MAGNETIC RESONANCE IN MEDICINE **21**, 39–48 (1991)

Coil Optimization for MRI by Conjugate Gradient Descent

ERIC C. WONG,* A. JESMANOWICZ, AND JAMES S. HYDE

Biophysics Section, Department of Radiology, Medical College of Wisconsin, Milwaukee, Wisconsin 53226

Received April 30, 1990; revised June 29, 1990

A flexible iterative algorithm is presented for optimizing gradient and radio frequency coils for MRI. It is based on a model using discrete current elements and direct Biot–Savart calculation of the fields. An error function is defined over a region of interest (ROI) and is minimized by conjugate gradient descent. The choice of error function allows optimization of the field uniformity, the inductance, and the efficiency of the coil in any combination. Neither the coil nor the ROI is restricted to any particular geometry. A 40-turn cylindrical z-gradient coil of radius a and length $4a$, designed for a ROI of radius $0.7a$ and length $2a$ has an average error in the gradient fields generated of 0.85%, an inductance of $0.014a$ mH/cm, and an efficiency of $6.65a^{-2}$ Gcm/A. A 16-turn birdcage-like RF coil of radius 5 cm, designed for a ROI of radius 4 cm has an average error of 0.79%.

INTRODUCTION

In several aspects of magnetic resonance imaging, magnetic fields with a specified distribution in space are required. These include gradient fields, the main static field, shim fields, and RF fields. In the past, coils made to produce such fields have been designed by analytic methods. Romeo and Hoult (*1*) used a spherical harmonic expansion of the fields to produce coils that approximate the lower order harmonics. These were then combined to approximate the desired fields. Compton (*2*) defined an error function to describe the uniformity of the gradient fields and used a least squares approach to derive a current distribution that would minimize this function.[1]

* Candidate for the combined MD and PhD degrees in the Medical Scientist Training Program at the Medical College of Wisconsin.

[1] This work was brought to our attention by one of the referees and describes a scheme for gradient coil optimization that is similar to the one presented here. An error function that is identical to our Eq. [1] is used to describe the uniformity of the gradient field, but this function is minimized by a different technique. A current is assigned to every point on the surface of the coil, and the fields generated are expressed as a function of these currents. Because the fields are a linear function of these currents, the error function can be minimized by a least squares approach without iteration. The optimized current distribution must then be approximated by discrete current elements. We chose an iterative approach which adjusts the positions of discrete current elements for two reasons. First, the coil that is constructed should give almost exactly the fields that are calculated because there is no step that involves the approximation of a current distribution with discrete current elements. This is especially important for configurations with smaller numbers of current elements for which the approximation of a current distribution would be less accurate. Also, for configurations like cylindrical transverse gradients, it is more straightforward using our technique to guarantee the existence of current return paths for all current elements and to include them in the field calculations.

0740-3194/91 $3.00

For cylindrical geometries, Turner (*3*) inverted a Fourier–Bessel expansion of the fields to derive a continuous current distribution flowing over a cylinder that would produce the desired fields on the surface of the region of interest (ROI). More recently, Turner (*4*) and Schweikert *et al.* (*5*) have introduced methods using Lagrange multipliers that can minimize the inductance or the power-to-field ratio and simultaneously guarantee that the fields will equal the desired fields at a number of points in the ROI.

These techniques have some theoretical and practical limitations: (a) In the methods that use Lagrange multipliers, the number of adjustable parameters in the coil must equal or exceed the number of points in the ROI at which the fields are specified. If the number of points at which the fields are specified is too small, then the field between these points can deviate greatly from the desired fields. This presents a problem for coils with smaller numbers of adjustable parameters. (b) For those methods which derive a distributed or continuous current distribution, such as those of Compton (*2*) and Turner (*3, 4*), the current distribution must be approximated by discrete wires or current paths. For small numbers of wires, the approximation to the calculated current distribution is poor, resulting in errors in the fields. (c) For current distributions expressed in terms of continuous functions, it is difficult to place explicit constraints on the overall size of the coil. This is important because the size of the coil is often a primary consideration in the coil design.

