



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/236016/>

Version: Accepted Version

Article:

El Arar, E.-M., Fasi, M., Filip, S.-I. et al. (2025) Probabilistic Error Analysis of Limited-Precision Stochastic Rounding. *SIAM Journal on Scientific Computing*, 47 (5). B1227-B1249. ISSN: 1064-8275

<https://doi.org/10.1137/24m1681458>

This is an author produced version of an article published in *SIAM Journal on Scientific Computing*, made available via the University of Leeds Research Outputs Policy under the terms of the Creative Commons Attribution License (CC-BY), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.





HAL
open science

Probabilistic error analysis of limited-precision stochastic rounding

El-Mehdi El Arar, Massimiliano Fasi, Silviu-Ioan Filip, Mantas Mikaitis

► **To cite this version:**

El-Mehdi El Arar, Massimiliano Fasi, Silviu-Ioan Filip, Mantas Mikaitis. Probabilistic error analysis of limited-precision stochastic rounding. 2025. hal-04665809v3

HAL Id: hal-04665809

<https://hal.science/hal-04665809v3>

Preprint submitted on 24 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Probabilistic error analysis of limited-precision stochastic rounding*

El-Mehdi El arar[†] Massimiliano Fasi[‡] Silviu-Ioan Filip[†]
Mantas Mikaitis[‡]

Abstract. Classical probabilistic rounding error analysis is particularly well suited to stochastic rounding (SR), and it yields strong results when dealing with floating-point algorithms that rely heavily on summation. For many numerical linear algebra algorithms, one can prove probabilistic error bounds that grow as $\mathcal{O}(\sqrt{nu})$, where n is the problem size and u is the unit roundoff. These probabilistic bounds are asymptotically tighter than the worst-case ones, which grow as $\mathcal{O}(nu)$. For certain classes of algorithms, SR has been shown to be unbiased. However, all these results were derived under the assumption that SR is implemented exactly, which typically requires a number of random bits that is too large to be suitable for practical implementations. We investigate the effect of the number of random bits on the probabilistic rounding error analysis of SR. To this end, we introduce a new rounding mode, limited-precision SR. By taking into account the number r of random bits used, this new rounding mode matches hardware implementations accurately, unlike the ideal SR operator generally used in the literature. We show that this new rounding mode is biased and that the bias is a function of r . As r approaches infinity, however, the bias disappears, and limited-precision SR converges to the ideal, unbiased SR operator. We develop a novel model for probabilistic error analysis of algorithms employing SR. Several numerical examples corroborate our theoretical findings.

Keywords. Rounding error analysis, stochastic rounding, low-precision, floating-point arithmetic, inner product

MSC classification 65G50, 65F30.

1 Introduction

A finite-precision number system cannot represent all reals, and values that cannot be represented must be approximated by a nearby machine-representable number. This process, known as *rounding*, ensures that partial results that are not representable can be used in subsequent limited-precision computations. Rounding operators usually make a choice between two rounding candidates, one larger and one smaller than the value that has to be rounded, and the rules for making that choice are called *rounding modes*.

Most rounding modes available in computer hardware are deterministic, meaning that for the same value, they will always return the same rounding candidate. For example, round-to-nearest (RN) will always choose a machine-representable number that minimizes the distance from the

*Version of February 24, 2025. **Funding:** This work was funded by the HOLIGRAIL project (ANR-23-PEIA-0010).

[†]Université de Rennes, Inria, CNRS, IRISA, F-35000 Rennes, France (el-mehdi.el-arar@inria.fr, silviu.filip@inria.fr).

[‡]School of Computing, University of Leeds, Woodhouse Lane, Leeds LS2 9JT, UK (m.fasi@leeds.ac.uk, m.mikaitis@leeds.ac.uk).

input. Here, we are concerned with a non-deterministic rounding mode: mode 1 *stochastic rounding* (SR) [6]. SR guarantees that rounding errors are random, which makes them more likely to cancel out. This rounding mode is particularly beneficial in alleviating *stagnation*, a numerical phenomenon that causes small addends to be lost to rounding in long chains of sums.

SR can be traced back to the 1950s [15], but it has seen a resurgence of interest in recent years. It has been shown to be particularly well-suited to machine learning computations [19, 22], and it is starting to become available in low-precision floating-point hardware: both the Tesla D1 [33] and AMD MI300 GPUs [1, sect. 7.2] support SR down-conversion of `binary32` values to 8-bit floating-point representations, whereas the Graphcore Intelligence Processing Unit (IPU) [18, sect. 7] offers the option to stochastically round (intermediate) `binary32` accumulation results in their Colossus IPU dot product operations to `binary16`. Starting with the Blackwell architecture, NVIDIA GPUs support SR down-conversion from `binary32` to 16, 8, 6 and 4-bit values [30, sect. 9.7.9.21]. The NeuronCore-v2 architecture powering AWS Trainium chips [13] offers the possibility of using either RN or SR when doing `binary16` or `bfloat16` arithmetic (at least additions) by setting a global rounding flag.

Indeed, artificial intelligence has motivated much of the recent research on SR, as this rounding mode has the potential to provide higher accuracy in various deep learning and optimization applications. When training neural networks with deterministic rounding modes, for example, the accumulation of gradients during parameter updates may be prone to stagnation, as the gradients become smaller in magnitude as training progresses. Gupta et al. [19] have shown that SR can alleviate this issue, and this has been further explored in subsequent work [35, 37].

SR has found some important uses in applications outside machine learning. It can alleviate stagnation in the solution of partial differential equations (PDEs), via Runge–Kutta finite difference methods in low precision [7], and of ordinary differential equations (ODEs), via the Euler, midpoint, and Heun methods [14]. Paxton et al. [31] have studied the effectiveness of low-precision arithmetic for climate simulations, focusing in particular on the effects of RN and SR in Lorenz systems, chaotic systems of ODEs, and PDEs related to climate modeling. They have found that SR can effectively mitigate rounding errors across various applications, and these results provide evidence that SR may be relevant to next-generation climate models. For a comprehensive overview of the uses of SR, we refer the reader to [6, sect. 8].

Theoretical results have confirmed the beneficial effects of SR observed in numerical applications. Assuming precision- p arithmetic, Parker [32] has shown that SR is unbiased. Connolly et al. [5] have proven that SR satisfies the mean independence property, whereby one can construct martingales. Using the Azuma–Hoeffding inequality, they introduce a model that leads to probabilistic bounds on the relative error of a large class of linear algebra algorithms. Consider, for example, the evaluation of the inner product between two length- n vectors in precision- p floating-point arithmetic. The deterministic worst-case bound on the roundoff error that can occur during this computation grows as nu , where $u = 2^{-p}$. Connolly et al. [5] have obtained a bound in $\mathcal{O}(\sqrt{n \ln(n)u})$, and by using an alternative construction for the martingale, Ipsen and Zhou [25] have been able to derive a tighter one in $\mathcal{O}(\sqrt{nu})$. This construction has been extended to the Horner method for polynomial evaluation [10]. A more recent approach relies exclusively on the variance of the error and the Bienaymé–Chebyshev inequality [11]. The advantages of each approach are discussed in [9, sect. 4.3].

The rounding errors produced by SR are not always unbiased; examples of this are two algorithms for computing the sample variance [12]. Nevertheless, by using the aforementioned alternative martingale construction [11, 25], one can derive probabilistic error bounds in $\mathcal{O}(\sqrt{nu})$ for these two algorithms. Dexter et al. [8] show that, when rounding tall-and-thin matrices to low precision, if SR is used, then the smallest singular value of the rounded matrix is bounded away from zero with high probability.

SR implementations, both in hardware and in software, tend to use a summation-based algo-

rithm [6, sect. 7]. In a precision- p floating-point number system (see Section 2 for details), this amounts to adding randomly generated bits to the trailing bits of the significand to be rounded. Ideally, these random bits should cover all the trailing bits to be rounded off, but this may be unfeasible, as the exact significand may have a non-terminating binary expansion. Therefore, the number r of random bits to use must be limited in practice. This introduces a certain bias in the behavior of the implementations, with too small values of r negating the expected benefits derived from using SR.

Ali, Filip, and Sentieys [2] study this empirically in the context of deep neural network training, where they use SR for inner product computations in matrix multiplication. In the context of ODE solvers, experiments varying r were performed for fixed-point arithmetic, where SR was used in the multiplication of fixed-point numbers [23, sect. 5c].

As far as we are aware, however, theoretical work on SR does not put constraints on r , and we are the first to consider the effects of limited precision from a theoretical point of view. To fully understand the behavior of practical implementations of SR, however, it is imperative to understand the impact of r from an error analysis perspective. Therefore, we extend the analysis of standard algorithms under SR by taking into account the number of random bits used to perform the rounding. We call this new rounding mode *limited-precision stochastic rounding*, and we denote it by $\text{SR}_{p,r}$. In this notation, $p - 1$ denotes the number of fraction bits available in the format to which the operator rounds, and similarly, we denote by SR_p the classical SR operator, for which $r = \infty$. We take the unit roundoff for precision- p arithmetic to be $u_p := 2^{1-p}$, in line with previous work [11]. Our main contributions are outlined below.

- In Section 3, we show that the unbiased and mean independence properties of SR_p break down for $\text{SR}_{p,r}$.
- In Sections 3 and 4, we present a model whereby we can analyze algorithms under $\text{SR}_{p,r}$ and compute probabilistic error bounds proportional to $\sqrt{n}u_p + nu_{p+r}$, where p is the working precision, r the number of random bits, and $u_p = 2^{1-p}$.
- In Section 4, we use our model to analyze the recursive summation and inner product algorithms.
- In Section 5, we suggest a theoretically sound rule of thumb for choosing r in recursive summation and inner product computations: setting $r = \lceil (\log_2 n)/2 \rceil$ offers, with high probability, a good tradeoff between accuracy and operator complexity.
- In Section 6, we showcase the value of our analysis by means of numerical experiments focused on varying r .

