Bounded Nonlinear Optimization to Fit a Model of Acoustic Foams

MA5630 – Numerical Optimization

Andrew Willemsen
Yinan Zuo

4/29/2011

1. Introduction

Open cell porous materials, such as polymer foams, dissipate acoustic waves through viscous and thermal losses as the sound moves through the foam. The dissipation of acoustic energy is typically described in terms of its surface impedance. Several models have been developed relating the surface impedance of porous materials to a set of macroscopic morphological parameters describing the geometry and connectedness of the cells within the porous material. However, many of these morphological properties cannot be easily measured by direct experiments. Instead, these properties can be determined through model fitting methods. These methods initially measure the surface impedance of the material directly, and then the morphological properties are estimated so that the model prediction of surface impedance matches observation. This is most accurately achieved by using optimization techniques to minimize the square of the difference between the predicted and observed surface impedance.

The purpose of this project is to examine a number of techniques for performing this model-fitting optimization. Five model parameters will be estimated by optimization algorithms. All optimization implemented optimization algorithms used in this project are pre-built functions in the Matlab Optimization Toolbox. The goal is to determine which methods most accurately and efficiently optimize the model parameters. Additionally, the effects of initial conditions of the optimization problem and the selection measured data used to build the cost function will be examined. Finally, the estimated parameters will be compared to measured data to determine how closely the optimized model matched the actual foam properties.

2. Problem Background

2.1 Model of Surface Impedance for Acoustic Foam

The acoustic properties of foam can be described in terms of its surface impedance. The surface impedance is dependent on a number of material properties of the foam material and is thus modeled based on these same material properties [1]. These material properties are flow resistivity (σ), porosity (φ), tortuosity (α_∞), viscous pore shape parameter (c), and thermal pore shape parameter (c'). The model for surface impedance is given by

$$Z(\sigma,\phi,\alpha_\infty,\Lambda,\Lambda',\omega) = -j\frac{\sqrt{\rho(\sigma,\phi,\alpha_\infty,\Lambda,\Lambda',\omega)K(\sigma,\phi,\alpha_\infty,\Lambda,\Lambda',\omega)}}{\varphi Z_0}\cot\left(\omega L\sqrt{\frac{\rho(\sigma,\phi,\alpha_\infty,\Lambda,\Lambda',\omega)}{K(\sigma,\phi,\alpha_\infty,\Lambda,\Lambda',\omega)}}\right) \quad (1)$$

Equation 1 shows that the model is not only dependent on the 5 material properties previously discussed, but is also dependent on frequency, ω. The j is the model is the imaginary number and Z_0 and L are constants describing the dimensions of the foam and the properties of air. Although, the surface impedance is complex and frequency-dependent, the material properties are real-valued and independent of frequency. The functions ρ and K detail the relationship between surface impedance and the material properties of the foam and are given by

$$\rho(\sigma,\phi,\alpha_\infty,\Lambda,\Lambda',\omega) = \rho_0\alpha_\infty\left(1+\frac{\varphi\sigma}{j\omega\rho_0\alpha_\infty}\sqrt{1+j\frac{\omega\rho_0\alpha_\infty c^2}{2\sigma\varphi}}\right) \quad (2)$$

$$K(\sigma, \phi, \alpha_\infty, \Lambda, \Lambda', \omega) = \gamma P_0 \left(\gamma - \frac{\gamma - 1}{1 + \frac{\sigma \varphi c'^2}{j\omega N_{Pr} \rho_0 \alpha_\infty} \sqrt{1 + j \frac{\omega N_{Pr} \rho_0 \alpha_\infty}{2\sigma \varphi c'^2}}} \right)^{-1} \tag{3}$$

In Equations 2 and 3, the terms ρ_0, γ, P_0, and N_{pr} are all constants describing the properties of air.

2.2 Cost Function and Constraints

While it is relatively simple to measure the surface impedance of a material, directly measuring the five material properties used to model surface impedance is often difficult and requires specialized equipment. Instead these parameters are often determined through model fitting techniques using measured values of surface impedance [1]. In this project, optimization algorithms were used to estimate the model parameters in order to minimize the difference between the measured and modeled surface impedance. Ideally, the difference between the modeled and observed surface impedance values would be zero, however, since there is noise in the measurement this is unlikely to occur. Since the difference between the measured and modeled impedance is complex-valued, which would cause difficulties in using common optimization techniques, the square of the magnitude of this complex-value is instead minimized here. Measured impedance data is available at many different frequencies and this was utilized when forming the cost function for the optimization problem. Since the material properties used in the model are constant at all frequencies, data at all frequencies can be used to estimate these properties. Thus, the sum of squares of the difference between the measured and modeled impedance was used in the cost function. This approach helps to eliminate the effect of measurement noise/error when estimating the model parameters through optimization since the parameters will fit over a wide frequency range rather than at only a single frequency. The cost function examined in this paper is given as

$$F(\boldsymbol{x}) = \sum_{i=1}^{N} \left| Z_i^{model}(\boldsymbol{x}, \omega_i) - Z_i^{meas}(\omega_i) \right|^2 \tag{4}$$

In this equation, N is the number of frequencies included in the cost function. The 5 material properties in the model are now designated by the vector **x**, which this has five elements. It is also noted that for simplicity, the squared values of the properties c and c' are used as function parameters here.

Since the five model parameters are physical quantities, their values are limited to values which are physically feasible [1]. Thus, a number of constraints are introduced into the optimization problem. Each of the first four parameters is constrained by simple upper and lower bounds. However, the fifth parameter, c'^2, has the constraint that it must be less than or equal to the parameter c^2. However, this linear inequality constraint can be avoided due to the lower bounds of both c^2 and c'^2 both being the same value, 0.1. This is achieved by defining a new parameter, s, which is given as

$$s = \frac{c'^2 - 0.1}{c^2 - 0.1} \tag{5}$$

Solving this relationship for c'^2 and substituting this value into the equation defining the bounds of c'^2 yields simple bound constraints on the parameter *s*. Thus, by using the parameter *s* instead of the parameter c'^2, all function parameters now have simple upper and lower bounds. Making this substitution and defining the vector **x** as [σ, φ, α_∞, c^2, s], Equations 2 and 3 become

$$\rho(\mathbf{x},\omega_i)=\rho_0 x_3\left(1+\frac{x_1x_2}{j\omega\rho_0x_3}\sqrt{1+j\frac{\omega\rho_0x_3x_4}{2x_1x_2}}\right) \tag{6}$$

$$K(\mathbf{x},\omega_i)=\gamma P_0\left(1-(\gamma-1)\left(1+\frac{x_1x_2(x_4x_5-0.1x_5+0.1)}{j\omega N_{pr}\rho_0x_3}\sqrt{1+\frac{\omega N_{pr}\rho_0x_3}{2x_1x_2(x_4x_5-0.1x_5+0.1)}}\right)^{-1}\right)^{-1} \tag{7}$$

The bounds are defined as follows

$$\text{bound constraints}=\begin{cases}10^3\le x_1\le 10^6\\ 0.7\le x_2\le 1\\ 1\le x_3\le 4\\ 0.1\le x_4\le 10\\ 0\le x_5\le 1\end{cases} \tag{8}$$

Thus, the optimization problem is defined as the minimization of F defined by Equation 4 with respect to the vector **x**, such that the bound constraints defined by Equation 8 are satisfied.

2.3 Frequency Range Considerations

The last consideration when forming the optimization problem is choosing the frequency range over which the model is to be optimized. Although the model parameters are constant across all frequencies, the influence of each parameter on the value of the model is dependent on frequency. For example, at low frequencies, the modeled surface impedance value is highly dependent on the value of the parameter x_1 but have almost no dependence on the value of x_4. The same is not true at higher frequencies. This can lead to poor estimation of some of the parameters due to scaling issues. Parameters which do not influence the value of the model as much as other parameters could not be accurately estimated since the gradient with respect to those parameters is so much lower than for the other parameters. Ideally, a frequency range should be chosen over which all the model parameters have the same influence on the value of the model. However, due to the inherent scaling of the model, particularly with respect to x_1, this is not possible without altering the model parameters. Instead, a particular frequency range was chosen over which all five parameters influenced the model value to some degree, although not equally among all five parameters.

The frequency range of interest was split into three different “zones”. These zones are dependent on the location of the peak in the acoustic absorption curve for the material being studied. Absorption is computed directly from the surface impedance, which also exhibits this peak. Figure 1 illustrates the three frequency zones for a common acoustic foam. Zone 1 is the frequency range below the peak, Zone 2 is the frequency range directly around and including the peak, and Zone 3 is the frequency range after the peak.