Presented here is an iterative algorithm for optimizing coils that is based on a specified number of discrete current elements and involves direct Biot–Savart calculation of the fields. We chose this approach to coil design for several reasons: (a) The number of points that describe the ROI can be chosen arbitrarily and is not required to be smaller than the number of current elements. (b) Because we carry out the Biot–Savart calculation of the fields from discrete current elements, the actual fields produced resemble the calculated fields very closely for any number of wires. (c) Arbitrary constraints can easily be placed on the dimensions of the coil, and on the locations of individual current elements. (d) Because both the efficiency and the inductance of coils are largely dependent upon the number of turns of wire present, it is useful to ask what is the best possible coil configuration with a given number of turns. (e) By suitable choice of error functions, it is possible to optimize any combination of field uniformity, inductance, and efficiency. (f) The method can in principle be applied to any configuration in which the currents can be described by a number of discrete current elements.

In our approach, an error function is defined over a set of points in the ROI and is a function of the positions of the current elements. This error function is then minimized by conjugate gradient descent (*6*). In order to optimize inductance and/or efficiency in addition to field uniformity, expressions for these quantities are incorporated into the error function.

Conjugate gradient descent is an established method for finding minima in multidimensional spaces, and it displays rapid (quadratic) convergence in the vicinity of a minimum. It has been used previously in MRI for the optimization of RF pulses by Mao *et al.* (*7*) and others.

Analyses of the cylindrical *z*-gradient coil, and a birdcage-like RF coil have been carried out using this technique, yielding designs which give very homogeneous gradient and RF fields over large ROIs. For the z-gradient coil, combinations of gradient field uniformity, inductance, and efficiency have been optimized.

This is a theoretical paper which introduces our techniques for optimizing coils. We have built a z-gradient coil and an RF coil and have tested the accuracy of our calculations. We are currently designing and implementing new coils including cylindrical transverse and biplanar gradient coils, and will report on these and their applications in future publications.

ALGORITHM

The general algorithm is given here, and applications to the cylindrical z-gradient coil and an RF coil are described in the next sections.

If M points are chosen in the ROI over which to optimize the field, and initial conditions are set for N current elements, then an error function E can be defined as

$$E = \sum_{m=1}^{M} (B_m - \overline{B_m})^2, \tag{1}$$

where $\overline{B_m}$ is the desired field at the ith point in the ROI, and B_m is the calculated field at that point. $\overline{B_m}$ can be specified explicitly, or if a uniform field is desired, it can be reset to the average field after each iteration.

Let $\mathbf{P}_i$ be an N-dimensional vector representing the positions of the N current elements after the ith iteration. Let $\mathbf{D}_i$ be the direction of search at the ith iteration, and $\mathbf{G}_i$ the gradient of E with respect to $\mathbf{P}_i$. $\mathbf{P}$ is adjusted as follows (*6*),

$$\mathbf{D}_0 = -\mathbf{G}_0, \tag{2}$$

$$\mathbf{D}_{i+1} = -\mathbf{G}_{i+1} + \beta_i\mathbf{D}_i, \tag{3}$$

$$\beta_i = \frac{|\mathbf{G}_{i+1}|}{|\mathbf{G}_i|}, \tag{4}$$

$$\mathbf{P}_{i+1} = \mathbf{P}_i + \alpha\mathbf{D}_i. \tag{5}$$

At each iteration, α is chosen to minimize E along the direction of search. Thus the problem becomes one of minimizing E as a function of a single variable (α). This is carried out as follows: Starting with the α of the previous iteration, the error E is calculated, and the coefficient α is successively doubled or halved until the minimum of E is passed. The nearest three (α, E) pairs are then fit with a second order polynomial, and the calculated minimum is used to determine α.

