Accurate Computation of the Log-Sum-Exp and Softmax Functions

Blanchard, Pierre and Higham,
Desmond J. and Higham, Nicholas J.

2019

MIMS EPrint: **2019.16**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

ISSN 1749-9097

ACCURATE COMPUTATION OF THE LOG-SUM-EXP AND SOFTMAX FUNCTIONS*

PIERRE BLANCHARD†, DESMOND J. HIGHAM‡, AND NICHOLAS J. HIGHAM§

Abstract. Evaluating the log-sum-exp function or the softmax function is a key step in many modern data science algorithms, notably in inference and classification. Because of the exponentials that these functions contain, the evaluation is prone to overflow and underflow, especially in low precision arithmetic. Software implementations commonly use alternative formulas that avoid overflow and reduce the chance of harmful underflow, employing a shift or another rewriting. Although mathematically equivalent, these variants behave differently in floating-point arithmetic. We give rounding error analyses of different evaluation algorithms and interpret the error bounds using condition numbers for the functions. We conclude, based on the analysis and numerical experiments, that the shifted formulas are of similar accuracy to the unshifted ones and that the shifted softmax formula is typically more accurate than a division-free variant.

Key words. log-sum-exp, softmax, floating-point arithmetic, rounding error analysis, overflow, underflow, condition number

AMS subject classifications. 65G50

1. Introduction.

In many applications, especially in a wide range of machine learning classifiers such as multinomial linear regression and naive Bayes classifiers [4], [23], [26], one needs to compute an expression of the form

$$y = f(x) = \log \sum_{i=1}^{n} \mathrm{e}^{x_i}, \tag{1.1}$$

where $x = [x_1, x_2, \ldots, x_n]^T \in \mathbb{C}^n$. The function $f : \mathbb{C}^n \to \mathbb{C}$ is often referred to as log-sum-exp or LSE. Its gradient $g : \mathbb{C}^n \to \mathbb{C}^n$, given by

$$g_j(x) = \frac{\partial}{\partial x_j} f(x) = \frac{\mathrm{e}^{x_j}}{\sum_{i=1}^{n} \mathrm{e}^{x_i}}, \quad j = 1\colon n, \tag{1.2}$$

is called softmax and is also a key function in classification algorithms [7, p. 355], [8, p. 78], [10]. It is often the case that both log-sum-exp and softmax are required simultaneously.

The most obvious danger in evaluating (1.1) and (1.2) is overflow. We are interested in IEEE arithmetic in the precisions half (fp16), single (fp32), and double (fp64) [16], as well as the bfloat16 half precision format [17]. Table 1.1 shows the key parameters of interest for these precisions: the unit roundoff u, the largest finite number $r_{\max}$, and the smallest positive normalized and subnormal floating-point numbers. If some x_i exceeds the relevant $\log r_{\max}$ value in Table 1.2 then overflow will occur.

*Version of September 8, 2019. **Funding**: This work was supported by Engineering and Physical Sciences Research Council grant EP/P020720/1, The MathWorks, and the Royal Society. The opinions and views expressed in this publication are those of the authors, and not necessarily those of the funding bodies.

†Department of Mathematics, University of Manchester, Manchester, M13 9PL, UK (`pierre.blanchard00@gmail.com`).

‡School of Mathematics, University of Edinburgh, Edinburgh, EH9 3FD, UK (`d.j.higham@ed.ac.uk`).

§Department of Mathematics, University of Manchester, Manchester, M13 9PL, UK (`nick.higham@manchester.ac.uk`).

Clearly, overflow is possible even for quite modestly sized x, especially for half and single precision.

Underflow is also possible. For example, for $n = 1$, if x_1 is a finite floating-point number with $x_1 < \log r_{\min}^{(s)}$ then[1] $\mathrm{fl}(f(x_1)) = \mathrm{fl}(\log(\mathrm{fl}(\mathrm{e}^{x_1}))) = \mathrm{fl}(\log 0) = -\infty$, whereas $f(x_1) = x_1$. For $n > 1$, underflow in the exponential evaluations is a problem when the sum of the terms that underflow is significant compared with the sum of the other terms; otherwise underflows are harmless. As well as avoiding harmful underflow, it is desirable to avoid generating subnormal numbers, which incur a performance penalty if handled in software[2]; see [12] or [22] for details of subnormal numbers.

A way to avoid overflow, and to attempt to avoid underflow and subnormal numbers, in evaluating log-sum-exp is to rewrite

$$y = \log \sum_{i=1}^{n} \mathrm{e}^{x_i} = \log \sum_{i=1}^{n} \mathrm{e}^{a} \mathrm{e}^{x_i - a} = \log \left(\mathrm{e}^{a} \sum_{i=1}^{n} \mathrm{e}^{x_i - a} \right).$$

If $a \in \mathbb{R}$ then

$$y = a + \log \sum_{i=1}^{n} \mathrm{e}^{x_i - a}. \tag{1.3}$$

Here and throughout, log denotes the principal logarithm: the logarithm whose imaginary part lies in $(-\pi, \pi]$. Equation (1.3) is not, in general, true for $a \in \mathbb{C}$ [1, Lem. 2.5]. The softmax can be expressed in a related form (for any a):

$$g_j = \frac{\mathrm{e}^{x_j - a}}{\sum_{i=1}^{n} \mathrm{e}^{x_i - a}}, \quad j = 1\colon n. \tag{1.4}$$

This shifting, typically with $a = \max_i x_i$, is a well known way to attempt to avoid overflow and underflow in the evaluation of f and g, described in many places, including on Wikipedia[3], in blog posts[4], and even in a YouTube video[5]. The functions `logsumexp` in SciPy 1.3.1 [18] and `LogSumExp` in R [25] both implement (1.3) with $a = \max_i x_i$. The function `softmax` in the MATLAB Deep Learning Toolbox (R2019a) [6] uses (1.4) with $a = \max_i x_i$.

An alternative to (1.4), which removes the denominator of (1.2) by subtracting log-sum-exp from the argument of exp in the numerator, is

$$g_j = \exp \left(x_j - \log \sum_{i=1}^{n} \mathrm{e}^{x_i} \right). \tag{1.5}$$

The conciseness of this division-free formula makes it attractive for implementing softmax when a log-sum-exp function is available. This formula is used in the SciPy 1.3.1 function `softmax`, in a MATLAB toolbox [20] associated with the book [2],

[1] $\log 0 = -\infty$ is the value recommended by the IEEE standard [16, p. 43].

[2] https://devblogs.nvidia.com/cuda-pro-tip-flush-denormals-confidence/, https://en.wikipedia.org/wiki/Denormal_number.

