

***Probabilistic Rounding Error Analysis of  
Householder QR Factorization***

Connolly, Michael P. and Higham, Nicholas J.

2022

MIMS EPrint: **2022.5**

Manchester Institute for Mathematical Sciences  
School of Mathematics

The University of Manchester

Reports available from: <http://eprints.maths.manchester.ac.uk/>

And by contacting: The MIMS Secretary  
School of Mathematics  
The University of Manchester  
Manchester, M13 9PL, UK

ISSN 1749-9097

# PROBABILISTIC ROUNDING ERROR ANALYSIS OF HOUSEHOLDER QR FACTORIZATION\*

MICHAEL P. CONNOLLY<sup>†</sup> AND NICHOLAS J. HIGHAM<sup>†</sup>

**Abstract.** When an  $m \times n$  matrix is premultiplied by a product of  $n$  Householder matrices the worst-case normwise rounding error bound is proportional to  $mnu$ , where  $u$  is the unit roundoff. We prove that this bound can be replaced by one proportional to  $\sqrt{mnu}$  that holds with high probability if the rounding errors are mean independent and of mean zero. The proof makes use of a matrix concentration inequality. In particular, this result applies to Householder QR factorization. The same square rooting of the error constant applies to two-sided transformations by Householder matrices and hence to standard QR-type algorithms for computing eigenvalues and singular values. It also applies to Givens QR factorization. These results complement recent probabilistic rounding error analysis results for inner-product based algorithms and show that the square rooting effect is widespread in numerical linear algebra. Our numerical experiments, which make use of a new backward error formula for QR factorization, show that the probabilistic bounds give a much better indicator of the actual backward errors and their rate of growth than the worst-case bounds.

**Key words.** floating-point arithmetic, backward error analysis, backward error, probabilistic rounding error analysis, Givens QR factorization, Householder QR factorization, matrix concentration inequality

**AMS subject classifications.** 65G50, 65F05

**1. Introduction.** It is well known that backward error bounds from worst-case rounding error analyses can greatly overestimate the backward error. Recently, it was proved that for inner product-based algorithms with backward error bounds of the form  $f(n)u$ , where  $n$  is the problem dimension and  $u$  is the unit roundoff, a bound proportional to  $\sqrt{f(n)}u$  holds with high probability under suitable assumptions on the rounding errors [15]. These results apply to matrix multiplication, Cholesky factorization, LU factorization, and the solution of triangular systems, but not to algorithms based on orthogonal transformations.

A QR factorization of  $A \in \mathbb{R}^{m \times n}$  ( $m \geq n$ ) is a factorization  $A = QR$  where  $Q \in \mathbb{R}^{m \times m}$  is orthogonal and  $R \in \mathbb{R}^{m \times n}$  is upper trapezoidal. A Householder matrix is a matrix of the form

$$P = I - \frac{2}{v^*v}vv^*, \quad 0 \neq v \in \mathbb{R}^m.$$

The vector  $v$  can be chosen so that in the product  $y = Px$  all elements of  $y$  except the first are zero. By applying a sequence of such Householder matrices to  $A$  we can reduce it to upper trapezoidal form:  $P_n \dots P_2 P_1 A = R$ . We then have  $Q = (P_n \dots P_2 P_1)^*$ .

The standard rounding error analysis for Householder QR factorization is summarized in the following result [12, Thm. 19.4]. We define

$$(1.1) \quad \gamma_n = \frac{nu}{1 - nu}$$

---

\*Version of February 15, 2022.

**Funding:** This work was supported by the Royal Society, Engineering and Physical Sciences Research Council grant EP/P020720/1, and the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.

<sup>†</sup>Department of Mathematics, University of Manchester, Manchester, M13 9PL, UK (michael.connolly-3@manchester.ac.uk, nick.higham@manchester.ac.uk).

and use the 2-norm  $\|x\|_2 = (x^*x)^{1/2}$ , the corresponding subordinate matrix norm, and the Frobenius norm  $\|A\|_F = \text{trace}(A^*A)^{1/2}$ .

**THEOREM 1.1.** *Let  $\widehat{R} \in \mathbb{R}^{m \times n}$  be the computed upper trapezoidal QR factor of  $A \in \mathbb{R}^{m \times n}$  ( $m \geq n$ ) obtained via the Householder QR algorithm. There exists an orthogonal  $Q \in \mathbb{R}^{m \times m}$  such that*

$$A + \Delta A = Q\widehat{R},$$

where

$$(1.2) \quad \|\Delta a_j\|_2 \leq \gamma_{cmn} \|a_j\|_2, \quad j = 1:n,$$

with  $c$  a small integer constant.

The columnwise bounds (1.2) yield the normwise backward error bound

$$(1.3) \quad \|\Delta A\|_F \leq \gamma_{cmn} \|A\|_F.$$

The purpose of this work is to show that the worst-case backward error bounds (1.2) and (1.3) for Householder QR factorization can be replaced by probabilistic bounds in which the dimension-dependent constants are proportional to the square roots of the original constants under suitable assumptions on the rounding errors.

We use the standard model of floating-point arithmetic, which assumes that the elementary operations are correctly rounded, so that

$$(1.4) \quad \text{fl}(x \text{ op } y) = (x \text{ op } y)(1 + \delta), \quad |\delta| \leq u.$$

In a rounding error analysis products of  $1 + \delta$  terms arise, and their distance from 1 can be bounded using the following lemma [12, Lem 3.1].

**LEMMA 1.2.** *If  $|\delta_i| \leq u$  and  $\rho_i = \pm 1$  for  $i = 1:n$ , and  $nu < 1$ , then*

$$(1.5) \quad \prod_{i=1}^n (1 + \delta_i)^{\rho_i} = 1 + \theta_n, \quad |\theta_n| \leq \gamma_n.$$

An analogous probabilistic result [7, Thm. 4.6] is given in terms of the constant

$$(1.6) \quad \tilde{\gamma}_n(\lambda) = \exp\left(\frac{\lambda\sqrt{nu} + nu^2}{1 - u}\right) - 1 = \lambda\sqrt{nu} + O(u^2).$$

with  $\lambda > 0$ . We note that random variables  $\delta_1, \delta_2, \dots, \delta_n$  are *mean independent* if  $\mathbb{E}(\delta_k | \delta_{k-1}, \dots, \delta_1) = \mathbb{E}(\delta_k)$  for  $k = 2:n$ , where  $\mathbb{E}$  denotes expectation.

**LEMMA 1.3.** *Let  $\delta_1, \delta_2, \dots, \delta_n$  be mean independent random variables of mean zero such that  $|\delta_i| \leq u$  for all  $i$ , and let  $\rho_i = \pm 1$ ,  $i = 1:n$ . Then for any constant  $\lambda > 0$ ,*

$$(1.7) \quad \prod_{i=1}^n (1 + \delta_i)^{\rho_i} = 1 + \theta_n, \quad |\theta_n| \leq \tilde{\gamma}_n(\lambda),$$

holds with probability at least  $p(\lambda) = 1 - 2\exp(-\lambda^2/2)$ .

