Scientific Computing, Fall 2019
Assignment IV: Nonlinear Equations and Optimization

Aleksandar Donev

Courant Institute, NYU, donev@courant.nyu.edu

Posted Oct 20th, 2019
Due by Sunday **Nov 3rd**, 2019

For grading purposes the maximum is considered to be 80 points.

1 [35 points] Newton-Raphson Method in One Dimension

Consider finding the three roots of the polynomial

$$f(x) = 816x^3 - 3835x^2 + 6000x - 3125,$$

which happen to all be real and all contained in the interval $[1.4, 1.7]$ [due to Cleve Moler].

1.1 [5pts] The roots

Plot this function on the interval $[1.4, 1.7]$ and find all of the zeros of this polynomial using the MATLAB function $fzero$ (report how you did this) [*Hint: The roots values can be obtained in MATLAB using the built-in function roots but Maple tells us the roots are 25/16, 25/17 and 5/3*].

1.2 [10 pts] Newton's Method

[2.5pts] Implement Newton's method (no safeguards necessary) for finding the roots of $f(x)$ and test it with some initial guess in the interval $[1.4, 1.7]$.

[7.5pts] Verify that the order of convergence is quadratic, as predicted by the theory from class:

$$\lim_{k\to\infty} \frac{|e^{k+1}|}{|e^k|^2} \to C = \left| \frac{f''(\alpha)}{2f'(\alpha)} \right|.$$

[*Hint: Due to roundoff errors and the very fast convergence, the error quickly becomes comparable to roundoff, so one must be careful not to use very large k. The errors must be dominated by truncation errors and not roundoff errors for the above theory to apply!*]

1.3 [20pts] Robustness

[10pts] Starting from many (say 100) guesses in the interval $[1.4, 1.7]$ [*Hint: In MATLAB, you can create a grid of 100 points with $x_0 = linspace(1.4, 1.7, 100)$*], run 100 iterations of Newton's method and plot the value to which it converges, if it does (make your own criterion for when to stop and report what you did), as a function of the initial guess. If the initial guess is sufficiently close to one of the roots α, i.e., if it is within the *basin of attraction* for root α, it should converge to α. What is the basin of attraction for the middle root ($\alpha = 25/16$) based on the plot?

[5pts] The theory from the lecture suggested an estimate for the width of each basin of attraction around a given root α of the form:

$$\left| x^0 - \alpha \right| \le \left| \frac{f''(\alpha)}{2f'(\alpha)} \right|^{-1}.$$

Compare these estimates to what is actually observed numerically.

[5pts] Is the estimate particularly bad for one of the roots, and if so, can you think of reasons why?

2 [55 pts] Nonlinear Least-Squares Fitting

In homework 2 you considered fitting a data series (x_i, y_i), $i = 1, \ldots, m$, with a function that that depends linearly on a set of unknown fitting parameters $\boldsymbol{c} \in \mathbb{R}^n$. Consider now fitting data to a nonlinear function of the fitting parameters, $y = f(x; \boldsymbol{c})$. The least-squares fit is the one that minimizes the squared sums of errors,

$$\boldsymbol{c}^\star = \arg\min_{\boldsymbol{c}} \sum_{i=1}^{m} [f(x_i; \boldsymbol{c}) - y_i]^2 = \arg\min_{\boldsymbol{c}} \|\boldsymbol{f} - \boldsymbol{y}\|_2^2, \tag{1}$$

where $\boldsymbol{f}(\boldsymbol{c}) \equiv f(\boldsymbol{x}; \boldsymbol{c})$ is a vector of m function values, evaluated at the data points, for a given set of the parameters $\boldsymbol{c}$. We can think of this as an overdetermined system of nonlinear equations for $\boldsymbol{c}$,

$$f(\boldsymbol{c}) = \boldsymbol{y}, \tag{2}$$

although of course in the end this is an optimization problem rather than solving a square system of equations (remember that the two are closely related).

We will consider here fitting an exponentially-damped sinusoidal curve with four unknown parameters (amplitude c_1, decay c_2, period c_3, and phase c_4, respectively),

$$f(x; \boldsymbol{c}) = c_1 e^{-c_2 x} \sin(c_3 x + c_4), \tag{3}$$

to a synthetic data set. Note: The method you will use is related to Newton's method for optimization with a trust-region safeguard, but slightly different.

2.1 [5pts] Synthetic Data

Generate synthetic data by generating $m = 100$ points randomly and uniformly distributed in the interval $0 \leq x \leq 10$ by using the *rand* function. Compute the actual function

$$f(x; \boldsymbol{c}) = e^{-x/2} \sin(2x), \tag{4}$$

and then add perturbations with absolute value on the order of $\epsilon = 10^{-2}$ to the y values (use the *rand* or the *randn* function),

$$y_i = e^{-x_i/2} \sin(2x_i) + \epsilon \cdot rand(). \tag{5}$$

Compare the synthetic data to the actual function on the same plot to make sure your synthetic data closely (but not exactly!) follows the relation (4).