2 Notation and definitions

We recall basic definitions and properties from probability theory [28] and floating-point arithmetic [21, 29]. We also formally introduce SR and its limited-precision variant.

2.1 Probability

Let X be a random variable. Throughout this paper, $\mathbb{E}(X)$ denotes the expected value of X and $\mathbb{V}(X)$ denotes its variance. The conditional expectation of X given Y is denoted by $\mathbb{E}(X | Y)$.

The following result provides a bound on the probability of a random variable falling within a specific number of standard deviations from its mean.

Lemma 1 (Bienaymé–Chebyshev inequality). *Let X be a random variable. If $\mathbb{E}(X)$ and $\mathbb{V}(X)$ are finite, then for any real number $\alpha > 0$, one has*

$$\mathbb{P}(|X - \mathbb{E}(X)| \leq \alpha \sqrt{\mathbb{V}(X)}) \geq 1 - \frac{1}{\alpha^2}.$$

To analyze how rounding errors accumulate, we will model them as random variables, and we will use probability to make precise statements about how rounding errors interact.

The linear relationship between two random variables X and Y is captured by their *covariance*

$$\text{Cov}(X, Y) = \mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y).$$

If $\text{Cov}(X, Y) = 0$, then the two random variables are *uncorrelated*, which means that there is no linear dependence between them. The lack of linear correlation is not sufficient in our analysis, and we need to rely on the following property instead.

Definition 1 (Mean independence). The random variable X is *mean independent* of the random variable Y if $\mathbb{E}(X | Y) = \mathbb{E}(X)$. The sequence of random variables X_0, X_1, \dots is *mean independent* if $\mathbb{E}(X_k | X_0, X_1, \dots, X_{k-1}) = \mathbb{E}(X_k)$ for all k .

As the following result demonstrates, being mean independent is a stronger property than being uncorrelated but weaker than being truly independent.

Proposition 1. *Let X and Y be random variables.*

1. *If X and Y are independent, then X is mean independent of Y .*
2. *If X is mean independent of Y , then X and Y are uncorrelated.*

The reciprocals of these two implications are false.

When X is fully determined by Y , the conditional expectation [3, sect. 34] satisfies the following property.

Proposition 2. *Let X and Y be random variables defined on the same probability space. If X is entirely determined by Y , meaning there exists a measurable function g such that $X = g(Y)$, then the conditional expectation of X given Y satisfies $\mathbb{E}(X | Y) = X$.*

Definition 2 ([28, p. 295]). A sequence of random variables M_0, \dots, M_n is a *martingale* with respect to the sequence X_0, \dots, X_n if, for all k ,

- M_k is a function of X_0, \dots, X_k ,
- $\mathbb{E}(|M_k|)$ is finite, and
- $\mathbb{E}(M_k | X_0, \dots, X_{k-1}) = M_{k-1}$.

Lemma 2 (Azuma–Hoeffding inequality, [28, p. 303]). *Let M_0, \dots, M_n be a martingale with respect to a sequence X_0, \dots, X_n . We assume that there exist $0 < a_k$ such that $-a_k \leq M_k - M_{k-1} \leq a_k$ for $k = 1, \dots, n$. Then, we have*

$$\mathbb{P}\left(|M_n - M_0| \leq \sqrt{\sum_{k=1}^n a_k^2} \sqrt{2 \ln(2/\lambda)}\right) \geq 1 - \lambda,$$

where $0 < \lambda < 1$.

The mean independence property is essential to improve the error analysis of algorithms performed using stochastic rounding. It allows one to obtain a martingale (Definition 2), and then by applying Lemma 2, derive probabilistic bounds on the error that grows in $\mathcal{O}(\sqrt{nu_p})$.

2.2 Floating-point arithmetic

Any nonzero real number $x \in \mathbb{R}$ can be represented as

$$x = (-1)^s \cdot 2^e \cdot m, \quad s \in \{0, 1\}, \quad e \in \mathbb{Z}, \quad m \in [1, 2). \quad (1)$$

This representation, which we shall call *scientific notation*, is unique for any $x \neq 0$.

Let $p > 0$. The set $\mathcal{F} \subset \mathbb{R}$ is a normalized precision- p binary floating-point number system if it contains all real numbers x that can be written as

$$x = (-1)^s \cdot 2^e \cdot m,$$

where the sign s is either 0 or 1, the exponent e is an integer in some range $[e_{\min}, e_{\max}]$, with $e_{\min} < e_{\max}$, and the significand $m \in [1, 2)$ is a real number with at most p significant binary digits (bits). Customary choices of p are defined in the IEEE 754 standard for floating-point arithmetic [24], the OCP 8-bit floating-point specification [27], and the IEEE P3109 standard¹ on arithmetic formats for machine learning.

For $x \in \mathbb{R}$, we denote the smallest precision- p floating-point number no smaller than x by $\lceil x \rceil_p$, and the largest floating-point number no greater than x by $\lfloor x \rfloor_p$. In other words, we have

$$\lceil x \rceil_p = \min\{y \in \mathcal{F} : y \geq x\}, \quad \lfloor x \rfloor_p = \max\{y \in \mathcal{F} : y \leq x\},$$

and by definition, $\lfloor x \rfloor_p \leq x \leq \lceil x \rceil_p$, with equality throughout if and only if $x \in \mathcal{F}$. A real number $x \neq 0$ has two possible representations in floating-point arithmetic, $\lfloor x \rfloor_p$ or $\lceil x \rceil_p$, which coincide if $x \in \mathcal{F}$. The operation that maps x to the chosen candidate \hat{x} is called rounding, and the rounded quantity satisfies

$$\hat{x} = x(1 + \delta), \quad (2)$$

where the relative error $\delta = (\hat{x} - x)/x$ is such that $|\delta| < 2^{1-p} = u_p$. For round-to-nearest (RN) with any tie-breaking rule, we have the tighter bound $|\delta| \leq u_p/(2 + u_p)$ [29, Thm. 2.3].

For $x \in \mathbb{R}$, we define the *unit in the last place* as $\text{ulp}_p(x) = 2^{e-p+1}$, where e is the normalised exponent of x . Since $x = (-1)^s \cdot 2^e \cdot m$ with $1 \leq m < 2$, we have that $2^e \leq |x| < 2^{e+1}$ and

$$\text{ulp}_p(x) = 2^e u_p \leq |x| u_p. \quad (3)$$

Let $x, y \in \mathcal{F}$ and $\text{op} \in \{+, -, \times, \div\}$. We assume the standard model of floating-point arithmetic, whereby the error in one elementary operation is bounded, and we have

$$\text{fl}_p(x \text{ op } y) = (x \text{ op } y)(1 + \delta), \quad \text{fl}_p(\sqrt{x}) = \sqrt{x}(1 + \delta), \quad (4)$$

where $|\delta|$ is bounded by $u_p/2$ for RN with any tie-breaking rule and by u_p for directed and stochastic rounding modes.

2.3 Stochastic rounding

Definition 3 (Stochastic rounding). Let $x \in \mathbb{R}$. The precision- p stochastic rounding of x to \mathcal{F} is the Bernoulli random variable

$$\text{SR}_p(x) = \begin{cases} \lceil x \rceil_p, & \text{with probability } q(x), \\ \lfloor x \rfloor_p, & \text{with probability } 1 - q(x), \end{cases} \quad q(x) = \frac{x - \lfloor x \rfloor_p}{\text{ulp}_p(x)}. \quad (5)$$

If $x \notin \mathcal{F}$, then $\text{ulp}_p(x) = \lceil x \rceil_p - \lfloor x \rfloor_p$, which is the distance between the two floating-point numbers enclosing x . Substituting this quantity in (5) when $x \notin \mathcal{F}$ allows us to recover previous definitions of SR_p from the literature [6].

¹<https://sagroups.ieee.org/p3109wgpublic/>




FIGURE 1. Quantities used in the definitions (5) and (6).

Remark 1. This definition corresponds to what is called *mode 1 SR* or *SR-nearness*. *Mode 2 SR* consists of taking $q(x) = 1/2$ if $x \notin \mathcal{F}$ and $q(x) = 0$ otherwise. While not our focus here, mode 2 SR can be an effective tool for doing what is called *stochastic arithmetic* [34, 16, 17], allowing one to detect instabilities in numerical routines and to provide accuracy estimates of computed results.

The quantities in this definition are represented pictorially in Figure 1. Note that if $x \in \mathcal{F}$, then $q(x) = 0$ and $\text{SR}_p(x) = \lfloor x \rfloor_p = x$ with probability 1. More generally, for $x \in \mathbb{R}$ we have

$$\begin{aligned} \mathbb{E}(\text{SR}_p(x)) &= q(x)\lceil x \rceil_p + (1 - q(x))\lfloor x \rfloor_p \\ &= q(x)(\lceil x \rceil_p - \lfloor x \rfloor_p) + \lfloor x \rfloor_p = x. \end{aligned}$$

The definition of $q(x)$ in (5) assumes that x is known with infinite precision. We now want to investigate how the behavior of SR_p changes when an infinitely precise x is not available or not used for efficiency reasons, and we only have $\text{fl}_{p+r}(x)$, the binary representation of x truncated to the first $p+r$ binary digits, for some positive integer r . We use the letter r to denote the additional bits of precision as this quantity corresponds to the number of random bits used to perform SR, as we will see later. This situation is also represented pictorially in Figure 1.

We note that fl_{p+r} could use a rounding mode other than truncation. All directed rounding modes satisfy $\text{fl}_{p+r}(x) = x(1 + \beta)$, where β is a (deterministic) relative error such that $|\beta| \leq 2^{1-p-r} = u_{p+r}$. For RN, the bound on $|\beta|$ can be reduced by half, but this does not significantly affect our analysis.

