MAGNITUDE-CONSTRAINED SEQUENCE DESIGN WITH APPLICATION IN MRI

Marcus Björk and Petre Stoica

Department of Information Technology, Uppsala University, Uppsala, Sweden

ABSTRACT

In this paper we present an algorithm for sequence design with magnitude constraints. We formulate the design problem in a general setting, but also illustrate its relevance to parallel excitation MRI. The formulated non-convex design optimization criterion is minimized locally by means of a cyclic algorithm, consisting of two simple algebraic sub-steps. Since the algorithm truly minimizes the criterion, the obtained sequence designs are guaranteed to improve upon the estimates provided by a previous method, which is based on the heuristic principle of the Iterative Quadratic Maximum Likelihood algorithm. The performance of the proposed algorithm is illustrated in two numerical examples.

Index Terms— Sequence design, MRI, Optimization

1. INTRODUCTION

Sequence, or waveform, design aims to generate sequences with specific desired properties, such as a certain spectral content, or good correlation properties. There is a wide range of applications, for example, in communications, active sensing, and MRI [1, 2, 3]. Typically, the signal to be designed is characterized by means of an optimization problem. Solving the problem globally can be difficult when the criterion is non-convex; however, in some cases a local minimization algorithm is sufficient to find a good solution. Indeed, different local optima correspond to possible candidates for a signal, and since the problem is usually solved offline, it is possible to generate several signals and choose the best among these based on the criterion.

In this paper, we derive a cyclic algorithm that locally solves a class of sequence design problems where a constraint on the magnitude of the designed complex-valued signal is enforced. This formulation has applications in MRI [4], but could also apply to other fields where small low-cost amplifiers are used. Typically, such amplifiers are single stage and are not equipped with feedback control [5]. This can cause nonlinear distortions of the signal in cases where the magnitude of the signal is varying rapidly [6]. By penalizing such variations in the design, the resulting sequences can be amplified and transmitted with higher fidelity.

This work was partially funded by the European Research Council (ERC) Advanced Grant, number 247035.

2. PROBLEM

In general terms, the criterion to be minimized can be formulated as

$$f(\mathbf{x}) = \|\mathbf{d} - \mathbf{A}\mathbf{x}\|^2_{\mathbf{W}} + \lambda\|\mathbf{R}|\mathbf{x}|\|^2, \quad (1)$$

where $\|\cdot\|$ is the Euclidean norm, $|\cdot|$ denotes the elementwise magnitude, $\mathbf{d} \in \mathbb{C}^M$ is the desired signal, $\mathbf{x} \in \mathbb{C}^N$ is the signal to be designed, $\mathbf{A} \in \mathbb{C}^{M\times N}$ and $\mathbf{R} \in \mathbb{C}^{P\times N}$ are arbitrary linear transformation matrices, and $\mathbf{W} \in \mathbb{C}^{M\times M}$ is a positive semidefinite weighting matrix. The regularization term contains a magnitude vector, which makes this function non-convex in general. The minimization of (1) with respect to $\mathbf{x}$ can be done in several ways; however, for large problems it is necessary to find an efficient method with low computational complexity.

The algorithm used in [4] for minimizing (1) is similar to the heuristic Iterative Quadratic Maximum Likelihood (IQML) algorithm and is not guaranteed to converge, nor is it a true minimization algorithm for the criterion [7]. However, IQML does typically converge to a vector fairly close to a minimizer of the stated criterion.

