Problem Set 5 - Solution

Aden Forrow

Unless otherwise specified, you may use MATLAB to assist with computations. Please provide a print-out of the code used and its output with your assignment.

1. Positive definiteness.

   (a) Without doing any calculations, determine whether the following matrices are positive definite (and explain your reasoning):

$$A_1 = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \qquad A_2 = \begin{bmatrix} -1 & 8 & 3 \\ 1 & 4 & -7 \\ -2 & -3 & 6 \end{bmatrix}. \qquad A_3 = \begin{bmatrix} 4 & 0 & 0 \\ 0 & -2 & 0 \\ 0 & 0 & 9 \end{bmatrix},$$

   Answer: None of them are positive definite. A_1 has two identical columns, so it is not full rank and has an eigenvalue of 0 (see part (b)). A_2 has a negative entry in the upper left corner, so it fails the upper left determinant test. A_3 is diagonal, so its eigenvalues are on the diagonal; one of them is -2, which is negative.

   (b) Show that if a $n \times n$ symmetric matrix B is not full rank, it cannot be positive definite.

   Answer: If B is not full rank, it has a nontrivial null space: there exists a nonzero vector x such that $Bx = 0$. This x is an eigenvector with eigenvalue 0. Since at least one of B's eigenvalues isn't positive, B is not positive definite.

   (c) Show that if C is an $n \times n$ matrix such that $u^\top C u > 0$ for every $u \neq 0$, then all eigenvalues of C are positive.

   Answer: Suppose λ is an eigenvalue of C with with eigenvector x. Then

$$x^\top C x = x^\top (\lambda x) = \lambda x^\top x = \lambda \|x\|^2.$$

   Since both $\|x\|^2$ and $x^\top C x$ are positive, λ must also be positive.

(d) Show that if the symmetric $n \times n$ matrices D_1 and D_2 are positive definite, then so is $D_1 + D_2$.

Answer: It's sufficient to show $u^\top (D_1 + D_2)u > 0$ for any $u \neq 0$. For such a u,

$$u^\top (D_1 + D_2)u = u^\top D_1 u + u^\top D_2 u.$$

Both $u^\top D_1 u$ and $u^\top D_2 u$ are positive if D_1 and D_2 are positive definite, so their sum is also positive.

2. SVD. **This problem is to be done by hand.** Find the singular value decomposition of the following matrices:

(a)

$$A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 1 \\ 0 & 1 \end{bmatrix},$$

Answer: We can start either by finding U and Σ from the eigenvalue decomposition of $AA^\top$ or by finding V and Σ from the eigenvalue decomposition of $A^\top A$. In this case, since $AA^\top$ is 4×4 while $A^\top A$ is 2×2, the latter will be much easier.

$$A^\top A = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 3 & 2 \\ 2 & 3 \end{bmatrix}$$

The eigenvalues $\lambda_i = \sigma_i^2$ satisfy

$$\begin{aligned} 0 &= \det(A^\top A - \lambda_i I) \\ &= (3 - \lambda_i^2) - 4 \\ &= \lambda_i^2 - 6\lambda_i + 5 \\ &= (\lambda_i - 5)(\lambda_i - 1) \end{aligned}$$

The eigenvalues of $A^\top A$ are $\lambda_1 = 5$ and $\lambda_2 = 1$, so the nonzero singular values of A are $\sigma_1 = \sqrt{5}$ and $\sigma_2 = 1$.
The eigenvector x_1 for $\lambda_1 = 5$ is in the null space of

$$A^\top A - 5I = \begin{bmatrix} -2 & 2 \\ 2 & -2 \end{bmatrix}.$$

We see that the columns sum to zero, so $x_1 = (1, 1)^\top$ is an eigenvector. Similarly, for $\lambda_2 = 1$,

$$A^\top A - I = \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix}.$$

The columns are equal, so $x_2 = (1, -1)^\top$ is an eigenvector.
We normalize these eigenvectors to get columns of V:

$$v_1 = \frac{x_1}{\|x_1\|} = \frac{1}{\sqrt{2}}\begin{bmatrix}1\\1\end{bmatrix} \qquad v_2 = \frac{x_2}{\|x_2\|} = \frac{1}{\sqrt{2}}\begin{bmatrix}-1\\1\end{bmatrix}$$

