Final Exam

1. For each of the following, indicate whether the statement is true or false and explain your answer. The explanation may be a few words or a few sentences.

a. We have a large $n \times n$ unsymmetric matrix, A. Computing the LU factors takes no more than five times as much time as performing back and forward substitution.

b. A function $f(x)$, for $x \in R^n$, is called *strictly convex* if the Hessian matrix is positive definite for every x. If f is strictly convex then f can have at most one local minimum.

c. If a function $f(x)$ is a unique local minimum then f is convex and the Hessian matrix will have an LL^t factorization for every x.

d. We are solving the initial value problem $\dot{x}(t) = f(x)$, $x(0)$ given. We wish to compute $x(T)$ for some T. In exact arithmetic it is possible to achieve any given accuracy by taking time step Δt small enough.

e. If I had a $2GHz$ computer that performed exact arithmetic, there would be no incentive to use a higher order method, such as the fourth order Runge Kutta method, rather than forward Euler.

f. We have a normal computer using double precision floating point arithmetic and a smooth function $f(x)$ that can be evaluated to full machine precision for any x. We want the integral $I = \int_0^1 f(x)dx$ as accurately as possible. The greatest accuracy achievable using Simpson's rule and the trapezoid rule are about the same.

2. I need to compute the values of $f(x) = e^x - 1$ to high relative accuracy when x is close to zero. I will use double precision arithmetic and need 10 correct digits as long as x is not a denormalized (also called subnormal) number. The code will be roughly

```
#define EPS ??? (the question)
...
double f, x;
... ( get an x value )
if ( abs( x ) > EPS )
   f = exp(x) - 1;
else
   f = x + (1.d0/2)*x*x + (1.d0/(2*3))*x*x*x + ... + (n-th term)
```

Assume that `exp(x)` computes e^x to full double precision accuracy. Double precision arithmetic has 51 fraction bits.

A. What is the largest EPS so that the `if` branch gives 10 digits whenever it is executed? It is OK to be off by up to factor of 2 in EPS.

B. How many Taylor series terms do you need in the `else` branch to get 10 digits of accuracy whenever $|x| < EPS$. It is OK to be off by one.

3. A computer code of some kind was used to calculate a function of x. Some of the computed values are listed below.

A. Estimate the condition number of the problem of computing $f(x)$ when $x = 12.846$.

B. In single precision arithmetic (23 fraction bits), how many of the computed digits of $f(12.846)$ are likely to be correct?

x	f
12.845	112.65
12.846	0.2496
12.847	-105.71
12.848	-201.42

4. We need to compute $A = \int_0^1 f(x)/\sqrt{|x-a|}dx$ where $f(x)$ is smooth and easy to evaluate but we may have $a \in [0,1]$.

A. If we apply the trapezoid rule or Simpson's rule to the function $g(x) = f(x)/\sqrt{|x-a|}$, will we achieve more than first order accuracy?

B. Suppose we have a panel $I = [x_*, x_* + \Delta x]$ that may or may not contain a and we approximate $f(x)$ by its two term Taylor approximation $f_*(x) = f(x_*) + f'(x_*)(x - x_*)$ in the panel. For what p is

$$\int_I f(x)/\sqrt{|x-a|}x = \int_I f_*(x)/\sqrt{|x-a|}dx + O(\Delta x^p) ,$$

for any a and and x_* as $\Delta x \to 0$?

C. Suggest a method for computing A that uses only the mesh point values $f_k = f(k\Delta x)$ with $\Delta x = 1/n$ for some n that achieves greater than first order accuracy.

5. We are using Newton's method to solve a nonlinear system of n equations. The Jacobian matrix is reasonably conditioned at the solution. The residuals after iterations 4 and 5 are respectively .234 and .0447. Estimate the residual after iteration 6.

6. We have a probability density $f(x) = \dfrac{\cos(x)}{sqrtx}$ for $0 < x < \pi/2$ and $f(x) = 0$ otherwise. Suggest a rejection strategy for sampling this random variable. Hint: suppose U is a standard uniform random variable such as is supplied by the computer random number generator. What is the probability density function for $X = U^2$?

7. We are using Monte Carlo to estimate a number, A. We are unable to create a random variable, Y whose expected value is exactly A. However, there is a "step size", h, and a family of random variables whose bias has an asymptotic error expansion $E[Y_h] \sim$

$A + h^{1/2}A_1 + hA_2 + h^{3/2}A_3 + \cdots$ We have a procedure `float randY( float h)` so that successive calls produce independent samples $(Y_h^{(1)}, Y_h^{(2))}, \ldots)$ of Y_h. Sketch a Monte Carlo code that uses `randY(h)` and `randY(2*h)` to produce an estimate of A with bias $O(h)$ rather than $O(h^{1/2})$.