Definition 4 (limited-precision stochastic rounding). Let $x \in \mathbb{R}$. The limited-precision stochastic rounding using r random bits ($\text{SR}_{p,r}$) of x to \mathcal{F} is the Bernoulli random variable

$$\text{SR}_{p,r}(x) = \begin{cases} \lceil x \rceil_p, & \text{with probability } q_r(x), \\ \lfloor x \rfloor_p, & \text{with probability } 1 - q_r(x), \end{cases} \quad q_r(x) = \frac{\text{fl}_{p+r}(x) - \lfloor x \rfloor_p}{\text{ulp}_p(x)}. \quad (6)$$

Remark 2. $\text{SR}_{p,r}$ and SR_p both output p -bit precision results. The main difference is that $\text{SR}_{p,r}$ uses r random bits during the operation before the truncation step, whereas SR_p essentially assumes an infinite value for r . Note that $\text{fl}_{p+r}(x)$ is represented with $p+r$ bits of precision.

Remark 3. Depending on how x is produced, one can determine the value of r such that $\text{SR}_{p,r}(x) = \text{SR}_p(x)$. For instance, if $x = ab$ is the multiplication of two precision- p values a and b , then x can be represented exactly as a precision- $2p$ value, meaning that $r = p$ random bits suffice. The situation is a bit different when $x = a \pm b$, since after normalization, x can have a significand with up to $e_{\max} - e_{\min} + p - 1$ bits, meaning that at most $r = e_{\max} - e_{\min}$ random bits are needed, which is indeed too large for common-use formats such as those specified in the IEEE 754 standard. If $x = a/b$, then x can have a non-terminating binary representation, meaning that we need an infinite number of random bits to perform SR_p . More generally, if x is a precision- q number with $q > p$, then it suffices to take $r = q - p$.

In actual hardware, the choice of r varies among architectures. Whereas the Tesla D1, AMD MI300 GPU, and Graphcore Colossus IPU take r to be the minimal value needed to implement

SR down-conversion exactly (i.e., it respects Definition 3 when the input is a `binary32` value), for the NVIDIA Blackwell architecture this is the case only for down-conversion to 16-bit formats. For 8, 6 and 4-bit down-conversion, a value of $r = 16$ seems to always be used (see [30, sect. 9.7.9.21, Figs. 38–40]), which is not sufficient to respect Definition 3. There does not seem to be any public information regarding what value of r is used in the Amazon Trainium chips.

3 Properties of limited-precision stochastic rounding

In this section, we demonstrate various properties of $\text{SR}_{p,r}$ and, in particular, show that the mean independence property fails to hold. However, we prove a theorem that bounds the bias of the accumulated errors on algorithms with this rounding mode. Furthermore, using a technique similar to the one introduced in [12], we present a model that allows us to compute probabilistic error bounds for algorithms using $\text{SR}_{p,r}$.

Unlike SR_p , the limited-precision $\text{SR}_{p,r}$ operator is biased, since

$$\begin{aligned}\mathbb{E}(\text{SR}_{p,r}(x)) &= q_r(x)\lceil x \rceil_p + (1 - q_r(x))\lfloor x \rfloor_p \\ &= q_r(x)(\lceil x \rceil_p - \lfloor x \rfloor_p) + \lfloor x \rfloor_p \\ &= \text{fl}_{p+r}(x).\end{aligned}$$

Consequently, the bias $\mathbb{E}(\text{SR}_{p,r}(x) - x)$ is given by $\text{fl}_{p+r}(x) - x$. Let δ be such that $\text{SR}_{p,r}(x) = x(1 + \delta)$. Then

$$\mathbb{E}(\delta) = \frac{\text{fl}_{p+r}(x) - x}{x} = \beta. \quad (7)$$

Remark 4. With each error δ obtained with $\text{SR}_{p,r}$, we have a corresponding deterministic error β that arises from the computation of the function fl_{p+r} . Moreover, for $x \notin \mathcal{F}$ and a large number of random bits r , $\text{fl}_{p+r}(x)$ tends to x , which means that the rounding operator $\text{SR}_{p,r}$ will become less and less biased as r increases.

We now turn to bound the variance of the error in computation. Let x be the result of an elementary operation, and let $\text{SR}_{p,r}(x) = x(1 + \delta)$ that value rounded with $\text{SR}_{p,r}$. If $x \in \mathcal{F}$, $\delta = 0$, $\mathbb{E}(\text{SR}_{p,r}(x)) = x$, and $\mathbb{V}(x) = 0$. If not, $\mathbb{E}(\text{SR}_{p,r}(x)) = \text{fl}_{p+r}(x)$, and

$$\begin{aligned}\mathbb{V}(\text{SR}_{p,r}(x)) &= \mathbb{E}(\text{SR}_{p,r}(x)^2) - \text{fl}_{p+r}(x)^2 = \lceil x \rceil_p^2 q_r(x) + \lfloor x \rfloor_p^2 (1 - q_r(x)) - \text{fl}_{p+r}(x)^2 \\ &= q_r(x)(\lceil x \rceil_p^2 - \lfloor x \rfloor_p^2) - (\text{fl}_{p+r}(x)^2 - \lfloor x \rfloor_p^2) \\ &= q_r(x)\text{ulp}_p(x)(\lceil x \rceil_p + \lfloor x \rfloor_p) - (\text{fl}_{p+r}(x) - \lfloor x \rfloor_p)(\text{fl}_{p+r}(x) + \lfloor x \rfloor_p) \\ &= q_r(x)\text{ulp}_p(x)(\lceil x \rceil_p + \lfloor x \rfloor_p) - q_r(x)\text{ulp}_p(x)(\text{fl}_{p+r}(x) + \lfloor x \rfloor_p) \\ &= q_r(x)\text{ulp}_p(x)(\lceil x \rceil_p - \text{fl}_{p+r}(x)) \\ &= \text{ulp}_p(x)^2 q_r(x)(1 - q_r(x)).\end{aligned}$$

Using (3), we obtain $\mathbb{V}(\text{SR}_{p,r}(x)) \leq x^2 \frac{u_p^2}{4}$, and

$$\mathbb{V}(\delta) = \mathbb{V}\left(\frac{\text{SR}_{p,r}(x) - x}{x}\right) \leq \frac{u_p^2}{4}. \quad (8)$$

The bound (8) is analogous to that for SR_p discussed in [11, sect. 3] and, interestingly, does not depend on the number of random bits r used.

Let us consider an execution of k elementary operations performed in floating-point arithmetic using $\text{SR}_{p,r}$. The following result shows that the mean independence property, which holds for SR_p and is shown in [5, lem. 5.2], does not extend to $\text{SR}_{p,r}$. Moreover, the conditional expectation of the k th rounding error generated by an operation $a \text{ op } b$ using $\text{SR}_{p,r}$ is equal to the k th error produced when applying the rounding function fl_{p+r} to the infinitely precise result c of the operation $a \text{ op } b$.

Lemma 3. *Let a and b be the result of $k - 1$ scalar operations performed using limited-precision SR, and let $\delta_1, \dots, \delta_{k-1}$ be the rounding errors these operations have produced. Let $c = a \text{ op } b$ for $\text{op} \in \{+, -, \times, \div\}$, let β_k be the deterministic relative rounding error defined by $\text{fl}_{p+r}(c) = c(1 + \beta_k)$ corresponding to the truncation of c to $p+r$ bits of precision, and let $\hat{c} = \text{SR}_{p,r}(c) = (a \text{ op } b)(1 + \delta_k)$. Then the random variable δ_k satisfies $\mathbb{E}(\delta_k \mid \delta_1, \dots, \delta_{k-1}) = \beta_k$.*

Proof. The proof is analogous to that of [5, Lem. 5.2]: the only difference is in the use of the probability $q_r(x)$ in (6) instead of $q(x)$. If $c = 0$, δ_k and β_k are equal to 0 and the result is immediate. Otherwise, the rounding error $\delta_k = (\hat{c} - c)/c$ is a random variable that depends on $\delta_1, \dots, \delta_{k-1}$ and has distribution

$$\delta_k = \begin{cases} (\lceil c \rceil_p - c)/c, & \text{with probability } q_r(c), \\ (\lfloor c \rfloor_p - c)/c, & \text{with probability } 1 - q_r(c), \end{cases} \quad q_r(c) = \frac{\text{fl}_{p+r}(c) - \lfloor c \rfloor_p}{\text{ulp}_p(c)}.$$

Moreover, $(\lceil c \rceil_p - c)/c$ and $(\lfloor c \rfloor_p - c)/c$ are themselves random variables that are determined by $\delta_1, \dots, \delta_{k-1}$, which implies that

$$\mathbb{E}\left(\frac{\lceil c \rceil_p - c}{c} \mid \delta_1, \dots, \delta_{k-1}\right) = \frac{\lceil c \rceil_p - c}{c}, \quad \mathbb{E}\left(\frac{\lfloor c \rfloor_p - c}{c} \mid \delta_1, \dots, \delta_{k-1}\right) = \frac{\lfloor c \rfloor_p - c}{c}.$$

Therefore, we obtain

$$\begin{aligned} \mathbb{E}(\delta_k \mid \delta_1, \dots, \delta_{k-1}) &= q_r(c) \left(\frac{\lceil c \rceil_p - c}{c}\right) + (1 - q_r(c)) \left(\frac{\lfloor c \rfloor_p - c}{c}\right) \\ &= q_r(c) \left(\frac{\lceil c \rceil_p - c}{c} - \frac{\lfloor c \rfloor_p - c}{c}\right) + \frac{\lfloor c \rfloor_p - c}{c} \\ &= q_r(c) \left(\frac{\lceil c \rceil_p - \lfloor c \rfloor_p}{c}\right) + \frac{\lfloor c \rfloor_p - c}{c} \\ &= \frac{\text{fl}_{p+r}(c) - \lfloor c \rfloor_p}{c} + \frac{\lfloor c \rfloor_p - c}{c} \\ &= \frac{\text{fl}_{p+r}(c) - c}{c} = \beta_k. \end{aligned}$$

□

Remark 5. From Remark 4, if we assume that we have a large number of random bits r , the β_1, \dots, β_k will be equal to zero, and the $\delta_1, \dots, \delta_k$ satisfy the mean independence property, which implies that the limited-precision SR operator $\text{SR}_{p,r}$ coincides with the SR operator SR_p .