[3] https://en.wikipedia.org/wiki/LogSumExp

[4] For example, https://hips.seas.harvard.edu/blog/2013/01/09/computing-log-sum-exp/, http://bayesjumping.net/log-sum-exp-trick/, and https://jblevins.org/log/log-sum-exp. And similarly for the softmax: https://timvieira.github.io/blog/post/2014/02/11/exp-normalize-trick/.

[5] https://youtu.be/-RVM21Voo7Q

Table 1.1

Parameters for bfloat16 and IEEE fp16, fp32, and fp64 arithmetics, to three significant figures: unit roundoff u, smallest positive (subnormal) number $r_{\min}^{(s)}$, smallest positive normalized number $r_{\min}$, and largest finite number $r_{\max}$. In Intel's bfloat16 specification, subnormal numbers are not supported, so $r_{\min}^{(s)} = r_{\min}$ [17].

	u	$r_{\min}^{(s)}$	$r_{\min}$	$r_{\max}$
bfloat16	3.91×10^{-3}	9.18×10^{-41}	1.18×10^{-38}	3.39×10^{38}
fp16	4.88×10^{-4}	5.96×10^{-8}	6.10×10^{-5}	6.55×10^{4}
fp32	5.96×10^{-8}	1.40×10^{-45}	1.18×10^{-38}	3.40×10^{38}
fp64	1.11×10^{-16}	4.94×10^{-324}	2.22×10^{-308}	1.80×10^{308}

Table 1.2

Logarithms of key parameters in Table 1.1, to three significant figures.

	$\log r_{\min}^{(s)}$	$\log r_{\min}$	$\log r_{\max}$
bfloat16	-92.2	-87.3	88.7
fp16	-16.6	-9.70	11.0
fp32	-103	-87.3	88.7
fp64	-744	-708	710

and in the internal function `softmax` in the MATLAB Statistics and Machine Learning Toolbox (R2019a) [24]; in each case the log-sum-exp term is computed by (1.3) with $a = \max_i x_i$. The formula (1.5) can also be found in codes posted in online communities such as Stack Exchange.

The accuracy properties of the formulas above are not clear. In particular, when $a = x_{\max} < 0$, y in (1.3) is computed as a sum of two terms of opposite sign, so there could potentially be damaging subtractive cancellation.

In this work we analyze the unshifted and shifted formulas and (1.5) in order to determine which choices of formulas give the best accuracy and reliability. In particular, we carry out a rounding error analysis of algorithms for the evaluation and relate the error bounds to the conditioning of f and g. We show that the shifted formulas have broadly similar error bounds to the unshifted ones, and so are entirely appropriate for practical use. We find, however, that the alternative softmax formula (1.5) has a less favorable error bound than the shifted formula and tends to produce larger errors in practice.

We begin, in the next section, by investigating the conditioning of the log-sum-exp and softmax functions. In section 3 we give detailed rounding error analyses of the basic formulas. In section 4 we analyze the shifted formulas and (1.5) and compare their error bounds with those for unshifted formulas. Numerical experiments are given in section 5 to test the accuracy of the evaluations and also to examine how the sum of the computed softmax vector entries compares with the exact value 1. Conclusions are given in section 6.

From this point on, we assume that the x_i are real and we write

$$x_{\max} = \max_i x_i, \qquad x_{\min} = \min_i x_i. \tag{1.6}$$

We will use the standard model of floating-point arithmetic [12, sec. 2.2]

$$\mathrm{fl}(a \mathbin{\mathrm{op}} b) = (a \mathbin{\mathrm{op}} b)(1+\delta), \quad |\delta| \le u, \quad \mathrm{op} \in \{+,-,\times,/\}. \tag{1.7}$$

2. Condition number. Before considering algorithms for computing log-sum-exp and softmax we investigate the conditioning of these functions, that is, the sensitivity of $f(x)$ and $g(x)$ in (1.1) and (1.2) to small perturbations in x.

We define the condition number of f in the usual way (see, e.g., [13, chap. 3]), by

$$\operatorname{cond}(f,x) := \lim_{\epsilon\to 0} \sup_{\|e\|\le\epsilon\|x\|} \frac{|f(x+e)-f(x)|}{\epsilon|f(x)|}.$$

This definition implies that

$$\frac{|f(x+e)-f(x)|}{|f(x)|} \le \operatorname{cond}(f,x)\frac{\|e\|}{\|x\|} + o(\|e\|), \tag{2.1}$$

so that $\operatorname{cond}(f,x)$ measures the worst-case relative change in f corresponding to a small relative change in x. It is easy to show that for the ∞-norm,

$$\operatorname{cond}_\infty(f,x) = \frac{\|\nabla f(x)\|_1\|x\|_\infty}{|f(x)|} = \frac{\|x\|_\infty}{|f(x)|} = \frac{\max_i |x_i|}{|\log\sum_i \mathrm{e}^{x_i}|}, \tag{2.2}$$

since $\|\nabla f(x)\|_1 = 1$ by (1.2).

We identify two extreme cases. First, the condition number is infinite for $x_i \equiv -\log n$, because $f(x) = 0$. Hence when $x_i \approx -\log n$ for all i the condition number must be large. Second, if $\max_i x_i = \max_i |x_i|$ then $|f(x)| \ge \max_i |x_i|$ by (4.2) below, so $\operatorname{cond}_\infty(f,x) \le 1$ and the problem is perfectly conditioned.

A forward stable algorithm for computing log-sum-exp is one for which the relative error of the computed result is bounded by $p(n)\operatorname{cond}(f,x)u$, for some low degree polynomial p. Ideally, we would like the algorithm that we use to be forward stable. To see whether it is reasonable to expect forward stability, consider the case $n = 1$. Then $f(x) = \log \mathrm{e}^x = x$, so $\operatorname{cond}(f,x) = 1$: the problem is perfectly conditioned. When we compute f using standard library functions we can expect to obtain relative errors in the computed exponential and logarithm bounded by u [5], [21], [22, Chap. 10], that is,

$$\widehat{y} = \mathrm{fl}(f(x)) = \log(\mathrm{e}^x(1+\delta_1))(1+\delta_2), \quad |\delta_1|,|\delta_2| \le u. \tag{2.3}$$

The term $1+\delta_2$ just causes a small relative perturbation of the output, so we have

$$\widehat{y} \approx \log(\mathrm{e}^x(1+\delta_1)) = x + \log(1+\delta_1) = x + \delta_1 + O(\delta_1^2).$$

Hence, since $y = x$,

$$\frac{|y-\widehat{y}|}{|y|} \lesssim \frac{u}{|x|} + O(u^2). \tag{2.4}$$

This relative error bound is much larger than u for $|x| \ll 1$, even though the problem is perfectly conditioned. So it is not reasonable to expect an algorithm to be unconditionally forward stable in floating-point arithmetic. For this trivial computation, backward error and forward error are the same, so we also conclude that we cannot expect to obtain an algorithm that is unconditionally backward stable.

