set of initial parameters that fulfill all constraints is available at all times. This ensures that even if all optimizations in the second step of the proposed algorithm fail, at least one solution is available to carry out the MR scan.

The second stage optimization itself is a concatenation of three sub-optimizations based on ES methods. The first ES – which we term phase update – maintains all values of the initial solution, s_{init} , except for optimizing the target phase with a solution that reduces the overall RMSE. It thus returns $s_{init} = (\lambda, \mathbf{w}, \boldsymbol{\varphi}_{\text{target}}^*, k_x, k_y)$, where the asterisk indicates an updated variable. The second ES optimizes the spoke locations (k_x, k_y) and the regularization parameter λ and keeps the target phase pattern fixed. It is referred to as location update (LU) and returns $s_{init} = (\lambda^*, \mathbf{w}, \boldsymbol{\varphi}_{\text{target}}, k_x^*, k_y^*)$. Finally, the VE method is carried out repeatedly to improve $\boldsymbol{\varphi}_{\text{target}}$ and \mathbf{w} until no further reduction of RMSE is achieved. The loop is terminated if the reduction of RMSE is less than 1% in subsequent iterations. Step two of the proposed algorithm is looped with the set of initial values taken from the previous iteration until it converges with the same threshold as above.

It should be noted that all ES optimizations of the proposed algorithm in the second stage use the MLS cost function (8) to generate the channel weights. However, this has been modified to include a large penalty, as suggested in [40], for all solutions that violate the imposed power or SAR constraints. A value of 99 was arbitrarily chosen for the penalty, which is well above any RMSE values encountered during the optimization. The cost function including the penalty is implemented as

$$\text{RMSE} = \begin{cases} \sqrt{\frac{\sum_{i=1}^{N_r} (|\mathbf{A}\mathbf{w}|_i - |\mathbf{m}|_i)^2}{N_r}} & \text{constraints fulfilled} \\ 99 & \text{otherwise} \end{cases} \quad (10)$$

Note also, that the overall algorithm structure allows for two varieties to be implemented. The algorithm can either be terminated after a single execution of the second stage in order to reduce computation time, known as the 'fast ES' method or, alternatively, looping over the second stage is feasible with each loop being initialized with the solutions from the prior iteration. This 'looped ES' with iterative executions of the second stage is referred to as 'the proposed algorithm' in the following.

The VE method optimizes for the optimal weights $\mathbf{w}_{\text{MLS}}^*$ by looping over individual least squares problems as described above. It is, however, possible to solve the problem for the weights directly by treating (\mathbf{w}, k_x, k_y) as optimization parameters, without

using the variable exchange method. In this case, the target phase φ_{target} and the regularization parameter λ are not needed, as the constraints can be modeled explicitly in the objective function. Using the non-convex MLS objective function in (7), this problem can be solved by non-linear optimization methods such as ES. We refer to this as a non-VE MLS optimization to distinguish it from the VE MLS loop. Notice that we use eq. (10) as the cost function both for the VE and non-VE strategies. The only difference lies in the optimization method employed to obtain the optimal weights $\mathbf{w}_{\text{MLS}}^*$ and spoke locations k_x^*, k_y^* .

Although other algorithms are feasible, for the sake of simplicity and computation speed, the proposed algorithm uses the simple $(N_{\text{parent}} + 1)$ -ES [32] during the phase update φ_{target} , (i), and the CMA-ES method [37] during the location update for k_x^*, k_y^* and λ^* , (ii). In order to adapt the algorithm to include an additional spoke, steps one and two are repeated. All ES implementations use the open-source software package Pagmo [41]. Its CMA-ES implementation follows the description given in [37]. It is possible to change the learning rates for the CMA-ES covariance matrix update steps as well as the time horizon of the evolution path, however, we kept those values unchanged in our implementation as they were automatically assigned. We set the initial step size parameter σ_0 to a value of 0.01, which is smaller than the default value of 0.5. The convergence tolerances were kept at the default values of 1×10^{-6} . Only the number of generations and the random seed can be set for the $(N_{\text{parent}} + 1)$ -ES, and these were left at the default value. The population size and the generation number are hyper-parameters, but even low values, such as 50 for both arguments, are often sufficient. For the (k_x, k_y, λ) update we inserted 10 random picks from the set of saved solutions of the grid search to be used to initialize the population. If the chosen population size of the ES in the second step is larger than 11, the remaining population is initialized with random solutions by the library employed [41]. The complete algorithm is implemented in the python programming language (Python Software Foundation, Delaware, USA) and is summarized in pseudo code in the appendix.

2.4 Performance evaluation

Table 1 contains an overview of the different optimization schemes investigated in this manuscript. The proposed algorithm is compared with published spoke design algorithms, namely the region-growing approach described in [29], sequential quadratic programming (SQP) and quadratically-constrained quadratic-program (QCQP) as suggested in [7].

In [17], several strategies for solving the MLS problem in parallel transmit pulse design are compared and the use of active sets (A-S) – which were also used in [6] – or alternatively sequential quadratic programming (SQP) with initial solutions found with the VE method (also called the Gerchberg-Saxton algorithm [43]), is advocated. As results obtained with the A-S and SQP approaches are comparable, we used a trust-region based SQP variant [44] as implemented in Scipy [45] as the non-VE optimization to optimize spoke locations (k_x, k_y) and weights \mathbf{w}_{MLS} that replaces the second step of the proposed algorithm for comparison. The required Hessian for the SQP is approximated with a Broyden-Fletcher-Goldfarb-Shanno (BFGS) [46] based

Table 1: Overview of RF pulse design algorithms compared in this study

Design Algorithm	Method	Description
proposed algorithm (looped ES)	VE	Step 1: Grid-Search + Step 2 opt. of \mathbf{w}_{MLS} with VE MLS method by opt. of its input arguments $\boldsymbol{\varphi}_{\text{target}}, \mathbf{w}, k_x, k_y, \lambda$ (Fig. 2)
Published algorithms		
SQP	non-VE	Grid-Search + opt. of (k_x, k_y) and \mathbf{w}_{MLS} with SQP
QCQP	non-VE	Grid-Search + opt. of \mathbf{w}_{MLS} with QCQP solver for given fixed $\boldsymbol{\varphi}_{\text{target}}$
Region growing	VE	opt. of \mathbf{w}_{MLS} with VE MLS method, fixed (k_x, k_y) , region growing update of $\boldsymbol{\varphi}_{\text{target}}$ with $\boldsymbol{\varphi}_{\text{MLS}}$
Variations of step 2 in the proposed algorithm		
fast ES	VE	proposed algorithm with only single iteration of step 2
CMA-ES	non-VE	Grid-Search + opt. of (k_x, k_y) and \mathbf{w}_{MLS} with non-VE CMA-ES
Newuoa [42]	non-VE	Grid-Search + opt. of (k_x, k_y) and \mathbf{w}_{MLS} with non-VE Newuoa
fast ES + SQP	hybrid	fast ES + opt. of (k_x, k_y) and \mathbf{w}_{MLS} with non-VE SQP
fast ES + QCQP	hybrid	fast ES + opt. of \mathbf{w}_{MLS} with QCQP, given the ES-opt. $\boldsymbol{\varphi}_{\text{target}} = \boldsymbol{\varphi}_{\text{ESopt}}$
fast ES + CMA-ES	hybrid	fast ES + opt. of (k_x, k_y) and \mathbf{w}_{MLS} with non-VE CMA-ES
proposed algorithm + CMA-ES	hybrid	proposed algorithm + opt. of (k_x, k_y) and \mathbf{w}_{MLS} with non-VE CMA-ES
Grid search without further optimization		
Grid-Search	VE	opt. of \mathbf{w}_{MLS} with VE MLS method by opt. of its inputs (k_x, k_y, λ) on grid, random initial $\boldsymbol{\varphi}_{\text{target}}$

update strategy, as implemented in Scipy, and with default arguments. A Jacobian is estimated with a 2-point finite difference estimation as implemented in the Scipy library. This implementation is labelled 'SQP' in Table 1.

