

QR VERSUS CHOLESKY: A PROBABILISTIC ANALYSIS

MARK LIRA, RAM IYER, A. ALEXANDRE TRINDADE, AND VICTORIA HOWLE

Abstract. Least squares solutions of linear equations $Ax = b$ are very important for parameter estimation in engineering, applied mathematics, and statistics. There are several methods for their solution including QR decomposition, Cholesky decomposition, singular value decomposition (SVD), and Krylov subspace methods. The latter methods were developed for sparse A matrices that appear in the solution of partial differential equations. The QR (and its variant the RRQR) and the SVD methods are commonly used for dense A matrices that appear in engineering and statistics. Although the Cholesky decomposition is backward stable and known to have the least operational count, several authors recommend the use of QR in applications. In this article, we take a fresh look at least squares problems for dense A matrices with full column rank using numerical experiments guided by recent results from the theory of random matrices. Contrary to currently accepted belief, comparisons of the sensitivity of the Cholesky and QR solutions to random parameter perturbations for various low to moderate condition numbers show no significant difference to within machine precision. Experiments for matrices with artificially high condition numbers reveal that the relative difference in the two solutions is on average only of the order of 10^{-6} . Finally, Cholesky is found to be markedly computationally faster than QR – the mean computational time for QR is between two and four times greater than Cholesky, and the standard deviation in computation times using Cholesky is about a third of that of QR. Our conclusion in this article is that for systems with $Ax = b$ where A has full column rank, if the condition numbers are low or moderate, then the normal equation method with Cholesky decomposition is preferable to QR.

Key words. Least squares problems, QR decomposition, Cholesky decomposition, random matrix, statistics.

1. Introduction

The solution of linear equations of the type $Ax = b$, where $A \in \mathbb{R}^{m \times n}$ is fundamental to problems in science, engineering, applied mathematics and statistics. However, depending on the area, the problems have different features. For instance, linear PDEs in applied mathematics are characterized by a sparse matrix A with a large value of n (typically at least in the thousands), whereas classical parameter estimation problems in engineering and statistics are characterized by a dense matrix A with moderately large value for n (in the tens or hundreds). Furthermore, problems in engineering tend to be minimum norm and least squares if there is periodicity in the data, or have $m \geq n$ and $\text{rank}(A) = n$. In recent years, the area of smart materials and structures have yielded linear compact operator equations, which upon discretization result in least squares problems of moderately sized A matrix [1, 2, 3, 4].

Methods for the solution of linear equations include QR decomposition, Cholesky decomposition, singular value decomposition (SVD), Krylov subspace and Multigrid methods. Krylov subspace methods such as the generalized minimal residual method (GMRES) [5, 6] and the Lanczos method were developed for sparse A matrices that appear in the solution of partial differential equations [7]. Multigrid methods are useful in solving discretized differential equations [7]. The QR method of Francis [8, 9] and the singular value decomposition (SVD) methods are commonly

Received by the editors July 31, 2015.

2000 *Mathematics Subject Classification.* 65F99.

used for dense A matrices that appear in engineering and statistics [3]. Although the Cholesky decomposition is backward stable [10], several authors recommend the use of QR in applications [14, 15].

Assuming $\text{rank}(A) = n$, the Cholesky method for the solution of $Ax = b$ involves the formulation of the normal equations $A^T Ax = A^T b$, decomposing $A^T A = LL^T$ where L is a lower triangular matrix, and then solving for x using forward and backward substitutions. The basic QR method involves the solution of $Rx = Q^T b$.

A second class of applications where the Cholesky method might find favor are the minimum norm – least squares problems. Consider a linear system $Ax = b$, where $A \in \mathbb{R}^{m \times n}$ and $\text{rank}(A) = p < \min\{m, n\}$. We assume that p is unknown. A variant of the QR — the rank revealing QR (RRQR) [11, 12, 13] — may be used to find p and obtain a thin QR decomposition of A . Suppose $A = QR$ where Q is a $m \times p$ matrix with orthonormal columns and R is an upper-triangular $p \times n$ matrix. The normal equation then reduces to $(RR^T)v = Q^T b$ and $x = R^T v$.