The softmax function has condition number

$$\operatorname{cond}(g,x) := \lim_{\epsilon\to 0} \sup_{\|e\|\le\epsilon\|x\|} \frac{\|g(x+e)-g(x)\|}{\epsilon\|g(x)\|},$$

which is given explicitly by

$$\operatorname{cond}(g,x) = \frac{\|G(x)\|\,\|x\|}{\|g(x)\|}.$$

Here, the $n \times n$ matrix $G(x) = (\partial g_i/\partial x_j)$ is the Jacobian of g and $\|\cdot\|$ denotes any vector norm and the corresponding subordinate matrix norm. Now

$$\frac{\partial g_i}{\partial x_j} = \begin{cases} \dfrac{-\mathrm{e}^{x_i}\mathrm{e}^{x_j}}{\left(\displaystyle\sum_{k=1}^{n}\mathrm{e}^{x_k}\right)^2}, & i \neq j, \\[2ex] \dfrac{\mathrm{e}^{x_i}\displaystyle\sum_{k=1}^{n}\mathrm{e}^{x_k} - \mathrm{e}^{2x_i}}{\left(\displaystyle\sum_{k=1}^{n}\mathrm{e}^{x_k}\right)^2}, & i = j. \end{cases}$$

We have, for each i,

$$\sum_{j=1}^{n}\left|\frac{\partial g_i}{\partial x_j}\right| = \frac{2\mathrm{e}^{x_i}\displaystyle\sum_{\substack{j=1\\ j\neq i}}^{n}\mathrm{e}^{x_j}}{\left(\displaystyle\sum_{k=1}^{n}\mathrm{e}^{x_k}\right)^2} \le 1,$$

that is, $\|G(x)\|_\infty \le 1$. Hence

$$\operatorname{cond}_\infty(g,x) \le \frac{\|x\|_\infty}{\|g(x)\|_\infty} \le n\|x\|_\infty,$$

because $\|g\|_\infty \ge n^{-1}\|g\|_1 = n^{-1}$. We note in passing that G is the Hessian of f and can be shown to be symmetric positive semidefinite for all x [3, p. 74].

We also note that shifting, as in (1.3) and (1.4), does not change the functions so does not change their condition numbers; likewise for (1.5). These reformulations may, of course, affect the accuracy of the floating-point evaluation.

3. Basic algorithms and error analysis. Algorithm 3.1 gives a naive implementation of (1.1) and (1.2).

ALGORITHM 3.1. *Given $x \in \mathbb{R}^n$, this algorithm computes $f(x) = \log\sum_{i=1}^{n}\mathrm{e}^{x_i}$ and the gradient $g(x) = \nabla f(x)$.*

```
s = 0
for i = 1:n
    w_i = exp(x_i)
    s = s + w_i
end
f = log(s)
for i = 1:n
    g_i = w_i/s
end
```

What can be said about the accuracy of this algorithm when it is implemented in floating-point arithmetic? To answer this question we carry out a rounding error analysis. Throughout this section, we assume that there is no overflow or underflow.

First, we consider the error in evaluating the sum of nonnegative terms

$$s = \sum_{i=1}^{n} \mathrm{e}^{x_i} \equiv \sum_{i=1}^{n} w_i.$$

Evaluating $w_i = \mathrm{e}^{x_i}$ yields a computed result satisfying

$$\widehat{w}_i = \mathrm{e}^{x_i}(1 + \delta_1), \tag{3.1}$$

where, as noted in Section 2, we can expect the relative error from the exponential evaluation to satisfy $|\delta_1| \le u$. Therefore

$$|\widehat{w}_i - w_i| \le w_i u.$$

Write the (exact) sum of computed quantities as

$$\widetilde{s} = \sum_{i=1}^{n} \widehat{w}_i.$$

The rounding error analysis in [11], [12, sec 4.2] shows that the computed sum $\widehat{s}$ satisfies

$$|\widetilde{s} - \widehat{s}| \le u \sum_{i=1}^{n-1} |t_i| + O(u^2),$$

where $t_i = \sum_{j=1}^{i+1} \widehat{w}_j$, so that, since $\widehat{w}_i \ge 0$,

$$|\widetilde{s} - \widehat{s}| \le u(n-1)(\widehat{w}_1 + \widehat{w}_2) + u \sum_{i=3}^{n} (n+1-i)\widehat{w}_i + O(u^2).$$

Writing $s - \widehat{s} = s - \widetilde{s} + \widetilde{s} - \widehat{s}$, we obtain

$$\begin{aligned} |s - \widehat{s}| &\le \sum_{i=1}^{n} |\widehat{w}_i - w_i| + |\widetilde{s} - \widehat{s}| \\ &\le u \sum_{i=1}^{n} w_i + u \sum_{i=1}^{n} (n+1-i)\widehat{w}_i + O(u^2) \\ &= \sum_{i=1}^{n} (n+2-i) w_i + O(u^2), \end{aligned} \tag{3.2}$$

since $\widehat{w}_i = w_i + O(u)$. Hence

$$\widehat{s} = s + \Delta s, \quad |\Delta s| \le (n+1)us + O(u^2). \tag{3.3}$$

Then the computed log-sum-exp is

$$\begin{aligned} \widehat{y} &= \mathrm{fl}(\log \widehat{s}) = \log(\widehat{s})(1+\epsilon), \quad |\epsilon| \le u, \\ &= \log(s + \Delta s)(1+\epsilon) \\ &= \left(\log s + \frac{\Delta s}{s} + O(u^2)\right)(1+\epsilon) \\ &= y(1+\epsilon) + \frac{\Delta s}{s} + O(u^2). \end{aligned} \tag{3.4}$$

Using (3.3) we obtain

$$|y - \widehat{y}| \le u|y| + (n+1)u + O(u^2),$$

which gives the following result.

THEOREM 3.2 (Basic log-sum-exp algorithm). *In the absence of overflow and underflow, the computed log-sum-exp $\widehat{y}$ from Algorithm 3.1 satisfies*

$$\left|\frac{y - \widehat{y}}{y}\right| \le \left(1 + \frac{n+1}{|y|}\right) u + O(u^2). \tag{3.5}$$

Comparing this bound with $\mathrm{cond}(f, x)u$ in (2.2) we see that it is larger by the factor $(|y| + n + 1)/\|x\|_\infty$. But $|y| \le \|x\|_\infty + \log n$ by (4.2) below, so this factor is bounded by $1 + (n + 1 + \log n)/\|x\|_\infty$. Hence we have forward stability as long as $\|x\|_\infty \gtrsim 1$, but for $\|x\|_\infty \ll 1$ the bound does not guarantee forward stability. This is consistent with the bound (2.4) for the case $n = 1$.