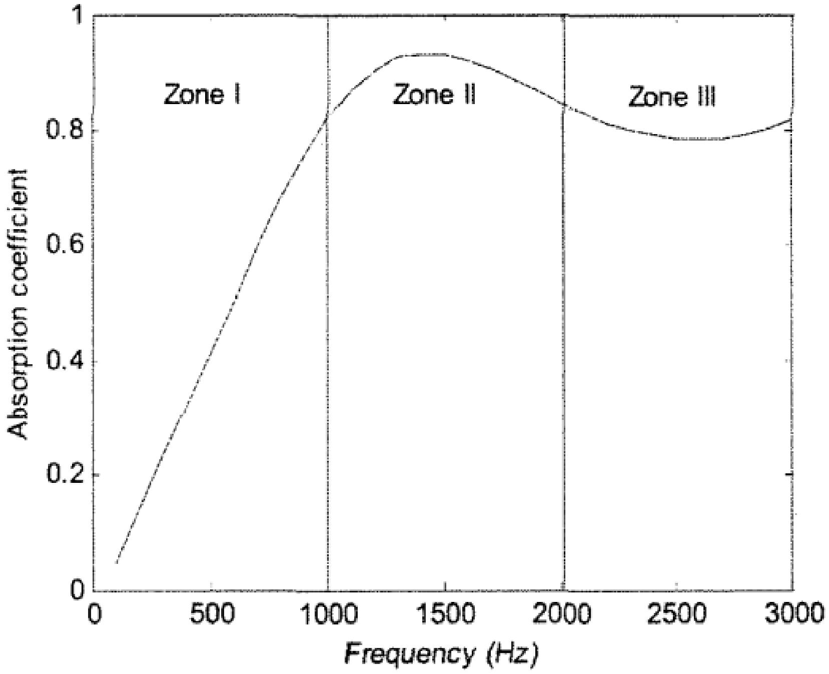

Figure 1: Three frequency range zones for acoustic foams [1]

The degree of influence each parameter had on the value of the modeled surface impedance was described as either a primary influence or a secondary influence. The primary influence(s) is what most controls the value of the model. The secondary influences also effect the model value but to a much lesser degree than the primary influence. Finally, any parameter that influenced the value of the model to a degree less than the secondary influences was deemed to have no significant influence. The influence of each of the five parameters in all three frequency zones was determined. This was based on graphical comparisons of 2-D planes of the cost function (three parameters held constant). The parameters x_1 and x_3 were found to have primary and secondary influences, respectively, for all three frequency zones. Thus, these two parameters were compared to the remaining three parameters to determine their degree of influence. Figure 2 shows the relationship between the parameters x_1 and x_3 in Zone 2, although it is similar for all three zones. It can be seen that the contours of the function are almost exactly perpendicular to the x_1 axis, with a long shallow valley along to the x_3 axis. Figures 3 through 5 illustrate the influence of the x_2 parameter over the three frequency zones. These show that x_2 has a primary influence in Zones 2 and 3 and a secondary influence in Zone 1. Figures 6 through 8 show that the x_4 parameter has a secondary influence in Zones 2 and 3, but no influence in Zone 1. Parameter x_5 has the opposite behavior of x_4, having a secondary influence in Zones 1 and 2, but no influence in Zone 3, as shown in Figures 9 through 11.

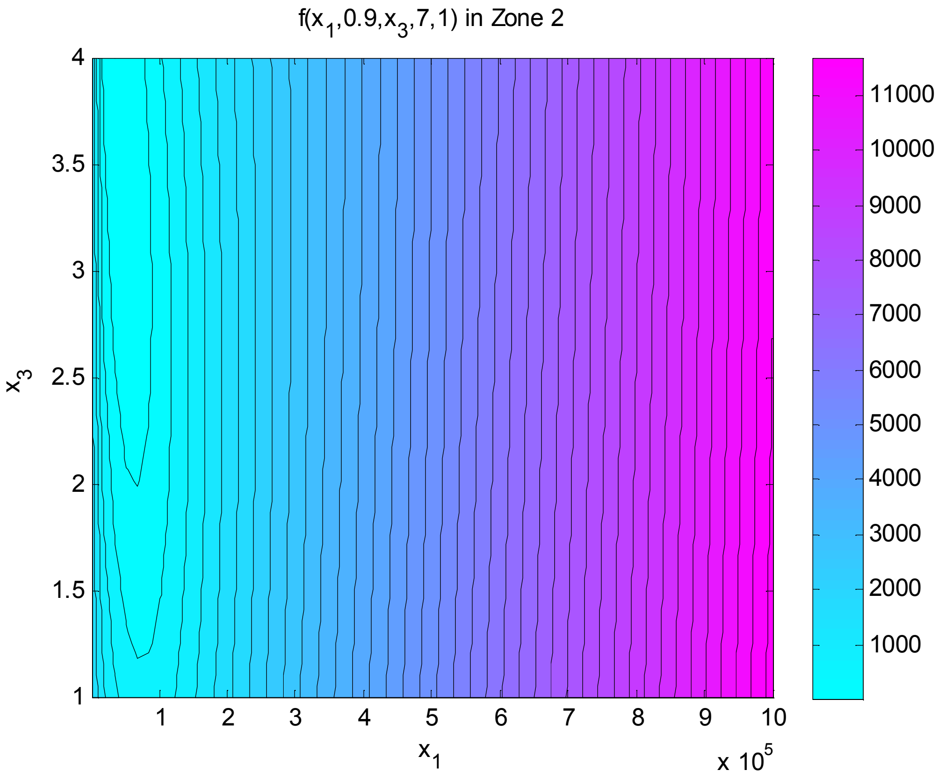

Figure 2: Primary and secondary influences of the x_1 and x_3 parameters, respectively in Zone 2.

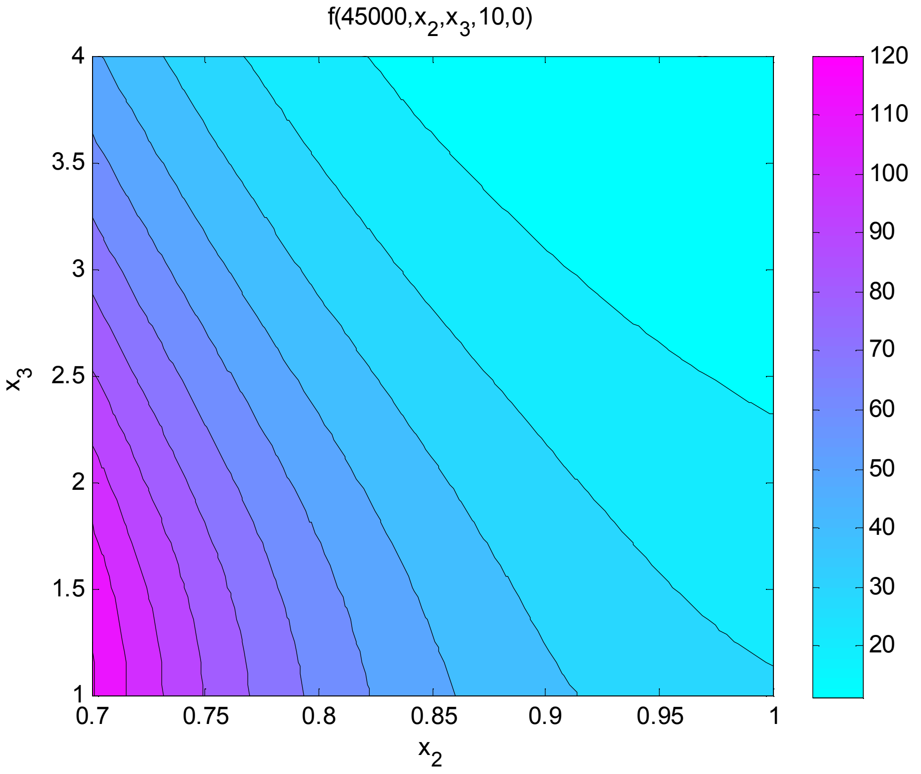

Figure 3: Secondary influence of x_2 in comparison to x_3 in Zone 1.

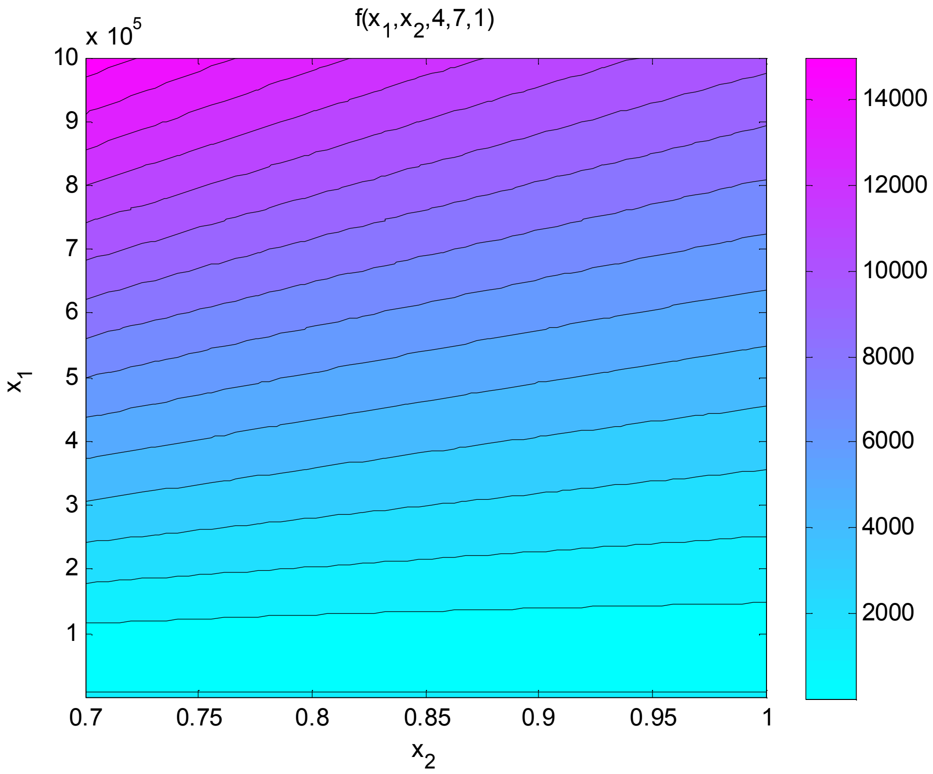

Figure 4: Primary influence of x_2 in comparison to x_1 in Zone 2.