- (i) One method for solving for x , which we refer to as QRC, computes a Cholesky factorization of the reduced normal equations. The matrix RR^T is a non-singular $p \times p$ matrix. Therefore, we may compute a Cholesky factorization $LL^T = RR^T$ and proceed to solve for v using forward and backward substitutions. Once v is found, x is computed.
- (ii) Another procedure to solve for x is the complete orthogonal factorization method (COF) [14], in which a QR factorization of R^T is computed. Suppose $R^T = US$, where U is a $n \times p$ matrix with orthonormal columns and S is a nonsingular and upper-triangular $p \times p$ matrix. Then $x = USv$ and the normal equation yields $S^T S v = Q^T b$. We may solve for $z = Sv$ from $S^T z = Q^T b$ and then find $x = Uz$.

In [14, 15] one finds a sensitivity analysis of the normal equation method, computing the sensitivity of the system to perturbations. The analysis looks at the upper bounds, which are not indicative of the behavior of the normal equation method for low to moderate conditioned systems. A perturbation analysis for the QR decomposition can be found in [16]. A related analysis is found in [17]. An error analysis of the Cholesky method is done in [18] and for positive semidefinite matrices in [10]. In Trefethen and Bau [15], an artificial example is constructed to show that the QR method should be considered to be superior to the normal equation method. The argument presented is that the normal equation method is susceptible to larger errors in the solution if the condition number κ_2 is at least as large as $1/\sqrt{\epsilon}$, where ϵ is the machine precision. Golub and Van Loan [14] state that the normal equation method is less accurate than a stable QR approach, though when the systems are ill-conditioned with large residuals, both methods are apt to produce comparable inaccurate results, which is a somewhat different statement than that of Trefethen. On the other hand, Higham[10] states that the Cholesky decomposition is one of the most numerically stable of all matrix algorithms, but the normal equation method is guaranteed to be backward stable only for well-conditioned matrices [19]. Trefethen [15] asserts that the SVD method is the only fully stable algorithm for solving rank-deficient problems. For high condition number systems (that is, $1/\sqrt{\epsilon} \leq \kappa(A) \leq 1/\epsilon$), it is possible for the solution of the normal equations to be highly erroneous for some vectors \bar{b} . An example verifying this is presented in Trefethen [15].

From the above discussion, it may be gathered that the authors were very concerned about backward stability *for all* matrices, and one can categorically say that for condition numbers greater than $\frac{1}{\sqrt{\epsilon}}$, the QR method is preferable to the

Cholesky method. However, for low to moderate condition numbers for the matrix A , there are no results that show one method to be superior over the other in terms of accuracy. Moreover, these references did not have the benefit of insights gained from recent developments in the theory of random matrices. The main point from this more recently developed theory [20] is that the probability of a matrix such as the one presented in Trefethen [15] (page 137) appearing in applications is much less than machine precision (to be precise, it is less than 2.5026×10^{-44} – this upper bound may be calculated using inequality (2) below [20]). In Trefethen’s example, 100 data points were fit with a 15-th degree polynomial. Such high order polynomials are known to produce very high variance fits in regression theory. Therefore, a commonly used statistical strategy is to trade-off variance and use much lower order polynomials that have slightly higher bias [21]. In other words, the example chosen to illustrate the superiority of the QR method in [15] is not likely to occur in real-world applications. As will be demonstrated, our conclusion in this article is that for low to moderate condition numbers the normal equation method with Cholesky decomposition is preferable to QR.

To resolve the QR versus Cholesky issue, we take recourse to probability and statistics. The theory of random matrices provides an answer to the question “how likely are very high condition number matrices to arise in engineering practice?”. According to Theorem 6.1 of [22], the probability density function of $\kappa(A)/n$ for an $n \times n$ matrix with independent and identically distributed (IID) Gaussian random entries converges in distribution to $(2/x^2 + 4/x^3) \exp\{-2/x - 2/x^2\}$. Elementary computations therefore lead to the following tail probability approximation for the condition number of a large square matrix of dimension n ,

$$(1) \quad P(\kappa(A) > y) \approx 1 - \exp\left\{-\frac{2n}{y} - \frac{2n^2}{y^2}\right\}.$$