The only piece still missing is U. We can find the first two columns of U using the relationship $Av_i = \sigma_i u_i$ between right and left singular vectors.

$$u_1 = \frac{Av_1}{\sigma_1} = \frac{1}{\sqrt{5}}\begin{bmatrix}1&0\\1&1\\1&1\\0&1\end{bmatrix}\begin{bmatrix}\frac{1}{\sqrt{2}}\\\frac{1}{\sqrt{2}}\end{bmatrix} = \frac{1}{\sqrt{10}}\begin{bmatrix}1\\2\\2\\1\end{bmatrix}$$

$$u_2 = \frac{Av_2}{\sigma_2} = \begin{bmatrix}1&0\\1&1\\1&1\\0&1\end{bmatrix}\begin{bmatrix}-\frac{1}{\sqrt{2}}\\\frac{1}{\sqrt{2}}\end{bmatrix} = \frac{1}{\sqrt{2}}\begin{bmatrix}-1\\0\\0\\1\end{bmatrix}$$

Putting what we have together is enough for a reduced SVD of A:

$$A = \hat{U}\hat{\Sigma}\hat{V}^\top = \begin{bmatrix}\frac{1}{\sqrt{10}}&-\frac{1}{\sqrt{2}}\\\frac{2}{\sqrt{10}}&0\\\frac{2}{\sqrt{10}}&0\\\frac{1}{\sqrt{10}}&\frac{1}{\sqrt{2}}\end{bmatrix}\begin{bmatrix}5&0\\0&1\end{bmatrix}\begin{bmatrix}\frac{1}{\sqrt{2}}&\frac{1}{\sqrt{2}}\\-\frac{1}{\sqrt{2}}&\frac{1}{\sqrt{2}}\end{bmatrix}$$

To get a full SVD of A, we need to add two more columns to $\hat{U}$ to make it square. They must be normalized and orthogonal to the other u_i's. An easy choice for u_3 is to take advantage of the zeros in u_2:

$$u_3 = \frac{1}{\sqrt{2}}\begin{bmatrix}0\\-1\\1\\0\end{bmatrix}.$$

Next, since u_4 must be orthogonal to u_2 and u_3, $u_{14} = u_{44} = a$ and $u_{24} = u_{34} = b$. Orthogonality with u_1 is then

$$0 = u_1^\top u_4 = \frac{1}{\sqrt{10}}\begin{bmatrix}1&2&2&1\end{bmatrix}\begin{bmatrix}a\\b\\b\\a\end{bmatrix} = \frac{1}{\sqrt{10}}(2a + 4b)$$

So $a = -2b$, and

$$u_4 = \frac{1}{\sqrt{10}}\begin{bmatrix}-2\\1\\1\\-2\end{bmatrix},$$

with the factor $1/\sqrt{10}$ there so u_4 is normalized. If we put these together with $\hat{U}$ and add extra zero rows to $\hat{\Sigma}$, we get the full SVD,

$$A = U\Sigma V^\top = \begin{bmatrix} \frac{1}{\sqrt{10}} & -\frac{1}{\sqrt{2}} & 0 & -\frac{2}{\sqrt{10}} \\ \frac{2}{\sqrt{10}} & 0 & -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{10}} \\ \frac{2}{\sqrt{10}} & 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{10}} \\ \frac{1}{\sqrt{10}} & \frac{1}{\sqrt{2}} & 0 & -\frac{2}{\sqrt{10}} \end{bmatrix} \begin{bmatrix} 5 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}$$

(We could also have found u_3 and u_4 by looking for vectors in the null space of $\hat{U}^\top$ or by choosing any vectors independent of the other u_i's and using Gram-Schmidt to make them orthogonal.)