Turning to the evaluation of the softmax function g from its definition (1.2), by (3.1) we have

$$\widehat{g}_j = \frac{\mathrm{e}^{x_j}(1+\delta_1)}{\widehat{s}}(1+\delta_2), \quad |\delta_2| \le u,$$

where δ_2 accounts for the division, and so by (3.3),

$$\widehat{g}_j = \frac{\mathrm{e}^{x_j}}{s(1+\eta)}(1+\delta_1)(1+\delta_2), \quad |\eta| \le (n+1)u + O(u^2).$$

Therefore

$$\widehat{g}_j = g_j(1+\theta), \quad |\theta| \le (n+3)u + O(u^2).$$

This bound guarantees a relative error of order at most nu in every component of g. We weaken the bound into a normwise bound for the next theorem.

THEOREM 3.3 (Basic softmax algorithm). *In the absence of overflow and underflow, the computed softmax $\widehat{g}$ from Algorithm 3.1 satisfies*

$$\frac{\|g - \widehat{g}\|_\infty}{\|g\|_\infty} \le (n+3)u + O(u^2). \tag{3.6}$$

While the error bounds of Theorem 3.2 and 3.3 have a very satisfactory form, they provide no useful information when $n \gtrsim 1/u$, and for fp16 this happens for n as small as 2048. We note, however, that the n terms, which come from the summation, are pessimistic. It is shown by Higham and Mary [14, Thm. 3.1] that, under a probabilistic model of rounding errors, n in the error bound for summation can be replaced by a small constant multiple of $\sqrt{n}$ with high probability, and the same holds for the bounds of Theorem 3.2 and 3.3.

Next, consider the alternative formula (1.5), which we rewrite here:

$$g_j = \exp\left(x_j - \log \sum_{i=1}^{n} \mathrm{e}^{x_i}\right) = \exp(x_j - y). \tag{3.7}$$

With $y = f(x)$ evaluated in floating-point arithmetic by Algorithm 3.1, we obtain

$$\begin{aligned}
\widehat{g}_j &= (1+\delta)\exp\big[(x_j-\widehat{y})(1+\epsilon)\big], \quad |\delta|,|\epsilon| \le u, \\
&= (1+\delta)\exp\big[(x_j-y+(y-\widehat{y}))(1+\epsilon)\big] && (3.8)\\
&= (1+\delta)g_j\exp[(x_j-y)\epsilon+(y-\widehat{y})(1+\epsilon)\big] \\
&= (1+\delta)g_j\big[(1+(x_j-y)\epsilon+(y-\widehat{y})(1+\epsilon)+O(u^2)\big] \\
&= (1+\theta)g_j, && (3.9)
\end{aligned}$$

where, using Theorem 3.2,

$$|\theta| \le (|y|+|x_j-y|+n+2)u+O(u^2).$$

We summarize this result as follows.

Theorem 3.4 (Alternative softmax algorithm). *In the absence of overflow and underflow, the computed $\widehat{g}$ from (3.7) with the log-sum-exp computed by Algorithm 3.1 satisfies*

$$\frac{\|g-\widehat{g}\|_\infty}{\|g\|_\infty} \le \Big(|y|+\max_j|x_j-y|+n+2\Big)u+O(u^2). \tag{3.10}$$

From (4.2) and (4.3) below, using the notation (1.6), we have

$$|y|+\max_j|x_j-y| \le |x_{\max}|+|x_{\max}-x_{\min}|+2\log n.$$

Hence (3.10) is less favorable than (3.6) when $x_{\max}-x_{\min} \gg n$ or $|x_{\max}| \gg n$. The analysis therefore suggests that (1.2) should be preferred to (1.5).

To give an intuitive explanation for the potential inaccuracy in (3.7), we refer to the steps leading to (3.9). A large absolute error in the argument of the final exp may lead to a large relative error in the result. This effect can be traced back to the appearance of $x_j - y$ in (3.8).

4. Algorithms with shifting. Now we consider the use of shifts in the log-sum-exp and softmax evaluations in order to avoid overflow and reduce the chance of harmful underflow.

Recall the definition (1.6) of $x_{\max}$ and $x_{\min}$. Overflow in the exponential evaluations in (1.3) is certainly avoided if we take $a = x_{\max}$, as we then have $x_i - a \le 0$ and hence $0 \le \mathrm{e}^{x_i-a} \le 1$ for all i. We can rewrite (1.3) as

$$y = x_{\max} + \log\Bigg(1+\sum_{\substack{i=1\\ i\ne k}}^{n}\mathrm{e}^{x_i-x_{\max}}\Bigg), \tag{4.1}$$

where $x_k = x_{\max}$. From this expression we see that

$$x_{\max} \le y \le x_{\max}+\log n. \tag{4.2}$$

It follows that when $x_{\max} \ge 0$, the sum "$x_{\max} + \log(\cdot)$" that produces y cannot suffer cancellation.

Note that for $n = 1$, (4.1) trivially provides the exact result $y = x_{\max}$, in contrast to the basic formula (1.1).

For later use, we note that (4.2) implies that, for any j,

$$|y - x_j| \le |x_{\max} - x_j| + \log n \le |x_{\max} - x_{\min}| + \log n. \tag{4.3}$$

The log term in (4.1) has the form $\log(1+z)$, where $z \ge 0$. If z is very small then $1+z$ will round to 1 and the logarithm will evaluate as zero, even though $\log(1+z) \approx z \ne 0$. To avoid this loss of information we will use the function $\mathrm{log1p}(z) = \log(1+z)$ provided in, for example, C, MATLAB, and Numpy. These functions guarantee an accurate result for small z (which can be achieved with a simple formula based on log [9], [12, Prob. 1.5]).

These considerations lead to Algorithm 4.1.

Algorithm 4.1 (log-sum-exp and softmax with shift). *This algorithm computes* $f(x) = \log \sum_{i=1}^n \mathrm{e}^{x_i}$ *and the gradient* $g(x) = \nabla f(x)$ *for* $x \in \mathbb{R}^n$.

$[a, k] = \max_i x_i$ % $a = x_k = \max_i x_i$
$s = 0$
for $i = 1\colon n$
    $w_i = \exp(x_i - a)$
    if $i \ne k$, $s = s + w_i$, end
end
$f = a + \mathrm{log1p}(s)$
for i = 1:n
    $g_i = w_i/(1+s)$
end

Note that while it is important to avoid forming $1+s$ for the f-evaluation, for g we can safely form $1+s$ because if s is small it has little influence on g.