Throughout this work we will use the following model of rounding errors, which was also used in [7] and [16].

MODEL 1.4 (probabilistic model of rounding errors). *Let the computation of interest generate rounding errors  $\delta_1, \delta_2, \dots$  (in that order) satisfying (1.4). The  $\delta_k$  are mean independent random variables of mean zero.*

For the round to nearest rounding mode (the default in IEEE arithmetic) the error bounds obtained under Model 1.4 often reflect the behavior of the actual errors obtained, but in some circumstances the model is not valid because the rounding errors are highly dependent or have nonzero mean, and the bound can then be violated [7, sec. 7], [8, sec. 8], [15, sec. 4.2]. For stochastic rounding [7], [8], the rounding errors *always* satisfy the model. For further discussion of the applicability of the model see [7], [15].

In our probabilistic rounding error analysis it is more difficult than in worst-case rounding error analysis to keep track of constants and to bound terms of all orders in  $u$ . We will express our bounds in terms of  $c_1, c_2, \dots$ , which denote generic integer constants of modest size, and we will state the bounds to first order in  $u$ . We will nevertheless give explicit bounds for the probabilities with which the error bounds hold, and these show that the analysis does not introduce exponential worsening of probabilities. All our probabilistic results can informally be stated as saying that the given backward error bound holds with  $\lambda$  a modest constant with high probability.

We begin in section 2 with a probabilistic analysis of the rounding errors in constructing a Householder vector. In section 3 we state a matrix concentration inequality that we use in section 4 in our probabilistic rounding error analysis of the application of products of Householder matrices to a matrix and of Householder QR factorization. We discuss the implications of the analysis in section 5, including for two-sided transformations and for Givens QR factorization. In section 6 we give a new backward error formula for QR factorization that we use in section 7, where we carry out numerical experiments to compare the results of the analysis with the practical behavior. Conclusions are given in section 8.

**2. Construction of a Householder vector.** We first give a probabilistic error result for the construction of the Householder vectors needed in QR factorization. We begin by considering the evaluation of the vector 2-norm. We denote by  $\text{fl}(\text{expr})$  the computed result of evaluating the expression  $\text{expr}$ .

LEMMA 2.1. *Under Model 1.4, the computed 2-norm of  $x \in \mathbb{R}^n$  satisfies*

$$(2.1) \quad \text{fl}(\|x\|_2) = \|x\|_2(1 + \eta), \quad |\eta| \leq c_1 \tilde{\gamma}_n(\lambda),$$

*with probability at least  $p_1(\lambda, n) = 1 - 2n \exp(-\lambda^2/2)$ .*

*Proof.* By standard inner-product analysis [7, Thm. 4.8], [15, Thm. 3.1],

$$(2.2) \quad \text{fl}(x^*x) = x^*x(1 + \theta_n), \quad |\theta_n| \leq \tilde{\gamma}_n(\lambda),$$

holds with probability at least  $p_1(\lambda, n)$ . Then  $\text{fl}(\|x\|_2) = \|x\|_2 \sqrt{1 + \theta_n}(1 + \delta)$ , with  $|\delta| \leq u$ . For  $\sqrt{1 + \theta_n} = 1 + \alpha$ , we certainly have  $|\alpha| \leq |\theta_n|$ . Then  $(1 + \alpha)(1 + \delta) = 1 + \beta$ , with  $|\beta| \leq |\theta_n| + u + |\theta_n|u$ . Hence provided  $|\theta_n| \leq \tilde{\gamma}_n(\lambda)$ , we have  $|\beta| \leq \tilde{\gamma}_n(\lambda) + u + \tilde{\gamma}_n(\lambda)u \leq c_1 \tilde{\gamma}_n(\lambda)$ . The bound holds with probability at least the probability of (2.2) holding.  $\square$

We now derive a probabilistic version of the worst-case error bound for construction of a Householder vector [12, Lem. 19.1].

LEMMA 2.2. Let  $x \in \mathbb{R}^n$ . Consider the construction of  $\beta \in \mathbb{R}$  and  $v \in \mathbb{R}^n$  such that  $Px = \sigma e_1$ , where  $P = I - \beta vv^*$  is a Householder matrix and  $\beta = 2/(v^*v)$ , by

$$\begin{aligned} v &= x, \\ s &= \text{sign}(x_1)\|x\|_2, \\ v_1 &= v_1 + s, \\ \beta &= 1/(sv_1). \end{aligned}$$

Under Model 1.4 the computed  $\hat{\beta}$  and  $\hat{v}$  satisfy  $\hat{v}(2:n) = v(2:n)$  and

$$\hat{v}_1 = v_1(1 + \eta), \quad \hat{\beta} = \beta(1 + \eta),$$

where  $|\eta| \leq c_3\tilde{\gamma}_n(\lambda) + O(u^2)$  with probability at least  $p_1(\lambda, n)$ .

*Proof.* From Lemma 2.1 we have that  $\hat{s} = s(1 + \Delta s)$ , with  $|\Delta s| \leq c_1\tilde{\gamma}_n(\lambda)$  holds with probability at least  $p_1(\lambda, n)$ . Define  $w = v_1 + s$ , so that

$$\begin{aligned} \hat{w} &= (v_1 + \hat{s})(1 + \delta), \quad |\delta| \leq u \\ &= v_1 + s + \delta(v_1 + s) + s\Delta s(1 + \delta) =: w + \Delta w_1. \end{aligned}$$

Because  $v_1$  and  $s$  have the same sign,  $|s| \leq |v_1 + s|$ , and so

$$\begin{aligned} |\Delta w_1| &= \left| (v_1 + s) \left( \delta + \frac{s}{v_1 + s} \Delta s(1 + \delta) \right) \right| \\ &\leq |w| (u + c_1\tilde{\gamma}_n(\lambda) + c_1\tilde{\gamma}_n(\lambda)u) \leq c_2\tilde{\gamma}_n(\lambda)|w|. \end{aligned}$$

So  $\hat{w} = w(1 + \Delta w)$ , where  $|\Delta w| \leq c_2\tilde{\gamma}_n(\lambda)$  holds with probability at least  $p_1(\lambda, n)$ . We then have, using (1.4) and (1.5),

$$\hat{\beta} = \text{fl}(1/(\hat{s}\hat{w})) = \frac{1 + \theta_2}{s(1 + \Delta s)w(1 + \Delta w)} = \frac{1}{sw}(1 + \xi),$$

where  $|\xi| \leq c_3\tilde{\gamma}_n(\lambda) + O(u^2)$ .  $\square$

If we redefine the Householder matrix as  $P = I - vv^*$  with  $\|v\|_2 = \sqrt{2}$  then Lemma 2.2 amounts to the result

$$(2.3) \quad \hat{v} = v + \Delta v, \quad |\Delta v| \leq c_4\tilde{\gamma}_n(\lambda)|v|,$$

where the bound holds with probability at least  $p_1(\lambda, n)$ .