Remark 6. Note that, for $k \geq 2$, the value of β_k depends on that of the random variables $\delta_1, \dots, \delta_{k-1}$. Therefore, β_k is a random variable despite having been produced by a deterministic rounding mode. From Lemma 3, we have

$$\mathbb{E}(\beta_k) = \mathbb{E}(\mathbb{E}(\delta_k \mid \delta_1, \dots, \delta_{k-1})) = \mathbb{E}(\delta_k). \quad (9)$$

Remark 7. Equation (9) shows that under $\text{SR}_{p,r}$, the rounding error of an operation, δ , and the truncation error applied for this operation, β , have the same expected value.

The following theorem helps in establishing lower and upper bounds on the expectation of the accumulated error $\prod_{k=1}^n (1 + \delta_k)$ in an algorithm, where $\delta_1, \delta_2, \dots, \delta_n$ are random errors obtained using $\text{SR}_{p,r}$.

Theorem 1. Let $\delta_1, \delta_2, \dots, \delta_n$ and $\beta_1, \beta_2, \dots, \beta_n$ be random variables such that $\mathbb{E}(\delta_1) = \mathbb{E}(\beta_1)$ and for all $2 \leq k \leq n$, $\mathbb{E}(\delta_k \mid \delta_1, \dots, \delta_{k-1}) = \beta_k$. Suppose there exists a constant $B > 0$ such that $|\beta_k| \leq B$ for all $1 \leq k \leq n$. Then

$$(1 - B)^n \leq \mathbb{E} \left(\prod_{k=1}^n (1 + \delta_k) \right) \leq (1 + B)^n.$$

Proof. The proof is by induction on n . For $n = 1$, we have $\mathbb{E}(1 + \delta_1) = 1 + \mathbb{E}(\delta_1) = 1 + \mathbb{E}(\beta_1)$, and

$$(1 - B) \leq 1 + \beta_1 \leq (1 + B).$$

For the inductive step, let $Q_n = \prod_{k=1}^n (1 + \delta_k)$, and assume that the inequalities hold for Q_{n-1} . Since $Q_n = Q_{n-1}(1 + \delta_n)$ and $\mathbb{E}(\delta_n \mid \delta_1, \dots, \delta_{n-1}) = \beta_n$, the law of total expectation $\mathbb{E}(X) = \mathbb{E}(\mathbb{E}(X \mid Y))$ gives

$$\begin{aligned} \mathbb{E}(Q_n) &= \mathbb{E}(Q_{n-1}(1 + \delta_n)) \\ &= \mathbb{E}(\mathbb{E}(Q_{n-1}(1 + \delta_n) \mid \delta_1, \dots, \delta_{n-1})) \\ &= \mathbb{E}(Q_{n-1} \mathbb{E}((1 + \delta_n) \mid \delta_1, \dots, \delta_{n-1})) \\ &= \mathbb{E}(Q_{n-1}(1 + \beta_n)). \end{aligned}$$

Moreover

$$\begin{cases} \mathbb{E}(Q_{n-1}(1 + \beta_n)) \leq \mathbb{E}(Q_{n-1})(1 + B) \leq (1 + B)^n, \\ \mathbb{E}(Q_{n-1}(1 + \beta_n)) \geq \mathbb{E}(Q_{n-1})(1 - B) \geq (1 - B)^n. \end{cases}$$

We thus have

$$(1 - B)^n \leq \mathbb{E} \left(\prod_{k=1}^n (1 + \delta_k) \right) \leq (1 + B)^n.$$

□

Remark 8. Let $\delta_1, \delta_2, \dots, \delta_n$ and $\beta_1, \beta_2, \dots, \beta_n$ as in lemma 3. Then, by applying theorem 1 we have:

$$(1 - u_{p+r})^n \leq \mathbb{E} \left(\prod_{k=1}^n (1 + \delta_k) \right) \leq (1 + u_{p+r})^n. \quad (10)$$

Remark 9. When using SR_p , [5, Lem 6.1] shows that $\mathbb{E} \left(\prod_{k=1}^n (1 + \delta_k) \right) = 1$. For $\text{SR}_{p,r}$, taking $r \rightarrow \infty$ in (10) implies that $u_{p+r} \rightarrow 0$, which allows us to recover the same result as SR_p .

In order to prove our main result, we need a technical lemma to rewrite a product of binomials as a sum of monomials. In the result, we denote by $\mathcal{P}(S)$ the power set of the set S , that is, the set whose elements are all possible subsets of S . We recall that $\mathcal{P}(S)$ can be defined recursively, since

$$\begin{aligned} \mathcal{P}(\emptyset) &= \{\emptyset\}, \\ \mathcal{P}(S \cup \{\ell\}) &= \{T \cup \{\ell\} : T \in \mathcal{P}(S)\} \cup \mathcal{P}(S). \end{aligned} \quad (11)$$

Lemma 4. Let $m, n \in \mathbb{N}$, with $m \leq n$, let $\mathcal{I} = \{k \in \mathbb{N} : m \leq k \leq n\}$, and let $x_k, y_k \in \mathbb{R}$, for all $k \in \mathcal{I}$. Then, we have

$$\prod_{k=m}^n (x_k + y_k) = \prod_{k=m}^n x_k + \sum_{\substack{K \in \mathcal{P}(\mathcal{I}) \\ K \neq \mathcal{I}}} \left(\prod_{i \in K} x_i \prod_{j \in \mathcal{I} \setminus K} y_j \right). \quad (12)$$

Proof. First, we prove by induction on the size of the index set \mathcal{I} that

$$\prod_{k=m}^n (x_k + y_k) = \sum_{K \in \mathcal{P}(\mathcal{I})} \left(\prod_{i \in K} x_i \prod_{j \in \mathcal{I} \setminus K} y_j \right), \quad \mathcal{I} = \{k \in \mathbb{N} : m \leq k \leq n\}. \quad (13)$$

For $K = \emptyset$, we use the fact that $\prod_{i \in K} x_i = 1$. It is easy to verify that (13) holds for $m = n$. For the inductive step, assume that (13) holds. Without loss of generality, we extend the set by incrementing n , but one could prove the result analogously by decrementing m . We have

$$\begin{aligned} \prod_{k=m}^{n+1} (x_k + y_k) &= (x_{n+1} + y_{n+1}) \prod_{k=m}^n (x_k + y_k) \\ &= (x_{n+1} + y_{n+1}) \sum_{K \in \mathcal{P}(\mathcal{I})} \left(\prod_{i \in K} x_i \prod_{j \in \mathcal{I} \setminus K} y_j \right), \quad \mathcal{I} = \{k \in \mathbb{N} : m \leq k \leq n\} \\ &= x_{n+1} \sum_{K \in \mathcal{P}(\mathcal{I})} \left(\prod_{i \in K} x_i \prod_{j \in \mathcal{I} \setminus K} y_j \right) + y_{n+1} \sum_{K \in \mathcal{P}(\mathcal{I})} \left(\prod_{i \in K} x_i \prod_{j \in \mathcal{I} \setminus K} y_j \right) \\ &= \sum_{K \in \mathcal{P}(\mathcal{I})} \left(\prod_{i \in K \cup \{n+1\}} x_i \prod_{j \in \mathcal{I} \setminus K} y_j \right) + \sum_{K \in \mathcal{P}(\mathcal{I})} \left(\prod_{i \in K} x_i \prod_{j \in \mathcal{I} \cup \{n+1\} \setminus K} y_j \right). \end{aligned}$$

Note that the sum on the left is over the elements of $\{T \cup \{n+1\} : T \in \mathcal{P}(\mathcal{I})\}$, and that on the right is over the elements of \mathcal{P} . Using the inductive step of (11) with $\ell = n+1$, we obtain

$$\prod_{k=m}^{n+1} (x_k + y_k) = \sum_{K \in \mathcal{P}(\mathcal{I}')} \left(\prod_{i \in K} x_i \prod_{j \in \mathcal{I}' \setminus K} y_j \right), \quad \mathcal{I}' = \{k \in \mathbb{N} : m \leq k \leq n+1\}.$$

This establishes (13) for \mathcal{I}' . To obtain (12) from (13), it suffices to extract the term $K = \mathcal{I}$, which corresponds to the product

$$\prod_{k=m}^n x_k.$$

□

Equation (7) shows that $\text{SR}_{p,r}$ is biased and Lemma 3 demonstrates that $\text{SR}_{p,r}$ does not satisfy the mean independence property. This complicates the direct theoretical analysis of $\text{SR}_{p,r}$, since unbiasedness and mean independence are principal tools used to derive error bounds for algorithms using SR. However, it has been shown [12] that, for some algorithms, even when the use of SR leads to a biased result, it is possible to obtain probabilistic bounds in $\mathcal{O}(\sqrt{n}u_p)$. Building on similar techniques, we propose the following general model to study algorithms under $\text{SR}_{p,r}$.

Lemma 5. *Let $\delta_1, \delta_2, \dots, \delta_n$ be random errors produced by a sequence of elementary operations using $\text{SR}_{p,r}$, and let $\beta_1, \beta_2, \dots, \beta_n$ be their corresponding errors incurred by fl_{p+r} . Then, the random variables $\alpha_k = \delta_k - \beta_k$ for $1 \leq k \leq n$ are mean independent, that is to say,*

$$\mathbb{E}(\alpha_k \mid \alpha_1, \dots, \alpha_{k-1}) = \mathbb{E}(\alpha_k) = 0.$$

Moreover, for all $1 \leq i \leq n$,

$$\prod_{k=i}^n (1 + \delta_k) = \prod_{k=i}^n (1 + \alpha_k) + \mathcal{B}_i, \quad (14)$$

where $\mathcal{B}_i = \sum_{\substack{K \in \mathcal{P}(\mathcal{I}_i) \\ K \neq \mathcal{I}_i}} \left(\prod_{i \in K} (1 + \alpha_i) \prod_{j \in \mathcal{I}_i \setminus K} \beta_j \right)$ and $\mathcal{I}_i = \{k \in \mathbb{N} : i \leq k \leq n\}$ such that

$$|\mathcal{B}_i| \leq \gamma_{n-i+1}(u_p + u_{p+r}) - \gamma_{n-i+1}(u_p), \quad (15)$$

where $\gamma_m(x) = (1 + x)^m - 1$.