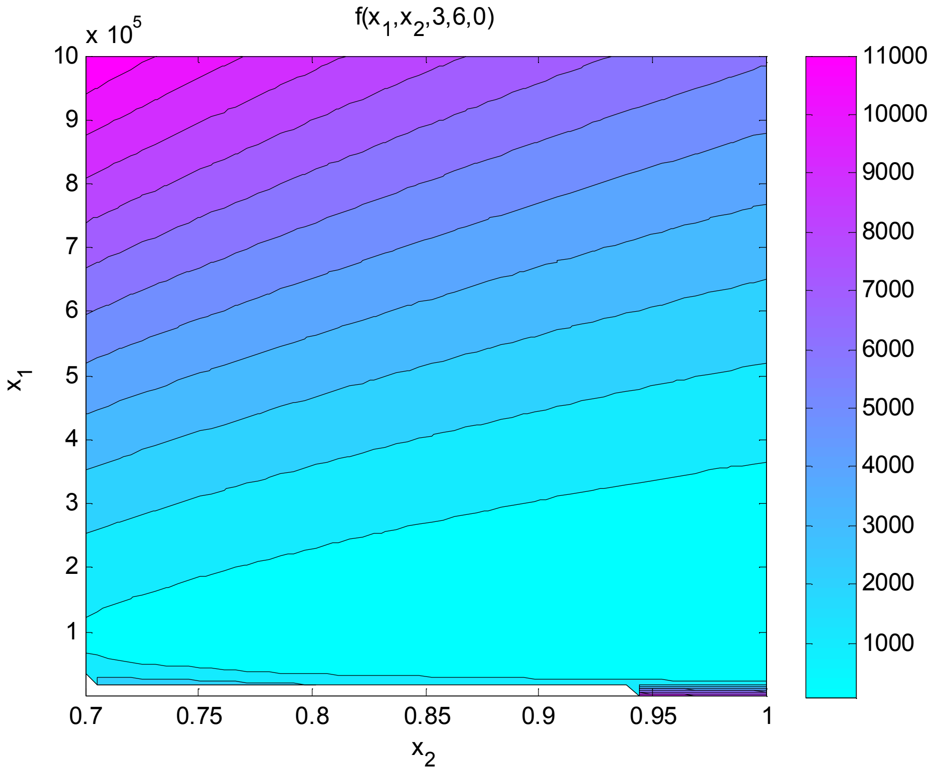

Figure 5: Primary influence of x_2 in comparison to x_1 in Zone 3.

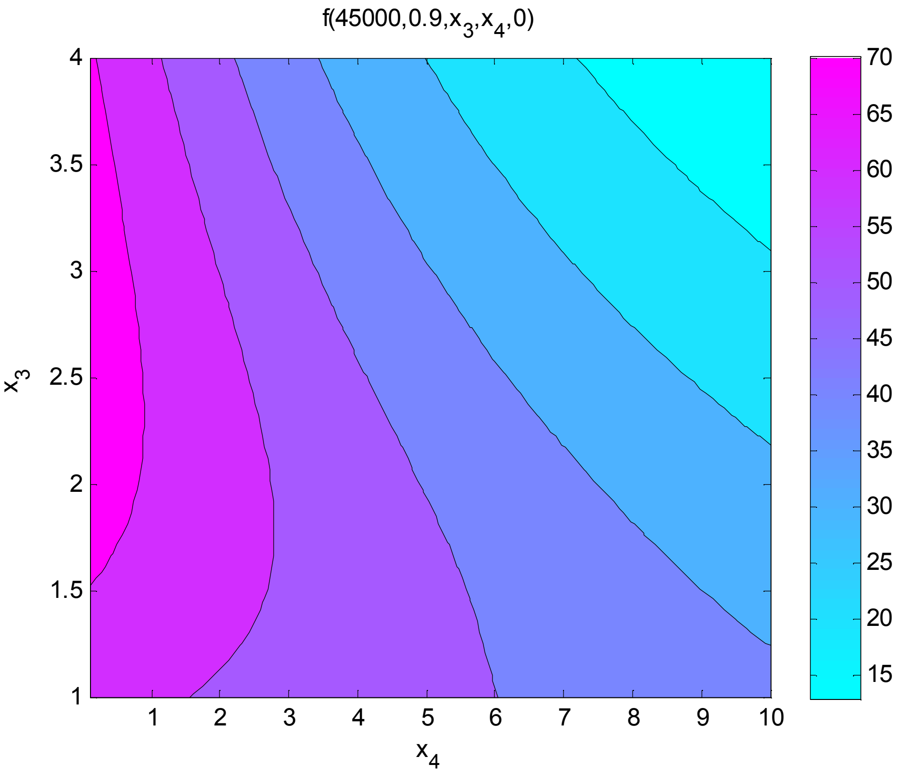

Figure 6: Negligible influence of x_4 in comparison to x_3 in Zone 1.

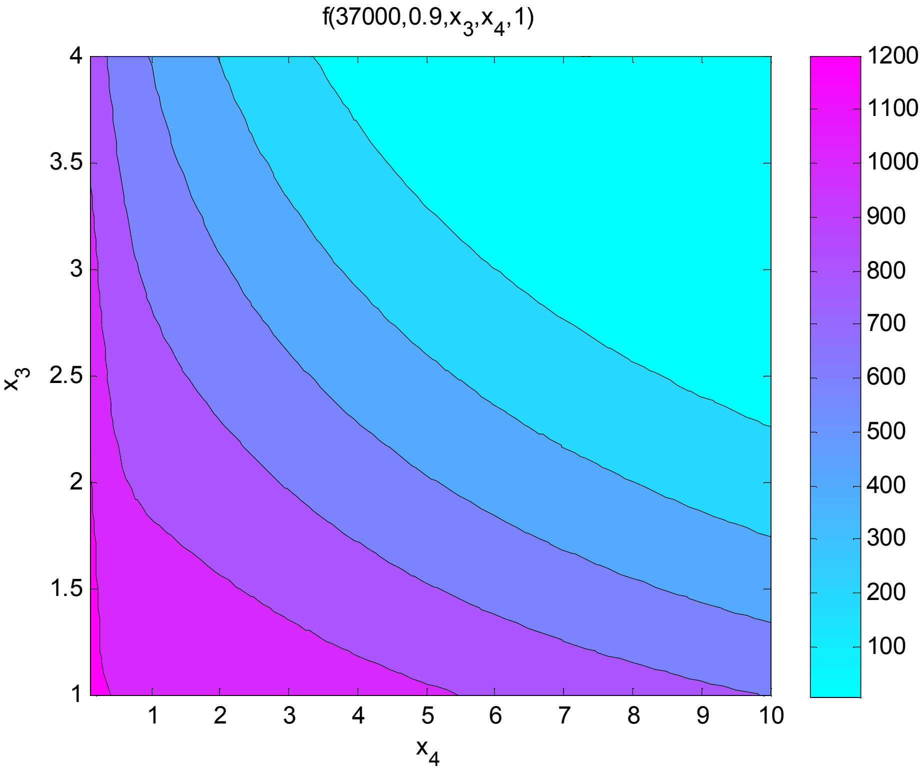

Figure 7: Secondary influence of x_4 in comparison to x_3 in Zone 2.

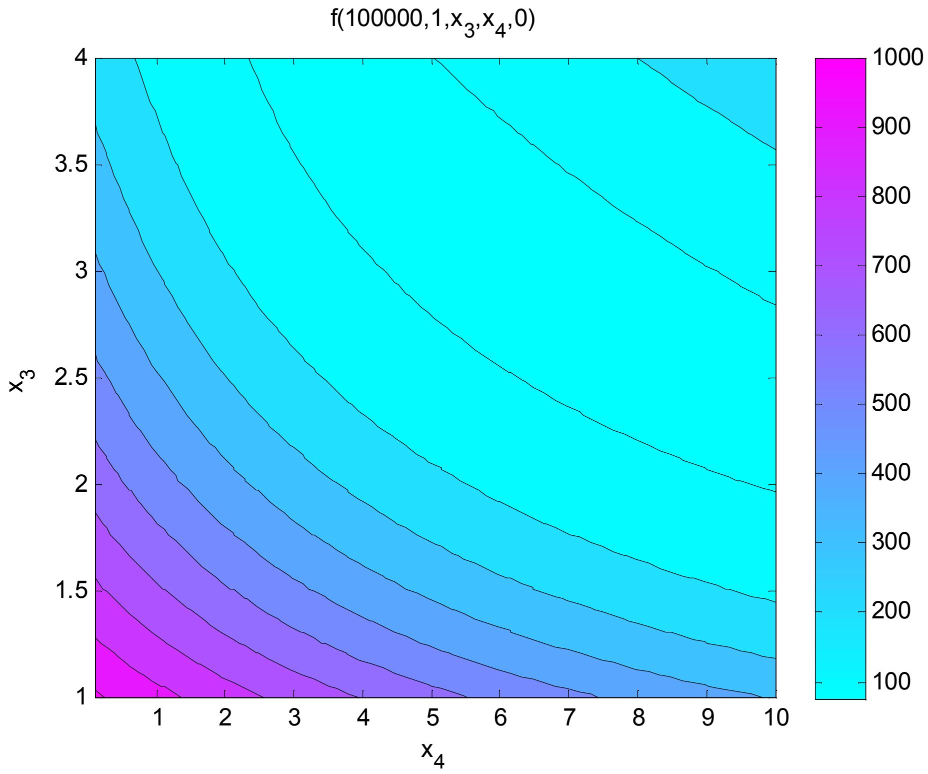

Figure 8: Secondary influence of x_4 in comparison to x_3 in Zone 3.