(b)

$$B = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix},$$

(Hint: Use the result from part (a))

Answer: Notice that $B = A^\top$ from the previous part, so

$$B = (U\Sigma V^\top)^\top = V\Sigma^\top U^\top,$$

using the same matrices as before. Explicitly,

$$B = \begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} \sqrt{5} & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{10}} & \frac{2}{\sqrt{10}} & \frac{2}{\sqrt{10}} & \frac{1}{\sqrt{10}} \\ -\frac{1}{\sqrt{2}} & 0 & 0 & \frac{1}{\sqrt{2}} \\ 0 & -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ -\frac{2}{\sqrt{10}} & \frac{1}{\sqrt{10}} & \frac{1}{\sqrt{10}} & -\frac{2}{\sqrt{10}} \end{bmatrix}$$

(c)

$$C = \begin{bmatrix} 6 & 3 \\ 8 & 4 \end{bmatrix}.$$

Answer: This time, we'll start by finding $CC^\top$:

$$CC^\top = \begin{bmatrix} 6 & 3 \\ 8 & 4 \end{bmatrix} \begin{bmatrix} 6 & 8 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} 45 & 60 \\ 60 & 80 \end{bmatrix}$$

Then

$$\begin{aligned} 0 =& \det(C^\top C - \lambda_i I) \\ =& (45-\lambda_i)(80-\lambda_i) - 3600 \\ =& \lambda_i^2 - 125\lambda_i \\ & . \end{aligned}$$

The eigenvalues of $CC^\top$ are $\lambda_1 = 125$ and $\lambda_2 = 0$, so the nonzero singular values of C are $\sigma_1 = 5\sqrt{5}$ and $\sigma_2 = 0$.
The eigenvector x_1 for $\lambda_1 = 125$ is in the null space of

$$CC^\top - 125I = \begin{bmatrix} -80 & 60 \\ 60 & -45 \end{bmatrix}.$$

We see that the first column is $-\frac{4}{3}$ times the second, so $x_1 = (1, 4/3)^\top$ is an eigenvector. Similarly, for $\lambda_2 = 0$,

$$CC^\top - 0I = \begin{bmatrix} 45 & 60 \\ 60 & 80 \end{bmatrix}.$$

The second column is $\frac{4}{3}$ times the first, so $x_2 = (4/3, -1)^\top$ is an eigenvector.
We normalize these eigenvectors to get columns of U (not V, since we used $CC^\top$ not $C^\top C$:

$$u_1 = \frac{x_1}{\|x_1\|} = \frac{1}{\sqrt{1+16/9}} \begin{bmatrix} 1 \\ 4/3 \end{bmatrix} = \begin{bmatrix} 3/5 \\ 4/5 \end{bmatrix}$$

$$u_2 = \frac{x_2}{\|x_2\|} = \frac{1}{\sqrt{1+16/9}} \begin{bmatrix} 4/3 \\ -1 \end{bmatrix} \begin{bmatrix} 4/5 \\ -3/5 \end{bmatrix}$$

We can find v_1 similarly to how we found u_1 in part (a):

$$v_1 = \frac{C^\top u_1}{\sigma_1} = \frac{1}{5\sqrt{5}} \begin{bmatrix} 6 & 8 \\ 3 & 4 \end{bmatrix} \begin{bmatrix} 3/5 \\ 4/5 \end{bmatrix} = \frac{1}{5\sqrt{5}} \begin{bmatrix} 50/5 \\ 25/5 \end{bmatrix} = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

Since $\sigma_2 = 0$, we can't use that method to find v_2 from u_2; instead, we look for a vector in the null space of C, since $Cv_1 = \sigma_1 u_1 = 0$. The first column of C is twice the second column, so $x = (-12)$ is in the null space. Normalized, this is

$$v_2 = \frac{1}{\sqrt{5}} \begin{bmatrix} -1 \\ 2 \end{bmatrix}$$

Putting everything together,

$$C = U\Sigma V^\top = \begin{bmatrix} 3/5 & 4/5 \\ 4/5 & -3/5 \end{bmatrix} \begin{bmatrix} 5\sqrt{5} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 2/\sqrt{5} & 1/\sqrt{5} \\ -1/\sqrt{5} & 2/\sqrt{5} \end{bmatrix}.$$

3. Error. Consider the matrix:

$$A = \begin{bmatrix} 1.2143 & 0.8054 \\ 0.8128 & 0.5391 \end{bmatrix}.$$

(a) Set `b1 = [1.2143; 0.8128]` and compute `x = A \ b1` in MATLAB.