Proof. From (9), we can conclude that $\mathbb{E}(\alpha_k) = \mathbb{E}(\delta_k) - \mathbb{E}(\beta_k) = 0$. Moreover, β_k is entirely determined by $\delta_1, \dots, \delta_{k-1}$, in particular $\alpha_1, \dots, \alpha_{k-1}$. Lemma 3 shows that $\mathbb{E}(\delta_k | \delta_1, \dots, \delta_{k-1}) = \beta_k$, in particular, $\mathbb{E}(\delta_k | \alpha_1, \dots, \alpha_{k-1})$ is also given by β_k . We thus have

$$\begin{aligned} \mathbb{E}(\alpha_k | \alpha_1, \dots, \alpha_{k-1}) &= \mathbb{E}(\delta_k - \beta_k | \alpha_1, \dots, \alpha_{k-1}) \\ &= \mathbb{E}(\delta_k | \alpha_1, \dots, \alpha_{k-1}) - \mathbb{E}(\beta_k | \alpha_1, \dots, \alpha_{k-1}) \\ &= \beta_k - \beta_k = 0. \end{aligned}$$

By applying Lemma 4 with $x_k = 1 + \alpha_k$ and $y_k = \beta_k$ we obtain

$$\begin{aligned} \prod_{k=i}^n (1 + \alpha_k + \beta_k) &= \prod_{k=i}^n (1 + \alpha_k) + \sum_{\substack{K \in \mathcal{P}(\mathcal{I}_i) \\ K \neq \mathcal{I}_i}} \left(\prod_{i \in K} (1 + \alpha_i) \prod_{j \in \mathcal{I}_i \setminus K} \beta_j \right) \\ &= \prod_{k=i}^n (1 + \alpha_k) + \mathcal{B}_i. \end{aligned}$$

Moreover

$$\begin{aligned} |\mathcal{B}_i| &= \sum_{\substack{K \in \mathcal{P}(\mathcal{I}_i) \\ K \neq \mathcal{I}_i}} \left(\prod_{i \in K} (1 + \alpha_i) \prod_{j \in \mathcal{I}_i \setminus K} \beta_j \right) \\ &\leq \sum_{\substack{K \in \mathcal{P}(\mathcal{I}_i) \\ K \neq \mathcal{I}_i}} (1 + u_p)^{|K|} u_{p+r}^{n-i+1-|K|} && \text{because } |\alpha_k| \leq u_p \text{ and } |\beta_k| \leq u_{p+r} \\ &= \sum_{k=1}^{n-i+1} \binom{n-i+1}{k} (1 + u_p)^k u_{p+r}^{n-i+1-k} \\ &= (1 + u_p + u_{p+r})^{n-i+1} - (1 + u_p)^{n-i+1} \\ &= \gamma_{n-i+1}(u_p + u_{p+r}) - \gamma_{n-i+1}(u_p). \end{aligned}$$

□

4 Error analysis of algorithms with limited-precision SR

We are now ready to apply our results on limited-precision SR in order to analyze two common algorithms, recursive summation and inner product of vectors of floating-point numbers. We use Theorem 1 to compute bounds on the biases of these algorithms. Furthermore, we use Lemma 5 to compute probabilistic bounds of the relative errors of these algorithms under $\text{SR}_{p,r}$. These bounds are established using two methods:

- martingales (Definition 2) and the Azuma–Hoeffding ineq. (Lemma 2).
- the variance bound proved in [11, Lem. 3.1] and the Bienaymé–Chebyshev ineq. (Lemma 1).

We show that these bounds are proportional to $\sqrt{n}u_p + nu_{p+r}$.

An important observation is that all theorems proved in this section naturally converge to their established counterparts in the literature. More precisely, by letting r tend to ∞ in our bounds, $u_{p+r} \rightarrow 0$, which recovers the results previously established for SR_p , confirming the consistency of our study with existing theory.

4.1 Recursive summation

Let $a \in \mathbb{R}^n$. We will now perform a roundoff error analysis for recursively computing the sum

$$y = \sum_{i=1}^n a_i \quad (16)$$

using precision- p floating-point arithmetic with $\text{SR}_{p,r}$. In our analysis, we will rely on the condition number of the sum in (16), defined by

$$\kappa(a) = \frac{\sum_{i=1}^n |a_i|}{|\sum_{i=1}^n a_i|}, \quad (17)$$

and on the error function

$$\gamma_n(u) = (1+u)^n - 1 = nu + \mathcal{O}(u^2) \text{ for } nu \ll 1. \quad (18)$$

If we denote $\hat{s}_k = \text{SR}_{p,r}(\hat{s}_{k-1} + a_k)$ for $k = 2, \dots, n$, we have

Limited-precision SR	Exact computation
$\hat{s}_1 = a_1$	$s_1 = a_1$
$\hat{s}_2 = (\hat{s}_1 + a_2)(1 + \delta_1)$	$s_2 = s_1 + a_2$
$\hat{s}_k = (\hat{s}_{k-1} + a_k)(1 + \delta_{k-1})$	$s_k = s_{k-1} + a_k$
$\hat{y} = \hat{s}_n$	$y = s_n$

It follows that

$$\hat{y} = \sum_{i=1}^n \left(a_i \prod_{k=\max\{i,2\}}^n (1 + \delta_{k-1}) \right). \quad (19)$$

Note that for $2 \leq k \leq n$, one has $|\delta_{k-1}| \leq u_p$ and $|\mathbb{E}(\delta_{k-1})| = |\mathbb{E}(\beta_{k-1})| \leq u_{p+r}$, where β_{k-1} is defined analogously to (7) as

$$\beta_{k-1} = \frac{\text{fl}_{p+r}(\hat{s}_{k-1} + a_k) - (\hat{s}_{k-1} + a_k)}{\hat{s}_{k-1} + a_k}. \quad (20)$$

In the following theorem, we give a bound on the summation bias in Equation (19).

Theorem 2. *The quantity \hat{y} in (19) satisfies*

$$\frac{|\mathbb{E}(\hat{y}) - y|}{|y|} \leq \kappa(a) \gamma_{n-1}(u_{p+r}), \quad (21)$$

where $\kappa(a)$ and $\gamma_n(u_{p+r})$ are defined in (17) and (18), respectively.

Proof. From (19), we have

$$\begin{aligned}
 |\mathbb{E}(\hat{y}) - y| &= \left| \mathbb{E} \left(\sum_{i=1}^n a_i \prod_{k=\max\{i,2\}}^n (1 + \delta_{k-1}) \right) - \sum_{i=1}^n a_i \right| \\
 &= \left| \sum_{i=1}^n a_i \left(\mathbb{E} \left(\prod_{k=\max\{i,2\}}^n (1 + \delta_{k-1}) \right) - 1 \right) \right| && \text{by linearity} \\
 &\leq \sum_{i=1}^n |a_i| \left| \mathbb{E} \left(\prod_{k=\max\{i,2\}}^n (1 + \delta_{k-1}) \right) - 1 \right| && \text{by triangle inequality} \\
 &\leq \sum_{i=1}^n |a_i| \left((1 + u_{p+r})^{n-\max\{i,2\}+1} - 1 \right) && \text{by Theorem 1} \\
 &\leq \sum_{i=1}^n |a_i| \left((1 + u_{p+r})^{n-1} - 1 \right) \\
 &= \sum_{i=1}^n |a_i| \gamma_{n-1}(u_{p+r}).
 \end{aligned}$$

We thus have

$$\frac{|\mathbb{E}(\hat{y}) - y|}{|y|} \leq \kappa(a) \gamma_{n-1}(u_{p+r}).$$

□

We now turn to give a probabilistic bound for the relative error of the summation (19) under $\text{SR}_{p,r}$ using martingales and Lemma 2.

Theorem 3. *For any $0 < \lambda < 1$, the quantity \hat{y} in (19) satisfies*

$$\frac{|\hat{y} - y|}{|y|} \leq \kappa(a) \left(\sqrt{u_p \gamma_{2(n-1)}(u_p)} \sqrt{\ln(2/\lambda)} + \gamma_{n-1}(u_p + u_{p+r}) - \gamma_{n-1}(u_p) \right), \quad (22)$$

with probability at least $1 - \lambda$.

Proof. For all $1 \leq i \leq n$, denote $\mathcal{I}_i = \{k \in \mathbb{N} : \max\{i, 2\} \leq k \leq n\}$. Lemma 5 shows that the random variables $\alpha_1, \alpha_2, \dots, \alpha_{n-1}$, with $\alpha_j = \delta_j - \beta_j$ are mean independent and

$$\prod_{k=\max\{i,2\}}^n (1 + \delta_{k-1}) = \prod_{k=\max\{i,2\}}^n (1 + \alpha_{k-1} + \beta_{k-1}) = \prod_{k=\max\{i,2\}}^n (1 + \alpha_{k-1}) + \mathcal{B}_i.$$

We thus have

$$\begin{aligned}
 |\hat{y} - y| &= \left| \sum_{i=1}^n a_i \left(\prod_{k=\max\{i,2\}}^n (1 + \delta_{k-1}) - 1 \right) \right| \\
 &= \left| \sum_{i=1}^n a_i \left(\prod_{k=\max\{i,2\}}^n (1 + \alpha_{k-1}) + \mathcal{B}_i - 1 \right) \right| && (23) \\
 &\leq \left| \sum_{i=1}^n a_i \left(\prod_{k=\max\{i,2\}}^n (1 + \alpha_{k-1}) - 1 \right) \right| + \left| \sum_{i=1}^n a_i \mathcal{B}_i \right| \\
 &= |M| + |A|,
 \end{aligned}$$

where $M = \sum_{i=1}^n a_i (\prod_{k=\max\{i,2\}}^n (1 + \alpha_{k-1}) - 1)$, and $A = \sum_{i=1}^n a_i \mathcal{B}_i$. Since the α_{k-1} are mean independent and bounded in magnitude by u_p , M can form a martingale as previously established in the literature [11, 25]. We now obtain a probabilistic bound on this martingale that is proportional to $\sqrt{n}u_p$. We will only outline our chosen construction here, referring the reader to [9, chap. 4] for a summary of possible methods.