The IEEE double precision, as used by MATLAB[®], $\epsilon = 2.2204 \times 10^{-16}$, and so $1/\sqrt{\epsilon} = 6.7109 \times 10^7$. Hence, for $n = 820$, we find that $P(\kappa(A) > 1/\sqrt{\epsilon}) \approx 2.44 \times 10^{-5}$. However, nearly every A matrix that appears in classical parameter estimation problems is not square and has more rows than columns (there is more data than parameters). The probability distribution of the condition number of a Gaussian random $n \times m$ matrix A satisfies [20]:

$$(2) \quad P(\kappa(A) > y) < \frac{1}{\sqrt{2\pi}} \left(\frac{Cn}{(|n-m|+1)y} \right)^{|n-m|+1},$$

where $5.3 \leq C \leq 6.414$. For $m = 820$ and $n = 516$, numbers from an application problem in [3], $P(\kappa(A) > 1/\sqrt{\epsilon}) \approx 10^{-1830}$. It is the exponent $|n-m|+1$ in inequality (2) that works in our favor. For example, for an A matrix with $m = 820$ and $n = 818$, $P(\kappa(A) > 1/\sqrt{\epsilon}) \approx 7.1 \times 10^{-13}$, which for such a small change in n compares dramatically with the value of 2.44×10^{-5} for a 820×820 matrix. Although problems with condition numbers as high as $1/\sqrt{\epsilon}$ are extremely rare, they can nevertheless be solved at least twice as fast by using Cholesky factorization, while retaining the same error probability as the QR method, as will be shown later in this article. The Cholesky factorization method performs even better for problems with smaller condition numbers than $1/\sqrt{\epsilon}$.

Next, we evaluate the asymptotic flop-counts for the solution of minimum-norm, least squares problems. A reference for flop-count analysis is Golub and Van Loan [14]. The COF method involves two thin-QR decompositions (using either Householder or Givens), and one forward substitution step. The asymptotic flop count for the Householder QR decomposition is $2mp^2 - \frac{2}{3}p^3$, while that for the Givens

QR decomposition is $3mp^2 - p^3$. The total flop count including the matrix multiplications and forward substitution is $2(2mp^2 - \frac{2}{3}p^3) + 2pn + \frac{1}{2}p^2 \approx 4mp^2 - \frac{4}{3}p^3$ for Householder QR, and $2(3mp^2 - p^3) + 2pn + p^2 \approx 6mp^2 - 2p^3$ for Givens QR.

The QRC method involves 1 thin QR decomposition, 1 Cholesky factorization, 1 forward, and 1 backward substitution. The asymptotic flop count for the QRC method using Householder QR decomposition is $2mp^2 - \frac{2}{3}p^3 + \frac{1}{3}p^3 + p^2 + np \approx 2mp^2 - \frac{1}{3}p^3$. The asymptotic flop count for the QRC method using Givens QR decomposition is $3mp^2 - p^3 + \frac{1}{3}p^3 + p^2 + np \approx 3mp^2 - \frac{2}{3}p^3$.

The asymptotic flop count for the COF method is at least as large as the QRC method if $p < 2m$ if the Householder QR is employed. If the Givens QR is employed, then the condition for the asymptotic flop count of COF method being at least as large as that for the QRC method is $p \leq \frac{3}{4}m$. Both conditions are satisfied because we have $p \leq n \leq m$.

This analysis shows that for minimum-norm least squares problems or least squares problems with full column rank, numerical solution methodologies that incorporate a final Cholesky decomposition step instead of a QR step have lower operational count.

2. Methodology and Results

In the literature [14, 15] one finds examples comparing the QR and Cholesky methods for high condition numbered systems. One also finds an upper bound for the sensitivity under perturbations for the normal equation [14], but not for the QR method. For low to moderately high condition numbered systems, the main unanswered question is: are the two methods comparable in terms of sensitivity to perturbations? A secondary question concerns the speed of computation of the two solutions. We provide some answers to these questions in this section.

Consider the system of equations $Ax = b$ where A has full-column rank. Let the solution obtained by using the normal equations and the Cholesky method be denoted x_{CH} , and that obtained by using the QR method be denoted x_{QR} . To compute the relative errors from the QR and Cholesky algorithms in this full-rank case, we compute the solution for x in the system of equations $Ax = b$, and the solution for $x + \delta x$ in the perturbed system $(A + \delta A)(x + \delta x) = b + \delta b$. The relative error for the Cholesky algorithm is then defined as $e_{CH} = \|\delta x_{CH}\|/\|x_{CH}\|$. The relative error e_{QR} for the QR algorithm is similarly defined. Specific numerical experiments (Experiments 1 and 2 below) were then conducted using the commercial software MATLAB[®].