Answer: See MATLAB script.

$$x = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

(b) Repeat the process but with a vector `b2` obtained from `b1` by rounding each value to three decimal places. Why do the two answers differ by so much?

Answer: See MATLAB script.

$$x = \begin{bmatrix} -32280 \\ 48670 \end{bmatrix}.$$

This is nowhere near the previous x; it's off by four orders of magnitude, even though `b2` is almost exactly `b1`. Computing the condition number of A,

$$c(A) = 3.1 \times 10^8,$$

shows why. A is extremely ill conditioned, magnifying errors by up to eight orders of magnitude.

(c) Now compute `x = B \ b1` where B is obtained from A by rounding each value to three decimal places. As you will see, our solution is also sensitive to small changes in the matrix A.

Answer: See MATLAB script.

$$x = \begin{bmatrix} -1.71 \\ 4.09 \end{bmatrix}.$$

It's not as bad as part (b), but this x is still quite far from x in part (a).

(d) Errors from ill-conditioned problems are not unique to solving $Ax = b$. As an example, compute $\tan(10^{100})$ three ways

i. Google: search for tan(10^100)

Answer: Google thinks $\tan(10^{100}) = 3.9231\ldots$.

ii. MATLAB: `tan(10 ^ 100)`

Answer: Computed this way, MATLAB thinks $\tan(10^{100}) = -0.4116\ldots$.

iii. MATLAB: Use `vpa('tan(1e100)', p)` to compute with p digits of precision for $p = 5$, 10, 25, 50, and 100.

Answer: See MATLAB script. We find the following

p	5	10	25	50	100
$\tan(10^{100})$	-0.81	-4.5	-2.8	19	0.401

That's six wildly different estimates - they agree on neither sign nor magnitude.

Explain the results you see. Why is this a hard problem to solve?

Answer: None of these values come close to agreeing with each other. The problem is that tangent itself varies from $-\infty$ to ∞ over each interval of width π, and even a small relative change in 10^{100} (including the error in representing 10^{100} as a floating point number) puts you in a completely different place in a different interval. To accurately place 10^{100} in the periodic cycle of tangent and get an answer remotely close to right, you need to keep track of all 100 digits of 10^{100}.

4. Condition number. Suppose the condition number of an invertible $n \times n$ matrix A is κ.

   (a) Find the condition number of $A^\top A$ in terms of κ. What does this imply about solving the normal equations by inverting $A^\top A$?

   Answer: If the SVD of A is $A = U\Sigma V^\top$, then

   $$A^\top A = V\Sigma^T U^\top U\Sigma V^\top = V\Sigma^2 V^\top.$$

   The singular values of $A^\top A$ are the diagonal entries of Σ^2, which are the singular values of A squared. So

   $$c(A^\top A) = \frac{\sigma_{max}(A)^2}{\sigma_{min}(A)^2} = \kappa^2.$$

   The condition number of $A^\top A$ will be larger than the condition number of A, much larger if A is ill-conditioned. This means that inverting $A^\top A$ will be much less stable and more sensitive to errors than dealing with A itself. Hence the usefulness of solving the normal equations via QR factorization.

   (b) Find the condition number of A^{-1} in terms of κ.

   Answer: If the SVD of A is $A = U\Sigma V^\top$, then

   $$V\Sigma^{-1}U^\top A = V\Sigma^{-1}U^\top U\Sigma V^\top = V\Sigma^{-1}\Sigma V^\top = VV^\top = I.$$

   So the SVD of A^{-1} is

   $$A^{-1} = V\Sigma^{-1}U^\top,$$

   up to some reodering of the singular values. Since Σ is diagonal, its inverse is also diagonal, with diagonal entries $1/\sigma_i$. Then

   $$c(A^{-1}) = \frac{\sigma_{max}(A^{-1})}{\sigma_{min}(A^{-1})} = \frac{1/\sigma_{min}(A)}{1/\sigma_{max}(A)} = \frac{\sigma_{max}(A)}{\sigma_{min}(A)} = \kappa.$$

   A and A^{-1} have the same condition number.