**3. Matrix concentration inequalities.** The proof of Lemma 1.3 in [7] makes use of the scalar Azuma–Hoeffding inequality [17, Thm. 13.4]. Here, we will use a matrix version of that result due to Tropp [18, Thm. 7.1]. Our notation is as follows: we write (in this section only),  $X \leq Y$ , where  $X$  and  $Y$  are symmetric, to mean that  $Y - X$  is positive semidefinite; the expectation of a matrix is defined componentwise; and we denote by  $\lambda_{\max}$  the largest eigenvalue of a symmetric matrix.

THEOREM 3.1. Let  $X_1, \dots, X_n$  be a sequence of random symmetric  $d \times d$  matrices. If  $\mathbb{E}(X_k \mid X_{k-1}, \dots, X_1) = 0$  and  $X_k^2 \leq A_k^2$  almost surely for all  $k$ , where  $A_1, \dots, A_n$  is a fixed sequence of symmetric  $d \times d$  matrices, then for all  $t \geq 0$ ,

$$(3.1) \quad \mathbb{P} \left( \lambda_{\max} \left( \sum_{k=1}^n X_k \right) \geq t \right) \leq d \exp(-t^2/(8\sigma^2)),$$

where

$$(3.2) \quad \sigma^2 = \left\| \sum_{k=1}^n A_k^2 \right\|_2.$$

We need a version of Theorem 3.1 for nonsymmetric matrices. Define

$$(3.3) \quad \phi(B) = \begin{bmatrix} 0 & B \\ B^* & 0 \end{bmatrix}$$

to be the symmetric dilation of the rectangular matrix  $B$ . It is well known that

$$(3.4) \quad \lambda_{\max}(\phi(B)) = \|B\|_2.$$

The required result can be obtained by applying the theorem to the symmetric dilation of the sequence  $X_1, \dots, X_n$  and using (3.4) [18, Rem. 7.3]. We also rescale by defining  $\lambda = t/(2\sqrt{2}\sigma)$ .

**THEOREM 3.2.** *Let  $X_1, \dots, X_n$  be a sequence of random  $d_1 \times d_2$  matrices. If  $\mathbb{E}(X_k \mid X_{k-1}, \dots, X_1) = 0$  for all  $k$ , and  $X_k X_k^* \leq A_k A_k^*$  and  $X_k^* X_k \leq A_k^* A_k$  almost surely for all  $k$ , where  $A_1, \dots, A_n$  is a fixed sequence of  $d_1 \times d_2$  matrices, then for all  $\lambda \geq 0$ ,*

$$(3.5) \quad \mathbb{P} \left( \left\| \sum_{k=1}^n X_k \right\|_2 \geq 2\sqrt{2}\sigma\lambda \right) \leq (d_1 + d_2) \exp(-\lambda^2),$$

where

$$(3.6) \quad \sigma^2 = \max \left( \left\| \sum_{k=1}^n A_k A_k^* \right\|_2, \left\| \sum_{k=1}^n A_k^* A_k \right\|_2 \right) \leq \sum_{k=1}^n \|A_k\|_2^2.$$

This is the key result that we need in the next section.

**4. Application of a sequence of Householder matrices.** Our probabilistic rounding error analysis for Householder QR factorization follows a similar strategy as for the worst-case analysis in [12, sec. 19.3]: we analyze the application of a sequence of general Householder matrices to a vector, then the application of the same sequence to a matrix, and finally specialize to the Householder matrices used in QR factorization.

For a single application of an  $m \times m$  Householder matrix our backward error analysis gives an error constant proportional to  $m^{1/2}$ . A straightforward inductive argument for the application of  $n$  Householder matrices would lead to a backward error bound with a constant proportional to  $nm^{1/2}$ , which is unsatisfactory as it is linear in  $n$ , as for the worst-case analysis. Our key observation is that by applying the matrix concentration inequality in Theorem 3.2 we can obtain a bound with a constant of order  $n^{1/2}m^{1/2}$ .

**LEMMA 4.1.** *Consider the (given) Householder matrices*

$$P_i = I - v_i v_i^* \in \mathbb{R}^{m \times m}, \quad v_i^* v_i = 2, \quad i = 1:r,$$

and the product

$$a_{r+1} = P_r \dots P_2 P_1 a, \quad a = a_1 \in \mathbb{R}^m,$$

computed as  $a_{j+1} = P_j a_j = a_j - (v_j^* a_j) v_j$ ,  $j = 1:r$ . Under Model 1.4, the computed  $\hat{a}_{r+1}$  satisfies

$$(4.1) \quad \hat{a}_{r+1} = (P_r + \Delta P_r) \dots (P_1 + \Delta P_1) a,$$

where

$$(4.2) \quad \|\Delta P_j\|_2 \leq c_5 \tilde{\gamma}_m(\lambda)$$

holds for all  $j$  with probability at least  $p_2(\lambda, m, r) = 1 - 2rm \exp(-\lambda^2/2)$ .

*Proof.* Define the scalars  $s_j = v_j^* a_j$ . From standard rounding error analysis [12, sec. 3.1], [15, Proof of Thm. 3.1], the computed  $\hat{s}_j = v_j^* D_j \hat{a}_j$ , where

$$D_j = \text{diag}(d_k^{(j)}), \quad d_k^{(j)} = \prod_{i=1}^m (1 + \delta_{k,i}^{(j)}), \quad |\delta_{k,i}^{(j)}| \leq u, \quad i, k = 1:m.$$

Then we form

$$\hat{a}_{j+1} = \text{fl}(\hat{a}_j - \hat{s}_j v_j) = (I + \Lambda_1^{(j)})(\hat{a}_j - (I + \Lambda_2^{(j)})\hat{s}_j v_j),$$

where

$$\Lambda_k^{(j)} = \text{diag}(\epsilon_{k,i}^{(j)}), \quad |\epsilon_{k,i}^{(j)}| \leq u, \quad k = 1, 2.$$