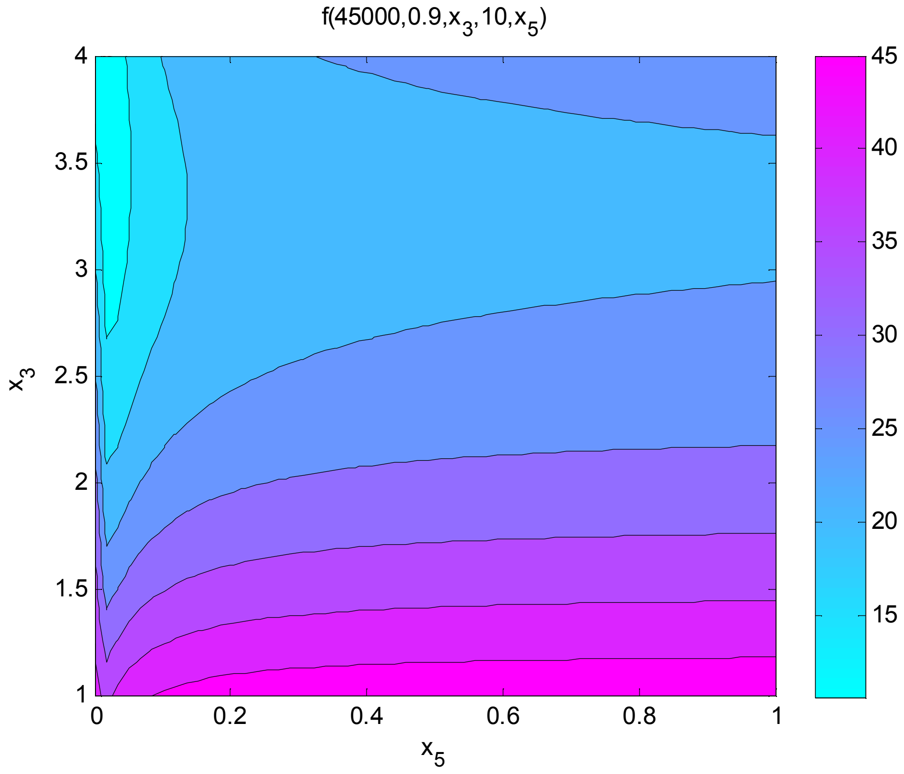

Figure 9: Secondary influence of x_5 in comparison to x_3 in Zone 1.

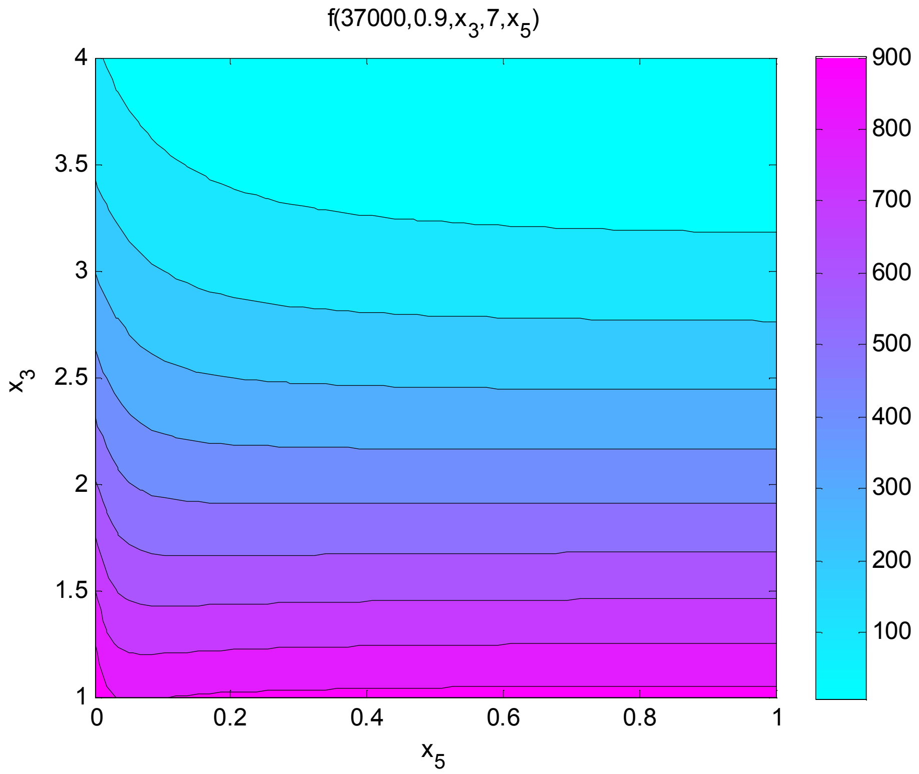

Figure 10: Secondary influence of x_5 in comparison to x_3 in Zone 2.

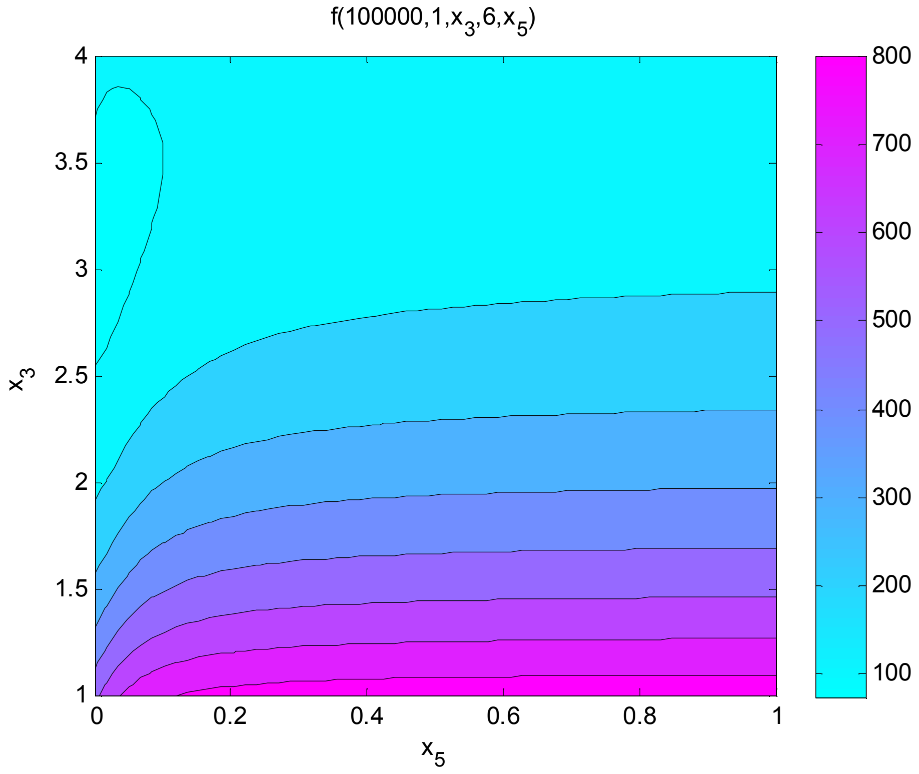

Figure 11: Negligible influence of x_5 in comparison to x_3 in Zone 3.

Based upon this analysis, the combination of both Zones 1 and 2 was selected for the optimization analysis. In this frequency range, all five parameters have at least a secondary influence on the value of the model. Additionally, the peak value of surface impedance is in this range. At the peak value, the measurement is least effected by measurement noise, which will help to improve the estimate of the model parameters resulting from the optimization. It should be noted that in this frequency range there is still a major scaling problem with the x_1 parameter. Its value and gradient are much larger than for the other parameters. This can cause some problems with accurately estimating the remaining four parameters. Since the literature discussing the optimization of this model does not discuss scaling the parameters, it was left un-scaled for this project as well. The effects of this are seen later in the discussion of results. In the frequency range covered in Zone 1 and 2, the surface impedance was measured at 246 different frequencies so that the cost function is the summation of 246 residuals.