Experiment 1. *Matrices A with $m = 100$ rows and $n = 90$ columns were generated, with entries comprising IID random draws from a standard normal distribution. Vectors b had IID entries drawn from $U[0, 1]$, a uniform distribution on $[0, 1]$. Perturbations δA and δb comprise IID entries drawn from $\mathcal{N}(0, t)$, a Gaussian distribution with mean zero and standard deviation t . For each of the three Gaussian standard deviation values of $t = 1, 0.1, 10^{-10}$, 1,000 trials were conducted (3,000 trials in total). One pair of relative errors, e_{CH} and e_{QR} , were then computed for each trial.*

The results of Experiment 1 appear in Tables 1 and 2. For each value of t , Table 1 shows a statistical five-number summary (minimum, 1st quartile, median, 3rd quartile, maximum) for the 1,000 differences in the pairs of relative errors, $e_{CH} - e_{QR}$. These summaries clearly show that the two relative errors, e_{CH} and e_{QR} , are practically identical to within machine precision ($\epsilon \approx 10^{-16}$). Moreover,

this finding seems to be insensitive of the value of t , which controls the magnitude of the perturbations in δA and δb .

TABLE 1. Statistical summaries for the 1,000 differences in the pairs of relative errors $e_{CH} - e_{QR}$ obtained from Experiment 1. Three sets of 1,000 trials were conducted, each set corresponding to a different value of t .

t	Minimum	1st Quartile	Median	3rd Quartile	Maximum
1.0E-10	-9.4E-14	-6.3E-15	-6.0E-17	6.0E-15	1.4E-13
1.0E-01	-1.2E-13	-7.2E-15	4.9E-16	7.1E-15	8.7E-14
1.0E-00	-1.3E-13	-4.4E-15	-5.6E-17	4.5E-15	5.7E-14

Table 1 therefore suggests that the QR and Cholesky methods are comparable in terms of sensitivity to perturbations. Table 2 provides another facet of this comparison, by shedding light on the speed of computation of each solution, x_{CH} and x_{QR} . The table shows statistical summaries of measures of centrality (mean and median) and of dispersion (standard deviation) for each set of 1,000 trials in Experiment 1. The main message from this table is that Cholesky is about twice as fast as QR, regardless of t . We also notice more variability in the QR solution time; its standard deviation is approximately three times that of the corresponding Cholesky time.

TABLE 2. Statistical summaries for the 1,000 compute times of each pair of solutions, x_{CH} and x_{QR} , obtained from Experiment 1. Three sets of 1,000 trials were conducted, each set corresponding to a different value of t .

t	QR times in seconds			CH times in seconds		
	mean	median	std. dev.	mean	median	std. dev.
1.0E-10	4.72E-04	4.37E-04	8.15E-05	2.41E-04	2.29E-04	2.76E-05
1.0E-01	4.83E-04	4.40E-04	8.71E-05	2.46E-04	2.29E-04	2.95E-05
1.0E-00	4.68E-04	4.35E-04	7.88E-05	2.40E-04	2.29E-04	2.64E-05

The condition numbers of the matrices were not controlled for in Experiment 1. Equation (1) assures us that for a 100×90 Gaussian random matrix A , $P(\kappa(A) > 1/\sqrt{\epsilon}) < 8.5 \times 10^{-68}$. Therefore, in order to investigate the statements by Golub and Van Loan [14] and Trefethen [15] concerning the performance of Cholesky and QR for extreme condition numbers, we conduct Experiment 2 where matrices A are chosen with $\kappa(A)$ forced to be close to $1/\sqrt{\epsilon}$.

Experiment 2. Vectors b were generated from $\mathcal{U}[0, 1]$ as in Experiment 1. Matrices A of size $100 \times n$ with singular values $\sigma_k = k^4 + r_k$ were constructed, where $n \in \{1, 2, 5, 10, \dots, 90\}$, r_k is a random number drawn from $\mathcal{U}[0, 1]$, and $k = 1, \dots, n$. This was achieved via the SVD of $A = USV$, where S is the diagonal matrix of singular values, and the entries of U are IID from $\mathcal{N}(0, t)$. The columns of U were then orthonormalized using Gram-Schmidt to a tolerance of 10^{-10} . V was created similarly. To gauge the size of the difference between the two solutions, x_{QR} and x_{CH} , for such high condition numbers, we compute the norm of the difference in the two solutions relative to the norm of the QR solution,

$$d_{CH,QR} = \frac{\|x_{QR} - x_{CH}\|}{\|x_{QR}\|},$$

and 1,000 trials are run for each value of n (12,000 trials in total). Each set of 1,000 trials spanned a variety of values for the standard deviation of the Gaussian entries in U , ranging over the set: $t = 10^{-1}, 10^{-3}, \dots, 10^{-15}$.