A convex optimization is described in [7] where a VE MLS generated initial target phase, φ_{target} , is used to subsequently solve a quadratically-constrained quadratic-program (QCQP) with an interior point algorithm to optimize the weights \mathbf{w}_{MLS} , which itself is also a non-VE method. Thus, step two of the proposed algorithm is replaced by QCQP optimization. We implemented the convex QP in cvxpy [47] with the open source ECOS solver to solve convex second-order cone programs [48]. The best case RMSE scenario of the quadratic convex optimization was performed constraints and is labeled 'QCQP' in Table 1.

Finally, in order to compare the RMSE performance with an algorithm that employs fixed spoke locations, we carried out a comparison with the region-growing approach described in [29], which is a pure VE method.

Variations of the proposed algorithms – which differ in the optimization strategies employed in the second stage following the initial grid search – are also investigated in addition to the published spoke design methods. The first variant reduces computational time by only executing step two once without looping ('fast ES'). We continue to look at strategies to solve the MLS problem in a non-VE way by replacing step two with published ES methods. These include the direct application of CMA-ES and Newuoa [42] to solve for optimal channel weights. Finally, we investigated hybrid approaches that utilize the strengths of both, VE and non-VE methods, to solve the spoke optimization problem. The first two use the fast ES approach described above and add a final SQP or CMA-ES optimization. The last variant then adds a final CMA-ES (non-VE) optimization to the algorithm proposed in this manuscript to investigate whether further progress is possible by adding additional optimization steps.

As published spoke designs do not carry out the initial grid search essential to the proposed algorithm, we investigated the impact of the grid search alone by comparing its performance against the methods described above. All performance evaluations are based on multi-slice B_0 and B_1^+ maps acquired from five healthy volunteers, as described in the next section. Special care is taken to evaluate the influence of both algorithmic steps of the proposed method independently. The comparisons are carried out as a performance analysis of RMSE for a two-spoke, 16-channel excitation.

2.5 Experimental Setup

All experiments were carried out on a 9.4T MAGNETOM scanner (Siemens Healthineers, Erlangen, Germany) using the work flow described in [29, 49], and one of the spoke design algorithms listed above. In addition, images were acquired with the RF coil operating in CP-mode. All numerical design was carried out on a standard laptop computer. The 9.4T scanner provides 16 transmit channels (1kW each) connected to the 16-channel pTX coil described in [50]. The safety of the coil was assured as described in [51] and 125 VOPs were extracted from electromagnetic simulations of the coil using a vendor-supplied Q-matrix compression software. The selected compres-

sion results in approximately 7 % overestimation of SAR. The scanner operates under software version VB17 (step 2.3) and has a maximum gradient capability of 80 mT/m and a maximum slew rate of 330 mT/m/ms.

B_1^+ maps were acquired with a dual refocusing acquisition mode (DREAM) sequence [52] with Fourier phase encoding of the transmit channels [53]. The parameters for this sequence were: FoV = 256 mm x 224 mm, 17 slices, slice thickness = 4 mm, TR/TE1/TE2/TI = 6.8 ms/2.22 ms/4.44 ms/7.1 ms, flip angle = 7° and matrix 64 x 64. Static field maps were generated from a slab selective double echo gradient echo sequence (FoV = 200 mm x 200 mm, 44 slices, slice thickness = 4 mm, TR/TE1/TE2 = 30 ms/1 ms/3.21 ms, flip angle = 8° and matrix 50 x 50).

The target excitation was set to achieve a homogeneous single slice excitation over the phantom or brain with a target flip angle of 40° using two spokes. The base excitation sub-pulses were sinc-pulses (slice thickness = 2 mm, time bandwidth product = 2.7). Phantom and in vivo excitation was measured with the pre-saturated turbo flash sequence as detailed in [29]. Gradient echo images were acquired with the following parameters: FoV = 210 mm x 210 mm, 12 slices, slice thickness = 2 mm, flip angle = 40°, TE/TR = 14 ms/400 ms, matrix 512 x 512.

Phantom measurements used a homogeneous spherical phantom (diameter: 150 mm) filled with tap water and doped with 0.5% NaCl and 3% agar. In vivo data were acquired from five healthy volunteers in order to test the robustness of our approach. All in vivo data were acquired after having obtained written informed consent, in accordance with the Declaration of Helsinki, and after a positive evaluation of the study protocol through the local IRB.

3 Results

Fig. 1 shows the non-convex optimization landscapes of the MLS objective (8) as a function of the (k_x, k_y) -position of the first spoke and a fixed regularization parameter λ for one slice in the human brain. The spoke location selected by the proposed method is marked with a blue circle.

Fig. 3 and Fig. 4 show the performance evaluation of the proposed algorithm, based on 66 measured B_1^+ and B_0 maps from 5 different subjects. Fig. 3 analyses the effect of sampling density of the initial grid search on RMSE. A very coarse sampling results in solutions with a higher RMSE. However, good performance without significant improvements with further increases in sampling densities can already be observed with moderate resolutions of the search grid. Thus, to benefit computational speed, a moderate resolution of 8 to 10 steps in each direction is sufficient for a good performance overall.

Fig. 4 (a) shows a performance comparison of the proposed algorithms with published algorithms and derivations of the algorithm described here when designing two-spoke single slice excitations. Table 2 lists the corresponding numerical values of the achieved RMSE as well as the accompanying local SAR burden. Clearly, a joint trajectory and pulse weight optimization outperforms solutions obtained using fixed spoke locations. Compared to the two published methods that employ a joint trajectory (at

Table 2: RMSE and local SAR values for several RF pulse design algorithms averaged over 66 slices acquired in 5 different subjects