The gradient $\mathbf{G}$ is obtained by direct Biot–Savart calculation at each iteration. The nth component of this gradient is the derivative of E with respect to the position of the nth wire:

$$G_n = \frac{\partial E}{\partial P_n} = 2 \sum_{m=1}^{M} (B_m - \overline{B_m}) \frac{\partial B_m}{\partial P_n}. \tag{6}$$

An average percent error (AE) in the fields over the ROI is defined as

$$\text{Average Error} = \frac{1}{M}\sum_{m=1}^{M}\frac{|B_m - \overline{B_m}|}{|\overline{B_m}|} \times 100\% \qquad [7]$$

and is the quantity that is presented in the examples below because its meaning is intuitive.

For optimizing field uniformity and inductance simultaneously for a given number of turns of wire, the error function is replaced by $E \times L$, where E is the error function of Eq. [1], and L is the inductance of the coil. This new error function is minimized by the conjugate gradient descent method described above. The error function and its derivatives are again obtained by Biot–Savart calculation. The inductance is calculated by integrating the magnetic flux through each loop due to itself and all other loops in the coil.

Our choice of error functions for the optimization of field uniformity and inductance is somewhat arbitrary, and implies that E and L are of equal importance. If an unequal weighting of E and L is desired, one could minimize $E(1 + \epsilon L)$ where ϵ is a selectable parameter, or $E^n \times L^m$, where n and m are selectable.

In our experience, the algorithm converges to a stable minimum of the error function in all cases. In some cases, however, this minimum is not the local minimum that is closest in the parameter space to the initial configuration. The reason for this is that the gradient of the error function is large at the beginning of the search, and the step size α is relatively coarse. This allows the configuration to jump into adjacent local minima. If the nearest local minimum is desired, a steepest descent (5) method with small step size can be used for the first few iterations. In this method, the search direction is directly down the gradient $-\mathbf{G}_i$, and the step size is fixed. For the z-gradient coil application, the closest local minimum usually gives the lowest error function after convergence and is the configuration that is reported. For the RF coil application, more than one local minimum was never observed.

The algorithm generally converges to a minimum in 20–200 iterations. Computing time for 200 iterations for a 24-turn z-gradient coil, optimizing $E \times L$, was 4 min. Calculations were performed on a NeXT computer.

Z-GRADIENT COIL

This algorithm has been applied to the design of cylindrical z-gradient coils. Two important design parameters for such coils are the uniformity of the gradient field and the inductance. For a given size coil, the inductance is strongly dependent upon the number of turns of wire used. This dependence lies between linear and quadratic, depending upon the spacing between the wires. Therefore, it is useful to ask what is the best uniformity that can be achieved using a given number of turns. This is a problem that can be addressed by our approach.

The ROI used was a cylinder of radius 0.7 a and length 2 a where a is the radius of the coil. The current distribution was composed of coaxial current loops and was assumed to be antisymmetric about $z = 0$. The fields were only calculated in one quadrant of the ROI because of the symmetry of the coil. The gradient fields were calculated at 48 points in the ROI, at spacings of 0.1 a in the radial (r) direction, and

TABLE 1

Average Error in Gradient Field for Cylindrical Z-Gradient Coils

Number of turns	Length of coil	Average error (%)
16	3.20 a	5.17
24	3.05 a	3.22
32	3.53 a	1.36
40	3.84 a	0.86
48	3.60 a	0.77

0.2 a in the axial (z) direction. The error functions were weighted by the radius of the points in order to provide volume-weighted errors.

The variables to be adjusted were the positions along the z-axis of N current loops. Initially, the loops were evenly spaced from $z = 0.5\ a$ to $z = 1.5\ a$, and they were constrained to $0 < z < 2\ a$. At the end of each iteration, the value of the average gradient was calculated, and the desired uniform gradient was reset to this value. Each turn carries one unit of current.