Hence

$$\begin{aligned} \hat{a}_{j+1} &= (I + \Lambda_1^{(j)})(\hat{a}_j - (I + \Lambda_2^{(j)})v_j^* D_j \hat{a}_j \cdot v_j) \\ &= (I + \Lambda_1^{(j)})(I - (I + \Lambda_2^{(j)})v_j v_j^* D_j) \hat{a}_j \\ &= \left[ I - (I + \Lambda_2^{(j)})(v_j v_j^* + v_j v_j^* (D_j - I)) + \Lambda_1^{(j)}(I - (I + \Lambda_2^{(j)})v_j v_j^* D_j) \right] \hat{a}_j \\ (4.3) \quad &= (P_j + \Delta P_j) \hat{a}_j, \end{aligned}$$

where

$$(4.4) \quad \Delta P_j = -\Lambda_2^{(j)} v_j v_j^* - \underbrace{(I + \Lambda_2^{(j)})v_j v_j^* (D_j - I) + \Lambda_1^{(j)}(I - (I + \Lambda_2^{(j)})v_j v_j^* D_j)}_{(*)},$$

and so

$$\begin{aligned} \|\Delta P_j\|_2 &\leq 2u + 2(1+u)\|D_j - I\|_2 + u(1 + 2(1+u)\|D_j\|_2) \\ &= 3u + 2(1+u)\|D_j - I\|_2 + 2u(1+u)\|D_j\|_2. \end{aligned}$$

From Lemma 1.3 we have that for any constant  $\lambda > 0$  and any particular  $j$ ,

$$(4.5) \quad d_k^{(j)} = 1 + \psi_j^{(k)}, \quad |\psi_j^{(k)}| \leq \tilde{\gamma}_m(\lambda),$$

holds for any particular  $k$  and  $j$  with probability at least  $p(\lambda) = 1 - 2 \exp(-\lambda^2/2)$ , or equivalently it fails to hold with probability at most  $1 - p(\lambda)$ . By the inclusion-exclusion principle [21, p. 39] (4.5) fails to hold for at least one of the pairs  $(j, k)$  with probability at most  $rm(1 - p(\lambda))$ , which means that it holds for all  $j$  and  $k$  with probability at least  $p_2(\lambda, m, r) = 1 - rm(1 - p(\lambda)) = 1 - 2rm \exp(-\lambda^2/2)$ .

Hence  $\|I - D_j\|_2 \leq \tilde{\gamma}_m(\lambda)$  holds for all  $j$  with probability at least  $p_2(\lambda, m, r)$  and hence

$$\|\Delta P_j\|_2 \leq 3u + 2(1+u)\tilde{\gamma}_m(\lambda) + 2u(1+u)(1 + \tilde{\gamma}_m(\lambda)) \leq c_5 \tilde{\gamma}_m(\lambda)$$

holds for all  $j$  with probability at least  $p_2(\lambda, m, r)$ .  $\square$

LEMMA 4.2. *With the same notation as in Lemma 4.1, assume that the bound (4.2) holds almost surely for all  $j$ . Then the computed  $\hat{a}_{r+1}$  satisfies*

$$(4.6) \quad \hat{a}_{r+1} = Q^*(a + \Delta a), \quad \|\Delta a\|_2 \leq c_6 \lambda \sqrt{r} \tilde{\gamma}_m(\lambda) \|a\|_2 + O(u^2)$$

with probability at least  $p_3(\lambda, m) = 1 - 2m \exp(-\lambda^2)$ .

*Proof.* We can rewrite (4.1) as

$$(4.7) \quad \begin{aligned} \hat{a}_{r+1} &= a_{r+1} + \left( \sum_{j=1}^r P_r \dots P_{j+1} \Delta P_j P_{j-1} \dots P_1 + O(u^2) \right) a \\ &= a_{r+1} + Q^* \left( \sum_{j=1}^r F_j + O(u^2) \right) a, \end{aligned}$$

where

$$(4.8) \quad F_j = P_1 \dots P_j \Delta P_j P_{j-1} \dots P_1.$$

Our aim is to show that  $\mathbb{E}(F_j \mid F_{j-1}, \dots, F_1) = 0$ , so that we can apply Theorem 3.2 with  $X_k = F_k$ . When we substitute the expressions for  $F_j$  and then  $\Delta P_j$  into this expression and use the linearity of the expectation we obtain a sum of expectations, each of which we need to show is zero. We will just consider two of the terms, which come from (\*) in (4.4), as the other terms are treated similarly.

The matrix  $D_j$  contains the rounding errors from the computation of  $v_j^* a_j$ , whereas all the other terms in  $F_j$  contain rounding errors from earlier computations.

For the first part of the term (\*) in (4.4) we have

$$(4.9) \quad \begin{aligned} \mathbb{E}(P_1 \dots P_j v_j v_j^* (I - D_j) P_{j-1} \dots P_1 \mid F_{j-1}, \dots, F_1) \\ = P_1 \dots P_j v_j v_j^* \mathbb{E}(I - D_j \mid F_{j-1}, \dots, F_1) P_{j-1} \dots P_1, \end{aligned}$$

since the  $v_j$  and  $P_j$  are constant. We need the general form of the law of total expectation (LTE),  $\mathbb{E}(\mathcal{X} \mid \mathcal{Y}) = \mathbb{E}(\mathbb{E}(\mathcal{X} \mid \mathcal{Z}) \mid \mathcal{Y})$  where " $\mathcal{Y} \subseteq \mathcal{Z}$ " [2, Thm. 34.4]. We define the sequence

$$(4.10) \quad \mathcal{Z}_s = \{ \delta_{k,i}^{(\ell)} : i, k = 1 : m, \ell = 1 : s \}, F_{j-1}, \dots, F_1.$$

By the LTE

$$\mathbb{E}(D_j \mid F_{j-1}, \dots, F_1) = \mathbb{E}(\mathbb{E}(D_j \mid \mathcal{Z}_{j-1}) \mid F_{j-1}, \dots, F_1) = I,$$

by [7, Lem. 6.4], in view of Model 1.4. Hence the expectation (4.9) is zero.

For the second part of the term (\*) in (4.4) we have

$$\begin{aligned} \mathbb{E}(P_1 \dots P_j \Lambda_2^{(k)} v_j v_j^* (I - D_j) P_{j-1} \dots P_1 \mid F_{j-1}, \dots, F_1) \\ = P_1 \dots P_j \mathbb{E}(\Lambda_2^{(j)} v_j v_j^* (I - D_j) \mid F_{j-1}, \dots, F_1) P_{j-1} \dots P_1. \end{aligned}$$

By the LTE,

$$\begin{aligned} \mathbb{E}(\Lambda_2^{(j)} v_j v_j^* (I - D_j) \mid F_{j-1}, \dots, F_1) &= \mathbb{E}(\mathbb{E}(\Lambda_2^{(j)} v_j v_j^* (I - D_j) \mid \mathcal{Z}_j) \mid F_{j-1}, \dots, F_1) \\ &= \mathbb{E}(\mathbb{E}(\Lambda_2^{(j)} \mid \mathcal{Z}_j) v_j v_j^* (I - D_j) \mid F_{j-1}, \dots, F_1) \\ &= 0, \end{aligned}$$



since  $\mathbb{E}(A_2^{(j)} \mid \mathcal{Z}_j) = 0$  by Model 1.4.