For the criterion in (1) the IQML algorithm can be described as follows. The vector $\mathbf{x}$ can be elementwise partitioned into its magnitude and phase as

$$x_k = |x_k|e^{i\phi_k}, \quad k = 1 \dots N. \quad (2)$$

By stacking the phases $\phi_k,\ k = 1 \dots N$, into a vector $\boldsymbol{\phi}$, we can form a criterion function

$$g(\mathbf{x}, \boldsymbol{\phi}) = \|\mathbf{d} - \mathbf{A}\mathbf{x}\|^2_{\mathbf{W}} + \lambda\|\mathbf{R}\mathrm{diag}(e^{-i\boldsymbol{\phi}})\mathbf{x}\|^2, \quad (3)$$

where $\mathrm{diag}(e^{i\boldsymbol{\phi}})$ is the matrix with diagonal elements $e^{i\phi_k}$. The criterion in (3) has the property that $g(\mathbf{x}, \boldsymbol{\phi}) = f(\mathbf{x})$, but for a fixed $\boldsymbol{\phi}$, the minimization becomes quadratic in $\mathbf{x}$, and can be solved by the least squares method. Then the phases can be updated as $\boldsymbol{\phi} = \angle\mathbf{x}$. These two steps are then iterated until some predefined stopping condition is satisfied.

Since IQML is not a minimizer of (1), it does not get stuck in local minima in the same way as a true minimization algorithm does. This property, together with the observation that IQML often converges rather rapidly, makes IQML a potential candidate for initialization of the local minimization algorithm described in Section 3. However, when there is no optimal vector $\mathbf{x}_{\mathrm{opt}}$ such that $\mathbf{A}\mathbf{x}_{\mathrm{opt}}$ is close enough to $\mathbf{d}$,

978-1-4799-2893-4/14/$31.00 ©2014 IEEE

IQML tends to have poor performance and might even never converge. An example of this type of behavior is shown in Section 5.

In [4], IQML is initialized by $\boldsymbol{\phi} = \mathbf{0}$, meaning that the first optimization step consists of solving the following least-squares problem

$$\underset{\mathbf{x}}{\text{minimize}} \ \|\mathbf{d} - \mathbf{A}\mathbf{x}\|_{\mathbf{W}}^2 + \lambda\|\mathbf{R}\mathbf{x}\|^2. \tag{4}$$

This provides a reasonably good initialization for the non-convex magnitude-constrained problem in (1).

3. MAGNITUDE-CONSTRAINED CYCLIC OPTIMIZATION (MACO)

Using (2), and defining $z_k = |x_k| \geq 0$, we can re-write the problem of minimizing (1) as:

$$\begin{aligned} \underset{\mathbf{z},\boldsymbol{\phi}}{\text{minimize}} \quad & \left\| \mathbf{d} - \sum_{k=1}^{N} \mathbf{a}_k z_k e^{i\phi_k} \right\|_{\mathbf{W}}^2 + \lambda\|\mathbf{R}\mathbf{z}\|^2, \\ \text{subject to} \quad & \mathbf{z} \geq \mathbf{0} \end{aligned} \tag{5}$$

where $\mathbf{a}_k$ is the k:th column of $\mathbf{A}$. Assuming $\mathbf{z}$ and $\{\phi_k\}_{k\neq p}$ are given, let

$$\mathbf{d}_p = \mathbf{d} - \sum_{\substack{k=1 \\ k\neq p}}^{N} \mathbf{a}_k z_k e^{i\phi_k}, \tag{6}$$

and observe that

$$\begin{aligned} &\|\mathbf{d}_p - \mathbf{a}_p z_p e^{i\phi_p}\|_{\mathbf{W}}^2 + \lambda\|\mathbf{R}\mathbf{z}\|^2 \\ &\quad = \|\mathbf{d}_p\|_{\mathbf{W}}^2 + z_p^2\|\mathbf{a}_p\|_{\mathbf{W}}^2 + \lambda\|\mathbf{R}\mathbf{z}\|^2 \\ &\qquad - 2\,\mathrm{Re}(z_p e^{-i\phi_p}\mathbf{a}_p^*\mathbf{W}\mathbf{d}_p) \\ &\quad = \text{const.} - 2z_p|\mathbf{a}_p^*\mathbf{W}\mathbf{d}_p|\cos(\arg(\mathbf{a}_p^*\mathbf{W}\mathbf{d}_p) - \phi_p), \end{aligned} \tag{7}$$

where the constant term is independent of $\boldsymbol{\phi}_p$. Then it follows that the minimizer with respect to ϕ_p is

$$\hat{\phi}_p = \arg(\mathbf{a}_p^*\mathbf{W}\mathbf{d}_p). \tag{8}$$

By cycling through the entire $\boldsymbol{\phi}$ vector we obtain an updated estimate, $\hat{\boldsymbol{\phi}}$, for the next iteration.