Algorithm	RMSE (%) (mean \pm std. dev.)		Loc. SAR (W kg ⁻¹) (mean \pm std. dev.)	
2 spokes				
Region growing	16.281 \pm	3.980	2.066 \pm	0.922
Grid-Search	13.659 \pm	3.995	2.604 \pm	1.078
fast ES	12.545 \pm	3.559	3.073 \pm	1.023
proposed algorithm	12.259 \pm	3.570	3.129 \pm	1.007
SQP	13.242 \pm	3.782	3.017 \pm	0.994
CMA-ES	12.492 \pm	3.673	3.614 \pm	1.110
Newuoa	12.842 \pm	3.684	3.366 \pm	1.076
fast ES + SQP	12.523 \pm	3.661	2.953 \pm	0.910
fast ES + CMA-ES	11.967 \pm	3.373	3.414 \pm	1.094
proposed algorithm + CMA-ES	11.869 \pm	3.471	3.335 \pm	0.963
QCQP	13.057 \pm	3.584	2.788 \pm	0.964
fast ES + QCQP	12.400 \pm	3.452	3.090 \pm	0.949
local SAR optimization				
2% RMSE relaxation	12.096 \pm	3.500	2.607 \pm	0.848
5% RMSE relaxation	12.445 \pm	3.600	2.273 \pm	0.829
15% RMSE relaxation	13.619 \pm	3.961	1.696 \pm	0.643
25% RMSE relaxation	14.820 \pm	4.294	1.496 \pm	0.557
40% RMSE relaxation	16.608 \pm	4.801	1.246 \pm	0.469
1, 3 and 4 spokes				
1 spoke MLS	19.251 \pm	3.532	2.321 \pm	1.023
fixed RG 3 spokes	15.353 \pm	3.717	1.900 \pm	0.987
fixed RG 4 spokes	15.916 \pm	4.359	2.152 \pm	1.135
Grid-Search 3 spokes	12.557 \pm	3.630	2.532 \pm	1.134
Grid-Search 4 spokes	11.869 \pm	3.777	2.348 \pm	0.898
SQP 3 spokes	11.557 \pm	3.626	3.181 \pm	1.007
SQP 4 spokes	10.976 \pm	3.868	3.565 \pm	0.978
fast ES 3 spokes	10.792 \pm	3.431	2.988 \pm	0.966
fast ES 4 spokes	9.630 \pm	3.423	3.300 \pm	1.055

least via an initial grid search) and pulse weight optimization (SQP and QCQP), our proposed algorithm shows an improvement in average RMSE and a comparable SAR burden. However, relaxing the RMSE constraint of the proposed algorithm as a trade-off for the SAR burden enables similar RMSEs as SQP and QCQP to be maintained with an almost 40 % decrease in SAR. Finally, a number of VE, non-VE and hybrid ES variations of the second step of the proposed algorithm all show similar RMSE and SAR performance.

Fig. 4 (b) compares algorithm performance when designing for three and four spoke pulses. Generally, longer pulse trains achieve higher excitation homogeneity and, in most cases, lower SAR burdens. However, longer excitation trains increase the minimum echo time and are more prone to off-resonance effects.

Fig. 5 shows simulation results and measured flip angle maps for two different spoke design algorithms based on measured B_1^+ and B_0 maps of the phantom described above. The left column shows the maps achieved with the algorithm used in [29] (region growing) and the right column shows the maps obtained with the proposed algorithm. Predicted and measured flip angle maps agree qualitatively and both show improved excitation homogeneity for the proposed algorithm, as indicated by the RMSE values given for each slice. Improved homogeneity is maintained for a spatially shifted phantom using the proposed algorithm, while the original spoke design algorithm from [29] tends to create excitation voids despite using a region growing algorithm to generate suitable target phases (Fig. S4 of supplementary material).

Fig. 6 shows simulated flip angle maps derived from measured B_0 and B_1^+ maps for a single slice in the human brain. In the first row, the effect of grid resolution in the first step of the proposed algorithm is analyzed. It can be seen that an excessively coarse sampling potentially results in regions with signal dropout, however, higher grid-resolutions can lead to improved results. Both, SQP (d) and QCQP (e) do not completely get rid of the excitation voids when they are initialized with non-optimum starting values. However, the algorithm presented here manages to achieve an excitation without excitation voids (f). Finally, the last row presents flip angle maps when relaxing the RMSE constraint in favour of a lower SAR burden (g).

Fig. 7 compares the in vivo flip angles achieved using a two-spokes region growing excitation according to [29] (first row) and our proposed method. Two optimizations were carried out with the proposed algorithm: the first targeted minimum RMSE (second row) and the second aimed at reducing the local SAR by relaxing the RMSE target by 25 % (third row). The figure also shows that the spoke pulses designed according to [29] violate peak power constraints. Flip angle maps for additional volunteers can be found in Supplementary Fig. S1.

Fig. 8 compares exemplary in vivo GRE images (voxel size: $0.4 \times 0.4 \times 2.0 \text{ mm}^3$) of a single healthy volunteer after CP excitation. Here, the proposed two-spokes minimum RMSE optimization and the proposed two-spokes method with a lower SAR burden were used, resulting in reduced excitation homogeneity. Additionally, it demonstrates the differences in the measured flip angle distributions by taking the ratios of the flip angle maps. The SAR burden and achieved RMSE for the GRE sequence are listed in Table 3 for the subject shown in Fig. 8 averaged over multiple slices.

Fig. 9 evaluates the influence of the target phase φ_{target} of VE based MLS designs

Table 3: Measured excitation properties for the gradient echo sequence detailed above. Values are given as mean \pm standard deviation.

	CP-Mode	Optimization for minimal RMSE	Optimization for low SAR burden
Measured RMSE (%)	30.85 \pm 6.35	14.40 \pm 2.38	14.48 \pm 2.29
Max. local SAR (W/kg)	3.07 \pm 1.34	2.34 \pm 0.34	1.62 \pm 0.41
Head SAR (W/kg)	1.13 \pm 0.46	1.26 \pm 0.25	0.91 \pm 0.22

on the RMSE. The top image shows RMSE achieved with the algorithm from [29] for each region growing step (blue points) that uses the MLS phase $\varphi_{\text{target}} = \varphi_{\text{MLS}}$ from the previous step as updated target phase for the next iteration. Using the proposed ES-optimized target phases $\varphi_{\text{target}} = \varphi_{\text{ESopt}}$ achieves lower RMSE during all region growing steps (green points) compared to MLS optimized phases. Thus, for constant spoke location and regularization, the importance of designing suitable target phase patterns is underlined. Note, that mean RMSE increases with the number of voxels as a larger region-of-interest needs to be homogenized. Fig. 9 bottom demonstrates the effect of the proposed looping of the VE MLS function to update the target phase with its returned MLS phase, $\varphi_{\text{target}} = \varphi_{\text{MLS}}$. While standard spoke designs terminate after a single optimization, we use the VE-computed MLS weights \mathbf{w}_{MLS} and phase pattern φ_{MLS} as initial values for a new optimization run. Repeating this until RMSE levels off can further improve the excitation fidelity.

Fig. 10 shows the convergence behavior of the proposed algorithm over the different successive optimization steps. It can be seen that subsequent steps allow a further reduction in excitation error or the SAR burden, if so desired. The algorithm was initially set to optimize for RMSE only. However, if the hard RMSE constraint is relaxed and further optimization carried out with (9) as a cost function, a further improvement in the SAR burden can be achieved at a minor expense of increased RMSE.

As far as computation time is concerned, a grid search including a single phase and location update can be performed in just a few seconds. This allows the number of evolution generations to be changed, with an approximate linear influence on the execution time. Running interleaved channel weight, (k_x, k_y, λ) -tuple and target phase updates until convergence achieved results in execution times of approximately 20 s and 30 s per slice for phantom and in vivo spoke calculations on a standard PC from 2013 (core i7 CPU, 16 GB of RAM), respectively.