Algorithm 4.1 avoids overflow. If underflow occurs in the exponential then it is in a term in the sum added to 1 in (4.1), so that term is negligible and the underflow is harmless. Note, in particular, that if $x_i \approx x < \log r_{\min}^{(s)}$ for all i then whereas Algorithm 3.1 returns $f = -\infty$, Algorithm 4.1 suffers no underflow and returns $f \gtrsim x_{\max}$.

The main question is how shifting affects the accuracy of the evaluations. We give a rounding error analysis to assess this question. The analysis is a generalization of that in the previous section for the unshifted algorithm.

We first examine the error in evaluating the sum of nonnegative terms

$$s = \sum_{\substack{i=1\\ i\ne k}}^{n} \mathrm{e}^{x_i - a} =: \sum_{\substack{i=1\\ i\ne k}}^{n} w_i. \tag{4.4}$$

Evaluating $w_i = \mathrm{e}^{x_i - a}$ yields a computed result satisfying

$$\widehat{w}_i = \mathrm{e}^{(x_i - a)(1+\delta_1)}(1+\delta_2), \quad |\delta_1| \le u,\ |\delta_2| \le u.$$

Therefore

$$\widehat{w}_i = \mathrm{e}^{x_i - a}\mathrm{e}^{(x_i - a)\delta_1}(1+\delta_2) = \mathrm{e}^{x_i - a}\big(1 + (x_i - a)\delta_1 + O\big(\delta_1^2\big)\big)(1+\delta_2),$$

and hence

$$|\widehat{w}_i - w_i| \le ((1 + a - x_i)u + O(u^2))w_i.$$

Assuming for notational simplicity that $k = n$, we can write the (exact) sum of computed quantities as

$$\widetilde{s} = \sum_{i=1}^{n-1} \widehat{w}_i.$$

The rounding error analysis in [11], [12, sec 4.2] shows that the computed sum $\widehat{s}$ satisfies

$$|\widetilde{s} - \widehat{s}| \le u \sum_{i=1}^{n-2} |t_i| + O(u^2),$$

where $t_i = \sum_{j=1}^{i+1} \widehat{w}_j$, so that, since $\widehat{w}_i \ge 0$,

$$|\widetilde{s} - \widehat{s}| \le u \sum_{i=1}^{n-1} (n-i)\widehat{w}_i + O(u^2).$$

Hence

$$\begin{aligned} |s - \widehat{s}| &\le \sum_{i=1}^{n-1} |\widehat{w}_i - w_i| + |\widehat{s} - \widetilde{s}| \\ &\le u \sum_{i=1}^{n-1} (1 + a - x_i) w_i + u \sum_{i=1}^{n-1} (n-i)\widehat{w}_i + O(u^2) \\ &= \sum_{i=1}^{n-1} \left(u\,(n-i) + u\,(1 + a - x_i)\right) w_i + O(u^2), \qquad (4.5) \end{aligned}$$

since $\widehat{w}_i = w_i + O(u)$. Hence

$$\left| \frac{\widehat{s} - s}{s} \right| \le (n + x_{\max} - x_{\min}) u + O(u^2), \qquad (4.6)$$

which guarantees an accurate computed sum as long as $n + x_{\max} - x_{\min}$ is not too large.

The final stage of the computation is to evaluate $y = x_{\max} + \log(1 + s)$ using the computed $\widehat{s}$, for which we have

$$\widehat{y} = \big(x_{\max} + \log(1 + \widehat{s})(1 + \delta_3)\big)(1 + \delta_4), \quad |\delta_3|, |\delta_4| \le u.$$

Here, we are assuming that the log1p function has the property

$$\mathrm{fl}(\mathrm{log1p}(s)) = \mathrm{log1p}(s)(1 + \delta), \quad |\delta| \le u.$$

Ignoring the innocuous δ_4 term and writing, by (4.6),

$$\widehat{s} = s(1 + \eta), \quad |\eta| \le (n + x_{\max} - x_{\min}) u + O(u^2), \qquad (4.7)$$

we have

$$\begin{aligned} \widehat{y} &= x_{\max} + \log(1 + s(1 + \eta))(1 + \delta_3) \\ &= x_{\max} + \log(1 + s + s\eta))(1 + \delta_3) \\ &= x_{\max} + \left(\log(1 + s) + \frac{s\eta}{1 + s} + O(u^2) \right)(1 + \delta_3), \end{aligned}$$

using a Taylor series expansion about $1 + s$ of the logarithm. Hence

$$\widehat{y} - y = \log(1 + s)\delta_3 + \frac{s\eta}{1 + s}(1 + \delta_3) + O(u^2).$$

Bounding η using (4.7) gives

$$|y - \widehat{y}| \le \log(1 + s)u + \frac{s}{1 + s}(n + x_{\max} - x_{\min})u + O(u^2) \tag{4.8}$$

or, as a relative error bound, since $s \ge 0$,

$$\left|\frac{y - \widehat{y}}{y}\right| \le \left(\frac{\log(1 + s) + n + x_{\max} - x_{\min}}{|y|}\right)u + O(u^2). \tag{4.9}$$

Simplifying the bound gives the next result.

THEOREM 4.2 (Shifted log-sum-exp algorithm). *The computed log-sum-exp $\widehat{y}$ from Algorithm 4.1 satisfies*

$$\left|\frac{y - \widehat{y}}{y}\right| = \left|\frac{y + n - x_{\min}}{y}\right|u + O(u^2). \tag{4.10}$$

The main question is how this result compares with Theorem 3.2 for the unshifted algorithm. The only difference in the bounds is that $|y| + n + 1$ in (3.5) is replaced by $|y + n - x_{\min}|$ here. Now $|y + n - x_{\min}| \gg |y| + n$ is possible only if $x_{\min} \ll 0$ and $x_{\min} \ll x_{\max}$, so let us assume that these two inequalities hold. The term $|y + n - x_{\min}|$ comes from bounding the term $(1 + a - x_i)w_i$, where w_i is defined in (4.4) and $x_i = x_{\min}$, and if $x_{\min} \ll 0$ then $w_i = \mathrm{e}^{x_i - a} = \mathrm{e}^{x_{\min} - x_{\max}} \ll 1$. Hence the potentially large constant is mitigated by the w_i term that it multiplies—something that is lost in the manipulations to achieve a readable bound. We conclude that shifting should have little effect on the accuracy.

We note that (4.10) is weaker than necessary when $s \ll 1$ (recall that $s \ge 0$), since we bounded $s/(1 + s)$ by 1 in going from (4.8) to (4.9). If $s \ll 1$ then (4.8) becomes

$$|y - \widehat{y}| \lesssim s(1 + n + x_{\max} - x_{\min})u + O(u^2).$$

Since $s \ll 1$ also implies $x_i \ll x_{\max}$ for $i \ne k$ and hence $y \approx x_{\max}$, we have

$$\frac{|y - \widehat{y}|}{|y|} \lesssim s\frac{|1 + n + y - x_{\min}|}{|y|}u + O(u^2),$$

which is a factor s smaller than (4.10).

