

CS 6210: HOMEWORK 3

Instructor: Anil Damle

Due: September 27, 2023

POLICIES

You may discuss the homework problems freely with other students, but please refrain from looking at their code or writeups (or sharing your own). Ultimately, you must implement your own code and write up your own solution to be turned in. Your solution, including plots and requested output from your code should be typeset and submitted via the Gradescope as a pdf file. Additionally, please submit any code written for the assignment. This can be done by either including it in your solution as an appendix, or uploading it as a zip file to the separate Gradescope assignment.

QUESTION 1:

Implement LU factorizations with no, partial, and complete pivoting for $A \in \mathbb{R}^{n \times n}$. Then use your code to address the following:

- (a) Demonstrate that your implementations can accurately solve well-conditioned square non-singular linear systems $Ax = b$ given A and b , and scales as $\mathcal{O}(n^3)$. Clearly outline how you test this and include your results.
- (b) Construct test cases where omitting any pivoting yields solutions with large relative residuals, and demonstrate this happens using your code.
- (c) Construct a test case where partial pivoting fails to yield a good solution but complete pivoting does, and demonstrate this happens using your code. Such examples exist, but are typically considered pathological and “rare” enough that simply using LU with partial pivoting (despite its poor worst case theoretical performance) is fine in practice.¹
- (d) Consider computing the LU factorization (all three ways) of the so-called Hilbert matrix of size n (defined as $H_{ij} = 1/(i+j-1)$ for $i, j = 1, \dots, n$). For moderate n what do you observe about $\|LU - H\|_2$ (adding permutation matrices as applicable)? What about the accuracy of a solution to $Hx = b$ as quantified by the relative residual? Explain your observations.
- (e) Since the Hilbert matrix H is symmetric positive definite (this is a fun proof to try if you are so inclined), presumably it would be natural to use a Cholesky decomposition rather than an LU factorization. Repeat the previous problem using a built-in Cholesky decomposition routine (or implement it yourself if you like). What do you observe?

¹One way to accomplish this is to appeal to the matrices that realize the worst case growth factor for LU with partial pivoting—they can be found in most of the textbooks listed for this course.

QUESTION 2 (UNGRADED, BUT INTERESTING):

In class we saw a backwards stability style result for pivoted LU factorization that incorporated the growth factor ρ . You may have noticed that we did not see any forward error results, i.e., we never said anything about $\tilde{L} - L$ or $\tilde{U} - U$. In fact, the forward error can be quite large. However, often what we ultimately care about is the solution to $Ax = b$. Devise a backwards error bound for solving $Ax = b$ with non-singular A using LU with partial pivoting followed by a sequence of triangular solves.² Your bound should explicitly incorporate the growth factor ρ . For this problem you may assume that in exact arithmetic there are no ties in the pivoting procedure and μ is sufficiently small such that the computed permutation matches the exact one.

One reason this problem is ungraded is that it has a very subtle point burried within it. While the goal of this question is for you to see how backwards error results can be “chained” together to get a result for the solution to the linear system, it requires bounds on $\|\tilde{L}\|$ and $\|\tilde{U}\|/\|A\|$.

In class we saw that when using partial pivoting we can ensure $\|L\| = \mathcal{O}(1)$ because all of the entries have magnitude bounded by 1. For \tilde{L} we can make a similar argument, though formally the easiest path is to ensure all the entires have magnitude less than $1 + \mu$ (which does not really change anything).

Unfortunately, bounding $\|\tilde{U}\|/\|A\|$ is more complex. In particular we (like many, but not all, books) defined the growth factor in terms of the exact U . Therefore, it is slightly delicate to use it in directly bounding $\|\tilde{U}\|/\|A\|$. For the purposes of this problem it suffices to derive a bound in terms of both the “exact” growth factor ρ and a computed growth factor $\tilde{\rho}$ (defined as the growth factor realized by \tilde{U}).

For those interested, this is actually a rather annoying sticking point in the analysis of these algorithms, and I am certainly happy to discuss it further. For example, to quote Higham (*Accuracy and Stability of Numerical Algorithms*; Section 9.3, page 165) when providing a theorem on backward error for solving $Ax = b$ via partially pivoted LU (Note, that in this book \hat{L} and \hat{U} are the computed LU factors):

“We hasten to admit to using an illicit manoeuvre in the derivation of this theorem: we have used bounds for \hat{L} and \hat{U} that strictly are valid only for the exact L and U .”

Higham goes on to comment (*Accuracy and Stability of Numerical Algorithms*; Section 9.14, page 189):

“The dilemma of whether to define the growth factor in terms of exact or computed quantities is faced by all authors; most make one choice or the other, and go on to derive, without comment, bounds that are strictly incorrect.”

So, there you have that.

²The computed solution \tilde{x} solves a linear system $(A + \delta A)\tilde{x} = b$, what can you say about $\|\delta A\|/\|A\|$?