   (c) What is the value of κ if A is an orthogonal matrix?

   Answer: If A is an orthogonal matrix, we can write its SVD as $A = A * I * I$. Its singular values are the diagonal entries of the identity matrix, which are all ones. Therefore $\kappa = \sigma_{max}/\sigma_{min} = 1/1 = 1$.

(d) What is the smallest possible value for κ? Construct a 3×3 matrix with this minimum condition number.

Answer: Since the largest singular value is at least as big as the smallest singular values, $\sigma_{max}/\sigma_{min}$ is at least 1. The simplest example of a matrix with condition number 1 is the identity matrix,

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

5. Initial Value Problem. Consider the following initial value problem for $x(t)$:

$$\frac{dx}{dt} = x^2$$

with initial condition $x(0) = -1$.

(a) Find the exact solution for $x(t)$. (Hint: Separate variables and integrate both sides). Compute the exact value of $x(1)$.

Answer: After separating variables,

$$\begin{aligned}\int \frac{1}{x^2} dx &= \int dt \\ -\frac{1}{x} &= t + C \\ x &= \frac{-1}{t + C}\end{aligned}$$

Since $x(0) = -1$, $C = 1$ and

$$x = \frac{-1}{t+1}$$

So $x(1) = -1/2$.

(b) We would like to setup this problem to be solved numerically with a time step Δt. Recall that $x_n = x(n\Delta t)$. Find an explicit formula for x_{n+1} as a function of x_n and Δt using a Forward Euler method.

Answer: For Forward Euler,

$$\begin{aligned}x_{n+1} &= x_n + \Delta t \left.\frac{dx}{dt}\right|_{x=x_n} \\ &= x_n + \Delta t x_n^2\end{aligned}$$

(c) Implement this explicit numerical method in MATLAB and find the approximation of $x(1)$ using $\Delta t = \frac{1}{10}$. Calculate the difference between the exact value and this approximate value.

Answer: See MATLAB script. Using this time step, we find

$$x(1) \approx -0.4817$$

with error

$$|-0.4817 - x(1)| = 0.0183.$$

(d) If we reduce the time step from $\Delta t = \frac{1}{10}$ to $\Delta t = \frac{1}{100}$, how much do you expect the error to be reduced?

Answer: Forward Euler schemes are first order accurate, so reducing the time step by a factor of 10 should reduce the error by about a factor of 10.

(e) Find the approximation of $x(1)$ with $\Delta t = \frac{1}{100}$. Calculate the difference between the exact value and this approximate value. By what factor is the error at $x(1)$ reduced by switching from $\Delta t = \frac{1}{10}$ to $\Delta t = \frac{1}{100}$?

Answer: See MATLAB script. This time we find

$$x(1) \approx -0.4983$$

with error

$$|-0.4983 - x(1)| = 0.0017.$$

The error was reduced by a factor of about 10.

(f) Derive an explicit formula for x_{n+1} as a function of x_n and Δt using a:

i. Backward Euler method,

Answer: For backward Euler,

$$\begin{aligned} x_{n+1} &= x_n + \Delta t \left.\frac{dx}{dt}\right|_{x=x_{n+1}} \\ &= x_n + \Delta t x_{n+1}^2 \end{aligned}$$

Rearranging, we have a quadratic equation for x_{n+1}:

$$\Delta t x_{n+1}^2 - x_{n+1} + x_n = 0.$$

Using the quadratic formula,

$$x_{n+1} = \frac{1 \pm \sqrt{1 - 4\Delta t x_n}}{2\Delta t}$$

Which sign do we pick? Well, as $\Delta t \to 0$, $\sqrt{1 - 4\Delta t x_n} \approx 1 - 2\Delta t x_n$. If we choose the positive sign, we get

$$x_{n+1} \approx \frac{2 - 2\Delta t x_n}{2\Delta t},$$

which goes to infinity as Δt goes to zero. That's unreasonable. Instead, we choose the negative sign, so

$$x_{n+1} = \frac{1 - \sqrt{1 - 4\Delta t x_n}}{2\Delta t},$$

which is approximately x_n if Δt is very small.

ii. Trapezoidal Rule method.