It can be seen that for the matrices in Experiment 2, the condition numbers are of the order of 10^7 (but because of the artificial construction the resulting matrices cannot be viewed as being truly random). The resulting scatter of $\log_{10}(d_{CH,QR})$ vs. \log_{10} of condition number in Figure 2, shows that the difference between x_{QR} and x_{CH} increases with condition number (and rank = n), but even for condition numbers as large as $1/\sqrt{\epsilon} = 10^8$ the value of the relative difference $d_{CH,QR}$ is of the order of 10^{-3} , with a median value of approximately 4×10^{-6} .

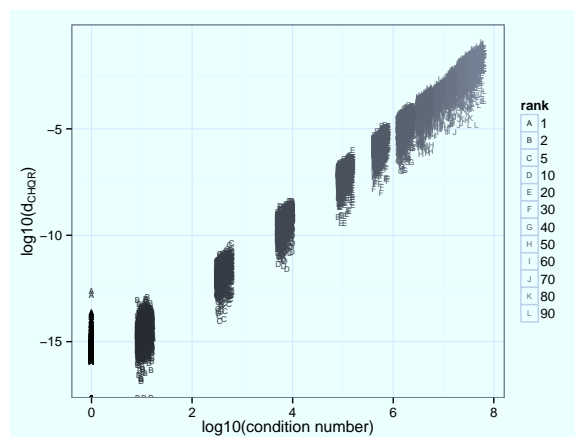


FIGURE 1. Log base 10 of the error of the CH solution relative to QR solution, $d_{CH,QR}$, as a function of log base 10 of the condition number, for full-rank matrices with 100 rows and various column sizes. The 12 groups of condition numbers are organized according to the rank of the matrix, each group identified by a distinct letter. Each of the 12 groups of points is based on 1,000 simulations. (A total of 104 values with $d_{CH,QR} = 0$ were discarded.)

The computation times for the two solutions, x_{CH} and x_{QR} , in Experiment 2 were also recorded. Figure 2 displays \log_{10} of the ratio of computation times as a function of \log_{10} of condition number and rank for a matrix A with 100 rows. Since very few points fall below zero (horizontal line) for $\log_{10}(\text{QR time}/\text{CH time})$, we conclude that the CH method overwhelmingly enjoys smaller computation times (about 4 times faster than QR on average) even in this extreme condition number setting. Moreover, this ratio of compute times is fairly insensitive to condition number and rank.

3. Conclusion

In this article, we compared the normal equation method using a Cholesky decomposition to the QR decomposition method for solving full-rank least squares problems, using numerical experiments motivated by probabilistic arguments. For generic matrices with low to moderate condition numbers, the accuracy and sensitivity of the solutions to parameter perturbations obtained by each of the two methods were found to be comparable to within a tolerance specified by machine

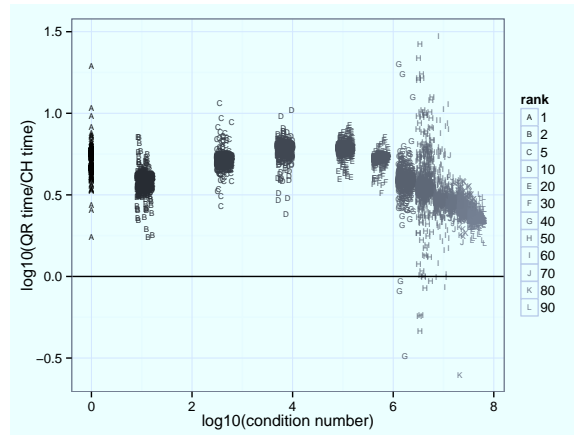


FIGURE 2. Log base 10 of ratio of computation times of the QR vs. CH solutions as a function of log base 10 of the condition number, for full-rank matrices with 100 rows and various column sizes. The 12 groups of condition numbers are organized according to the rank of the matrix, each group identified by a distinct letter. Each of the 12 groups of points is based on 1,000 simulations.

precision. Experiments for matrices with artificially high condition numbers, revealed that the norm of the difference in the two solutions relative to QR varies approximately from machine precision to an upper bound of 10^{-3} , with an average value of the order of 10^{-6} . Similar experiments showed Cholesky to be computationally faster with more consistent compute times across a broad spectrum of condition numbers; its execution times were found to be between two and four times shorter, with the standard deviation of Cholesky times being about a third of that for QR. Our conclusion is therefore that, for systems with $Ax = b$ where A has full column rank, if the condition numbers are low to moderate, then the normal equation method with Cholesky decomposition is preferable to QR.