3. Optimization Algorithms

In this project, two different pre-built optimization functions in Matlab's Optimization Toolbox were tested. These were the *fmincon* function for general constrained optimization and the *lsqnonlin* function for nonlinear least squares algorithms. These functions were tested to determine which function and which options (i.e. algorithms) within each function worked best at optimizing the cost function. For the *fmincon* function, the active set, sequential quadratic programming (SQP), and interior point algorithms were implemented. For the *lsqnonlin* function, only the trust-region reflective algorithm was implemented since it was the only one which allowed bound constraints. The default options for convergence tolerance were used for both functions. It was not possible to define the gradient of the cost function in functional form within Matlab, so the default option of estimating the gradient by finite differences was implemented as well. Where the Hessian was used, a BFGS update strategy was implemented by Matlab. Descriptions of the four optimization algorithms tested are given in this section.

3.1 Active Set Method

The active set method was the first of the algorithms implemented using Matlab's *fmincon* function [2,3]. The general formulation of minimizing functions subject to constraints on the variables is

$$\min_{x \in \mathbf{R}^n} f(x) \qquad \text{subject to} \quad c_i(x) = 0, i \in \mathcal{E} \qquad \text{and} \quad c_i(x) \geq 0, i \in \mathcal{I} \tag{9}$$

The active set consists of the equality constraints from $\mathcal{E}$ and the inequality constraints where $c_i(x) = 0$ is satisfied. The constraint is called active if $c_i(x) = 0$ and called inactive if $c_i(x) > 0$ is satisfied. Thus, the active set can be defined as

$$\mathcal{A}(x) = \mathcal{E} \cup \{i \in I \mid c_i(x) = 0\} \tag{10}$$

The Lagrange function of this algorithm is defined by

$$\mathcal{L}(x,\lambda) = f(x) - \sum_{i\in\mathcal{A}(x)} \lambda_i c_i(x) \tag{11}$$

As an algorithm of optimization, the active determines which constraints determine the final optimization and which do not. The algorithm starts from an initial point. For each step, the equality problem defined by the active set is solved approximately. Then, for the active set, the Lagrange multipliers are computed. Those constraints whose Lagrange multipliers are negative are removed and infeasible constraints are searched. Those above are repeated until the result of the optimization is good enough.

3.2 Sequential Quadratic Programming Method

The second *fmincon* algorithm used use the sequential quadratic programming method [2,3]. Rather than convert it to a sequence of unconstrained minimization problems, one attempts to solve a nonlinear program directly is SQP. It is an iterative method. Requirement for the problem is twice continuous differentiability of the function and continuous differentiability of its constraints. For significant nonlinearities in the constraints, SQP methods show their strength especially. The search direction of each iteration is defined by

$$\min_p \frac{1}{2} p^T \nabla^2_{xx} \mathcal{L}(x_k,\lambda_k) p + \nabla f(x_k)^T p + f_k$$

subject to

$$\begin{aligned} &\nabla c_i(x_k)^T p + c_i(x_k) = 0, i \in \boldsymbol{\mathcal{E}} \\ &\nabla c_i(x_k)^T p + c_i(x_k) \geq 0, i \in \boldsymbol{I} \end{aligned} \tag{12}$$

where

$$\mathcal{L}(x,\lambda) = f(x) - \sum_{i\in I\cup\mathcal{E}} \lambda_i c_i(x)$$

This problem can be solved by the algorithms for quadratic programming. By solving for p_k, the solution of kth step, and λ_{k+1}, the Lagrange multiplier, the new iterate can be found by $(x_k + p_k, \lambda_{k+1})$. The process is repeated until the optimization is good enough.

3.3 Interior Point Method

The third algorithm implemented by Matlab's *fmincon* function is the interior point algorithm [2,3]. This algorithm can be used for any general optimization problem with equality or inequality constraints. However, for the interior point method, the inequality constraints are made into equality constraints through the use of slack variables. These slack variables are updated at each iteration, and are scalar values greater than zero. Additionally, instead of minimizing the function *f*, this method minimizes a function *f* which is the original function with a "barrier function" dependent on a barrier parameter which is greater than zero and the slack variables *s*. The minimization problem is then rewritten as

$$\min_{x,s} f_\mu(x,s) \tag{13}$$

such that

$$c_E(x) = 0$$
$$c_I - s = 0$$

where $f_\mu(x,s) = f(x) - \mu \sum_{i=1}^{m} \ln s_i$

In this problem, c_E and c_I are the equality and inequality constraints, respectively, and m is the number of inequality constraints. As the algorithm converges to the global minimum, the parameter converges to zero, so that the minimum of the function f converges to the minimum of the function f. In Matlab, the steps for the vectors **x** and **s** and the Lagrange multipliers for the constraints at each iteration are computed by one of two different approaches. The first, default approach is to solve a system of KKT equations for this problem using linear algebra. The KKT system used in this subproblem is given by

$$\begin{bmatrix} H & 0 & A_E^T & A_I^T \\ 0 & SZ & 0 & -S \\ A_E & 0 & I & 0 \\ A_I & -S & 0 & I \end{bmatrix} \begin{bmatrix} p_x \\ p_s \\ -p_y \\ -p_z \end{bmatrix} = \begin{bmatrix} \nabla f - A_E^T y - A_I^T z \\ SZ - \mu e \\ c_E \\ c_I + s \end{bmatrix} \tag{14}$$

where H is the hessian of the Lagrangian, A_E and A_I are the Jacobian matrices of the equality and inequality constraints, respectively, y and z are the Lagrange multipliers for the equality and inequality constraints, respectively, S is a diagonal matrix of the slack variables, and Z is a diagonal matrix of the Lagrange multipliers of the inequality constraints. This symmetric system is solved by LDL factorization. The Hessian can be computed directly from the function by

$$H = \nabla^2 f + \sum_i y_i \nabla^2 c_{E,i} + \sum_i z_i \nabla^2 c_{I,i} \tag{15}$$

However, since the Hessian of the cost function was not originally formed for this project, the Hessian is instead estimated by a BFGS update.

The solution to the KKT equations can move to other stationary points in the function, so it will only move towards the minimum if the Hessian is positive definite. If the Hessian is indefinite for a particular iteration, Matlab uses an alternative method to compute the step. A conjugate gradient approach is used to minimize a quadratic approximation of the function f, adjusting both **x** and **s**. The Lagrange multipliers are first obtained by approximately solving the KKT equations in a least squares sense. The KKT equation solve is

$$\nabla f + \sum_i y_i \nabla c_{E,i} + \sum_i z_i \nabla c_{I,i} = 0 \tag{16}$$

The step for **x** and **s** are then determined from the quadratic approximation of f within a trust region given by

$$\min_{p_x, p_s} \nabla f^T p_x + \frac{1}{2} p_x^T H \, p_x + \mu e^T S^{-1} p_s + \frac{1}{2} p_s^T S^{-1} Z p_s \tag{17}$$

and subject to

$$c_I + A_I p_x + p_s = 0, c_E + A_E p_x = 0 \tag{18}$$

After the step is determined by either the KKT subproblem solver or the conjugate gradient trust-region approach, the algorithm checks for if the step is appropriate using a merit function. The barrier parameter   is then updated and the iterations continue until the first order optimality conditions are sufficiently near zero.

3.4 Nonlinear Least Squares – Trust-Region Reflection Method

The final optimization algorithm tested in this project was the nonlinear least squares algorithm using a trust-region reflective method for bound constrained problems. Least squares optimization is specifically used for model-fitting in which the cost function is a sum of squared residuals. The problem statement is restated as

$$\min_{\mathrm{x}} ||\mathrm{f(x)}||_2^2 \tag{19}$$

Here, *f* is a vector-valued function, each element being the residual between a measured value and modeled value (dependent on a set of modal parameters). Algorithms can take advantage of the special structure of the nonlinear least squares problem to minimize the function more efficiently [4]. For example, when computing the Newton step for a trust region algorithm (*s*), only the Jacobian of the vector-valued function (*J*) is needed and the method of preconditioned conjugate gradients is used to solve the normal equations given by

$$J^T J s = -J^T f \tag{20}$$

In Matlab's *lsqnonlin* optimization function, only the trust-region reflective algorithm can be used if bound constraints are applied to the problem. This algorithm is very similar to the trust-region method for unconstrained optimization with a few key modifications [3,5]. The first modification is in how the Newton step is computed from the quadratic model function. A vector *v* is defined for a particular value of *x* which is dependent on the gradient at the current location and the value of the upper and lower bounds for each element of *x*. This vector is defined as

$$v_i = \begin{cases} x_i - u_i & \text{if } g < 0 \text{ and } u_i < \infty \\ x_i - l_i & \text{if } g \geq 0 \text{ and } u_i > -\infty \\ -1 & \text{if } g < 0 \text{ and } u_i = \infty \\ 1 & \text{if } g \geq 0 \text{ and } u_i = -\infty \end{cases} \tag{21}$$

The values of the vector *v* are used to define a matrix *D* as follows

$$D(x) = \mathrm{diag}\left(|v(x)|^{-1/2}\right) \tag{22}$$

The first order necessary conditions at the minimum are then a modification of those for unconstrained minimization, based on the matrix *D* and are given by

$$D(x)^{-2}g(x) = 0 \tag{23}$$

The Newton step is then computed from the minimum of a modified quadratic approximation of the cost function given by

$$\min_{s}(D_k^{-1}g_k)^T s + \frac{1}{2}s^T(D_k^{-1}B_kD_k^{-1} + \mathrm{diag}(g_k)J_k^v)s \tag{24}$$

where J^v is the Jacobian of the vector *v*.