Answer: For the trapezoidal rule,

$$\begin{aligned} x_{n+1} =& x_n + \Delta t \frac{dx}{dt}\Big|_{x=(x_{n+1}+x_n)/2} \\ =& x_n + \frac{\Delta t}{4}(x_{n+1}+x_n)^2 \end{aligned}$$

Rearranging, we again have a quadratic equation for x_{n+1}:

$$\frac{\Delta t}{4}x_{n+1}^2 + (\frac{\Delta t}{2}x_n - 1)x_{n+1} + \frac{\Delta t}{4}x_n^2 + x_n = 0.$$

Using the quadratic formula,

$$x_{n+1} = \frac{1 - \frac{\Delta t}{2}x_n \pm \sqrt{\left(1-\frac{\Delta t}{2}x_n\right)^2 - 4\frac{\Delta t}{4}\left(\frac{\Delta t}{4}x_n^2 + x_n\right)}}{\Delta t/2}$$

Again, we don't want this to explode when $\Delta t \to 0$, so we choose the negative sign,

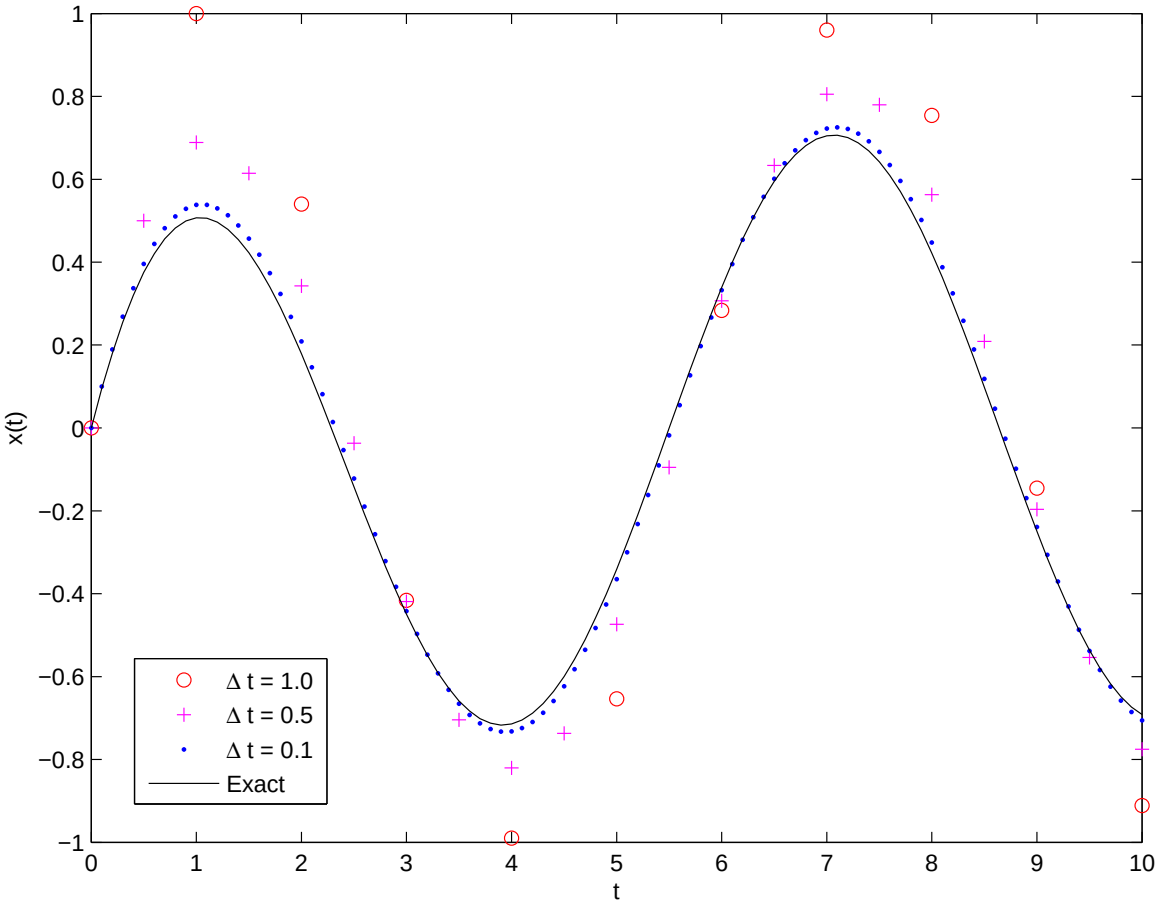
$$x_{n+1} = \frac{1 - \frac{\Delta t}{2}x_n - \sqrt{\left(1-\frac{\Delta t}{2}x_n\right)^2 - 4\frac{\Delta t}{4}\left(\frac{\Delta t}{4}x_n^2 + x_n\right)}}{\Delta t/2}$$