By Lemma 2, we obtain the bound

$$|M| \leq \left| \sum_{i=1}^n a_i \right| \sqrt{u_p \gamma_{2(n-1)}(u_p)} \sqrt{\ln(2/\lambda)}, \quad (24)$$

which holds with probability at least $1 - \lambda$. Let us bound the second term in (23). For all $1 \leq i \leq n$, because $\mathcal{I}_i \subset \mathcal{I}_2$, $|\mathcal{B}_i| \leq |\mathcal{B}_2|$ and Lemma 5 shows that

$$|\mathcal{B}_i| \leq |\mathcal{B}_2| \leq \gamma_{n-1}(u_p + u_{p+r}) - \gamma_{n-1}(u_p).$$

It follows that

$$|A| = \left| \sum_{i=1}^n a_i \mathcal{B}_i \right| \leq \sum_{i=1}^n |a_i| |\mathcal{B}_i| \leq \left(\sum_{i=1}^n |a_i| \right) (\gamma_{n-1}(u_p + u_{p+r}) - \gamma_{n-1}(u_p)). \quad (25)$$

Therefore, the inequality

$$\frac{|\hat{y} - y|}{|y|} \leq \kappa(a) \left(\sqrt{u_p \gamma_{2(n-1)}(u_p)} \sqrt{\ln(2/\lambda)} + \gamma_{n-1}(u_p + u_{p+r}) - \gamma_{n-1}(u_p) \right)$$

holds with probability at least $1 - \lambda$. □

In the following, we give a probabilistic bound on the relative error of the summation (19) under $\text{SR}_{p,r}$ using the bound on the variance proposed in [11, Lem. 3.1] and Lemma 1.

Theorem 4. *For any $0 < \lambda < 1$, the quantity \hat{y} in (19) satisfies*

$$\frac{|\hat{y} - y|}{|y|} \leq \kappa(a) \left(\sqrt{\gamma_{n-1}(u_p^2)/\lambda} + \gamma_{n-1}(u_p + u_{p+r}) - \gamma_{n-1}(u_p) \right), \quad (26)$$

with probability at least $1 - \lambda$.

Proof. This proof relies on (23). Since α_{k-1} are mean independent and satisfy $|\alpha_{k-1}| \leq u_p$, [11, Lem. 3.1] and Lemma 1 show that

$$|M| \leq \left(\sum_{i=1}^n |a_i| \right) \sqrt{\gamma_{n-1}(u_p^2)/\lambda} \quad (27)$$

holds with probability at least $1 - \lambda$. Using (25), we conclude that

$$\frac{|\hat{y} - y|}{|y|} \leq \kappa(a) \left(\sqrt{\gamma_{n-1}(u_p^2)/\lambda} + \gamma_{n-1}(u_p + u_{p+r}) - \gamma_{n-1}(u_p) \right),$$

with probability at least $1 - \lambda$. □

4.2 Computation of inner products

Let $a, b \in \mathbb{R}^n$. We will now extend the analysis of summation to the computation of the inner product

$$y = \sum_{i=1}^n a_i b_i. \quad (28)$$

Our analysis will rely on the condition number of the inner product $\kappa(a \circ b)$, where $\kappa(\cdot)$ is defined in (17) and $a \circ b = (a_1 b_1, \dots, a_n b_n)$ denotes the Hadamard product. If we denote

$$\begin{cases} \widehat{s}_{2k-1} = \widehat{s}_{2k-2} + \text{SR}_{p,r}(a_k b_k), \\ \widehat{s}_{2k} = \text{SR}_{p,r}(\widehat{s}_{2k-1}), \end{cases} \quad \text{for } k = 1, \dots, n,$$

we have

Limited-precision SR	Exact computation
$\widehat{s}_1 = a_1 b_1$	$s_1 = a_1 b_1$
$\widehat{s}_2 = \widehat{s}_1(1 + \delta_1)$	$s_2 = s_1$
$\widehat{s}_{2k-1} = \widehat{s}_{2k-2} + a_k b_k(1 + \delta_{2k-2})$	$s_{2k-1} = s_{2k-2} + a_k b_k$
$\widehat{s}_{2k} = \widehat{s}_{2k-1}(1 + \delta_{2k-1})$	$s_{2k} = s_{2k-1}$
$\widehat{y} = \widehat{s}_{2n}$	$y = s_{2n}$

It follows that

$$\widehat{y} = \sum_{i=1}^n a_i b_i (1 + \delta_{2i-1}) \prod_{k=i}^n (1 + \delta_{2(k-1)}), \quad (29)$$

with $\delta_0 = 0$. Note that for $1 \leq k \leq 2n - 1$, one has $|\delta_k| \leq u_p$ and $|\mathbb{E}(\delta_k)| = |\mathbb{E}(\beta_k)| \leq u_{p+r}$, where β_k is defined analogously to (7) as

$$\beta_{2k-1} = \frac{\mathfrak{fl}_{p+r}(\widehat{s}_{2k}) - (\widehat{s}_{2k-1})}{\widehat{s}_{2k-1}} \quad \text{and} \quad \beta_{2k-2} = \frac{\mathfrak{fl}_{p+r}(\widehat{s}_{2k-1}) - (\widehat{s}_{2k-2} + a_k b_k)}{\widehat{s}_{2k-2} + a_k b_k}. \quad (30)$$

In the following theorem, we give a bound on the inner product bias in Equation (29).

Theorem 5. *The quantity \widehat{y} in Equation (29) satisfies*

$$\frac{|\mathbb{E}(\widehat{y}) - y|}{|y|} \leq \kappa(a \circ b) \gamma_n(u_{p+r}), \quad (31)$$

where $\kappa(a \circ b)$ and $\gamma_n(u_{p+r})$ are defined in (17) and (18), respectively.

Proof. The proof follows the same structure as that of Theorem 2, with one additional error term resulting from the multiplication $a_i b_i$. \square

The following two theorems give probabilistic bounds on the relative error of the inner product (29) under $\text{SR}_{p,r}$ using martingales and Lemma 2 and the bound on the variance proposed in [11, Lem. 3.1] and Lemma 1, respectively.

Theorem 6. *For any $0 < \lambda < 1$, the quantity \widehat{y} in (29) satisfies*

$$\frac{|\widehat{y} - y|}{|y|} \leq \kappa(a \circ b) \left(\sqrt{u_p \gamma_{2n}(u_p)} \sqrt{\ln(2/\lambda)} + \gamma_n(u_p + u_{p+r}) - \gamma_n(u_p) \right), \quad (32)$$

with probability at least $1 - \lambda$.

Proof. The proof follows the same structure as that of Theorem 3, with one additional error term resulting from the multiplication $a_i b_i$. \square

Theorem 7. For any $0 < \lambda < 1$, the quantity \hat{y} in (29) satisfies

$$\frac{|\hat{y} - y|}{|y|} \leq \kappa(a \circ b) \left(\sqrt{\gamma_n(u_p^2)/\lambda} + \gamma_n(u_p + u_{p+r}) - \gamma_n(u_p) \right), \quad (33)$$

with probability at least $1 - \lambda$.

Proof. The proof follows the same structure as that of Theorem 4, with one additional error term resulting from the multiplication $a_i b_i$. \square

Remark 10. We have shown the applicability of Lemma 5 to summation and inner product. These results should hold for all algorithms considered for SR_p that satisfy a probabilistic bound in terms of $\mathcal{O}(\sqrt{n}u_p)$, in particular, algorithms with multi-linear errors [9, sect. 4]. We plan to explore this in future work.

5 Bound analysis

We analyze the bound in Theorem 6; the same analysis holds for Theorem 7 and the probabilistic bounds in Theorem 3 and Theorem 4.

Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be twice continuously differentiable at $(0, 0)$. By Taylor's theorem, we have

$$f(x, y) = f(0, 0) + \frac{\partial f}{\partial x}(0, 0)x + \frac{\partial f}{\partial y}(0, 0)y + \mathcal{O}(\|(x, y)\|_2). \quad (34)$$

By using (34), we obtain

$$\begin{aligned} f(u_p, u_{p+r}) &:= \gamma_n(u_p + u_{p+r}) = (1 + u_p + u_{p+r})^n - 1 \\ &= n(u_p + u_{p+r}) + \mathcal{O}(\|(u_p, u_{p+r})\|_2). \end{aligned} \quad (35)$$

Combining (35) with

$$\gamma_n(u_p + u_{p+r}) = (1 + u_p + u_{p+r})^n - 1 = n(u_p + u_{p+r}) + \mathcal{O}(\|(u_p, u_{p+r})\|_2)$$

and

$$\sqrt{u_p \gamma_{2n}(u_p)} \sqrt{\ln(2/\lambda)} = \sqrt{2n} \sqrt{\ln(2/\lambda)} u_p + \mathcal{O}(u_p^2),$$

we obtain

$$\begin{aligned} \frac{|\hat{y} - y|}{|y|} &\leq \kappa(a \circ b) \left(\sqrt{u_p \gamma_{2n}(u_p)} \sqrt{\ln(2/\lambda)} + \gamma_n(u_p + u_{p+r}) - \gamma_n(u_p) \right) \\ &= \kappa(a \circ b) \left(\sqrt{2n} \sqrt{\ln(2/\lambda)} u_p + n(u_p + u_{p+r}) - nu_p + \mathcal{O}(\|(u_p, u_{p+r})\|_2) \right) \\ &= \kappa(a \circ b) \left(\sqrt{2n} \sqrt{\ln(2/\lambda)} u_p + nu_{p+r} \right) + \mathcal{O}(\|(u_p, u_{p+r})\|_2). \end{aligned}$$

This result indicates that the bound on $\text{SR}_{p,r}$ with fixed λ is made up of two components: $\sqrt{n}u_p$, a probabilistic term that captures the random behavior of the algorithm, and nu_{p+r} , a deterministic one that captures the truncations performed in precision $p + r$. As the number of random bits r increases, the bound becomes tighter as the magnitude of u_{p+r} decreases. Moreover, the probabilistic bounds converge asymptotically to the probabilistic bound on SR_p , which is consistent with Remark 5.