The minimized AE was found to be strongly dependent upon the number of turns of wire specified, as shown in Table 1. Thus, there is a strong trade off between uniformity and inductance that must be taken into account for each application. Also shown in Table 1 is the overall length of the coils (maximum separation between current loops) in units of a, the radius of the coil. Figure 1 shows the wire distribution and the ROI for the 24-turn configuration.

For a coil with a given number of turns, the product of E and the inductance can be minimized explicitly. For the inductance calculations, the wire was assumed to have a radius of 0.02 a. When $E \times L$ is minimized, the resultant coils have an AE

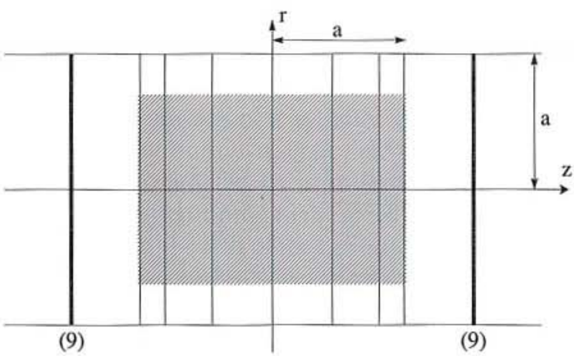

FIG. 1. Wire positions and ROI for 24-turn cylindrical z-gradient coil designed using the conjugate gradient descent algorithm. ROI is $|z| < a$, and $r < 0.7\ a$, and is shaded. Number in parentheses indicates the number of turns at that location. Other wires are single turn. Wire positions (z/a): 0.442, 0.798, 0.988, nine wires from 1.466 to 1.524.

that is 0–12% greater, but an inductance that is 21–38% lower than coils optimized for field uniformity alone. Results for various numbers of wires are shown in Fig. 2.

The efficiency, here defined as gradient per unit current is another important design parameter that can be incorporated into the optimization scheme. In the algorithm presented above for the z-gradient coil, the error function E was minimized without regard to efficiency, and the desired gradient was reset after each iteration to the calculated average gradient. This approach was taken because the gradient strength at which the error will be a minimum is not known beforehand.

The efficiency of the coil can be taken into account in several different ways. One way is to set it explicitly by fixing the desired gradient at a given value. Because the currents in the turns are fixed, specifying the desired gradient is equivalent to specifying the efficiency. This results in a coil with the best uniformity consistent with the specified efficiency. Another way is to use the AE (Eq. [7]) as the error function for the search. The AE is the error in the gradient fields divided by the desired gradient. If it is minimized while the magnitude of the desired gradient is allowed to vary, then the product of uniformity and efficiency is maximized. This is an arbitrary weighting of uniformity and efficiency, and as in the minimization of $E \times L$, other weightings are easily specified. The procedure is identical except that the gradient used to determine the direction of search is the gradient of the AE with respect to the wire positions.

Both of these approaches have been taken, and the results are shown in Fig. 3. The curve represents the AE of configurations for which the efficiency was fixed at various values. The AE is also shown for configurations in which E and AE were minimized