6. The exact solution of the differential equation

$$\frac{dx}{dt} = -x + \cos t, \tag{1}$$

with the initial condition $x(0) = x_0$ is $x(t) = (x_0 - \frac{1}{2})e^{-t} + \frac{1}{\sqrt{2}}\cos(t - \pi/4)$. Implement the Forward Euler method in MATLAB to solve this equation numerically. Use three different time steps: $\Delta t = 0.1$, $\Delta t = 0.5$ and $\Delta t = 1.0$. In a single plot, show the exact solution and your three numerical solutions for $x_0 = 0$.

Answer: See MATLAB script. Notice the steady improvement as the time step decreases.



7. Newton's method **(This problem is to be done by hand.)**

(a) Use two steps of Newton's method to approximate $x^* = \sqrt{3}$ starting from an initial guess $x_0 = 2$. How many correct digits are there at each step? How many iterations do you expect it to take to find x^* to sixteen digits (the limit of MATLAB's double-precision arithmetic)?

Answer: To use Newton's method to find $\sqrt{3}$, we need to choose a function that is zero at $x^* = \sqrt{3}$. The simplest is $f(x) = x^2 - 3$. This means $f'(x) = 2x$. As in class, we can fill in a table from left to right and top to bottom, using

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$$

to find the new estimate at each step.

n	x_n	$f(x_n)$	$f'(x_n)$	$f(x_n)/f'(x_n)$
0	2	1	4	0.25
1	1.75	0.0625	3.5	$1/56 \approx 0.01786$
2	1.73214	–	–	–

The answer to sixteen digits is $x^* = 1.732050807568877$. So we have no digits (or maybe one) at step 0, two at step 1, and four at step 2. Continuing the pattern of doubling the number of digits each step, we expect eight at step 3 and sixteen after only four steps of Newton's method. (You can verify this with the `newtons_method.m` code on the website.)

(b) Consider solving the system of equations

$$\frac{dx}{dt} = 1 - txy$$
$$\frac{dy}{dt} = 1 - ty^2 - x$$

with initial condition $(x(0), y(0)) = (1, 1)$. Set up the system of equations for the backwards Euler approximation to $(x(1), y(1))$ with a timestep of $\Delta t = 1$. Use the forward Euler approximation as your initial guess and improve it with one step of Newton's method.

Answer: For backwards Euler,

$$\begin{bmatrix} x(1) \\ y(1) \end{bmatrix} = \begin{bmatrix} x(0) \\ y(0) \end{bmatrix} + \Delta t \left[\begin{matrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{matrix} \right]\Bigg|_{x=x(1),y=y(1),t=1} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \begin{bmatrix} 1 - x(1)y(1) \\ 1 - y(1)^2 - x(1) \end{bmatrix}$$

After rearranging terms, this is

$$\begin{bmatrix} 0 \\ 0 \end{bmatrix} = \begin{bmatrix} x(1)y(1) + x(1) - 2 \\ y(1)^2 + x(1) + y(1) - 2 \end{bmatrix} = \begin{bmatrix} f_1(x(1), y(1)) \\ f_2(x(1), y(1)) \end{bmatrix}$$

Our initial guess for the solution is the forward Euler approximation,

$$\begin{bmatrix} x_0 \\ y_0 \end{bmatrix} = \begin{bmatrix} x(1) \\ y(1) \end{bmatrix} = \begin{bmatrix} x(0) \\ y(0) \end{bmatrix} + \Delta t \left[\begin{matrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{matrix} \right]\Bigg|_{x=x(0),y=y(0),t=t(0)} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