FIGURE 2. *Left: relative error of RN and $SR_{11,r}$ in IEEE-754 binary16 arithmetic ($p = 11$) for the recursive summation of n floating-point numbers drawn from a uniform distribution between 0 and 1. For each value of n , the reported relative error for $SR_{11,r}$ is the average value over 500 runs. Right: comparison of deterministic bound and probabilistic bounds (Theorem 4) with the associated random bits.*

Remark 11. A good rule of thumb is to pick a value of r that ensures that the term containing $\sqrt{nu_p}$ is not dominated by the one containing nu_{p+r} . Simplifying the two terms leads to $r \geq \lceil (\log_2 n)/2 \rceil$, where $\lceil x \rceil$ is the smallest integer greater or equal than $x \in \mathbb{R}$. Since our lemma 5 should be applicable to all algorithms with multi-linear errors, we expect this rule of thumb to work for all such algorithms. Section 6 looks at various numerical examples where such a value of r is indeed a sensible choice.

6 Numerical experiments

We perform a set of numerical experiments by considering computations that are prone to stagnation with RN, and we investigate the effect that the value of r has when SR is used. The first two sets of experiments use the `srfloat` C++ library², also available with Python bindings, which simulates $SR_{p,r}$ arithmetic as described in Section 2.3. Internally, the software uses `binary64` arithmetic, and it can handle values of p up to 53 and values of r up to $53 - p$ for a chosen p . The neural network training example relies on the `mptorch`³ library, a PyTorch extension with similar $SR_{p,r}$ simulation functionality for deep learning computations.

6.1 Summation

We first look at the effect of r when recursively summing n floating-point numbers, for n between 2 and 6,000. The addends are sampled uniformly at random from the interval $[0, 1]$, and in this experiment we use `binary16` arithmetic, for which $p = 11$. The results are reported in Figure 2. The left panel shows relative errors when using RN and various $SR_{11,r}$ rounding modes. While RN is superior for smaller n , as soon as the running sum becomes sufficiently large, new addends will tend to get absorbed, leading to stagnation. The biased errors that are introduced this way accumulate, leading to a significant increase in the overall error. While performing worse than RN in the beginning, the $SR_{11,r}$ configurations are less affected by stagnation as n increases. Taking r

²<https://github.com/sfilip/srfloat>

³<https://github.com/mptorch/mptorch>




FIGURE 3. Convergence profiles for 5,000 iterations of gradient descent on the Rosenbrock function. The parameter updates are performed using **binary64** arithmetic with RN and **binary16** arithmetic with RN and $SR_{11,r}$. The starting value for the iteration is $\mathbf{x}_0 = (0, 0)$ for the profiles on the left and $\mathbf{x}_0 = (0.5, 0.5)$ for those on the right. For both experiments, we average each $SR_{11,r}$ error over 500 different runs, and the learning rate is $t_k = 0.001$.

close to $\lceil (\log_2 6,000)/2 \rceil = 7$ gives good results, with larger values of r not showing any significant improvements.

The figure on the right corroborates the findings presented in Section 5 and shows that the probabilistic bounds are tighter than deterministic ones. We use the bound obtained with the variance bound and the Bienaymé–Chebyshev inequality because, with probability 0.9, this method ensures tight bounds for a larger n . For more details on the comparison between the bounds in theorem 3 and Theorem 4, we refer the reader to [9, sect 4.3]. The effect of r is also evident in this figure.

6.2 Rosenbrock function

The Rosenbrock function is a non-convex function defined by

$$f(x_1, x_2) = (1 - x_1)^2 + 100(x_2 - x_1^2)^2,$$

with a global minimum of 0, occurring at $\mathbf{x}^* = (1, 1)$. This function is often used to assess the performance of mathematical optimization algorithms. In this experiment, we look at the convergence of gradient descent to this minimum for RN and $SR_{p,r}$. The update rule of gradient descent is

$$\mathbf{x}_{k+1} = \mathbf{x}_k - t_k \nabla f(\mathbf{x}_k),$$

where \mathbf{x}_k is the current point, t_k is the learning rate, and $\nabla f(\mathbf{x}_k)$ is the gradient of the function at \mathbf{x}_k . We assume gradient computation using **binary64** with RN (gradients are then downcast to **binary16**) and apply $SR_{p,r}$ only during the parameter update process.

Figure 3 illustrates the effect of r on the convergence of the first 5,000 iterations of gradient descent for two different starting values. In the left panel, the starting value for the iteration is $\mathbf{x}_0 = (0, 0)$, for which $f(\mathbf{x}_0) = 1$ and $\nabla f(\mathbf{x}_0) = (-2, 0)$. This choice of \mathbf{x}_0 is close to the narrow, flat valley where the minimum lies, but the point is not too close to the minimum itself, which guarantees the need for a large number of iterations. In **binary16** arithmetic, RN and $SR_{11,3}$ are both significantly affected by rounding errors, with RN stagnating rather early on. For larger values of r , the SR results are much better, matching those of **binary64** arithmetic with RN.

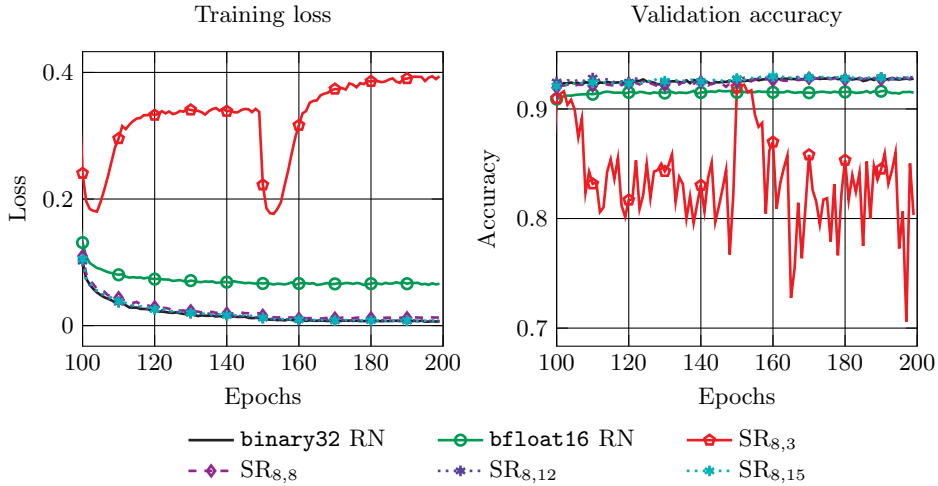


FIGURE 4. Evolution of training loss (left) and validation accuracy (right) for a ResNet32 image classification network on the CIFAR-10 dataset using various parameter and optimizer update configurations. In the baseline configuration, `binary32` arithmetic with RN is used for compute and the same format is used for storage. For the low-precision configurations, parameters are stored and updated using `bfloat16` arithmetic with either RN or $SR_{p,r}$.

In the right panel, the starting value for the iteration is $\mathbf{x}_0 = (0.5, 0.5)$, at which $f(\mathbf{x}_0) = 25.25$ and $\nabla f(\mathbf{x}_0) = (-50.5, 50)$. This choice of which \mathbf{x}_0 lies in a region with higher curvature but closer to the minimum. While all implementations we consider do make progress initially, as soon as the iterates get close to $(1, 1)$, `binary16` RN stagnates, whereas the limited-precision SR alternatives are not affected and continue to progress. As in the previous case, larger values of r recover the baseline `binary64` convergence profile.

In both cases, as soon as r is close to $\lceil \log_2(5,000)/2 \rceil = 7$ the accuracy improvement of $SR_{11,r}$ starts to plateau: the difference between $r = 6$ and $r = 7$ is somewhat visible, whereas $r = 7$ and $r = 8$ yield almost identical curves. This indicates that increasing r above 7 leads to diminishing returns.

6.3 Parameter update in deep neural network training

SR has found two main uses in modern deep learning scenarios: as a quantization procedure for network parameters and signals, in particular gradient signals [4, 19, 36], and to avoid stagnation during parameter updates [35, 37] when low-precision formats (16 and 8 bits) are used to compute and store them. The smaller memory footprint associated with lower precision formats is one of the key factors that has enabled the current boom of large language models.

Here, we focus on the second case and consider parameter updates during the training of a ResNet32 image classification model [20] on the CIFAR-10 [26] dataset. The training hyperparameters are identical to those used in [20, sect. 4.2]: minibatch gradient descent with 128 batch size, momentum set to 0.9, weight decay to 10^{-4} , 0.1 learning rate that gets divided by 10 after 32,000 and 48,000 iterations, and $n = 64,000$ iterations (200 epochs) of training. We use `bfloat16` arithmetic, for which $p = 8$, and the update rule for the network parameters \mathbf{x} in this configuration is

$$\begin{aligned} \mathbf{v}_{k+1} &= \circ(\mu \mathbf{v}_k + \mathbf{g}_k), \\ \mathbf{x}_{k+1} &= \circ(\mathbf{x}_k - t_k \mathbf{v}_{k+1}), \end{aligned}$$

where μ is the momentum term, \mathbf{v}_k is the velocity vector, \mathbf{g}_k is the gradient of the loss function

with respect to \mathbf{x}_k , and \circ is the rounding operator, which can round to `bfloat16` using either RN or $\text{SR}_{8,r}$, depending on the configuration.