4 Discussion

We investigated the application of ES to the design of spoke excitation pulses in MRI and compared the results to constraint design algorithms published in literature. We found that both designs resulted in similar RMSEs, but the design using ES was more robust with respect to the choice of appropriate starting values in both phantom and in vivo experiments. A high field strength of 9.4 T was employed as homogenizing the excitation flip angles increases in complexity with a shorter wavelength or higher field

strength. Performance was also tested with a 40° target flip angle, which is outside the strict definition of the small tip angle regime but is still below the 45° limit given in [54]. At the same time, the higher flip angle allows a more robust calculation of the SAR burden for the GRE sequence employed.

The improvements in excitation fidelity as a result of the proposed algorithm come at little additional cost in terms of time, as a single VE cost function evaluation only requires a computation time of only a few ms on a standard laptop CPU. Tse et al. [29] used a VE-based region-growing approach that employs an MLS optimized target phase pattern for each iteration to escape B_1^+ voids that correspond to the local minimal solutions of the search space. Using the proposed algorithm, we can show that further optimization of the initial target phase always outperforms the result of an initial fixed MLS target phase pattern (Fig. 9, top row) and more robustly avoids local minima as demonstrated in Fig. 6 and in the supplementary material for the spatially shifted phantom. Looping the VE method multiple times (proposed step 2.3 in Fig. 2) is a fast and efficient way to further update the phase until the RMSE converges (Fig. 9, bottom row) to even smaller values. The example shown converges after 1000 iterations, however, the number of possible iterations varies and depends on the starting conditions. Moreover, using a fixed target phase restricts control over the resulting constraints following the MLS optimization. In contrast, our approach discards any phase solution that exceeds the constraints and gives more control when steering the MLS solution towards reduced RMSE or reduced SAR solutions, while at the same time being flexible enough to include any kind of additional non-linear, non-convex constraint. Note that the use of meta-heuristic optimization techniques allows the algorithm to find a new target phase set outside the conventional local optimization in gradient-based optimization strategies. However, even with these more global optimization strategies, it is not guaranteed that the overall global optimum will be found as the final solution.

In our observations, the initial grid-search in the (k_x, k_y, λ) parameter space is an important feature of the proposed algorithm. It has been demonstrated that all implementations investigated here, as well as those reported in the literature, benefit from it, as demonstrated in Figs. 4 and 6. While the overall performance does improve with a higher grid-resolution (Fig. 4(a), diminished improvements are observed above a certain number. If required, the grid search is easily parallizable on multiple compute nodes for shorter execution times. Furthermore, as described previously [38], the optimal choice of λ depends on the transmit coil array configuration as well as the selected slice orientation and is thus not known a priori. Our approach requires only minimal a priori knowledge of the approximate order of magnitude of the regularization parameter. Although a grid search in transmit k -space has recently been reported [55], the implementation did not include the λ -optimization in this step. An alternative approach has also been proposed [56] which avoids the need for regularization by using a reduced search space reliant on semidefinite programming.

Fixing the potential spoke location candidates may still lead to sub-optimal results since better solutions might be missed, as demonstrated in the visualizations of the non-convex complex search landscapes in Fig. 1 and Figs. S3. It can be observed that the optimal spoke location varies in each slice, even within a single subject. For

example, in the fourth and fifth columns of Fig.S6 (representing neighboring slices), the chosen locations are almost at opposite locations in the (k_x, k_y) -plane. This is consistent with the result of Dupas et al. [6] who report that constraints have an impact on the spoke placement and that fixing the potential k -space locations may possibly miss some deeper local minima.

We examined the robustness of multiple optimization methods in the second step of the proposed algorithm by solving the regularized MLS problem using the VE loop described in [14] and by using non-VE (direct) optimization, e.g. as reported in [6, 17]. The challenge of all the approaches is to find good initial values (compare Gras et. al. [23]). The VE formulation profits from initialization with good candidate target phase patterns, φ_{target} , and regularization parameters, λ , which is not required when using non-VE optimizations with explicit constraints. However, VE and non-VE strategies achieve comparable results in RMSE as shown in Fig. 4(a). There is an indication that evolution strategies outperform both the QCQP [7] and the SQP [6, 17] implementations in the case where a sub-optimal starting point for the spoke design is selected, as demonstrated in Fig. 6. Overall, the presented algorithm using the $(N_{\text{parent}} + 1)$ -ES and CMA-ES optimization achieves stable, low RMSE spoke pulses, even for three or four spokes pulses (Fig. 4(b)). If computation time is considered, the loop in step two of the proposed algorithm can be executed with only one iteration only ("fast ES"), which results in only a minor performance degradation. Finally, using derivative-free methods, e.g. like the NEWUOA method [42], are fast and efficient alternatives.

Due to the high dimensionality of the phase search space and the non-convex cost function, we agree with [6] that no claim regarding global optimality can be made. With the VE method, however, we found that optimization of the (k_x, k_y) -locations alone leads to sub-optimal results without additionally optimizing the target phase pattern. When both (k_x, k_y) and the phase are optimized, the VE design is comparable in performance compared to non-VE methods that do not require phase optimizations. Based on the results presented here, a wide variety of the optimization strategies described in the literature can be employed in the proposed algorithm. They include, but are not limited to, differential evolution, particle swarm optimization, simulated annealing, and bound optimization by quadratic approximation (BOBYQA) [57].

The proposed algorithm allows other parameters, such as SAR, to be optimized in the same way as RMSE (Fig. 10). This has been demonstrated by adding a SAR optimization after the RMSE optimization loop. A trade-off between a moderate relaxation on RMSE performance and a lower SAR burden is feasible and has been demonstrated. Figs. 6 to 8 demonstrate that a small RMSE relaxation for the SAR optimization leads to only minimal differences in the reached flip angle patterns. The implication of this observation is that it might be possible to tailor spoke excitation schemes for specific sequence requirements. That is to say, if a high excitation homogeneity is required, e.g. in quantitative sequences, an RMSE optimization alone would be the best solution. Note that, with the proposed algorithm, this optimization will only return values within imposed SAR limits. In contrast, high SAR sequences may profit from adding the SAR optimization loop – at the cost of lower excitation homogeneity. It has been reported in [55] that SAR increases by up to 90% at even

higher field strengths of 10.5 T compared to 7 T making the ability to reduce SAR at only minor RMSE trade-offs an important step towards applicability at higher fields.

5 Conclusion

The suitability of using evolution strategies in combination with a coarse search for starting values has been demonstrated for combined spoke location and RF weight optimization. In particular, it has been shown that a sequential dual-target optimization that first allows excitation homogeneity to be improved and then subsequently reduces the SAR burden is feasible. Trade-offs between an increase in RMSE in the excitation homogeneity and SAR burden have been investigated and potentially mean that the proposed framework could be tailored for SAR limited sequences, such as MPRAGE or turbo spin-echo.

Acknowledgement

The authors thank A. W. Magill for scientific discussions during the start of this project and D.H.Y. Tse for providing several scripts used in this work. They acknowledge C. Rick for proofreading the manuscript.