By (4.2), (4.8), and the assumption of the lemma,

$$\|F_j\|_2^2 = \|\Delta P_j\|_2^2 \leq c_5^2 \tilde{\gamma}_m(\lambda)^2, \quad j = 1 : r,$$

almost surely. Hence we can take  $X_j = F_j$ ,  $A_j = c_5 \tilde{\gamma}_m(\lambda) I_{d_1, d_2}$ , and  $\sigma^2 = r c_5^2 \tilde{\gamma}_m(\lambda)^2$  in Theorem 3.2 and set  $E = \sum_{j=1}^r F_j$ , to obtain

$$(4.11) \quad \mathbb{P}(\|E\|_2 \geq 2\sqrt{2} c_5 \lambda \sqrt{r} \tilde{\gamma}_m(\lambda)) \leq 2m \exp(-\lambda^2),$$

or equivalently

$$(4.12) \quad \|E\|_2 \leq c_6 \lambda \sqrt{r} \tilde{\gamma}_m(\lambda),$$

holding with probability at least  $1 - 2m \exp(-\lambda^2)$ .

From (4.7) we have

$$\hat{a}_{r+1} = Q^* a + Q^* E a + O(u^2) = Q^* (a + \Delta a)$$

with  $\Delta a = E a + O(u^2)$ , and so

$$(4.13) \quad \|\Delta a\|_2 \leq \|E a\|_2 + O(u^2) \leq c_6 \lambda \sqrt{r} \tilde{\gamma}_m(\lambda) \|a\|_2 + O(u^2). \quad \square$$

Now we consider the backward error in applying a sequence of Householder matrices to a matrix.

LEMMA 4.3. *Consider the sequence of transformations*

$$A_{k+1} = P_k A_k, \quad k = 1 : r.$$

where  $A_1 = A \in \mathbb{R}^{m \times n}$ , each  $P_k \in \mathbb{R}^{m \times m}$  is a Householder matrix. Under Model 1.4 and the assumption of Lemma 4.2, the computed matrix  $\hat{A}_{r+1}$  satisfies

$$\hat{A}_{r+1} = Q^* (A + \Delta A),$$

where  $Q^* = P_r P_{r-1} \dots P_1$  and where

$$(4.14) \quad \|\Delta a_j\|_2 \leq c_6 \lambda \sqrt{r} \tilde{\gamma}_m(\lambda) \|a_j\|_2 + O(u^2), \quad j = 1 : n,$$

holds with probability at least  $p_4(\lambda, m, n) = 1 - 2mn \exp(-\lambda^2)$ .

*Proof.* Lemma 4.2 shows that for each  $j$  the bound (4.6) holds with  $a = a_j$  with probability at least  $p_3(\lambda, m)$ . The probability that it holds for all columns is given by 1 minus the probability that it fails for at least one, which is at least  $1 - n(1 - p_3(\lambda, m)) = 1 - 2mn \exp(-\lambda^2)$ .  $\square$

The probabilistic result for Householder QR factorization, analogous to the worst-case rounding error result Theorem 1.1, now follows.

THEOREM 4.4. *Let  $\hat{R} \in \mathbb{R}^{m \times n}$  be the computed upper trapezoidal QR factor of  $A \in \mathbb{R}^{m \times n}$  obtained via the Householder QR algorithm. Under Model 1.4 and the assumption of Lemma 4.2, there exists an orthogonal  $Q \in \mathbb{R}^{m \times m}$  such that*

$$(4.15) \quad A + \Delta A = Q \hat{R},$$

where

$$(4.16) \quad \|\Delta a_j\|_2 \leq c_6 \lambda \sqrt{n} \tilde{\gamma}_m(\lambda) \|a_j\|_2 + O(u^2), \quad j = 1 : n,$$

holds with probability at least  $p_4(\lambda, m, n) = 1 - 2mn \exp(-\lambda^2)$ .

TABLE 5.1  
Value of  $p_5(\lambda, m, n)$  for various choices of  $\lambda$  and  $m$ , with  $m = n$ .

$\lambda$	$m$			
	$10^2$	$10^4$	$10^6$	$10^8$
6.0	9.9970e-01	-2.0460e+00	-3.0459e+04	-3.0460e+08
7.0	1.0000e+00	9.9542e-01	-4.4795e+01	-4.5795e+05
8.0	1.0000e+00	1.0000e+00	9.7467e-01	-2.5228e+02
9.0	1.0000e+00	1.0000e+00	9.9999e-01	9.4846e-01
10.0	1.0000e+00	1.0000e+00	1.0000e+00	1.0000e+00
11.0	1.0000e+00	1.0000e+00	1.0000e+00	1.0000e+00
12.0	1.0000e+00	1.0000e+00	1.0000e+00	1.0000e+00

*Proof.* The theorem follows from Lemma 4.3 as long as we note two subtleties. First, the Householder matrices in the lemma are completely general, yet for QR factorization they are chosen to introduce zeros into vectors and we explicitly set those elements to zero rather than compute them. This essentially forces rows of  $\Delta P_i$  in (4.4) to be zero, so  $\|\Delta P_i\|_2$  is no larger than the analysis assumes.

The second subtlety is that for Householder QR factorization the Householder vector  $v_j$  in Lemma 4.1 depends on previous computed quantities and is computed itself, so is subject to rounding error. The fact that  $v_j$  is no longer a constant as regards the conditional probabilities can be dealt with by adding “ $v_j, v_{j-1}, \dots, v_1$ ” to  $\mathcal{Z}_s$  in (4.10). The key point is that  $D_j$ ,  $\Lambda_1^{(j)}$ , and  $\Lambda_2^{(j)}$  depend on rounding errors that occur later than those on which  $v_j$  depends. The error in  $v_j$  introduces an  $O(u^2)$  perturbation in (4.4), so it does not affect the bound (4.16).  $\square$

Ideally, we would remove the assumption in Lemmas 4.2 and 4.3 and Theorem 4.4 that the bound (4.2) holds almost surely for all  $j$ . We have effectively replaced a probability of  $1 - 2rm \exp(-\lambda^2/2)$  with a probability of 1. To remove this assumption we need a version of Theorem 3.2 that allows  $X_k X_k^* \leq A_k A_k^*$ , and  $X_k^* X_k \leq A_k^* A_k$  to hold with a certain probability less than 1 rather than almost surely. Such a result is not, to our knowledge, available in the literature on concentration inequalities, and deriving one requires further research that is beyond the scope of this paper.

The columnwise bound (4.16) implies a normwise one:

$$(4.17) \quad \|\Delta A\|_F \leq c_6 \lambda \sqrt{n} \tilde{\gamma}_m(\lambda) \|A\|_F + O(u^2).$$

The equivalent bound for the worst-case analysis is (1.3).

The conclusion from our analysis is that the constant  $mnu$  in the worst-case bound reduces to  $\sqrt{m}nu$  in the probabilistic bound.