Once the phase vector is updated, we have to solve the minimization problem in (5) with respect to $\mathbf{z} \geq 0$ for a fixed $\boldsymbol{\phi} = \hat{\boldsymbol{\phi}}$. This is a linearly constrained quadratic program (LCQP), which can be solved rather efficiently, for example by using interior-point methods [8]. However, for large dimensions it might be favorable to determine the $\{z_k\}$ one-by-one, as was done above for $\{\phi_k\}$. To see how this can be done let us rewrite the criterion in (5) as

$$\left\| \begin{bmatrix} \mathbf{d} \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{A}\mathrm{diag}(e^{i\boldsymbol{\phi}}) \\ -\sqrt{\lambda}\mathbf{R} \end{bmatrix} \mathbf{z} \right\|_{\tilde{\mathbf{W}}}^2 \triangleq \|\mathbf{c} - \mathbf{B}\mathbf{z}\|_{\tilde{\mathbf{W}}}, \tag{9}$$

where

$$\tilde{\mathbf{W}} = \begin{bmatrix} \mathbf{W} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_P \end{bmatrix}, \tag{10}$$

and $\mathbf{I}_P$ is the identity matrix of size $P \times P$. If we assume that $\boldsymbol{\phi}$ and $\{z_k\}_{k\neq p}$ are given, and define

$$\mathbf{c}_p = \mathbf{c} - \sum_{\substack{k=1 \\ k\neq p}}^{N} \mathbf{b}_k \hat{z}_k, \tag{11}$$

we can write (9) as

$$\begin{aligned} \|\mathbf{c}_p - \mathbf{b}_p z_p\|_{\tilde{\mathbf{W}}}^2 &= \|\mathbf{c}_p\|_{\tilde{\mathbf{W}}} + z_p^2\|\mathbf{b}_p\|_{\tilde{\mathbf{W}}}^2 \\ &\qquad - 2z_p\,\mathrm{Re}(\mathbf{b}_p^*\tilde{\mathbf{W}}\mathbf{c}_p) \\ &= \text{const.} + \|\mathbf{b}_p\|_{\tilde{\mathbf{W}}}^2 \left[z_p - \frac{\mathrm{Re}(\mathbf{b}_p^*\tilde{\mathbf{W}}\mathbf{c}_p)}{\|\mathbf{b}_p\|_{\tilde{\mathbf{W}}}^2} \right]^2, \end{aligned} \tag{12}$$

where $\mathbf{b}_p$ is the p:th column of $\mathbf{B}$, and the constant term is independent of z_p. The minimizer $\hat{z}_p \geq 0$ of (12) has the following simple expression:

$$\hat{z}_p = \begin{cases} \frac{\mathrm{Re}(\mathbf{b}_p^*\tilde{\mathbf{W}}\mathbf{c}_p)}{\|\mathbf{b}_p\|_{\tilde{\mathbf{W}}}^2} & \text{if } \mathrm{Re}(\mathbf{b}_p^*\tilde{\mathbf{W}}\mathbf{c}_p) > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{13}$$

This can be used to update $\hat{\mathbf{z}}$, element by element, in the same manner as for $\hat{\boldsymbol{\phi}}$. By iterating the two steps, (8) and (13), the criterion function in (5) will decrease monotonically, as each step minimizes a part of the criterion. Since the criterion is bounded from below, it follows that the algorithm will converge to a local minimum. The proposed MACO algorithm is summarized below:

Algorithm 1: MACO Sequence Design

Input: $\mathbf{A}, \mathbf{R}, \mathbf{d}, \mathbf{W}, \lambda$, initial guess of $\mathbf{z}$
repeat
  Step 1:
  for all p **do**
    Compute $\mathbf{d}_p$ using (6)
    Compute $\hat{\phi}_p$ using (8)
  end for
  Step 2:
  for all p **do**
    Compute $\mathbf{c}_p$ using (11)
    Compute $\hat{z}_p$ using (13)
  end for
until convergence
Output: Compute $\hat{\mathbf{x}}$ from $\hat{\mathbf{z}}$ and $\hat{\boldsymbol{\phi}}$ using (2)

The computations in (6) and (11) can be performed recursively to reduce the computational burden. We have

$$\mathbf{d}_p = \mathbf{d} - \sum_{\substack{k=1 \\ k\neq p}}^{N} \mathbf{a}_k e^{i\phi_k} z_k = \mathbf{d} - \mathbf{A}\mathbf{x} + \mathbf{a}_p x_p, \tag{14}$$

where x_p is the current estimate. After obtaining the updated estimate $\hat{x}_p$ we can express the next residual as

$$\mathbf{d}_{p+1} = \mathbf{d}_p - \mathbf{a}_p \hat{x}_p + \mathbf{a}_{p+1} x_{p+1}. \tag{15}$$

Because of this, $\mathbf{d} - \mathbf{A}\mathbf{x}$ only has to be computed once; although, to prevent accumulating numerical errors in the recursion, a full re-computation of the residual can be done at each iteration. A similar recursion holds for $\mathbf{c}_p$.

The algorithm can be initialized in several ways. A good guess is, typically, provided by solving (4). The other option considered here is to initialize the algorithm by IQML, given that it has converged properly. By using the estimate obtained from IQML as initialization, MACO is guaranteed to perform at least as well, while taking advantage of IQMLs potential ability to avoid some local minima.

4. APPLICATION TO MRI

In MRI, the problem is to design sequences used to excite, or tip, the magnetic field vector in a certain region of a subject. Typically, such parallel/multi-coil excitation pulses have rapidly varying magnitudes [9, 10]. As mentioned, the low-cost amplifiers commonly used in this application can distort these signals, which leads to artifacts in the resulting images. Therefore, the signal magnitude should be made smooth while trying to maintain a desired excitation pattern. The multi-coil problem can be stated as follows [4]

$$\underset{\{\mathbf{x}_i\}_{i=1}^{N_c}}{\operatorname{argmin}} \left\| \mathbf{d} - \sum_{i=1}^{N_c} \operatorname{diag}(\mathbf{s}_i)\tilde{\mathbf{A}}\mathbf{x}_i \right\|_{\mathbf{W}}^2 + \lambda \left\| \sum_{i=1}^{N_c} \tilde{\mathbf{R}}|\mathbf{x}_i| \right\|^2, \tag{16}$$

where N_c is the number of parallel transmit channels in the multi-coil array, $\mathbf{s}_i \in \mathbb{C}^M$ is the vectorized spatial sensitivity of coil i, and $\mathbf{x}_i$ is the corresponding complex-valued signal to be designed. By stacking the N_c signal vectors in one vector $\mathbf{x} = [\mathbf{x}_1^T \ \dots \ \mathbf{x}_{N_c}^T]^T$, and defining the matrices $\mathbf{A} = [\operatorname{diag}(\mathbf{s}_1)\tilde{\mathbf{A}} \ \dots \ \operatorname{diag}(\mathbf{s}_{N_c})\tilde{\mathbf{A}}]$, and $\mathbf{R} = \mathbf{I}_{N_c} \otimes \tilde{\mathbf{R}}$, where $\otimes$ is the Kronecker product, we get the problem in the same form as (1). The desired signal $\mathbf{d}$ is in this case a vectorized multi-dimensional excitation pattern in space. The matrix $\mathbf{A}$ corresponds to a Fourier-type matrix that captures the, possibly non-uniform, sampling trajectory in the spatial Fourier domain (k-space) over time. The regularization matrix $\mathbf{R}$ can, for example, be determined by using a linear approximation of the filtering occurring in the amplifier, and computing the expected distortion filter. However, this requires knowledge of, or direct measurements from, the amplifier used. The distortion of the amplifiers used in [4] was shown to be fairly accurately modeled by a first-order difference filter, that is