References

- [1] B. R. Godlewska, S. Clare, P. J. Cowen, and U. E. Emir, “Ultra-high-field magnetic resonance spectroscopy in psychiatry,” *Front Psychiatry*, vol. 8, p. 123, 2017.
- [2] K. Ugurbil, “Imaging at ultrahigh magnetic fields: History, challenges, and solutions,” *Neuroimage*, vol. 168, pp. 7–32, 2018.
- [3] P. F. Van de Moortele, C. Akgun, G. Adriany, S. Moeller, J. Ritter, C. M. Collins, M. B. Smith, J. T. Vaughan, and K. Ugurbil, “B(1) destructive interferences and spatial phase patterns at 7 t with a head transceiver array coil,” *Magn Reson Med*, vol. 54, no. 6, pp. 1503–18, 2005.
- [4] T. Vaughan, L. DelaBarre, C. Snyder, J. Tian, C. Akgun, D. Shrivastava, W. Liu, C. Olson, G. Adriany, J. Strupp, P. Andersen, A. Gopinath, P. F. van de Moortele, M. Garwood, and K. Ugurbil, “9.4t human mri: preliminary results,” *Magn Reson Med*, vol. 56, no. 6, pp. 1274–82, 2006.
- [5] S. Saekho, C. Y. Yip, D. C. Noll, F. E. Boada, and V. A. Stenger, “Fast-kz three-dimensional tailored radiofrequency pulse for reduced b1 inhomogeneity,” *Magn Reson Med*, vol. 55, no. 4, pp. 719–24, 2006.

- [6] L. Dupas, A. Massire, A. Amadon, A. Vignaud, and N. Boulant, “Two-spoke placement optimization under explicit specific absorption rate and power constraints in parallel transmission at ultra-high field,” *J Magn Reson*, vol. 255, pp. 59–67, 2015.
- [7] B. Guerin, M. Gebhardt, S. Cauley, E. Adalsteinsson, and L. L. Wald, “Local specific absorption rate (sar), global sar, transmitter power, and excitation accuracy trade-offs in low flip-angle parallel transmit pulse design,” *Magn Reson Med*, vol. 71, no. 4, pp. 1446–57, 2014.
- [8] J. Pauly, D. Nishimura, and A. Macovski, “A k-space analysis of small-tip-angle excitation. 1989,” *J Magn Reson*, vol. 213, no. 2, pp. 544–57, 2011.
- [9] U. Katscher, P. Bornert, C. Leussler, and J. S. van den Brink, “Transmit sense,” *Magn Reson Med*, vol. 49, no. 1, pp. 144–50, 2003.
- [10] Y. Zhu, “Parallel excitation with an array of transmit coils,” *Magn Reson Med*, vol. 51, no. 4, pp. 775–84, 2004.
- [11] 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 ismr challenge,” *Magnetic Resonance in Medicine*, vol. 78, no. 4, pp. 1352–1361, 2017.
- [12] X. Wu, G. Adriany, K. Ugurbil, and P.-F. Van de Moortele, “Correcting for strong eddy current induced B0 modulation enables two-spoke RF pulse design with parallel transmission: Demonstration at 9.4T in the human brain,” *PloS One*, vol. 8, no. 10, p. e78078, 2013.
- [13] K. Setsompop, V. Alagappan, B. Gagoski, T. Witzel, J. Polimeni, A. Potthast, F. Hebrank, U. Fontius, F. Schmitt, L. L. Wald, and E. Adalsteinsson, “Slice-selective RF pulses for in vivo B inhomogeneity mitigation at 7 tesla using parallel RF excitation with a 16-element coil,” *Magnetic Resonance in Medicine*, vol. 60, no. 6, pp. 1422–1432, 2008.
- [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. 908–15, 2008.
- [15] I. Graesslin, H. Homann, S. Biederer, P. Bornert, K. Nehrke, P. Vernickel, G. Mens, P. Harvey, and U. Katscher, “A specific absorption rate prediction concept for parallel transmission mr,” *Magn Reson Med*, vol. 68, no. 5, pp. 1664–74, 2012.
- [16] G. Eichfelder and M. Gebhardt, “Local specific absorption rate control for parallel transmission by virtual observation points,” *Magn Reson Med*, vol. 66, no. 5, pp. 1468–76, 2011.

- [17] A. Hoyos-Idrobo, P. Weiss, A. Massire, A. Amadon, and 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 Transactions on Medical Imaging*, vol. 33, no. 3, pp. 739–748, 2014.
- [18] C. Ma, D. Xu, K. F. King, and Z.-P. Liang, "Joint design of spoke trajectories and RF pulses for parallel excitation," *Magnetic Resonance in Medicine*, vol. 65, no. 4, pp. 973–985, 2011.
- [19] 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," *Magnetic Resonance in Medicine*, vol. 68, no. 5, pp. 1553–1562, 2012.
- [20] M. A. Cloos, M. Luong, G. Ferrand, A. Amadon, D. Le Bihan, and N. Boulant, "Local sar reduction in parallel excitation based on channel-dependent tikhonov parameters," *J Magn Reson Imaging*, vol. 32, no. 5, pp. 1209–16, 2010.
- [21] A. C. Zelinski, L. L. Wald, K. Setsompop, V. K. Goyal, and E. Adalsteinsson, "Sparsity-enforced slice-selective mri rf excitation pulse design," *Ieee Transactions on Medical Imaging*, vol. 27, no. 9, pp. 1213–1229, 2008.
- [22] X. Wu, S. Schmitter, E. J. Auerbach, K. Uğurbil, and P.-F. Van de Moortele, "Mitigating transmit b1 inhomogeneity in the liver at 7t using multispoke parallel transmit rf pulse design," *Quantitative Imaging in Medicine and Surgery*, vol. 4, no. 1, pp. 4–10, 2014.
- [23] V. Gras, A. Vignaud, A. Amadon, F. Mauconduit, D. Le Bihan, and N. Boulant, "In vivo demonstration of whole-brain multislice multispoke parallel transmit radiofrequency pulse design in the small and large flip angle regimes at 7 tesla," *Magnetic Resonance in Medicine*, vol. 78, no. 3, pp. 1009–1019, 2017.
- [24] 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. 278–85, 2012.
- [25] V. Gras, A. Vignaud, A. Amadon, D. Le Bihan, and N. Boulant, "Universal pulses: A new concept for calibration-free parallel transmission," *Magnetic Resonance in Medicine*, vol. 77, no. 2, pp. 635–643, 2017.
- [26] V. Gras, F. Mauconduit, A. Vignaud, A. Amadon, D. Le Bihan, T. Stöcker, and N. Boulant, "Design of universal parallel-transmit refocusing kt-point pulses and application to 3d t2-weighted imaging at 7t," *Magnetic Resonance in Medicine*, vol. 80, no. 1, pp. 53–65, 2018.
- [27] F. Glover, "Future paths for integer programming and links to artificial-intelligence," *Computers & Operations Research*, vol. 13, no. 5, pp. 533–549, 1986.