**5. Discussion.** We now discuss several aspects and implications of the error analysis.

**5.1. Choice of  $\lambda$ .** We begin with a brief discussion of the probabilities associated with the bounds in Lemma 4.1 and Theorem 4.4. As noted in previous work [7], [15], [16], these probabilities are typically pessimistic and setting  $\lambda = 1$  almost always provides bounds that hold in practice. Table 5.1 shows values of  $p_5(\lambda, m, n) = 1 - 2mn(\exp(-\lambda^2) + \exp(-\lambda^2/2))$ , which is a lower bound on the probabilities in Lemma 4.1 and Theorem 4.4. For certain values the probabilities are negative. However, for  $\lambda = 12$  and all the problem sizes shown in Table 5.1,  $p_5$  is extremely close to 1—in fact, within  $10^{-7}$  of it.

**5.2. Aggregated Householder transformations.** In practice, Householder transformations are usually aggregated in order to express the computation primarily in terms of matrix multiplication. A common form of aggregation is the  $WY$  representation [3], [12, sec. 19.5]. We represent the product  $Q_r = P_r P_{r-1} \dots P_1$  of  $b$  Householder matrices  $P_i = I - v_i v_i^*$  as

$$Q_r = I + W_b Y_b^*, \quad W_b, Y_b \in \mathbb{R}^{m \times b}.$$

This is done through the recurrence

$$W_1 = -v_1, \quad Y_1 = v_1, \quad W_i = [W_{i-1} \quad -v_i], \quad Y_i = [Y_{i-1} \quad Q_{i-1}^* v_i].$$

We partition  $A$  as

$$A = [A_1 \quad B], \quad A_1 \in \mathbb{R}^{m \times r}$$

and transform  $A_1$  to upper trapezoidal form by forming  $P_r \dots P_1 A_1$ , accumulating the product  $P_r \dots P_1 = I + W_r Y_r^*$ . The matrix  $B$  is then updated via  $B \leftarrow B + W_r (Y_r^* B)$ , and we repeat this process on the remaining rows of  $B$ .

We do not perform a full analysis of the aggregated algorithm. We simply note that the two core operations of the aggregated Householder QR factorization are the application of a sequence of Householder matrices, in forming  $P_r \dots P_1 A_1$ , and matrix multiplication in forming  $B \leftarrow B + W_r (Y_r^* B)$ . Both of these operations have been shown to have probabilistic bounds whose constants are the square roots of those in the worst-case bounds, so we can expect the same to be true for the aggregated algorithm.

**5.3. Mixed-precision QR factorization.** Yang, Fox, and Sanders [22, Thm. 4.1] consider a mixed precision Householder QR factorization algorithm in which the working precision is  $u_{\text{low}}$  and the inner products are computed at precision  $u_{\text{high}}$ . They obtain a normwise backward error bound of order  $n(u_{\text{low}} + m u_{\text{high}})$ . Our analysis is readily adapted for this algorithm by replacing  $\tilde{\gamma}_m = \tilde{\gamma}_m(u)$  in (4.2) by  $\tilde{\gamma}_m(u_{\text{high}}) + O(u_{\text{low}})$ , and the resulting probabilistic error bound is of order  $n^{1/2}(u_{\text{low}} + m^{1/2} u_{\text{high}})$ .

**5.4. Two-sided transformations.** Householder matrices and Householder QR factorization are tools in many algorithms, including the reduction of matrices to tridiagonal or Hessenberg form for the QR algorithm for eigenvalues and to bidiagonal form for the QR algorithm for singular values [11]. Can we use our results to obtain probabilistic backward error bounds with reduced constants for these reductions? For the reduction of  $A \in \mathbb{R}^{n \times n}$  to Hessenberg form  $H$  we have

$$H = P_{n-2} \dots P_1 A P_1^* \dots P_{n-2}^*.$$

The reduction is carried out in the order  $P_{n-2}(\dots(P_1 A P_1^*)\dots)P_{n-2}^*$ , because  $P_{j+1}$  depends on  $P_j \dots P_1 A P_1^* \dots P_j^*$ . However, for the purposes of the analysis we can consider applying all the transformations on the left before applying those on the right, since the order should not affect the error bound:

$$(5.1) \quad H = (P_{n-2} \dots P_1 A) P_1^* \dots P_{n-2}^* \equiv (PA) P^*.$$

Two applications of Lemma 4.3 give

$$\begin{aligned} \hat{H} &= (P(A + \Delta A_1) + \Delta A_2) P^* \\ &= P(A + \Delta A) P^*, \quad \Delta A = \Delta A_1 + P^* \Delta A_2. \end{aligned}$$

The probabilistic bounds on  $\|\Delta A_1\|_F$  and  $\|\Delta A_2\|_F$  from the lemma are proportional to  $nu$  and hence so is the bound for  $\|\Delta A\|_F$ . Lemma 4.3 does not incorporate the errors in forming the Householder vectors  $v_j$ , but as in the proof of Theorem 4.4 doing so only affects the second order term. The usual worst-case bound obtained by Wilkinson is proportional to  $n^2u$  [20, pp. 160–161] (we have accounted for the fact that Wilkinson assumes that inner products are accumulated at twice the working precision).

For the multishift Hessenberg QR iteration the bulge chasing is implemented using Householder matrices [5], [6], so Lemma 4.3 can be applied again for sufficiently large bulges, and for small bulges one can apply the worst-case bounds.

**5.5. Other forms of QR factorization.** The modified Gram–Schmidt algorithm applied to  $A \in \mathbb{R}^{m \times n}$  with  $m \geq n$  is known to be equivalent both mathematically and numerically to Householder QR factorization applied to the augmented matrix  $\begin{bmatrix} 0 \\ A \end{bmatrix}$  [4], [12, sec. 19.8]. Hence we can apply Theorem 4.4, and since the augmented matrix is  $(m+n) \times n$  the constant obtained will be proportional to  $\sqrt{(m+n)nu} \approx \sqrt{mnu}$  since  $m \geq n$ . (As explained in the above references, some additional work is needed to obtain a backward error result for modified Gram–Schmidt and the constant will be increased.)

In [1], a randomized process for computing the QR factorization of  $A \in \mathbb{R}^{m \times n}$  is presented. It uses a randomized Gram–Schmidt process, with the approach based on the dimension reduction technique of random sketching. The authors assume that the rounding errors are independent, zero-mean random variables. They also exploit mixed-precision arithmetic with two precisions  $u_{\text{fine}} < u_{\text{crs}}$ . The analogous result to Theorem 4.4 is [1, Thm. 3.2]  $\|A - \hat{Q}\hat{R}\|_F \leq cu_{\text{crs}}n^{3/2}\|A\|_F$ .