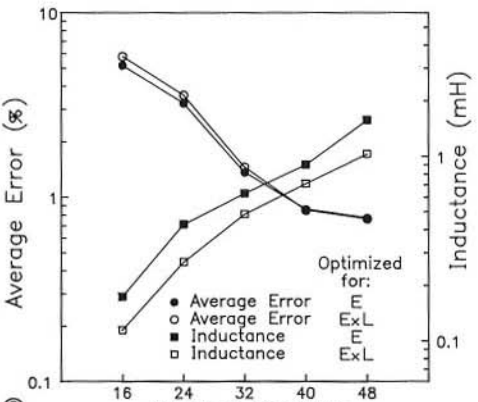

②

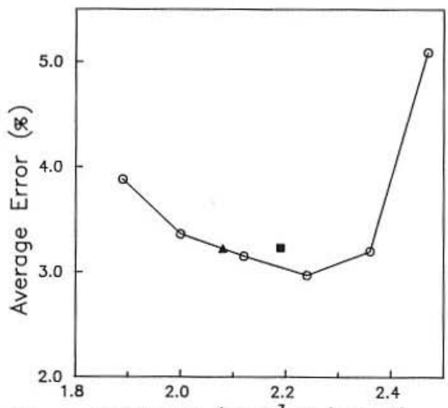

③

FIG. 2. Average error and inductance for z-gradient coils optimized for field uniformity only (closed symbols) and optimized for field uniformity and inductance (corresponding open symbols). Average error is defined by Eq. [7]. The ROI is $|z| < a$, and $r < 0.7\,a$. Inductance is shown for coils of radius 0.5 m.

FIG. 3. Average error vs efficiency for 24-turn z-gradient coils designed using the conjugate gradient descent algorithm. Average error is defined by Eq. [7]. Efficiency is defined as gradient per unit current, and is given for coils of radius 0.5 m. The ROI is $|z| < a$, and $r < 0.7\,a$. Triangle, square, and open circles are coils designed respectively by minimizing: E while the efficiency floats, AE while the efficiency floats, and E with the efficiency fixed at the values shown.

while the efficiency was allowed to "float." Note that when the AE is minimized, the algorithm settles in on a configuration with higher efficiency than when E is minimized. In general the configurations with higher efficiencies were more compact, but in all cases the loops were at $z < 1.7\ a$.

Gradient coils in which the turns of wire are driven in parallel, with varying currents in each turn, can also be considered. Such configurations have the potential for much lower inductance, but introduce the practical problem of constructing a coil with unequal currents. If the coil is wound in parallel, and resistors are used to limit the current in each turn, then a very large amount of heat must be dissipated by these resistors. However, if a configuration with a very small number of turns can be found that provides acceptable field uniformity, then it becomes possible to consider driving each wire or group of wires with a separate amplifier, thus avoiding the problem of dividing the currents.

For parallel coil configurations, the current carried by each wire can be adjusted by our algorithm. The algorithm is identical except that the gradient of E is calculated with respect to the currents. Optimization of the positions and currents of the wires can be done simultaneously.

In most of the configurations designed by our method, many of the turns of wire clustered at a location around $z = 1.5\ a$. This suggests a dual amplifier configuration in which a high current amplifier is used to drive one or a few turns around $z = 1.5\ a$, and a lower current amplifier drives other turns to increase the uniformity of the gradient.

Coil configurations with fixed but unequal currents, and also configurations for which the currents were allowed to vary are described in Table 2. For coil A the current ratios were fixed while the positions were allowed to move. For coils B–D only the current in the last turn was allowed to vary and for coil E, all the currents were allowed to vary. Coils A–D can be driven with two amplifiers as described above. It is thus possible with very few turns (and thus very low inductance) to design coils with good field uniformity.

CONSTANT-CURRENT BIRDCAGE COIL

Our technique has also been applied to the optimization of RF coils that are similar in principle to the birdcage design of Hayes *et al.* (*8*). It has been shown (*9–12*) that an axial current flow over the surface of a cylinder that varies sinusoidally with the azimuthal angle will produce a uniform magnetic field inside an infinitely long cylinder. For MRI, this is usually approximated by a finite number of equally spaced axial wires, each carrying a current proportional to the sine of the angle between the uniform field direction and a vector from the axis of the cylinder to the wire. A notable practical realization of this concept is the conventional birdcage coil (CBC) (*8*).

As the number of wires increases, the approximation to an ideal current distribution improves. Sotgiu and Hyde (*13*) have analyzed the uniformity of the field as a function of the number of wires.