Turning to the evaluation of the softmax function g from the shifted formula (1.4), we have, using (4.6),

$$\widehat{g}_j = \frac{\exp((x_j - a)(1 + \delta_1))(1 + \delta_2)(1 + \delta_3)}{s(1 + \eta)},$$

where δ_2 corresponds to the exponential evaluation and δ_3 to the division, and

$$|\delta_i| \le u,\ i = 1\colon 3, \qquad |\eta| \le (n + x_{\max} - x_{\min})u + O(u^2).$$

Therefore

$$\begin{aligned}\widehat{g}_j &= g_j \frac{\exp((x_j - a)\delta_1)(1+\delta_2)(1+\delta_3)}{1+\eta} \\ &= g_j(1+\theta), \quad |\theta| \le \big(n + 2 + 2(x_{\max} - x_{\min})\big)u + O(u^2).\end{aligned}$$

Hence we have obtained the following result.

THEOREM 4.3 (Shifted softmax algorithm). *The computed $\widehat{g}$ from Algorithm 4.1 satisfies*

$$\frac{\|g - \widehat{g}\|_\infty}{\|g\|_\infty} \le \big(n + 2 + 2(x_{\max} - x_{\min})\big)u + O(u^2). \tag{4.11}$$

Again, this is broadly commensurate with Theorem 3.3 for the unshifted evaluation, bearing in mind the comments following Theorem 4.2.

Finally, we consider (1.5) with the log-sum-exp computed by Algorithm 4.1. In floating-point arithmetic we have the same equation (3.8) as for the unshifted algorithm, but now with θ bounded by, using (4.10),

$$|\theta| \le (1 + |x_j - y| + |y + n - x_{\min}|)u + O(u^2).$$

We have obtained the following result.

THEOREM 4.4 (Alternative shifted softmax algorithm). *The computed $\widehat{g}$ from (1.5) with the log-sum-exp computed by Algorithm 4.1 satisfies*

$$\frac{\|g - \widehat{g}\|_\infty}{\|g\|_\infty} \le \Big(1 + \max_j |x_j - y| + |y + n - x_{\min}|\Big)u + O(u^2). \tag{4.12}$$

This is broadly similar to Theorem 3.4 for the unshifted alternative softmax algorithm.

5. Computational experiments. We now perform some experiments in a realistic setting, using MATLAB R2019a. The codes and data used for the experiments are available online[6].

Our aims are to examine the sharpness of the rounding error bounds and to give a pairwise comparison of the accuracy of the algorithms in floating-point arithmetic. Our data comes from a deep learning application. To generate the data, we first set up and trained an artificial neural network, using the MATLAB Deep Learning Toolbox [6]. More precisely, we trained a network to classify handwritten digit data from the widely used MNIST data set [19]. Here each data point is a grayscale 28×28 pixel image and there are ten categories: 0, 1, ..., 9. We used a network whose architecture has the following general form:

1. Image Input $28 \times 28 \times 1$ with normalization.
2. Convolution 8 $3 \times 3 \times 1$ stride [1 1] padding 'same'.
3. Batch Normalization 8 channels.
4. ReLU
5. Max Pool 2×2 stride [2 2] padding [0 0 0 0].
6. Convolution 16 $3 \times 3 \times 8$ stride [1 1] padding 'same' .
7. Batch Normalization 16 channels.
8. ReLU.

[6] https://github.com/higham/logsumexp-softmax-tests

9. Max Pool 2×2 stride [2 2] padding [0 0 0 0].
10. Convolution 32 $3 \times 3 \times 16$ stride [1 1] padding 'same'.
11. Batch Normalization 32 channels.
12. ReLU.
13. Fully Connected 10 layer.
14. Softmax.
15. Classification Output crossentropy.

This is the default architecture from [6], where further details may be found.

The network was trained on 7500 images (750 from each of the ten categories), with 2500 further images (250 from each of the ten categories) used for validation.

The network takes as input a 28×28 matrix corresponding to the pixels in the image and returns a nonnegative 10×1 vector whose ith component may be interpreted as the probability that the image came from category i. If we categorize according to the highest probability from the output, then the trained network misclassifed 27 of the 2500 validation images, corresponding to a 98.9% success rate.

The network uses single precision arithmetic, `fp32`. In our experiments, we are concerned only with floating-point arithmetic issues, and we treat the trained network as a means to produce a realistic data set. To do this, we extracted the 2500 single precision vectors from the validation set that were passed into the softmax layer and converted them to fp16 or bfloat16. We then used this data in our implementation of the softmax and log-sum-exp algorithms that we have studied in the previous sections.

To record errors in computed results we applied the basic algorithm, Algorithm 3.1, in single precision to provide a reference solution and used the `chop` function of [15] to simulate half precision arithmetic, in both the fp16 format and the bfloat16 format.

We first describe experiments in fp16. The components in the 2500 test vectors $x \in \mathbb{R}^{10}$ vary between about -19 and $+20$. As indicated in Table 1.2, e^x overflows in fp16 for $x \gtrsim 11$. Hence, in these tests, overflow is an issue for the basic log-sum-exp implementation in Algorithm 3.1: it generated an `Inf` for 475 of the 2500 test vectors. The shifted version of log-sum-exp in Algorithm 4.1 did not overflow. In the plots below, we do not include results for the cases where Algorithm 3.1 produced overflow.

First, we look at the log-sum-exp algorithms. In the upper left plot of Figure 5.1 we used the basic implementation of log-sum-exp, Algorithm 3.1. We scatter plot over the 2025 vectors where no overflow occurred. For each such vector, the horizontal coordinate is the leading term in the error bound of Theorem 3.2, scaled by u, that is, $1 + (n+1)/|y|$. Here, as shown in Table 1.1, $u = 4.88 \times 10^{-4}$ for fp16. The vertical coordinate is the actual scaled relative error $|\widehat{y} - y|/(u|y|)$. The plot also gives a reference line of slope 1 from the origin. We see that the bound is always satisfied and is reasonably sharp in many cases.

In the upper right plot of Figure 5.1 we show corresponding results for the shifted log-sum-exp implementation in Algorithm 4.1, using the bound from Theorem 4.2.

In the lower part of Figure 5.1 we scatter plot the floating-point errors for the basic and shifted algorithms. Here, for 1863 out of the 2025 cases (92%) the two errors were identical to all digits in the half precision computation. In more detail, over all the data points the ratio of the error in the basic log-sum-exp (horizontal axis) divided by the error in the shifted version (vertical axis) varied between 0.19 and 59, with a mean of 1.07 and a standard error of 0.03. This indicates that the two versions perform similarly, with the shift producing slightly better results.