Finally, we consider Givens QR factorization. Givens transformations operate on vectors of length 2, so there is no benefit to a probabilistic approach in analyzing the application of a single rotation. The backward error analysis [12, sec. 19.6] shows that each individual rotation introduces a backward error bounded by a small constant and that the  $m+n-2$  products of disjoint rotations needed for a QR factorization lead to a result of the same form as Theorem 1.1 but with

$$\|\Delta a_j\|_2 \leq \gamma_{c(m+n-2)}\|a_j\|_2, \quad j = 1:n.$$

A probabilistic analysis analogous to that in the proof of Lemma 4.1, using the matrix concentration inequality, leads to a result of the same form as Theorem 4.4 but with probabilistic bound

$$\|\Delta a_j\|_2 \leq c_7\sqrt{m+n}\|a_j\|_2 + O(u^2), \quad j = 1:n.$$

**6. Backward error for QR factorization.** In our numerical experiments in the next section we need to compare the probabilistic backward error bounds with the actual backward errors. How to compute the backward error matrix  $\Delta A = A - Q\hat{R}$  for a given  $\hat{R}$ , though, is not clear, since the orthogonal matrix  $Q$  in Theorems 1.1 and 4.4 is unknown. We will focus on the backward error measure

$$(6.1) \quad \mu(\hat{R}) = \min \left\{ \left( \sum_{j=1}^n \|d_j \Delta a_j\|_2^2 \right)^{1/2} : A + \Delta A = Q\hat{R}, Q \in \mathbb{R}^{m \times m}, Q^*Q = I_m \right\} \\ = \min(\|(A - Q\hat{R})D\|_F : Q \in \mathbb{R}^{m \times m}, Q^*Q = I_m, D = \text{diag}(d_j))$$

For  $D = \text{diag}(\|a_j\|_2^{-1})$  we have a columnwise backward error and for  $D = \|A\|_F^{-1}I$  the normwise backward error.

The next result shows how to compute  $\mu(\hat{R})$ . Recall that the polar decomposition of  $A \in \mathbb{R}^{n \times n}$  is a factorization  $A = UH$ , where  $U$  is orthogonal and  $H$  is symmetric positive semidefinite.

**THEOREM 6.1.** *Let  $A, B \in \mathbb{R}^{m \times n}$  and let  $D = \text{diag}(d_i) \in \mathbb{R}^{n \times n}$  be nonsingular. Then  $\min\{\|(A - QB)D\|_F : Q \in \mathbb{R}^{m \times m}, Q^*Q = I_m\}$  is obtained for  $Q = U^*$ , where  $U$  is the orthogonal polar factor of the matrix  $BD^2A^*$ .*

*Proof.* For  $F, G \in \mathbb{R}^{m \times n}$ , the orthogonal Procrustes problem has the form  $\min\{\|F - GW\|_F : W \in \mathbb{R}^{n \times n}, W^*W = I_n\}$  and any orthogonal polar factor of  $G^*F$  is a solution [13, Thm. 8.6]. Writing  $\|(A - QB)D\|_F = \|AD - QBD\|_F = \|DA^* - DB^*Q^*\|_F$  therefore gives the result on taking  $F = DA^*$  and  $G = DB^*$ .  $\square$

By Theorem 6.1,  $\mu(\hat{R}) = \|(A - Q\hat{R})D\|_F$ , where  $Q^*$  is an orthogonal polar factor of  $\hat{R}D^2A^*$ . If  $\hat{R}D^2A^* = U\Sigma V^*$  is a singular value decomposition then we can take  $Q^* = UV^*$ .

We note that one can express  $\mu(\hat{R})^2 = \|AD\|_F^2 + \|\hat{R}D\|_F^2 - 2\sum_{i=1}^n \sigma_i(\hat{R}D^2A^*)$ , where  $\sigma_i$  denotes the  $i$ th largest singular value. However, this formula suffers from severe cancellation, and rounding errors can cause it to evaluate as negative, so it is better to use the expression  $\|(A - Q\hat{R})D\|_F$ . In fact, even the latter expression does not necessarily give a result of the correct order of magnitude, so it is best to compute the backward error at twice the working precision. Hence in our experiments we take single precision as the working precision and compute the backward error in double precision.

**7. Numerical experiments.** For all our all numerical experiments we set the parameter  $\lambda = 1$  and set all constants  $c_i$  in the error bounds (worst-case or probabilistic) to be 1. We use MATLAB R2021b and take IEEE single precision as the working precision. Backward errors (6.1) are computed as described in section 6, in double precision.

We use round to nearest in all the experiments. We have tried stochastic rounding, which ensures that the assumptions of Model 1.4 are satisfied [7], and found that the numerical results reported are virtually identical to those for round to nearest in these experiments.

In sections 7.1, 7.3, and 7.4 we use random  $m \times n$  matrices with entries drawn uniformly from the interval  $[0, 1]$  and we sample 10 different matrices for each pair of dimensions.

**7.1. Householder QR factorization for random matrices.** In this experiment we test the backward error bound in Theorem 4.4. In order to study how the error grows with  $n$  and  $m$  independently, we first fix a value of  $n$  and vary  $m$  and then fix  $m$  and vary  $n$ . For each pair of dimensions we plot the maximum and mean normwise backward errors for Householder QR factorization along with associated worst-case and probabilistic bounds obtained from (1.3) and (4.17). The results are given in Figures 7.1 and 7.2. We see that the probabilistic bound  $\sqrt{mnu}$  proves a much better indicator than the worst-case bound  $mnu$  of the size of the error and its rate of growth.

**7.2. Householder QR factorization for SuiteSparse matrices.** We also consider Householder QR factorization of some real life matrices from the SuiteSparse Matrix Collection [9], [10]. We select all matrices from the collection with  $10 \leq m, n \leq 2 \times 10^3$  and  $m \geq n$ . This results in 842 matrices. We plot the same error quantities as for the random matrices. Some of the matrices in the collection have zero columns, so




FIG. 7.1. Normwise backward errors and bounds for Householder QR factorization for  $n = 10$  and various  $m$ , for  $m \times n$  matrices with elements sampled uniformly from  $[0, 1]$ .




FIG. 7.2. Normwise backward errors and bounds for Householder QR factorization for  $m = 10^4$  and various  $n$ , for  $m \times n$  matrices with elements sampled uniformly from  $[0, 1]$ .

we filter these out. There are a few cases where the reported error exceeds even the deterministic bound, which we suspect is an underflow issue; a similar observation is made in [15, sec. 4.5]. We also filter out these cases from the reported results, which results in 774 matrices. We show the observed backward errors in Figure 7.3. Again, the probabilistic bound is satisfied and proves closer to the observed error than the worst-case bound by several orders of magnitude.