2.2 [30 pts] Gauss-Newton Method

The basic idea behind the Gauss-Newton method is to make a linearization of the function $f(x_i; \boldsymbol{c})$ around the current estimate $\boldsymbol{c}_k$,

$$\boldsymbol{f}(\boldsymbol{c}) \approx \boldsymbol{f}(\boldsymbol{c}_k) + [\boldsymbol{J}(\boldsymbol{c}_k)](\boldsymbol{c} - \boldsymbol{c}_k) = \boldsymbol{f}(\boldsymbol{c}_k) + [\boldsymbol{J}(\boldsymbol{c}_k)]\Delta\boldsymbol{c}_k,$$

where the Jacobian $m \times n$ matrix is the matrix of partial derivatives $\partial f / \partial c$ evaluated at the data points:

$$\boldsymbol{J}(\boldsymbol{c}) = \nabla_{\boldsymbol{c}} \boldsymbol{f}(\boldsymbol{c}).$$

This approximation (linearization) transforms the non-linear problem (2) into a linear least-squares problem, i.e., an overdetermined linear system

$$[\boldsymbol{J}(\boldsymbol{c}_k)]\Delta\boldsymbol{c}_k = \boldsymbol{J}_k \Delta\boldsymbol{c}_k = \boldsymbol{y} - \boldsymbol{f}(\boldsymbol{c}_k) = \boldsymbol{y} - \boldsymbol{f}_k, \tag{6}$$

which you know how to solve from previous homeworks and lectures. The standard approach is to use the normal equations

$$(\boldsymbol{J}_k^T \boldsymbol{J}_k) \Delta\boldsymbol{c}_k = \boldsymbol{J}_k^T (\boldsymbol{y} - \boldsymbol{f}_k), \tag{7}$$

2

which does not lead to substantial loss of accuracy if one assumes that the original problem is well-conditioned (you are welcome to use the QR factorization or backslash if you prefer, just report what you did). Gauss-Newton's algorithm is a simple iterative algorithm of the form

$$\boldsymbol{c}_{k+1} = \boldsymbol{c}_k + \Delta \boldsymbol{c}_k,$$

starting from some initial guess $\boldsymbol{c}_0$. The iteration is terminated, for example, when the increment $\|\Delta \boldsymbol{c}_k\|$ becomes too small.

[10 pts] Implement Gauss-Newton's algorithm and see whether it works for the problem at hand, using an initial guess $\boldsymbol{c}_0$ that is close to the correct values.

[5 pts] Is the convergence to the answer quadratic or linear? Answer this based on empirical (numerical) results for the purpose of this homework. If you can give a theoretical argument that's great.

[5pts] If you start with $\boldsymbol{c}_0 = (1, 1, 1, 1)$, does the method converge to the correct answer? Play around a bit with initial guesses and see if the method converges most of the time, and whether it converges to the "correct" solution or other solutions.

[5 pts] If the synthetic data points have no error, i.e., if $\epsilon = 0$ in (5), how many digits of accuracy in $\boldsymbol{c}$ can you obtain? How many steps do you need to achieve this accuracy?

[5pts] Is this method the same or even similar to using Newton's method (for optimization) to solve the non-linear problem (1)? Try to think about this on your own but if you cannot figure it out it is OK to use sources like Wikipedia (report the source).

2.3 [20pts] Levenberg-Marquardt Algorithm

The Gauss-Newton algorithm is not very robust. It is not guaranteed to have even local convergence. A method with much improved robustness can be obtained by using a modified (regularized) version of the normal equations (7),

$$\left[\left(\boldsymbol{J}_k^T \boldsymbol{J}_k \right) + \lambda_k \operatorname{Diag} \left(\boldsymbol{J}_k^T \boldsymbol{J}_k \right) \right] \Delta \boldsymbol{c}_k = \boldsymbol{J}_k^T \left(\boldsymbol{y} - \boldsymbol{f}_k \right), \tag{8}$$

where $\lambda_k > 0$ is a damping parameter that is used to ensure that $\Delta \boldsymbol{c}_k$ is a descent direction, in the spirit of trust-region and quasi-Newton algorithms for optimization. Here $\operatorname{Diag}(\boldsymbol{A})$ denotes a diagonal matrix whose diagonal is the same as the diagonal of $\boldsymbol{A}$ [*Hint: The MATLAB call diag(diag(A)) can be used to obtain* $\operatorname{Diag}\boldsymbol{A}$, *as used in (8)*].

If λ_k is large, the method will converge slowly but surely (can you understand why?), while a small λ_k makes the method close to the Gauss-Newton Method, which converges rapidly if it converges at all. So the idea is to use a larger λ_k when far from the solution, and then decrease λ_k as approaching the solution. The actual procedure used to adjust the damping parameter can be found in the literature, and consists of doubling λ if things are not going well, or halving λ if things are going well. Here we study one simple strategy: Start with some sufficiently large initial value λ_1, and then reduce it by a factor of 2 each iteration, $\lambda_k = \lambda_{k-1}/2$.

[15pts] Implement a code that tries this method and try it for a value of the initial guess for which the Gauss-Newton method in part 2.2 above failed, i.e., for which $\lambda_1 = 0$ does not work. Try different initial values of λ_1 and see how large it has to be before the method converges.

[5pts] Do you notice any difference in the speed of convergence for different values of λ_1 (i.e., is the speed of convergence faster for smaller values or larger values)?