We now move on to the four softmax implementations. In Figure 5.2 we use the shifted softmax implementation from Algorithm 4.1, analysed in Theorem 4.3, as the basis for comparison. The upper left plot has the scaled error $\|\widehat{g} - g\|_\infty/(u\|g\|_\infty)$ from

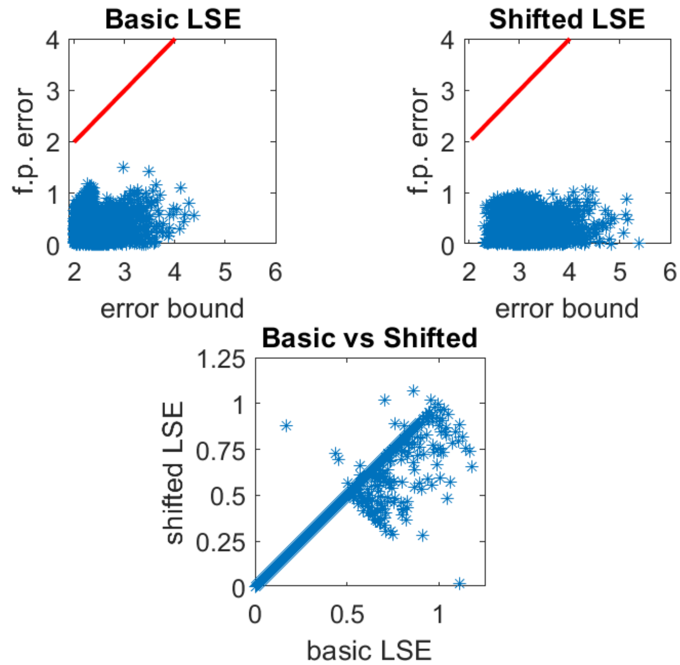

FIG. 5.1. *Scatter plots of errors and error bounds, scaled by unit roundoff, over* 2025 *vectors in* $\mathbb{R}^{10}$ *for log-sum-exp algorithms in fp16. See the text for a description of the axes. Upper left: basic implementation of log-sum-exp from Algorithm* 3.1*. According to the error analysis, all points should lie below the reference line* $y = x$ *(shown in red). Upper right: corresponding results for the shifted implementation of log-sum-exp in Algorithm* 4.1*. Lower: scaled error from Algorithm* 3.1 *versus scaled error from Algorithm* 4.1*.*

Algorithm 4.1 on the horizontal axis and the scaled error from the basic softmax in Algorithm 3.1 on the vertical axis. The upper right plot compares the shifted softmax against the alternative algorithm analyzed in Theorem 3.4. Similarly, the lower plot compares against the alternative shifted softmax algorithm analyzed in Theorem 4.4. We see that the softmax values obtained from Algorithms 3.1 and 4.1 have similar accuracy, whereas the alternative softmax versions based on the rewrite in (1.5) are typically less accurate.

The results in Figures 5.1 and 5.2 are consistent with our floating-point error analysis.

A further test is to compute the sum of each softmax vector, which should equal 1. In Figure 5.3 we compare the softmax sums for the basic algorithm (red circles) analyzed in Theorem 3.3 and the alternative version (blue crosses) analyzed in Theorem 3.4. Similarly, Figure 5.4 compares the shifted softmax algorithm analyzed in Theorem 4.3 and its alternative analyzed in Theorem 4.4. The order along the x-axis is arbitrary; it corresponds to the order in which the data vectors were generated. These figures provide further evidence that the alternative softmax algorithms are less accurate than the basic or shifted algorithms.

We also conducted the corresponding experiments in simulated bfloat16 arithmetic. Here, as indicated in Tables 1.1 and 1.2, the number range is increased at the expense of reduced precision. In this case there was no overflow in any of the algorithms. The results were very similar to those for fp16, so they are not shown here.

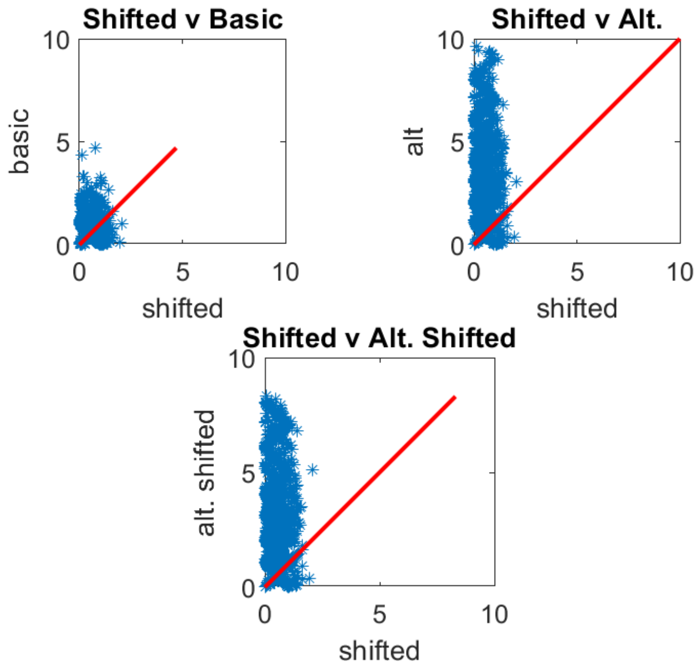

FIG. 5.2. *Scatter plots of errors, scaled by unit roundoff, for softmax algorithms in fp16. See the text for a description of the axes. Reference line is $y = x$.*

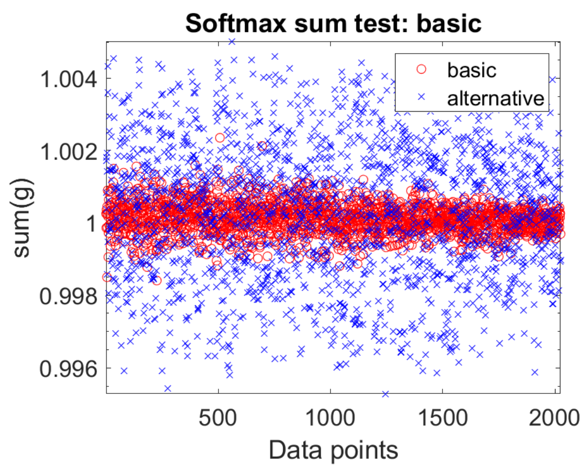

FIG. 5.3. *Sum of entries of computed softmax vector for Algorithm 3.1 (red circles), analyzed in Theorem 3.3, and the alternative (blue crosses) analyzed in Theorem 3.4.*