The results of this experiment are reported in Figure 4. For `bfloat16`, the training loss is degraded when using RN or $\text{SR}_{8,r}$ with very small r , but larger values of r match the baseline `binary32` results on both the training and the validation datasets: the validation accuracy of `binary32` is 92.85% but goes down to 91.68% for `bfloat16` with RN. For `bfloat16` with $\text{SR}_{p,r}$, the choice $r = 3$ leads to an unstable validation accuracy that hovers around 83%, but using $\text{SR}_{8,r}$ with $r \geq \lceil \log_2(64,000)/2 \rceil = 8$ recovers baseline accuracy, with 92.86%, 92.93%, and 92.99% for $r = 8$, $r = 12$, and $r = 15$, respectively. Diminishing improvements are again visible as r is increased beyond this threshold.

7 Conclusions

SR has garnered significant attention in various domains [6], as it can deliver improved accuracy compared with the deterministic rounding modes in the IEEE 754 standard [24]. We have investigated the number of random bits required to implement SR effectively. We introduced a new rounding mode, limited-precision SR, denoted by $\text{SR}_{p,r}$, and we showed that the bias that this rounding mode introduces depends on the value of r used. In Lemma 5, we presented a model that facilitates the theoretical analysis of algorithms using $\text{SR}_{p,r}$. With this lemma, we can analyze all algorithms that were previously studied for classical SR. Applying this lemma to recursive summation and inner product computation, we derived probabilistic error bounds proportional to $\sqrt{nu_p} + nu_{p+r}$.

Our findings suggest that $\text{SR}_{p,r}$ becomes less biased as the number of random bits r increases, and that for a large enough value it converges to the theoretical properties of SR_p . On the other hand, using a large value of r may not always be practical or necessary, as the computational overhead associated with generating and processing a large number of random bits can be significant. Therefore, it is crucial to strike a balance between the desired accuracy and the computational resources available when implementing SR in applications. To this end, our bounds suggest that choosing a value of r that is close to $\lceil (\log_2 n)/2 \rceil$ offers the best compromise in practice for summation and dot product operations. This has been verified through several numerical experiments: recursive summation, gradient descent on the Rosenbrock function, and parameter updates in the training of deep learning models.

References

- [1] Advanced Micro Devices. AMD instinct MI300 instruction set architecture, June 2024.
- [2] S. B. Ali, S.-I. Filip, and O. Sentieys. A stochastic rounding-enabled low-precision floating-point MAC for DNN training. In *Proceedings of the 2024 Design, Automation & Test in Europe Conference & Exhibition*, 2024, pages 1–6.
- [3] P. Billingsley. *Probability and Measure*. Wiley Series in Probability and Statistics. Wiley, 1995.
- [4] B. Chmiel, R. Banner, E. Hoffer, H. B. Yaacov, and D. Soudry. Accurate neural training with 4-bit matrix multiplications at standard formats. arXiv:2112.10769 [cs.LG], December 2021. Revised June 2024.
- [5] M. P. Connolly, N. J. Higham, and T. Mary. Stochastic rounding and its probabilistic backward error analysis. *SIAM J. Sci. Comput.*, 43(1):A566–A585, 2021.

- [6] M. Croci, M. Fasi, N. J. Higham, T. Mary, and M. Mikaitis. Stochastic rounding: Implementation, error analysis and applications. *Roy. Soc. Open Sci.*, 9(3), 2022.
- [7] M. Croci and M. B. Giles. Effects of round-to-nearest and stochastic rounding in the numerical solution of t. *IMA J. Numer. Anal.*, 43(3):1358–1390, 2022.
- [8] G. Dexter, C. Boutsikas, L. Ma, I. C. Ipsen, and P. Drineas. Stochastic rounding implicitly regularizes tall-and-thin matrices. *arXiv preprint arXiv:2403.12278*, 2024.
- [9] E.-M. El Arar. *Stochastic Models for the Evaluation of Numerical Errors*. PhD thesis, Université Paris-Saclay, 2023.
- [10] E.-M. El Arar, D. Sohier, P. de Oliveira Castro, and E. Petit. The positive effects of stochastic rounding in numerical algorithms. In *Proceedings of the 29th IEEE Symposium on Computer Arithmetic*, Institute of Electrical and Electronics Engineers, September 2022, pages 58–65.
- [11] E.-M. El Arar, D. Sohier, P. de Oliveira Castro, and E. Petit. Stochastic rounding variance and probabilistic bounds: A new approach. *SIAM J. Sci. Comput.*, 45(5):C255–C275, 2023.
- [12] E.-M. El Arar, D. Sohier, P. de Olivera Castro, and E. Petit. Bounds on non-linear errors for variance computation with stochastic rounding. *SIAM J. Sci. Comput.*, 46(5):B579–B599, 2024.
- [13] H. Fan, H. Zhou, G. Huang, P. Raman, X. Fu, G. Gupta, D. Ram, Y. Wang, and J. Huan. HLAT: High-quality large language model pre-trained on AWS trainium. *arXiv:2404.10630 [cs.CL]*, 2024.
- [14] M. Fasi and M. Mikaitis. Algorithms for stochastically rounded elementary arithmetic operations in IEEE 7. *IEEE Trans. Emerg. Topics Comput.*, 9(3):1451–1466, 2021.
- [15] G. E. Forsythe. Round-off errors in numerical integration on automatic machinery. *Bull. Amer. Math. Soc.*, 56:55–64, 1950.
- [16] S. Graillat, F. Jézéquel, and R. Picot. Numerical validation of compensated summation algorithms with sto. *Electron. Notes Theor. Comput. Sci.*, 317:55–69, 2015.
- [17] S. Graillat, F. Jézéquel, and R. Picot. Numerical validation of compensated algorithms with stochastic arith. *Appl. Math. Comput.*, 329:339–363, 2018.
- [18] Graphcore Limited. Mixed-precision arithmetic for AI: A hardware perspective, March 2023.
- [19] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan. Deep learning with limited numerical precision. In *Proceedings of the 32nd International Conference on Machine Learning*, F. Bach and D. Blei, editors, volume 37 of *Proceedings of Machine Learning Research*, PMLR, July 2015, pages 1737–1746.
- [20] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, Institute of Electrical and Electronics Engineers, December 2016.
- [21] N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. 2nd edition, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2002. xxx+680 pp. ISBN 0-89871-521-0.

- [22] M. Höhfeld and S. E. Fahlman. Probabilistic rounding in neural network learning with limited precision. *Neurocomputing*, 4(6):291–299, 1992.
- [23] M. Hopkins, M. Mikaitis, D. R. Lester, and S. Furber. Stochastic rounding and reduced-precision fixed-point arithmetic for solving neural ordinary differential equations. *Philos. Trans. R. Soc. A*, 378(2166), 2020.
- [24] *IEEE Standard for Floating-Point Arithmetic, IEEE Std 754-2019 (revision of IEEE Std 754-2008)*. Institute of Electrical and Electronics Engineers, Piscataway, NJ, USA, July 2019. 82 pp. ISBN 978-0-7381-5752-8.
- [25] I. C. F. Ipsen and H. Zhou. Probabilistic error analysis for inner products. *SIAM J. Matrix Anal. Appl.*, 41(4):1726–1741, 2020.
- [26] A. Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Technical report, University of Toronto, April 2009.
- [27] P. Micikevicius, S. Oberman, P. Dubey, M. Cornea, A. Rodriguez, I. Bratt, R. Grisenthwaite, N. Jouppi, C. Chou, A. Huffman, M. Schulte, R. Wittig, D. Jani, and S. Deng. OCP 8-bit floating point specification (OFP8). Technical report, Open Compute Project, June 2023. Revision 1.0.
- [28] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, Cambridge, UK, January 2005. ISBN 9780511813603.
- [29] J.-M. Muller, N. Brunie, F. de Dinechin, C.-P. Jeannerod, M. Joldes, V. Lefèvre, G. Melquiond, N. Revol, and S. Torres. *Handbook of Floating-Point Arithmetic*. 2nd edition, Birkhäuser, 2018. ISBN 978-3-319-76525-9.
- [30] NVIDIA Corporation. CUDA PTX ISA Release 8.7, January 2025.
- [31] E. A. Paxton, M. Chantry, M. Klöwer, L. Saffin, and T. Palmer. Climate modeling in low precision: Effects of both deterministic and stochastic rounding. *J. Climate*, 35(4):1215–1229, 2022.
- [32] P. D. Stott. Monte Carlo arithmetic: Exploiting randomness in floating-point arithmetic. Technical Report CSD-970002, Computer Science Department, University of California, Los Angeles, March 1997. 86 pp.
- [33] Tesla Dojo Technology. A guide to Tesla’s configurable floating point formats & arithmetic.
- [34] J. Vignes. Discrete stochastic arithmetic for validating results of numerical software. *Numer. Algorithms*, 37(1–4):377–390, 2004.
- [35] N. Wang, J. Choi, D. Brand, C.-Y. Chen, and K. Gopalakrishnan. Training deep neural networks with 8-bit floating point numbers. In *Proceedings of the 32nd Conference on Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, Curran Associates, December 2018, pages 7686–7695.
- [36] S. Wu, G. Li, F. Chen, and L. Shi. Training and inference with integers in deep neural networks. In *Proceedings of the International Conference on Learning Representations*, April 2018.
- [37] P. Zamirai, J. Zhang, C. R. Aberger, and C. De Sa. Revisiting BFloat16 training. arXiv:2010.06192 [cs.LG], October 2020. Revised March 2021.