The second modification the trust-region reflective method makes to the trust-region algorithm for unconstrained optimization is how it handles steps which cross one or more of the bounds. This is handled by reflecting the step. If the step crosses one or more bounds, and the index of the first bound crossed is i, then the ith element of the step is reflected ($p_i^{\text{reflect}} = -p_i^{\text{orignal}}$). The other elements of the step are not changed. Thus, if a bound is crossed, the step taken at that iteration is changed to move away from the bound in the dimension that is in violation of the bounds.

4. Results and Discussion

4.1 Initial Point and Global and Local Minima

The first look at the optimization routines was to determine the effect of the initial point on the solution found by the different minimization routines. This was done by evaluating the optimization problem repeatedly for different random initial conditions. The initial point with variables x_1, x_2, x_3, x_4, x_5 was chosen randomly within their bounds defined by the physical problem. A pentagon plot is applied for used to illustrate the the range of the five variables (anticlockwise starting from the top right of the pentagon for x_1). For each of the algorithms, once a starting point is initiated, a corresponding green point is plotted in the pentagon. Figures 12 and 13 show the initial point for 1000 different runs of each optimization routine and their resulting solutions. Figure 12 illustrates the initial conditions that converged to the global minimum of the function. The green points along each parameter axis indicate the initial conditions used, while the crosses indicate the solution the algorithms converged to. The global minimum found indicated in Figure 12 was determined by taking the solution with the lowest function value at the minimum for all 1000 runs of the routines. The four algorithms were consistent with each other in finding this global minimum. Figure 13, on the other hand, illustrates the initial conditions and resulting solutions for the four algorithms when a nonglobal minimum was found. The red points indicate the initial point values, while the crosses indicate the solutions. Only solutions which are not the global minimum are plotted in Figure 13.

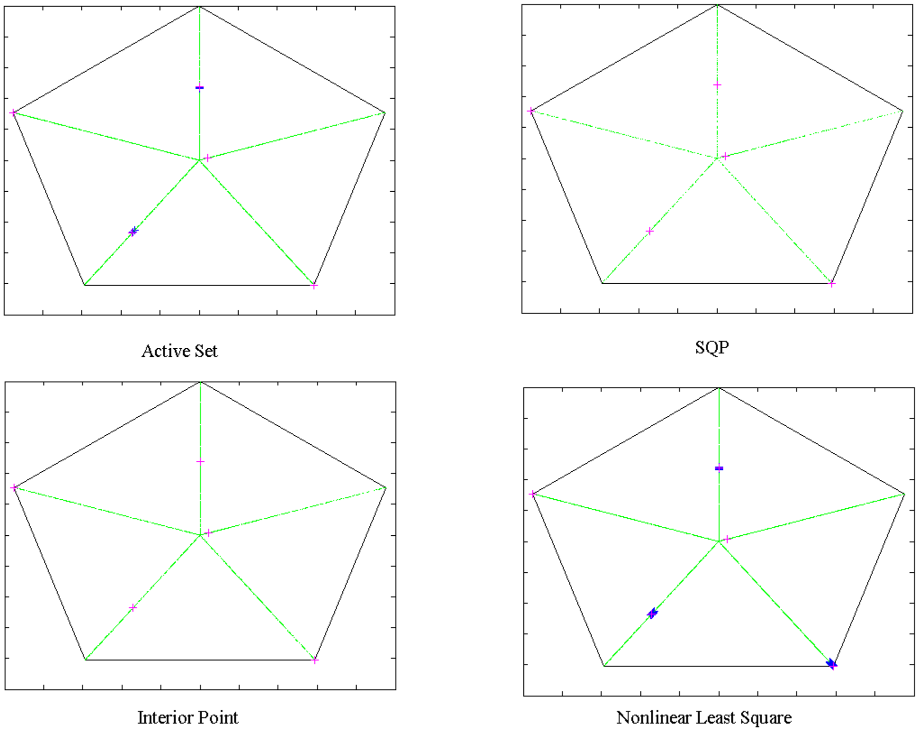

Active Set

SQP

Interior Point

Nonlinear Least Square

Figure 12: Initial points and optimization solution when global minimum found.

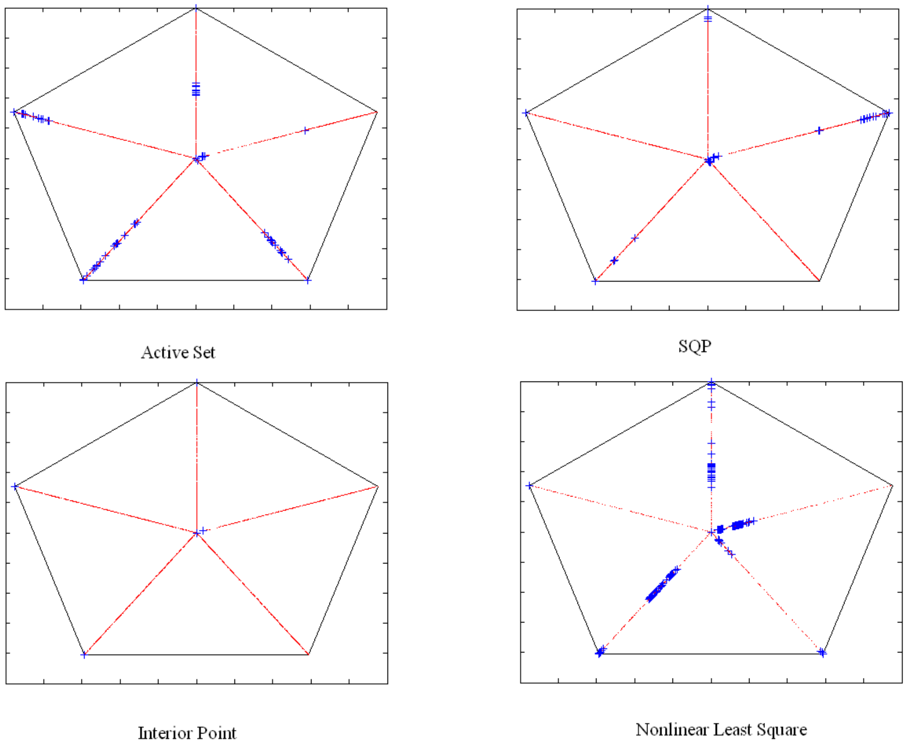

Active Set

SQP

Interior Point

Nonlinear Least Square

Figure 13: Initial points and optimization solution when only a local minimum found.

One can see from the figure that, for all of the algorithms, the optimization works correctly and each feasible initial value for each parameter can lead to the global minimum, if the other variables are chosen properly. However, the same can be said of convergence to one of the many local minimums. Any random initial point can lead to either the global minimum or the local minimum. However, as can be seen from the Figure 13, where there is a gap for the x_1 values, for certain range of x_1 values, the optimization always converges to the global minimum. This range is approximately 10^3 to 10^5. Thus, the variable x_1 plays a great role in the function value and also in the optimization process. Again this may be due to scaling issues with the parameter x_1.

4.2 Comparison of Algorithms

Using the same initial point (which was guaranteed to converge to the global minimum for all algorithms) the convergence properties, efficiency, and accuracy of the four optimization algorithms described in Section 3 were compared. A number of measurements of the algorithms' performance were compared including the final value of the function and the model parameters, the number of iterations, the number of function evaluations, and the value of the first order optimality criteria. Table 1 shows the results for the four algorithms.

Table 1: Performance criteria for four optimization algorithms.

Algorithm	**Function Value**	**Parameter Values**	**Iterations**	**First Order Optimality**	**Function Evaluations**
Active Set	33.4230	[46088, 0.8436, 4, 5.8401, 1]	31	7.8883E-03	191
SQP	33.4230	[46088, 0.8436, 4, 5.8401, 1]	36	2.8610E-06	231
Interior Point	33.4234	[46087, 0.8436, 3.9999, 5.8406, 0.9995]	40	2.0000E-06	253
Trust-Region Reflective	33.4254	[46262, 0.8432, 4, 5.8162, 1]	73	5.3090E-01	444

Based on the results shown in Table 1, several conclusions can be drawn on the performance of the various algorithms. The most efficient algorithm appears to be the active set algorithm used by *fmincon*. It converges to the minimum in the least number of iterations and requires the least number of function evaluations. This is expected since the bound constraints are fairly simple, and there are a relatively low number of possible active constraints. The active set method excels with problems such as this. However, the active set method did not produce the most accurate results, as both the SQP and interior point methods resulted in a much smaller value of the first order optimality at the global minimum. However, the final value of the cost function and model parameters for all three of the *fmincon* algorithms were almost identical (to the number of decimal places reported in Table 1). Since these estimates of the model parameters are based on a measurement of the surface impedance, which itself only has a certain degree of significant precision, all three *fmincon* algorithms results in solutions with acceptable accuracy to the precision necessary. It is also noted that while the active set and SQP methods result in two constraints being active, the interior point solution is only very close to the constraints, but still within the feasible region.