**7.3. Givens rotations.** Givens QR factorization is typically used for structured matrices, such as tridiagonal or upper Hessenberg matrices, but here we wish to see how the backward error of the factorization behaves for dense matrices. For random  $n \times n$  matrices we plot in Figure 7.4 the maximum and mean normwise backward errors, the worst-case error bound  $2nu$ , and the probabilistic error bound  $\sqrt{2nu}$ . The backward error grows at a rate very similar to that of the probabilistic bound.




FIG. 7.3. Normwise backward errors and bounds for 774  $m \times n$  matrices from the SuiteSparse collection.




FIG. 7.4. Normwise backward error and backward error bound for Givens QR factorization for  $n \times n$  matrices with elements sampled uniformly from  $[0, 1]$ .

**7.4. Reduction to Hessenberg form.** Finally, we consider Householder reduction matrix to Hessenberg form:  $A = QHQ^*$ . Figure 7.5 shows normwise backward errors, computed as  $\|A - \hat{Q}H\hat{Q}^*\|_F / \|A\|_F$ , where  $\hat{Q}$  is the computed product of Householder matrices (we do not have an explicit backward error formula such as that in Theorem 6.1 in this two-sided case), and the worst-case and probabilistic bounds,  $n^2u$  and  $nu$  respectively. We see that the backward error satisfies the probabilistic bound, and again the probabilistic bound is a better indicator than the worst-case bound of the rate of growth of the backward error with  $n$ .

**8. Conclusions.** In a classic 1961 paper, Wilkinson [19, p. 318] carries out rounding error analyses of LU factorization, Givens QR factorization, and Householder QR factorization. He notes that “The bounds we have obtained are in all cases strict upper bounds. In general, the statistical distribution of the rounding errors




FIG. 7.5. Normwise backward errors and bounds for reduction to Hessenberg form for  $n \times n$  matrices with elements sampled uniformly from  $[0, 1]$ .

will reduce considerably the function of  $n$  occurring in the relative errors. We might expect in each case that this function should be replaced by something which is no bigger than its square root and is usually appreciably smaller.” Recent probabilistic rounding error analysis has provided a rigorous foundation for Wilkinson’s statement for LU factorization and other inner product-based computations. Our work does the same for Householder QR factorization-based methods, as well as for Givens QR factorization.

A significant feature of the probabilistic backward error bounds is that they bound the likely rate of growth of the backward error as the problem dimensions increase. The rate of growth, along with blocking, exploiting architectural features of the hardware, and using other techniques to improve the accuracy of the computations, is what determines our ability to solve problems at extreme scale and possibly low precision in a numerically stable way [14].

**Acknowledgments.** We thank Joel Tropp, Neil Walton, and Marcus Webb for advice on concentration inequalities, and Theo Mary for comments on a draft manuscript.

## REFERENCES

- [1] Oleg Balabanov and Laura Grigori. **Randomized Gram-Schmidt process with application to GMRES**. ArXiv:2011.05090, November 2021. Revised January 2022.
- [2] Patrick Billingsley. *Probability and Measure*. Third edition, Wiley, New York, USA, 1995. ISBN 0-471-00710-2.
- [3] Christian H. Bischof and Charles F. Van Loan. The WY representation for products of Householder matrices. *SIAM J. Sci. Statist. Comput.*, 8(1):s2–s13, 1987.
- [4] Åke Björck and C. C. Paige. **Loss and recapture of orthogonality in the modified Gram-Schmidt algorithm**. *SIAM J. Matrix Anal. Appl.*, 13(1):176–190, 1992.
- [5] Karen Braman, Ralph Byers, and Roy Mathias. **The multishift QR algorithm. Part I: Maintaining well-focused shifts and level 3 performance**. *SIAM J. Matrix Anal. Appl.*, 23(4): 929–947, 2002.
- [6] Karen Braman, Ralph Byers, and Roy Mathias. **The multishift QR algorithm. Part II: Aggressive early deflation**. *SIAM J. Matrix Anal. Appl.*, 23(4):948–973, 2002.
- [7] Michael P. Connolly, Nicholas J. Higham, and Theo Mary. **Stochastic rounding and its proba-**



- bilistic backward error analysis. *SIAM J. Sci. Comput.*, 43(1):A566–A585, 2021.
- [8] Matteo Croci, Massimiliano Fasi, Nicholas J. Higham, Theo Mary, and Mantas Mikaitis. **Stochastic rounding: Implementation, error analysis, and applications**. MIMS EPrint 2021.17, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, October 2021. 28 pp. Revised January 2022. To appear in Roy. Soc. Open Sci.
  - [9] Timothy A. Davis. SuiteSparse Matrix Collection. <https://sparse.tamu.edu/>.
  - [10] Timothy A. Davis and Yifan Hu. **The University of Florida Sparse Matrix Collection**. *ACM Trans. Math. Software*, 38(1):1:1–1:25, 2011.
  - [11] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Fourth edition, Johns Hopkins University Press, Baltimore, MD, USA, 2013. xxi+756 pp. ISBN 978-1-4214-0794-4.
  - [12] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Second edition, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2002. xxx+680 pp. ISBN 0-89871-521-0.
  - [13] Nicholas J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2008. xx+425 pp. ISBN 978-0-898716-46-7.
  - [14] Nicholas J. Higham. **Numerical stability of algorithms at extreme scale and low precisions**. MIMS EPrint 2021.14, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, September 2021. 21 pp. To appear in Proc. Int. Cong. Math.
  - [15] Nicholas J. Higham and Theo Mary. **A new approach to probabilistic rounding error analysis**. *SIAM J. Sci. Comput.*, 41(5):A2815–A2835, 2019.
  - [16] Nicholas J. Higham and Theo Mary. **Sharper probabilistic backward error analysis for basic linear algebra kernels with random data**. *SIAM J. Sci. Comput.*, 42(5):A3427–A3446, 2020.
  - [17] Michael Mitzenmacher and Eli Upfal. *Probability and Computing. Randomization and Probabilistic Techniques in Algorithms and Data Analysis*. Cambridge University Press, Cambridge, UK, 2017. xx+467 pp. ISBN 978-1-107-15488-9.
  - [18] Joel A. Tropp. **User-friendly tail bounds for sums of random matrices**. *Found. Comput. Math.*, 12(4):389–434, 2012.
  - [19] J. H. Wilkinson. **Error analysis of direct methods of matrix inversion**. *J. ACM*, 8:281–330, 1961.
  - [20] J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Oxford University Press, Oxford, UK, 1965. xviii+662 pp. ISBN 0-19-853403-5 (hardback), 0-19-853418-3 (paperback).
  - [21] David Williams. *Weighing the Odds. A Course in Probability and Statistics*. Cambridge University Press, Cambridge, UK, 2001. xvii+547 pp. ISBN 0-521-00618-X.
  - [22] L. Minah Yang, Alyson Fox, and Geoffrey Sanders. **Rounding error analysis of mixed precision block Householder QR algorithms**. *SIAM J. Sci. Comput.*, 43(3):A1723–A1753, 2021.