We inquire here about a different approximation to the ideal current distribution. With the coil wound in series, the current in each wire is the same, and we adjust the azimuthal angle ϕ for each wire. The term "uniform current birdcage (UCB) coil" is

TABLE 2

Average Error in Gradient Field for Parallel-Wound *Z*-Gradient Coils

	Number of turns	Current ratios	Wire positions (z/a)	Average error (%)
A	8	1:1:1:9	0.437, 0.792, 0.970, 1.491	3.14
B	6	1:1:5.76	0.595, 0.981, 1.514	5.32
C	8	1:1:1:9.71	0.426, 0.794, 0.894, 1.466	2.59
D	10	1:1:1:1:11.17	0.414, 0.748, 0.886, 1.079, 1.534	1.71
E	6	1:1.36:7.19	0.572, 0.973, 1.545	5.09

introduced for this structure. Bolinger *et al.* (*14*) have described a UCB coil with spacings between wires determined by a geometrical construction in which intervals of equal length along a diameter are projected onto the circumference. This structure will be referred to as BB.

The rationale for studying the UCB is that it can provide additional options for the optimization of the free space Q. For very small coils at 1.5 T, for larger coils at lower fields, and for coils intended for use with nuclei of low magnetogyric ratio, body losses do not necessarily dominate coil losses. In these situations the highest possible free space Q commensurate with a specified degree of RF field homogeneity is required. The Qs of CBCs are not very high in our experience, and Qs of UCBs can be higher. Often the free space Qs are determined, at least in part, by nonideal characteristics of the capacitors (both dielectric loss and joulean losses in the plates of the capacitors). The higher inductance of a series-wound coil would reduce the capacitance required to resonate at a given frequency. The UCB thus provides additional design flexibility in management of the capacitive elements of the resonant structure.

In this application, the parameters to be adjusted are the azimuthal position of the axial wires. The wires are initially spaced to approximate a sinusoidal distribution as in the BB, and are assumed infinitely long. The points that define the ROI were spaced 5° apart in the ϕ direction, and 0.1 a apart in the radial direction. The radial extent of the ROI was 0.8 a. Again, the error function was weighted by the radius of the points to give volume-weighted errors. The calculated field on axis was considered to be the desired field.

The optimized configurations for 4- to 16-turn UCBs are given in Table 3 for a ROI of radius 0.8 a. The dependence of the AE upon the number of wires is shown in Fig. 4. Also shown in Fig. 4 for comparison are the errors for CBCs and BBs with the same numbers of wires. For the CBCs, the wires were evenly spaced, and the currents were assumed to vary sinusoidally. For the BBs, each wire carried the same current, and the wire positions were a geometrical approximation (*14*) of a sinusoidal distribution. For this ROI and this range of numbers of wires, the UCBs have slightly better field homogeneity than the CBCs, and the BBs.

DISCUSSION

This conjugate gradient descent algorithm has been shown to produce cylindrical z-gradient coils and constant current birdcage coils with very good field uniformity

TABLE 3

Wire Positions and Average Errors for Uniform Current Birdcage Coils

Number of wires	Wire positions in degrees	Average error (%)
4	25.39	11.68
8	13.40, 45.72	3.69
12	9.14, 28.89, 54.58	1.52
16	6.94, 21.51, 37.50, 59.70	0.79

over large ROIs. It has also been shown to allow for the optimization of inductance, efficiency, and parallel configurations.

An attractive feature of this algorithm is its flexibility. The desired fields can be prescribed arbitrarily over a ROI of any shape by specifying the set of points over which the fields are calculated and optimized. A field that varies irregularly can be prescribed by directly assigning a numerical value for the desired field at each point. Constraints on the positions of current elements are also easily applied. If a current element is moved because of a constraint that the designer wishes to apply, the algorithm will attempt to compensate for such arbitrary movements in further iterations. For instance, to guarantee that current elements do not overlap, the program can check for the overlap after each iteration, and directly move the elements apart. Similarly, limits can be placed on the positions of the current elements to restrict the overall dimensions of the coil. Because we use a discrete current model, and direct Biot–Savart field calculations, our method is very general. We feel that our method has