The trust-region reflective algorithm of the *lsqnonlin* performed the worst of the four algorithms. It was both the least efficient, requiring significantly more iterations and function evaluations to converge than the other methods, and the least accurate algorithm. This was not expected beforehand since the least squares algorithm takes advantage of the specialized definition of the problem to increase the efficiency of the optimization. However, since the global minimum was at a corner of two of the bound constraints, the trust-region reflective method had issues converging to the global minimum due to the reflection of the search direction away from these bound constraints. If the global minimum had been within the bounds, it is expected that the least squares trust-region reflective algorithm would have outperformed the other methods.

The convergence of the four algorithms was examined next. Figures 14 through 17 show both the value of the cost function and the value of the first order optimality term at each iteration for the four algorithms.

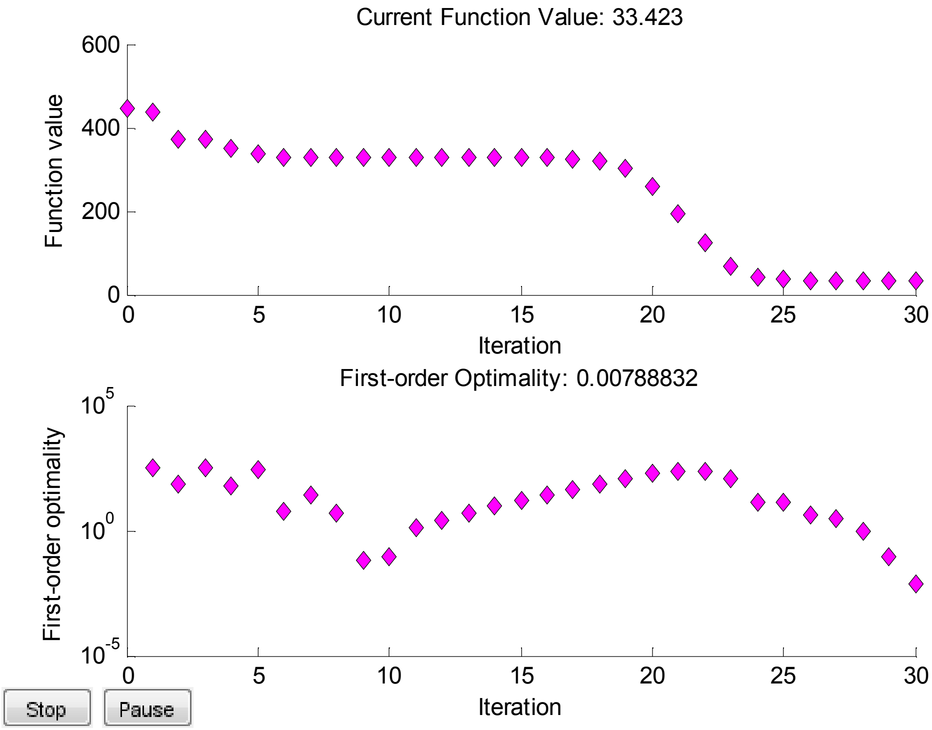

Figure 14: Convergence of the active set algorithm used by *fmincon*.

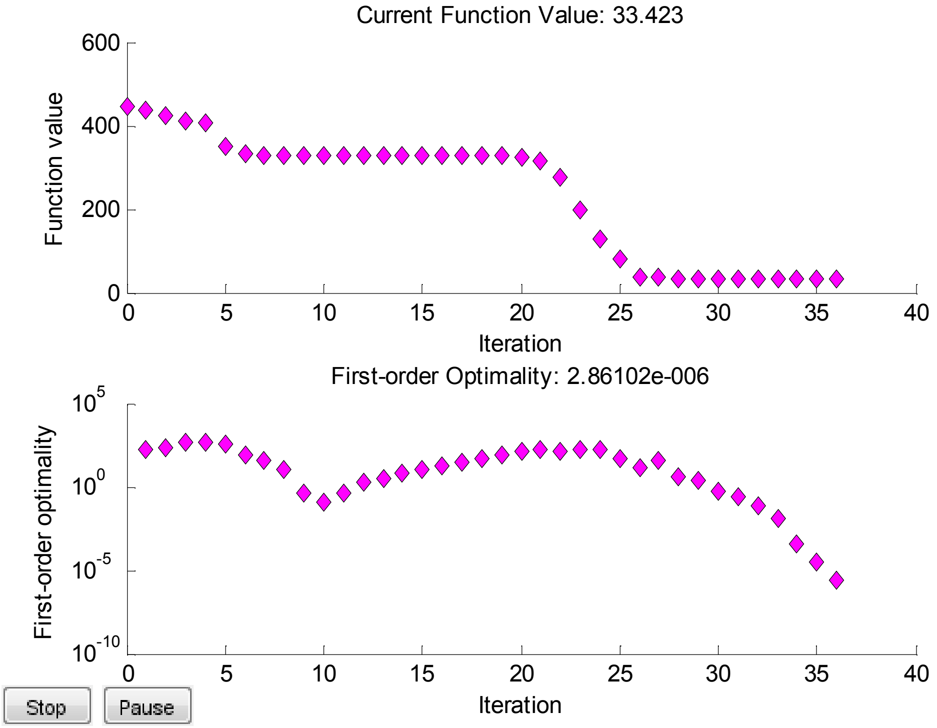

Figure 15: Convergence of the SQP algorithm used by *fmincon*.

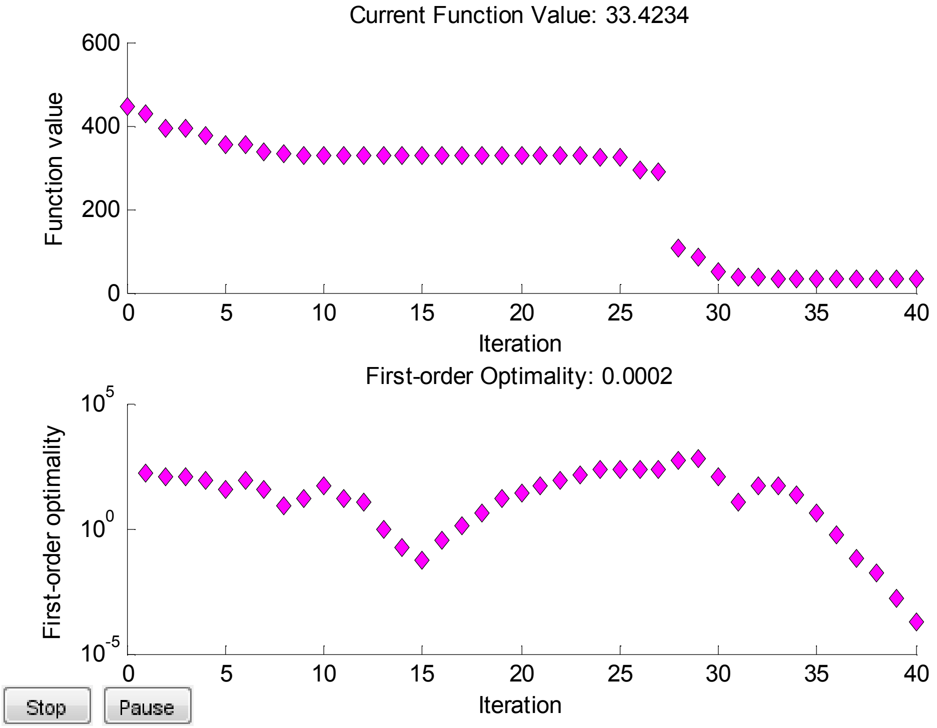

Figure 16: Convergence of the interior point algorithm used by *fmincon*.

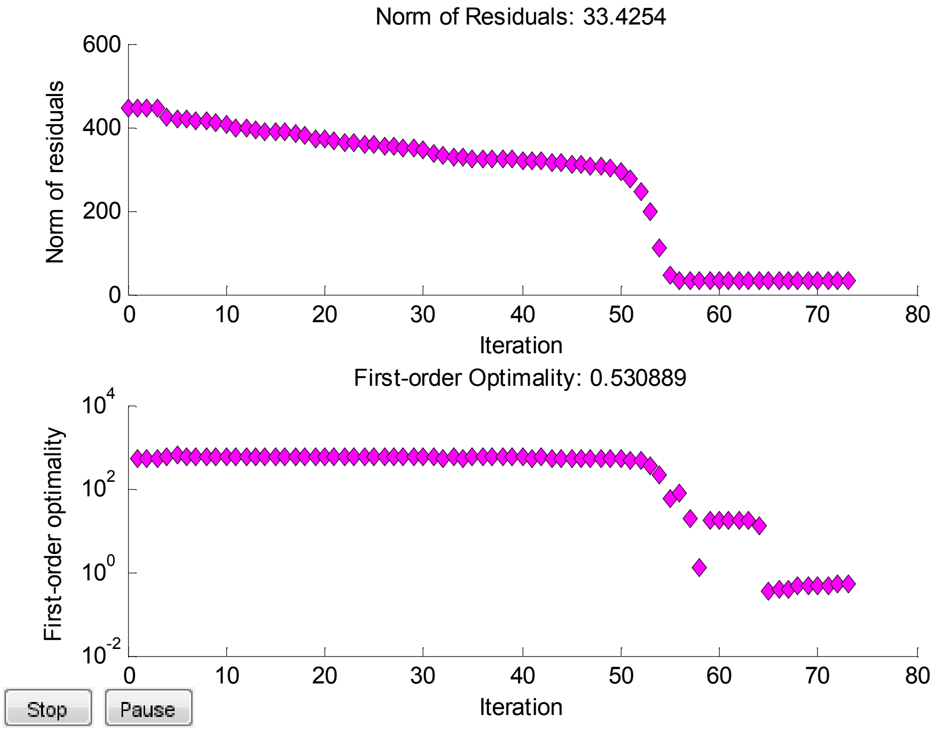

Figure 17: Convergence of the trust-region reflective algorithm used by *lsqnonlin*.