- [28] W. Grissom, C. Y. Yip, Z. Zhang, V. A. Stenger, J. A. Fessler, and D. C. Noll, "Spatial domain method for the design of rf pulses in multicoil parallel excitation," *Magn Reson Med*, vol. 56, no. 3, pp. 620–9, 2006.
- [29] D. H. Y. Tse, C. J. Wiggins, and B. A. Poser, "High-resolution gradient-recalled echo imaging at 9.4t using 16-channel parallel transmit simultaneous multislice spokes excitations with slice-by-slice flip angle homogenization," *Magnetic Resonance in Medicine*, vol. 78, no. 3, pp. 1050–1058, 2017.
- [30] P. W. Kassakian, "Convex approximation and optimization with applications in magnitude filter design and radiation pattern synthesis," Ph.D. in Engineering-Electrical Engineering and Computer Sciences, University of California, Berkeley, 2006.
- [31] D. Calvetti, S. Morigi, L. Reichel, and F. Sgallari, "Tikhonov regularization and the l-curve for large discrete ill-posed problems," *Journal of Computational and Applied Mathematics*, vol. 123, no. 1-2, pp. 423–446, 2000.
- [32] H.-G. Beyer and H.-P. Schwefel, "Evolution strategies - a comprehensive introduction," *Natural Computing*, vol. 1, no. 1, pp. 3–52, 2002.
- [33] J. Klockgether and H.-P. Schwefel, "Two-phase nozzle and hollow core jet experiment," in *Proc. 11th Symp. Engineering Aspects of Magnetohydrodynamics*, D. Elliott, Ed. California Institute of Technology, Pasadena CA, 1970, pp. 141–148.
- [34] N. Hansen, D. V. Arnold, and A. Auger, "Evolution Strategies," in *Handbook of Computational Intelligence*, 2015.
- [35] S.-K. Wang, J.-P. Chiou, and C.-W. Liu, "Non-smooth/non-convex economic dispatch by a novel hybrid differential evolution algorithm," *IET Generation, Transmission & Distribution*, vol. 1, pp. 793–803(10), September 2007.
- [36] Y. Diouane, S. Gratton, and L. N. Vicente, "Globally convergent evolution strategies for constrained optimization," *Computational Optimization and Applications*, vol. 62, no. 2, pp. 323–346, 2015.
- [37] N. Hansen, *The CMA Evolution Strategy: A Comparing Review*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 75–102.
- [38] A. Sbrizzi, H. Hoogduin, J. J. Lagendijk, P. Luitjen, G. L. Sleijpen, and C. A. van den Berg, "Time efficient design of multi dimensional rf pulses: application of a multi shift cgls algorithm," *Magn Reson Med*, vol. 66, no. 3, pp. 879–85, 2011.
- [39] C. Bauckhage, *NumPy / SciPy Recipes for Data Science: Regularized Least Squares Optimization*, 03 2015.
- [40] Z. Michalewicz and M. Schoenauer, "Evolutionary Algorithms for Constrained Parameter Optimization Problems," *Evolutionary Computation*, vol. 4, no. 1, pp. 1–32, Mar. 1996.

- [41] F. Biscani and D. Izzo, “esa/pagmo2: pagmo 2.11.3,” Sep. 2019.
- [42] M. J. D. Powell, “The NEWUOA software for unconstrained optimization without derivatives,” in *Large-Scale Nonlinear Optimization*, ser. Nonconvex Optimization and Its Applications, G. Di Pillo and M. Roma, Eds. Boston, MA: Springer US, 2006, pp. 255–297.
- [43] R. W. Gerchberg, “A practical algorithm for the determination of phase from image and diffraction plane pictures,” *Optik*, Jan. 1972.
- [44] M. Lalee, J. Nocedal, and T. Plantenga, “On the implementation of an algorithm for large-scale equality constrained optimization,” *SIAM Journal on Optimization*, vol. 8, pp. 682–706, 1998.
- [45] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and S. . . Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [46] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. New York, NY, USA: Springer, 2006.
- [47] S. Diamond and S. Boyd, “CVXPY: A Python-embedded modeling language for convex optimization,” *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.
- [48] A. Domahidi, E. Chu, and S. Boyd, “ECOS: An SOCP solver for embedded systems,” in *European Control Conference (ECC)*, 2013, pp. 3071–3076.
- [49] D. H. Y. Tse, C. J. Wiggins, D. Ivanov, D. Brenner, J. Hoffmann, C. Mirkes, G. Shajan, K. Scheffler, K. Uludag, and B. A. Poser, “Volumetric imaging with homogenised excitation and static field at 9.4 t,” *Magnetic Resonance Materials in Physics Biology and Medicine*, vol. 29, no. 3, pp. 333–345, 2016.
- [50] G. Shajan, M. Kozlov, J. Hoffmann, R. Turner, K. Scheffler, and R. Pohmann, “A 16-channel dual-row transmit array in combination with a 31-element receive array for human brain imaging at 9.4 t,” *Magn Reson Med*, vol. 71, no. 2, pp. 870–9, 2014.
- [51] J. Hoffmann, A. Henning, I. A. Giapitzakis, K. Scheffler, G. Shajan, R. Pohmann, and N. I. Avdievich, “Safety testing and operational procedures for self-developed radiofrequency coils,” *NMR in Biomedicine*, vol. 29, no. 9, pp. 1131–1144, 2016.
- [52] K. Nehrke and P. Bornert, “Dream—a novel approach for robust, ultrafast, multi-slice b(1) mapping,” *Magn Reson Med*, vol. 68, no. 5, pp. 1517–26, 2012.

- [53] D. H. Tse, M. S. Poole, A. W. Magill, J. Felder, D. Brenner, and N. Jon Shah, “Encoding methods for b1(+) mapping in parallel transmit systems at ultra high field,” *J Magn Reson*, vol. 245, pp. 125–32, 2014.
- [54] 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 Imaging*, vol. 28, no. 10, pp. 1548–59, 2009.
- [55] M. A. Ertürk, X. Wu, Y. Eryaman, P.-F. V. de Moortele, E. J. Auerbach, R. L. Lagore, L. DelaBarre, J. T. Vaughan, K. Uğurbil, G. Adriany, and G. J. Metzger, “Toward imaging the body at 10.5 tesla,” *Magnetic Resonance in Medicine*, vol. 77, no. 1, pp. 434–443, 2017.
- [56] D. O. Brunner and K. P. Pruessmann, “Optimal design of multiple-channel rf pulses under strict power and sar constraints,” *Magnetic Resonance in Medicine*, vol. 63, no. 5, pp. 1280–1291, 2010.
- [57] M. Powell, “The bobyqa algorithm for bound constrained optimization without derivatives,” *Technical Report, Department of Applied Mathematics and Theoretical Physics, University of Cambridge*, 2009.


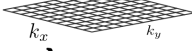
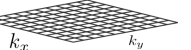



Figure 1: Optimization landscape for the first spoke location when using the 16-channel transmit coil at 9.4 T for one of the slices shown in Fig. 7. The Fig. shows the dependency of the RMSE (*top left*), peak RF voltage (*top right*), local SAR (*bottom left*) and global SAR (*bottom right*) on the spoke location for a fixed $\lambda = 0.357$. Transmit k-space is plotted from $-1/10$ cm to $1/10$ cm in the (k_x, k_y) -plane with a resolution of 250×250 . Also shown is the first found spoke location, indicated by a blue circle. Notice that the minimum in the RMSE landscape was not chosen, but a slightly different position that meets all constraints. The locations of the initial grid search for an 8×8 grid are shown as white dots.