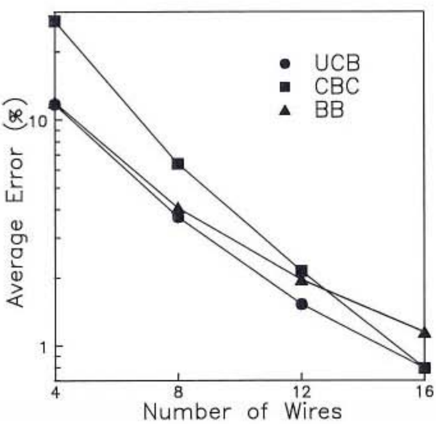

FIG. 4. Average error vs number of wires for uniform current birdcage coils. Average error is defined by Eq. [7]. UCB represents uniform current coils, CBC and BB represent conventional birdcage coils and Bolinger's birdcage coils, respectively, with the same number of turns. ROI is a cylinder of radius 0.8 a.

advantages over other design techniques, particularly in cases where the number of current elements is small, strict constraints are placed on the coil dimensions, or the geometry of the coil is not easily handled by other methods.

Potential applications of the algorithm include the design of windings to produce the main static field, gradient fields, and shim fields, and also the design of RF coils with specified sensitivity profiles.

Shielded gradient coils can be designed by distributing currents on two coaxial cylinders, and including the exterior of the coil in the ROI. A uniform gradient can be specified within the inner cylinder and zero field specified outside the coil.

We are currently applying our design methods to cylindrical transverse gradients. In this application, the current loops that make up the coil are broken down into small wire segments, and these segments are moved individually. Because the segments can move in both axial and azimuthal directions, an additional level of dimensionality is added to the parameter space, but the algorithm remains the same. We are also evaluating other numerical optimization techniques using this same discrete current element model.

Because the coil and the ROI are not required to be cylindrical, other configurations can be examined with this algorithm. Application to flat or biplanar geometries in which current is confined to one or two planes is easily implemented because the field perturbations that accompany the movement of current elements in a plane are easily calculated. The flexibility in the specification of the ROI and the ability to place arbitrary constraints on the positions of current elements also suggest the design of local gradient and/or RF coils for specific anatomic locations. Thus coils could be made both to fit the anatomic site and to produce the desired field profile.

ACKNOWLEDGMENTS

The authors would like to thank Dr. J. B. Kneeland and James D. Christensen for helpful discussion and suggestions. This work was supported by Grants CA41464 and RR01008 from the National Institutes of Health.

REFERENCES

1. F. Romeo and D. I. Hoult, *Magn. Reson. Med.* **1,** 44 (1984).
2. R. A. Compton, U.S. Patent no. 4,456,881 (1984).
3. R. Turner, *J. Phys. D* **19,** L147 (1986).
4. R. Turner, *J. Phys. E* **21,** 948 (1988).
5. K. H. Schweikert, R. Krieg, and F. Noack, *J. Magn. Reson.* **78,** 77 (1988).
6. L. Hasdorff, "Gradient Optimization and Nonlinear Control," Wiley, New York, 1976.
7. J. Mao, T. H. Mareci, K. N. Scott, and E. R. Andrew, *J. Magn. Reson.* **70,** 310 (1986).
8. C. E. Hayes, W. A. Edelstein, J. F. Schenck, O. M. Mueller, and M. Eash, *J. Magn. Reson.* **63,** 622 (1985).
9. R. A. Beth, *J. Appl. Phys.* **37,** 2568 (1966).
10. J. H. Coupland, *Nucl. Instrum. Methods* **78,** 181 (1970).
11. K. Halbach, *Nucl. Instrum. Methods* **78,** 185 (1970).
12. D. E. Lobb, *Nucl. Instrum. Methods* **64,** 251 (1968).
13. A. Sotgiu and J. S. Hyde, *Magn. Reson. Med.* **3,** 55 (1986).
14. L. Bolinger, M. G. Prammer, and J. S. Leigh, Jr., *J. Magn. Reson.* **81,** 162 (1988).