$$\mathbf{R} = \begin{bmatrix} 1 & -1 & 0 & \cdots & 0 \\ 0 & 1 & -1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 1 & -1 \end{bmatrix}_{N-1 \times N}, \tag{17}$$

which is the approximation we will consider here. For a more detailed explanation of how one can achieve a multi-dimensional excitation pattern in space from one or several scalar time series, see for example [11].

5. NUMERICAL EXAMPLES

5.1. Example 1: A simple design

Let $\mathbf{W}$ and $\mathbf{A}$ in (1) be identity matrices; then the optimal phases can be obtained in closed form as $\boldsymbol{\phi} = \angle\mathbf{d}$. The resulting problem for $\mathbf{z}$ is a convex LCQP that can be written as follows

$$\begin{aligned} \underset{\mathbf{z}}{\text{minimize}} \quad & \| |\mathbf{d}| - \mathbf{z} \|^2 + \lambda \|\mathbf{R}\mathbf{z}\|^2 \\ \text{subject to} \quad & \mathbf{z} \geq \mathbf{0} \end{aligned}. \tag{18}$$

Since the globally optimal solution of (18) can be computed, this special case can be used as benchmark to compare the IQML and MACO algorithms. In the example, $N = 100$ and $\lambda = 1$. Each element in $\mathbf{d}$ and the initialization $\mathbf{x}_0$, was generated from uniform distributions for both the phase (between 0 and 2π) and the magnitude (between 0 and 1), and the elements of $\mathbf{R} \in \mathbb{R}^{100 \times 100}$ were drawn from a zero-mean Gaussian distribution with unit variance.

Monte Carlo simulations were performed by generating 1000 random initializations, and using these to start each algorithm. The problem parameters, $\mathbf{d}$ and $\mathbf{R}$ were kept fixed in all simulations. The resulting mean criterion as a function of the iteration number is shown in Fig. 1, together with the spread in terms of two standard deviations. As can be seen, the proposed method converges to the optimal solution in less than 20 iterations for all initializations, while IQML does not converge at all. Even the initialization given by (4) resulted in a similar behavior (not shown). This shows that IQML will have poor performance in some cases, which is a partial motivation for the local minimization algorithm presented herein.

5.2. Example 2: An MRI design

To make this example simple to understand, we will consider the problem with a fully sampled rectangular grid in k-space, no weighting, and a single transmitter coil. Furthermore, $\lambda = 10$, and the desired 2D excitation pattern, $\mathbf{D} \in \mathbb{R}^{32 \times 32}$, is a 10×10 square passband with unit magnitude, centered in space, as shown in Fig. 3a. In 2D, the problem can be formulated as

$$\underset{\mathbf{X}}{\operatorname{argmin}} \left\| \mathbf{D} - \mathbf{F}\mathbf{X}\mathbf{F}^T \right\|_F^2 + \lambda \left\| \mathbf{R}|\operatorname{vec}(\mathbf{X})| \right\|^2, \tag{19}$$

where $\| \cdot \|_F$ denotes the Frobenius norm, $\mathbf{F} \in \mathbb{C}^{32 \times 32}$ is a inverse discrete Fourier transform matrix, and $\operatorname{vec}(\cdot)$ is the columnwise vectorizing operator. By letting $\mathbf{x} = \operatorname{vec}(\mathbf{X}) \in \mathbb{C}^{1024}$, $\mathbf{d} = \operatorname{vec}(\mathbf{D}) \in \mathbb{R}^{1024}$, and $\mathbf{A} = \mathbf{F} \otimes \mathbf{F} \in \mathbb{C}^{1024 \times 1024}$, we can re-write the problem in the same form as (16). Here,