For all spokes:

1st Step: grid-search

- initialize (k_x, k_y) Search-Grid 
- initialize regularization vector λ
- initialize random target phase φ_{target}
- $\forall (k_x, k_y) \in$  and $\forall \lambda \in \lambda$:
evaluate $\text{VE } \text{MLS}(\lambda, \mathbf{w}, \varphi_{\text{target}}, k_x, k_y)$
- Check if at least one solution fulfills all constraints (eq. 8/9)



if Yes: keep solution $(\lambda^*, \mathbf{w}, \varphi_{\text{target}}, k_x^*, k_y^*)$
with minimal RMSE (eq.9) for step 2
else: increase regularization λ and
repeat step 1 from beginning

2nd Step: improve solution

While RMSE progress is possible:

- optimize φ_{target} with ES, return $(\lambda, \mathbf{w}, \varphi_{\text{target}}^*, k_x, k_y)$
- optimize λ, k_x, k_y with ES, return $(\lambda^*, \mathbf{w}, \varphi_{\text{target}}, k_x^*, k_y^*)$
- while RMSE progress is possible:

optimize \mathbf{w} and φ_{target} by looping VE MLS
 $\mathbf{w} \leftarrow \mathbf{w}_{\text{MLS}}^*, \varphi_{\text{target}} \leftarrow \varphi_{\text{MLS}}$
 return $(\lambda, \mathbf{w}_{\text{MLS}}^*, \varphi_{\text{MLS}}^*, k_x, k_y)$

Result: final VE MLS evaluation with optimal

input parameters: $\text{MLS}(\lambda^*, \mathbf{w}^*, \varphi_{\text{target}}^*, k_x^*, k_y^*)$

returns optimal weights $\mathbf{w}_{\text{MLS}}^*$ for scanner




Figure 2: Schematic representation of the proposed optimization strategy. The first step uses a coarse grid-search that attempts to find suitable starting values for the subsequent optimization. It requires at least one solution to meet all constraints to guarantee convergence of the complete algorithm to a feasible solution. The second step uses ES initialized in addition to alternative starting values and allows mutation in order to avoid convergence to a single local minimum in the non-convex optimization landscape.




Figure 3: Dependence of RMSE on the grid resolution in the first step of the proposed algorithm. Results are based on 66 in vivo B_1^+ and B_0 maps of 5 subjects and show the median RMSE values as well as quartiles as boxes. To visualize the spread of the data, each data point is shown individually.




Figure 4: (a) Performance evaluation for the design algorithms listed in Table 1 for a two spoke excitation, and (b) evaluation for three and four spoke excitations. Results are based on 66 in vivo B_1^+ and B_0 maps of 5 subjects and show the median RMSE values as well as quartiles as boxes. To visualize the spread of the data, each data point is shown individually. For visual comparison, our proposed algorithm and one of the variants investigated (fast ES + CMA-ES) for computing two-spoke pulses are also shown.

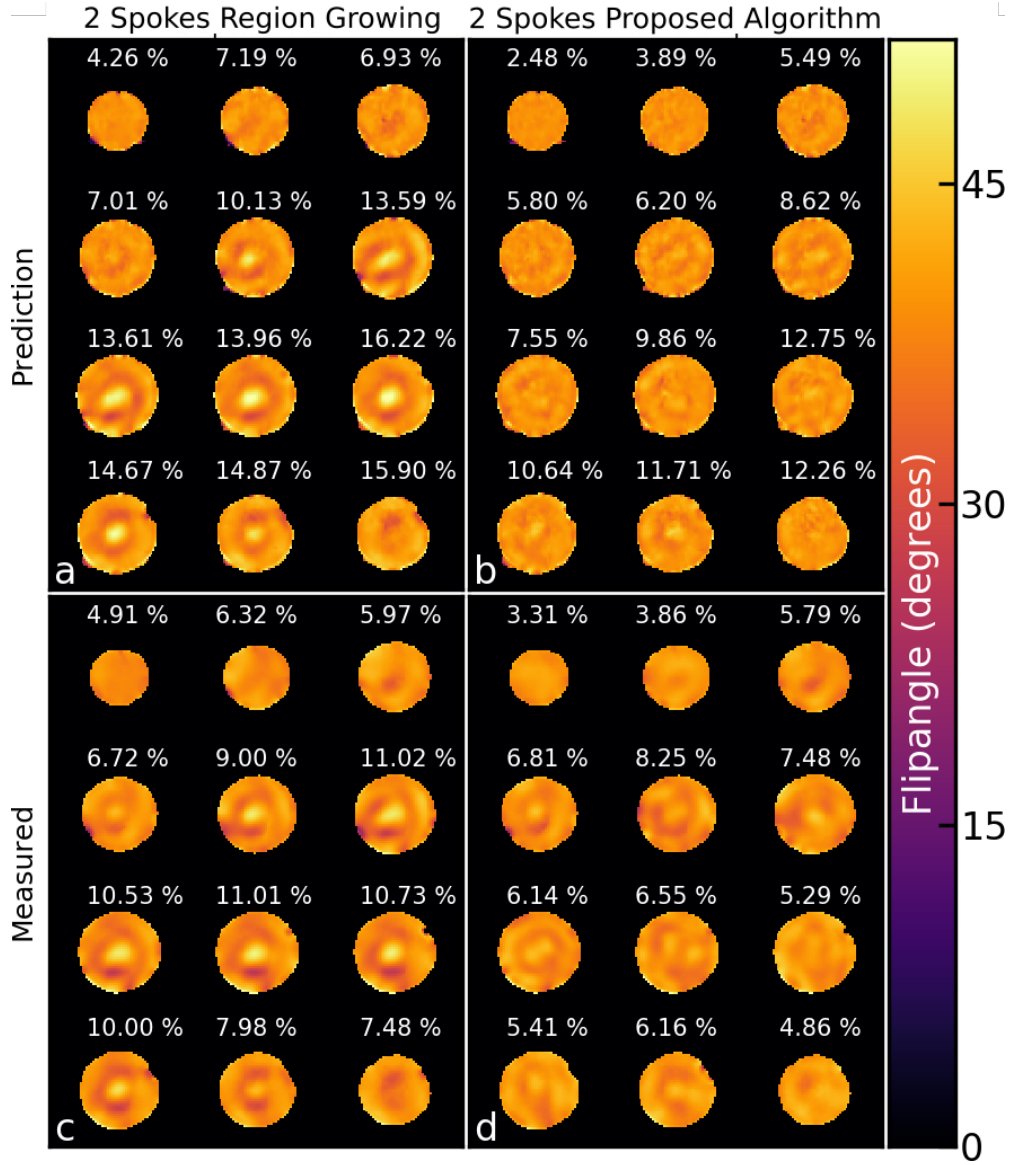


Figure 5: Phantom simulation and measurement results. The left column (*a,c*) shows the results of a MLS two-spoke optimization with fixed initial phase and spoke locations as described in [29], and the right column (*b,d*) shows our proposed approach to design two-spoke pulses. The first row shows the predicted results with RMSE percentages (describing the error with respect to the unit magnitude target excitation), and the second row shows the measured flip angle maps, also with RMSE percentages.