The Jacobian matrix we need is

$$J(x,y) = \begin{bmatrix} \frac{\partial f_1}{\partial x} & \frac{\partial f_1}{\partial y} \\ \frac{\partial f_2}{\partial x} & \frac{\partial f_2}{\partial y} \end{bmatrix} = \begin{bmatrix} y+1 & x \\ 1 & 2y+1 \end{bmatrix}$$

The new approximations (x_1, y_1) after one step of Newton's method satisfy

$$J(x_0,y_0) \begin{bmatrix} x_1 \\ y_1 \end{bmatrix} = J(x_0,y_0) \begin{bmatrix} x_0 \\ y_0 \end{bmatrix} - \begin{bmatrix} f_1(x_0,y_0) \\ f_2(x_0,y_0) \end{bmatrix}.$$

Plugging in numbers,

$$\begin{bmatrix} 2 & 2 \\ 1 & 5 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \end{bmatrix} = \begin{bmatrix} 2 & 2 \\ 1 & 5 \end{bmatrix} \begin{bmatrix} 2 \\ 1 \end{bmatrix} - \begin{bmatrix} 2 \\ 2 \end{bmatrix} = \begin{bmatrix} 4 \\ 5 \end{bmatrix}.$$

This is a linear system, which we can solve by Gaussian elimination. The augmented matrix is

$$\left[\begin{array}{cc|c} 2 & 2 & 4 \\ 1 & 5 & 5 \end{array}\right].$$

Swapping rows yields

$$\left[\begin{array}{cc|c} 1 & 5 & 5 \\ 2 & 2 & 4 \end{array}\right].$$

Subtract twice row 1 from row 2:

$$\left[\begin{array}{cc|c} 1 & 5 & 5 \\ 0 & -8 & -6 \end{array}\right].$$

So $y_1 = -6/(-8) = 3/4$ and $x_1 = 5 - 5y_1 = 5/4$. This is our final approximation to $x(1)$ and $y(1)$:

$$\begin{bmatrix} x(1) \\ y(1) \end{bmatrix} \approx \begin{bmatrix} 5/4 \\ 3/4 \end{bmatrix}$$

(The true solution is $(x(1), y(1)) = (1.53, 0.43)$, so this is better than the forward Euler approximation.)

(c) A common place Newton's method shows up is in finding minima of functions. Use it to find the minimum of

$$f(x, y) = x^2 + y^2 + xy - 3x - 3y$$

near $(x, y) = (0, 0)$. How many steps do you need? Why?
(Hint: what do you want to set to zero?)

Answer: To minimize $f(x, y)$, we set its gradient to zero. For this f,

$$\nabla f(x, y) = \begin{bmatrix} \frac{\partial f}{\partial x} \\ \frac{\partial f}{\partial y} \end{bmatrix} = \begin{bmatrix} 2x + y - 3 \\ 2y + x - 3 \end{bmatrix}$$

The Jacobian matrix we need is the Jacobian of $\nabla f(x, y)$, or the matrix of second derivatives of $f(x, y)$:

$$J(x, y) = \begin{bmatrix} \frac{\partial^2 f}{\partial x^2} & \frac{\partial^2 f}{\partial x \partial y} \\ \frac{\partial^2 f}{\partial x \partial y} & \frac{\partial^2 f}{\partial y^2} \end{bmatrix} = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$

. Using Newton's method,

$$J(x, y) \begin{bmatrix} x_1 \\ y_1 \end{bmatrix} = J(x, y) \begin{bmatrix} x_0 \\ y_0 \end{bmatrix} - \nabla f(x_0, y_0),$$

or

$$\begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \end{bmatrix} = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \end{bmatrix} - \begin{bmatrix} -3 \\ -3 \end{bmatrix} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}.$$

Here we can read off that $x_1 = 1$ and $y_1 = 1$. And that's it - we don't need any more steps. Since $f(x, y)$ is quadratic, its gradient is linear. Newton's method finds the exact zeros of linear functions in only one step.