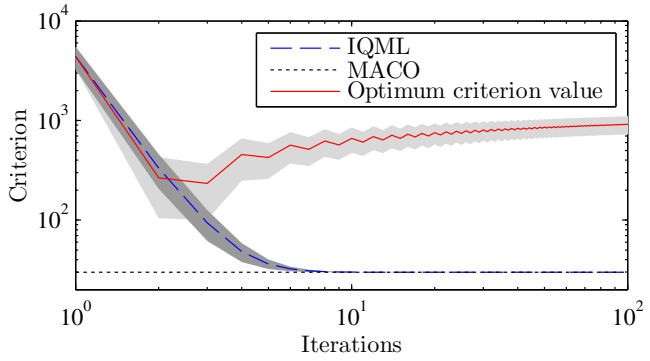

Fig. 1. The mean criterion for IQML and MACO versus the number of iterations, when applied to the simple design problem of (18) using 1000 random initializations. The light and dark gray fields show the spread of the criterion ($\pm 2\sigma$) for IQML and MACO, respectively.

$\mathbf{A}^*\mathbf{A}$ becomes diagonal, from which it follows that the optimal phase vector $\phi^\star$ is independent of $\mathbf{z}$. Therefore, MACO will reach the optimum in one iteration if step 2 is executed by solving the LCQP of (5). The corresponding algorithm, where the LCQP is solved by MOSEK, is denoted MACO-LCQP, and is included in the following for comparison. It should, however, be noted that solving an LCQP in step 2 might become intractable for large problems. In these cases, the elementwise update approach is preferable.

MACO, MACO-LCQP, and IQML were used to find the solution to this problem. The initial guess for all algorithms was obtained by solving the least-squares problem in (4). The convergence in terms of the criterion function versus the number of iterations is shown in Fig. 2. The magnitudes of the excitation patterns obtained after 30 iterations are shown in Fig. 3b-d. The resulting stopband and passband ripples, together with the sub-criteria for the fit (first term of (1)) and the magnitude smoothing (second term of (1)), are listed in Table 1. The stopband and passband ripples were defined as the maximum magnitude deviation from the desired excitation pattern in the respective areas.

In this example, the regularization is easier to handle than in the first example, and IQML converges. MACO-LCQP converges in one iteration, as expected, while the standard MACO has a slightly slower convergence rate. The time until convergence, with a tolerance of 10^{-6}, was 1041 s, 40 s, and 4 s, for IQML, MACO, and MACO-LCQP, respectively. At iteration 30, MACO closely approximates the MACO-LCQP solution, while IQML provides a smoother estimate with both lower fit and higher ripple values.

For smaller values of λ, that is, less smoothness imposed, IQML might outperform MACO for a given initialization as it does not get stuck in local minima. However, IQML would typically be used to initialize MACO in these cases, and therefore an improvement can still be expected.

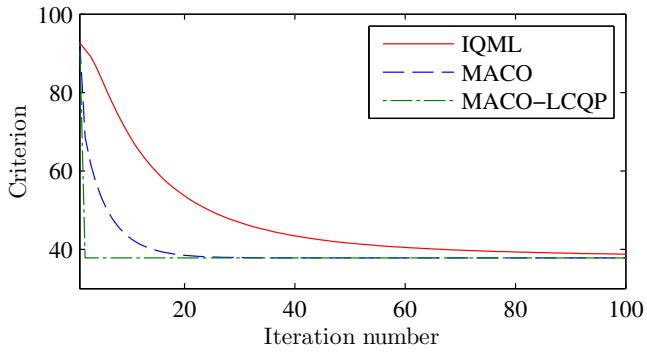

Fig. 2. Comparison of the criterion for IQML, MACO, and MACO-LCQP versus the number of iterations, when applied to the MRI example.