References

- [1] K. H. Hoffmann and G. H. Meyer, A least squares method for finding the Preisach hysteresis operator from measurements, *Numer. Math.* Vol. 55, pp. 695–710, 1989.
- [2] R. B. Gorbet and D. W. L. Wang and K. A. Morris, Preisach model identification of a two-wire SMA actuator, *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2161–2167, May, 1998, Leuven, Belgium.
- [3] R. Iyer and M. Shirley, Hysteresis Parameter Identification with Limited Experimental Data, *IEEE Trans. Magnetics*, Vol. 40, No. 5, pp 3227–3239, Sep. 2004.
- [4] X. Tan and J. Baras. Modeling and Control of hysteresis in magnetostrictive actuators, *Automatica*, Vol. 40, Iss. 9, pp. 1469–1480, Sep. 2004.
- [5] Y. Saad and M.H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.*, Vol. 7, pp. 856–869, 1986.
- [6] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd edition, Society for Industrial and Applied Mathematics, 2003.
- [7] J. Stoer and R. Bulirsch, *Introduction to Numerical Analysis*, 3rd edition, Springer-Verlag, New York Inc, 2002.
- [8] J. Francis, The QR transformation: a unitary analogue to the LR transformation - Part 1, *The Computer Journal*, pp. 265–271, 1961.
- [9] J. Francis, The QR transformation - Part 2, *The Computer Journal*, pp. 332–345, 1962.

- [10] N. J. Higham. Analysis of the Cholesky Decomposition of a Semi-Definite Matrix in M. G. Cox and S. J. Hammarling, eds., *Reliable Numerical Computation*, Oxford University Press, 1990, pp. 161–185.
- [11] T. Chan, Rank revealing QR factorizations, *Lin. Alg. Appl.*, Vol. 88/89, pp. 67–82, 1987.
- [12] Y. Hong and C.-T. Pan, Rank revealing QR factorizations, *Mathematics of Computation*, Vol. 58, No. 197, pp. 213–232, 1992.
- [13] M. Gu and S. Eisenstat, Efficient algorithms for computing a strong rank-revealing QR factorization, *SIAM J. Sci. Comput.*, Vol. 17, No. 4, pp. 848–869, Jul. 1996.
- [14] G. H. Golub and C. F. Van Loan, *Matrix Computations*, Third Edition, The John Hopkins University Press, 1996.
- [15] L. N. Trefethen and D. Bau III, *Numerical Linear Algebra*, Society for Industrial and Applied Mathematics, 1997.
- [16] G. W. Stewart, Perturbation Bound for the QR Factorization of a Matrix *SIAM J. Numerical Analysis*, Vol. 14, No 3, Jun. 1977.
- [17] C. C. Paige, Computer Solution and Perturbation Analysis of Generalized Linear Least Squares Problems, *Mathematics of Computation*, Vol. 33, No. 145, pp. 171–183, Jan. 1979.
- [18] J. Meinguet, Refined Error Analyses of Cholesky Factorization, *SIAM J. Numer. Anal.*, Vol. 20, No. 6, pp. 1243–1250, Dec. 1983.
- [19] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, Society of Industrial and Applied Mathematics, Philadelphia, 1996. ISBN 0-89871-355-2.
- [20] Z. Chen and J. Dongarra phCondition numbers of Gaussian Random Matrices, *SIAM J. Matrix Anal. & Appl.*, Vol. 27, No. 3, pp. 603–620, 2005.
- [21] F. E. Harrell Jr., *Regression Modeling Strategies*, Springer-Verlag, 2010.
- [22] A. Edelman, phEigenvalues and Condition numbers of Random Matrices, *SIAM J. Matrix Anal. Appl.*, Vol. 9, No. 4, pp. 543–560, Oct. 1988.

AFLCMC/WLZWE, 3001 Staff Drive, 2AI67A Tinker AFB, OK, 73145, USA
E-mail: `mark.lira@us.af.mil`

Department of Mathematics and Statistics, Texas Tech University, Lubbock, TX 79409, USA
E-mail: `ram.iyer@ttu.edu`

Department of Mathematics and Statistics, Texas Tech University, Lubbock, TX 79409, USA
E-mail: `alex.trindade@ttu.edu`

Department of Mathematics and Statistics, Texas Tech University, Lubbock, TX 79409, USA
E-mail: `victoria.howle@ttu.edu`