Figure 6: Simulated and normalized flip angle maps derived from measured B_1^+ and B_0 maps. (a, b, c) depict the results of grid-searches with resolutions 3×3 , 6×6 and 8×8 , respectively. $(d-f)$ show the result of using different spoke design algorithms when intentionally initialized with the values from the sub-optimal starting configuration (b) : (d) – SQP similar to [6], (e) – QCQP, and (f) – proposed algorithm. $(g-i)$ show the effect of subsequent SAR optimization using the proposed algorithm with RMSE relaxations of 2%, 5% and 15%, respectively.




Figure 7: Measured in vivo flip angle maps, comparing a fixed phase and (k_x, k_y) -location two-spokes excitation used in [29] without constraints (*first row*), and our proposed methods to reduce the RMSE (*second row*) and local SAR with RMSE relaxations (*third row*). A green circle indicates that all constraints are fulfilled, a red cross the violation of the maximum RF amplitude constraint V_{\max} of 175 V.




Figure 8: GRE images of an exemplary slice in the human head acquired with CP mode excitation (*left*), using the proposed spoke design algorithm to optimize for the best RMSE error (*middle*) and using the proposed spoke design algorithm with the target of reducing the SAR burden but relaxing the RMSE constraint by 5% (*right*). Ratios of the measured flip angle maps are shown for CP/RMSE, (*second from left*) and LSAR/RMSE (*second from right*)

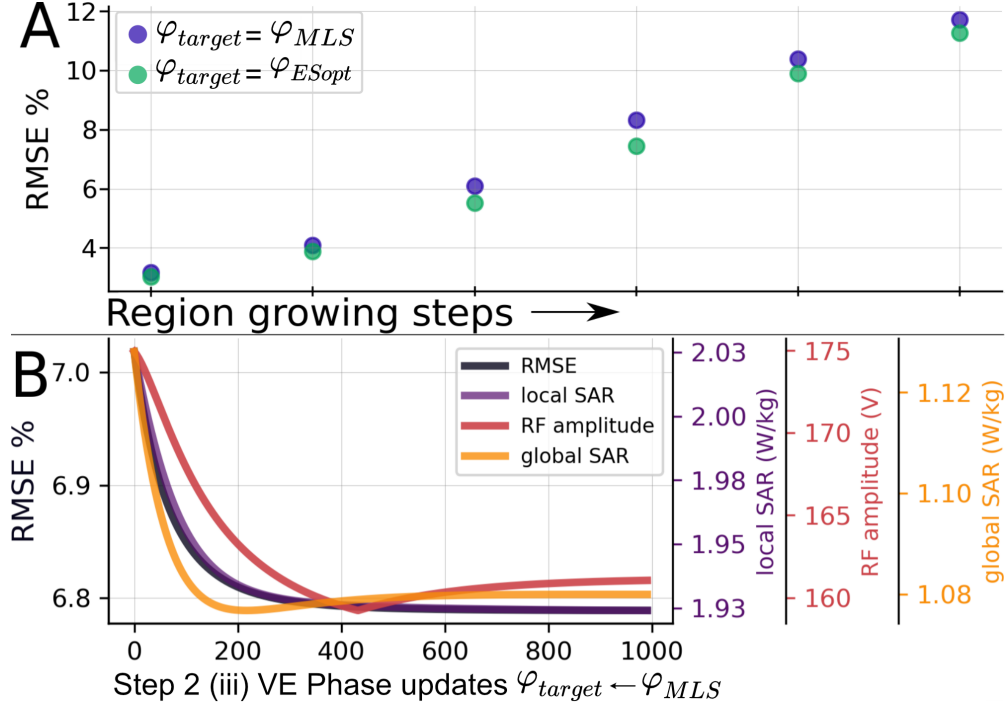


Figure 9: (A) RMSE as a function of the target phase. The blue dots indicate a target phase chosen according to the region growing for each region growing step (iteration) proposed in [29], while the RMSE obtained after updating the target phase by an ES step is superior for all region growing steps (green points). As the number of voxels in the mask increases with each region growing step the mean RMSE increases correspondingly. (B) shows that multiple runs of the variable exchange optimization with the updated target phases for each run have the potential to further improve the excitation fidelity at marginal cost on SAR burden and power requirements.

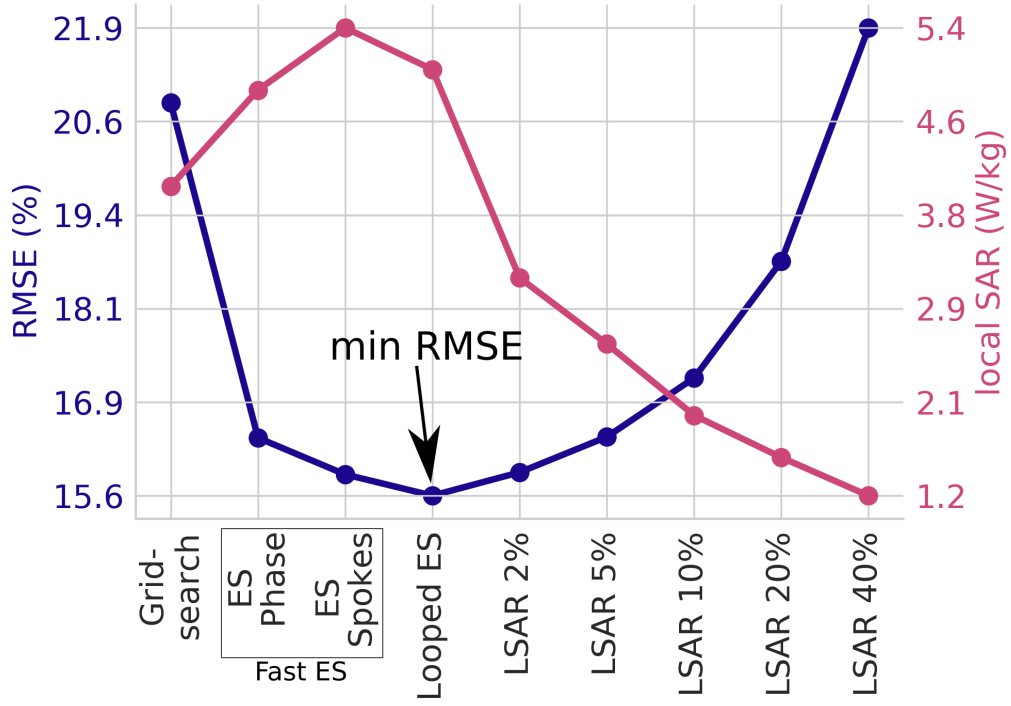


Figure 10: Performance of the proposed algorithm over the different optimization steps with respect to RMSE (left ordinate – blue curve) and SAR (right ordinate – red curve). The grid-search already provides a feasible solution that can be optimized further with respect to RMSE while staying within the imposed SAR and power constraints. The phase update and location update steps are implemented with evolution strategies while the final loop is a VE loop without an ES. Once a minimum RMSE solution has been found (at the tip of the arrow), a further SAR reduction is possible. The RMSE trade-off percentage for the local SAR optimization is given as labels on the x-axis.