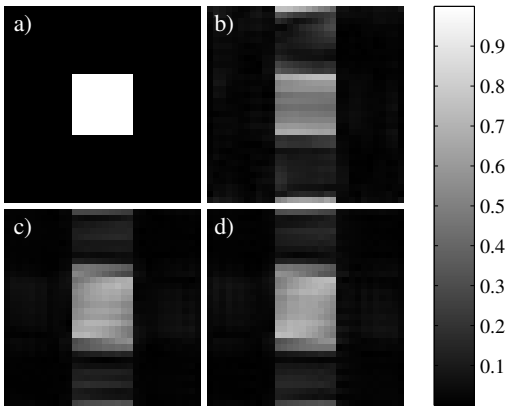

Fig. 3. a) The desired excitation pattern for the MRI example. Excitation patterns corresponding to the sequences designed by: b) IQML, c) MACO, and d) MACO-LCQP, obtained after 30 iterations.

Table 1. Ripples and sub-criteria at iteration 30 for the different methods in the case of the MRI example.

	MACO	MACO-LCQP	IQML
Passband ripple	0.56	0.53	0.60
Stopband ripple	0.56	0.57	0.75
Fit-term	26.9	26.7	35.9
Smoothness-term	1.10	1.11	1.04

6. CONCLUSION

We have derived a simple algorithm with low computational complexity, for solving least squares problems with magnitude constraints. The proposed MACO algorithm does not suffer from the potential convergence problems of IQML, and can further improve the results from IQML by truly minimizing the design criterion. The algorithm is useful for designing RF pulse excitation sequences in parallel MRI, which can be transmitted without compromising signal fidelity in the amplifier stage.

7. REFERENCES

[1] S. W. Golomb and G. Gong, *Signal design for good correlation: for wireless communication, cryptography, and radar*. Cambridge, U.K.: Cambridge University Press, 2005.

[2] H. He, J. Li, and P. Stoica, *Waveform design for active sensing systems: a computational approach*. Cambridge, U.K.: Cambridge University Press, 2012.

[3] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: physical principles and sequence design*. New York, NY, USA: Wiley-Liss, 1999.

[4] W. A. Grissom, A. B. Kerr, P. Stang, G. C. Scott, and J. M. Pauly, "Minimum envelope roughness pulse design for reduced amplifier distortion in parallel excitation," *Magnetic Resonance in Medicine*, vol. 64, no. 5, pp. 1432–1439, 2010.

[5] S. C. Cripps, *RF Power Amplifiers for Wireless Communications*, 2nd ed. Norwood, MA, USA: Artech House, Inc., 2006.

[6] J. C. Pedro and S. A. Maas, "A comparative overview of microwave and wireless power-amplifier behavioral modeling approaches," *IEEE Trans. Microw. Theory Techn.*, vol. 53, no. 4, pp. 1150–1163, 2005.

[7] P. Stoica, J. Li, and T. Söderström, "On the inconsistency of IQML," *Signal Processing*, vol. 56, no. 2, pp. 185–190, 1997.

[8] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.

[9] U. Katscher, P. Börnert, C. Leussler, and J. S. van den Brink, "Transmit SENSE," *Magnetic Resonance in Medicine*, vol. 49, no. 1, pp. 144–150, 2003.

[10] Y. Zhu, "Parallel excitation with an array of transmit coils," *Magnetic Resonance in Medicine*, vol. 51, no. 4, pp. 775–784, 2004.

[11] C. Y. Yip, J. A. Fessler, and D. C. Noll, "Iterative RF pulse design for multidimensional, small-tip-angle selective excitation," *Magnetic Resonance in Medicine*, vol. 54, no. 4, pp. 908–917, 2005.