6. Conclusions. The log-sum-exp and softmax functions both feature in many computational pipelines, so it is important to compute them accurately and to avoid generating infs or NaNs because of overflow or underflow. To this end, a shift is usually incorporated into the defining formulas, yielding (1.3) and (1.4). It is important to understand the effect of the shift on the accuracy of the computed result, especially when computations are carried out in a low precision such as bfloat16 or fp16, which

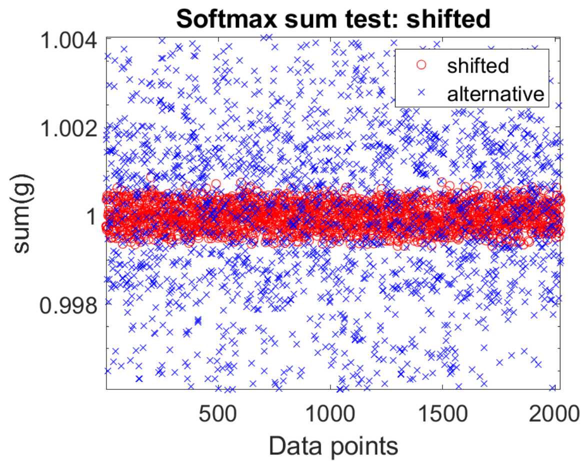

Fig. 5.4. *Sum of entries of computed softmax vector for Algorithm 4.1 (red circles), analyzed in Theorem 4.3, and the alternative (blue crosses) analyzed in Theorem 4.4.*

have the equivalent of only 3 or 4 decimal digits of precision.

Our rounding error analysis shows that shifting by the largest element of the input vector does not lessen the accuracy of the computed log-sum-exp and softmax. Underlying this pleasing fact is the phenomenon that any large coefficients caused by shifting are canceled by multiplication with small exponentials.

We obtained an explicit formula for the condition number of log-sum-exp and bounds for the condition number of softmax, and we were able to identify situations in which the log-sum-exp algorithms are guaranteed to be forward stable.

For the alternative and widely used softmax formula that avoids division, (1.5), we obtained larger error bounds than for the shifted formula (1.4). Since our numerical experiments confirm that larger errors are typically obtained in practice, we recommend using (1.4) instead of (1.5) to evaluate softmax.

In summary, Algorithm 4.1 is our recommendation for computing log-sum-exp and softmax. It avoids overflow, reduces the chance of harmful underflow, and generally produces results as accurate as those from the unshifted formulas.

REFERENCES

[1] Mary Aprahamian and Nicholas J. Higham. The matrix unwinding function, with an application to computing the matrix exponential. *SIAM J. Matrix Anal. Appl.*, 35(1):88–109, 2014.

[2] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag, New York, 2006. xx+738 pp. ISBN 978-0-387-31073-2.

[3] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004. xiii+716 pp. ISBN 0-521-83378-7.

[4] Giuseppe C. Calafiore, Stephane Gaubert, and Corrado Possieri. Log-sum-exp neural networks and posynomial models for convex and log-log-convex data. *IEEE Trans. Neural Networks and Learning Systems*, pages 1–12, 2019.

[5] Florent de Dinechin, Christoph Lauter, and Jean-Michel Muller. Fast and correctly rounded logarithms in double-precision. *RAIRO-Inf. Theor. Appl.*, 41(1):85–102, 2007.

[6] Deep Learning Toolbox. The MathWorks, Inc., Natick, MA, USA. http://www.mathworks.co.uk/products/deep-learning/.

[7] Bradley Efron and Trevor Hastie. *Computer Age Statistical Inference. Algorithms, Evidence, and Data Science*. Cambridge University Press, Cambridge, UK, 2016. xix+475 pp. ISBN 978-1-107-14989-2.

[8] Ian Goodfellow, Yoshua Bengio, and Aaaron Courville. *Deep Learning*. The MIT Press, Cambridge, MA, USA, 2016. xxii+775 pp. ISBN 978-0-262-03561-3.

[9] *HP-15C Advanced Functions Handbook*. Hewlett-Packard, Portable Computer Division, Corvallis, OR, USA, 1982. 221 pp. Part number 00015-90011 Rev. C.

[10] Catherine F. Higham and Desmond J. Higham. Deep learning: An introduction for applied mathematicians. *SIAM Review*, to appear, 2019.

[11] Nicholas J. Higham. The accuracy of floating point summation. *SIAM J. Sci. Comput.*, 14(4): 783–799, 1993.

[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 and Theo Mary. A new approach to probabilistic rounding error analysis. MIMS EPrint 2018.33, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, November 2018. 22 pp. Revised March 2019. To appear in SIAM J. Sci. Comput.

[15] Nicholas J. Higham and Srikara Pranesh. Simulating low precision floating-point arithmetic. MIMS EPrint 2019.4, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, March 2019. 18 pp. Revised July 2019. To appear in SIAM J. Sci. Comput.

[16] *IEEE Standard for Floating-Point Arithmetic, IEEE Std 754-2008 (revision of IEEE Std 754-1985)*. IEEE Computer Society, New York, 2008. 58 pp. ISBN 978-0-7381-5752-8.

[17] Intel Corporation. BFLOAT16—hardware numerics definition, November 2018. White paper. Document number 338302-001US.

[18] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. http://www.scipy.org/.

[19] Yann LeCun, Corinna Cortes, and Christopher J. C. Burges. The MNIST database of handwritten digits. Accessed June 17, 2019.

[20] Matlab code for machine learning algorithms in book PRML. https://github.com/PRML/PRMLT.

[21] Jean-Michel Muller. *Elementary Functions: Algorithms and Implementation*. Third edition, Birkhäuser, Boston, MA, USA, 2016. xxv+283 pp. ISBN 978-1-4899-7981-0.

[22] Jean-Michel Muller, Nicolas Brunie, Florent de Dinechin, Claude-Pierre Jeannerod, Mioara Joldes, Vincent Lefèvre, Guillaume Melquiond, Nathalie Revol, and Serge Torres. *Handbook of Floating-Point Arithmetic*. Second edition, Birkhäuser, Boston, MA, USA, 2018. xxv+627 pp. ISBN 978-3-319-76525-9.

[23] Kevin P. Murphy. *Machine Learning: a Probabilistic Approach*. Cambridge University Press, Cambridge, UK, 2012.

[24] Statistics and Machine Learning Toolbox. The MathWorks, Inc., Natick, MA, USA. https://uk.mathworks.com/products/statistics.html.

[25] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

[26] Christopher K. I. Williams and David Barber. Bayesian classification with Gaussian processes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 20(12):1342–1351, 1998.