The convergence behavior of the three *fmincon* algorithms was fairly similar. For all three, the first order optimality at first decreased at a near linear rate and then dropped in value at around ten iterations, although the function value did not decrease significantly. This initial drop may have been due to being near one of the many local minimums of the function. After this initial drop, the first order optimality then increased again back to a steady linear convergence. Near the end of the iterations, the first order optimality then sharply decreased at a much faster convergence rate corresponding with a large decrease in the function value. This was due to being near the global minimum, where the convergence is expected to be closer to quadratic. However, after this sharp decrease in the first order optimality, several more iterations occurred in which the first order optimality and the function value were fairly steady. At this point, the algorithms were making small steps very near to the true minimum until the convergence tolerance was met.

The convergence behavior of the least squares trust-region reflective algorithm was slightly different than the other algorithms. The first order optimality does not decrease significantly at all for the first 50 iterations. The norm of the residuals, however, does steady decrease in the first 50 iterations. After 50 iterations, there was a sharp decrease in the first order optimality corresponding to a sharp decrease in the norm of the residuals, similar to the sharp drop observed for the *fmincon* algorithms. Finally, at around 65 iterations, the first order optimality again levels off and does not decrease substantially, even if the iterations are allowed to continue. The norm of the residual also does not further decrease after this sharp drop. This is again due to the difficulty the trust-region reflective algorithm has when near the solution at the corner of two bound constraints. It is not possible for the algorithm to improve the solution any further.

4.3 Analysis of Estimated Parameters

The model parameters obtained from the optimization were compared to measured values to determine how well the optimization method works for obtaining correct material properties. Only measured values for the parameters x_1 and x_2 (flow resistivity and porosity, respectively) were available. The estimated value of flow resistivity was 46,088 Pa s/m. The measured value of flow resistivity for the modeled material is approximately 23,000 Pa s/m, but depending on conditions, could be as large as 40,000 Pa s/m. Considering the large range of possible values for flow resistivity (10^3 Pa s/m to 10^6 Pa s/m), this estimate is considered relatively accurate in comparison to the measurement.

The estimated value of x_2 (porosity) from the optimization procedure was 0.8436. The measured value of porosity for this material is 0.9486. The range of possible values for x_2 is much less than for x_1, so considering this, the estimated value is not as accurate. However, it is also possible that the particular specimen that the model is based upon has a lower porosity than the specimen the measured value of porosity is based on. Thus, the estimated value of porosity is still physically feasible. Had the optimization solution resulted in a porosity estimate on the bounds (0.7 or 1) this would not have been as good of a result since it is physically not possible for the porosity to be outside this range.

Lastly, the model parameters were substituted back into the original model of surface impedance, and the modeled surface impedance was compared to the measured surface impedance. The modeled and measured values should closely match if the optimization resulted in a good solution. Figure 18 shows the real and imaginary parts of the modeled and measured surface impedance over the entire available frequency range (including Zone 3 which the optimization was not based upon). There is a high degree of agreement between the measurement and modeled value of surface impedance. Both the real and imaginary parts of the modeled surface impedance almost exactly overlay on the measured value of surface impedance in the Zone 1 and Zone 2 frequency ranges. This is further supported by the value of the cost function at the minimum, which is approximately 33. Since the cost function is the sum and the squared residuals, and there are 246 residuals included in this summation, the residual at each frequency is very low. In Zone 3 (after the peak in the imaginary part of the surface impedance), there is still very good agreement between the measurement and model, although not as good as in Zones 1 and 2. Since there is good agreement in Zone 3 even though this frequency region was not directly used in the optimization procedure, this supports the conclusion that the global minimum found by the optimization algorithm reflects the true value of the material properties described by the five model parameters.

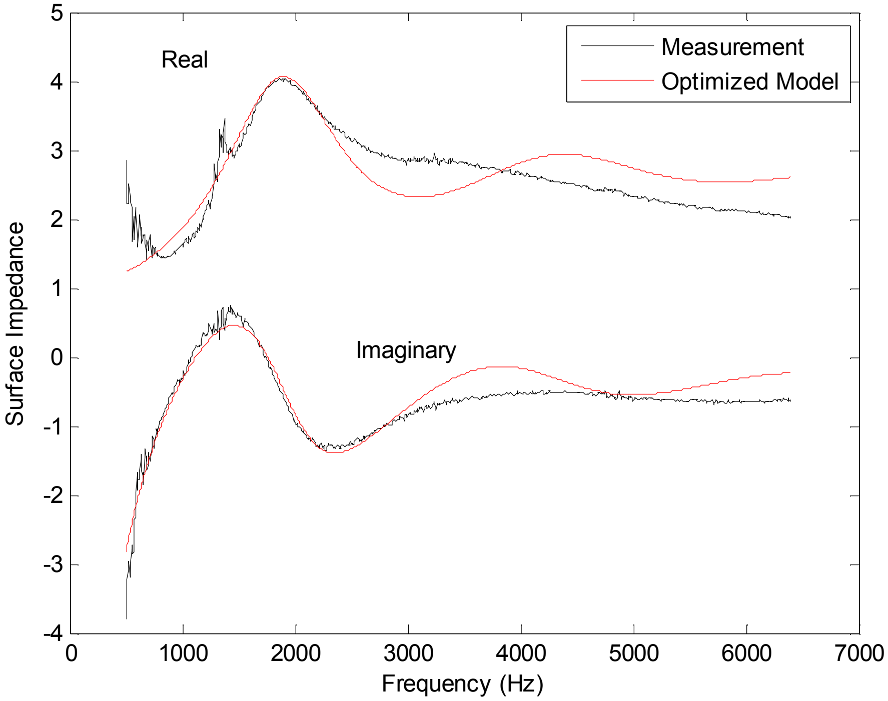

Figure 18: Comparison between the measured surface impedance and the modeled surface impedance based on the estimated model parameters obtained from optimization.

5. Conclusions

Based on the results presented above, the following conclusions were made in regards to the bounded optimization of the model for acoustic foams:

- An effective means to estimate the properties of acoustic foam is to minimize the residual between the modeled and measured value of surface impedance.
- Zone 1 and 2 is the best frequency range to optimize over since all parameters have at least a secondary influence on the value of the cost function. This allows for more accurate estimates by avoiding major scaling issues in the problem.
- There are still some scaling problems with the cost function that cannot be eliminated unless the cost function is modified. In particular, the scale of the parameter x_1 is much larger than any of the other parameters. This was not altered to remain consistent with the literature.
- Final solution is highly dependent on the initial conditions. Many local minimums within the feasible region of the cost function, and the optimization algorithms may converge to these based on the selection of starting parameters. However, any initial value in the range of all five parameters can result in convergence to either the global minimum or the local minimum. The results seem random.
- The only parameter for which certain initial conditions normally guarantee convergence to the global minimum is x_1. If the initial x_1 is in the range of approximately 10^3 to 10^5, then the algorithms normally always converge to the global minimum.
- For this particular problem:

  - Active set method was the fastest/most efficient. This is likely due to the well defined, straight bound constraints.
  - Least squares (trust region reflective method) was the least efficient and least accurate. This was likely due to the limitations of the trust-region reflective algorithm in dealing with a solution at the corner of multiple active constraints.
  - SQP and Interior Point methods resulted in the most accurate minimums and were nearly as efficient as active set.
  - Active set is the recommended method for estimating the parameters of this model since it is the most efficient and the precision of the solution is high enough to match or exceed the precision of the initial measurements used to form the cost function.
- The value of the parameters estimated by the optimization methods match well to the measured values of these same parameters. Additionally the value of the modeled surface impedance based on the estimated parameters closely matches the original measured value of the surface impedance. The residual between the measurement and model was acceptably low.

Further work on this problem could be focused on scaling the parameters so that all parameters have an equal influence on the value of the cost function and its gradient. It would also be of interest to examine these same algorithms when modeling a different material, whose solution may lie entirely within the bounds of the parameters. In this case, the nonlinear least squares trust-region reflective algorithm would be expected to perform the best.

6. References

1. Y. Atalla and R. Panneton, “Inverse acoustical characterization of open cell porous media using impedance tube measurements”, *Canadian Acoustics*, **33**(1), (2005).

2. J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer, New York, (2006).

3. Matlab Documentation, “Constrained Nonlinear Optimization Algorithms”, http://www.mathworks.com/help/toolbox/optim/ug/brnoxzl.html, (2011).

4. Matlab Documentation, “Least Squares (Model Fitting) Algorithms”, http://www.mathworks.com/help/toolbox/optim/ug/brnoybu.html, (2011).

5. T. F. Coleman and Y. Li, “An interior trust region approach for nonlinear minimization subject to bounds”, *SIAM J. Optimization*, **6**(2), 418-445, (1996).