

Widespread Underestimation of Sensitivity in Differentially Private Libraries and How to Fix It

Sílvia Casacuberta* Michael Shoemate† Salil Vadhan‡ Connor Wagaman§

Harvard University

July 21, 2022

Abstract

We identify a new class of vulnerabilities in implementations of differential privacy. Specifically, they arise when computing basic statistics such as sums, thanks to discrepancies between the implemented arithmetic using finite data types (namely, ints or floats) and idealized arithmetic over the reals or integers. These discrepancies result in the *sensitivity* of the implemented statistics — how much one individual’s data can affect the result — to be much higher than the sensitivity we expect. Consequently, essentially all differential privacy libraries fail to introduce enough noise to hide individual-level information as required by differential privacy, and we show that this may be exploited in realistic attacks on differentially private query systems. In addition to presenting these vulnerabilities, we also provide a number of solutions, which modify or constrain the way in which the sum is implemented in order to recover the idealized or near-idealized bounds on sensitivity.

*Email: scasacu@cs.harvard.edu. Supported by the Harvard Program for Research in Science and Engineering (PRISE) and a grant from the Sloan Foundation.

†Email: mshoemate@cs.harvard.edu. Supported by a grant from the Sloan Foundation, a grant from the US Census Bureau, and gifts from Apple and Facebook.

‡Email: vadhan@cs.harvard.edu. Supported by a grant from the Sloan Foundation, gifts from Apple and Facebook, and a Simons Investigator Award.

§Email: connorw@cs.harvard.edu. Supported by the Harvard College Research Program (HCRP).

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of our funders.

Contents

1	Introduction	4
1.1	Differential Privacy	4
1.2	Previous Work	6
1.3	Underestimation of Sensitivity	8
1.4	Attacks	11
1.5	Solutions	13
1.6	A Roadmap for Implementing Our Solutions	16
2	Preliminaries	18
2.1	Measuring Distances	18
2.2	Sensitivity of Functions	21
2.2.1	The Path Property	21
2.3	Differential Privacy	22
2.3.1	The Laplace Mechanism	22
2.4	Integers	23
2.4.1	Integer Values	23
2.4.2	Integer Arithmetic	23
2.5	Floating-Point Representation	24
2.5.1	Floating-Point Values	24
2.5.2	Rounding Modes	26
2.5.3	Floating-Point Arithmetic	27
3	An Overview of DP Libraries	27
4	Bounded Sum	29
4.1	Bounded Sum is a Building Block in DP Libraries	29
4.2	The Idealized Bounded Sum Function	30
4.3	Sensitivity of the Idealized Bounded Sum	30
4.3.1	Uses of the Idealized Sensitivity in Libraries of DP Functions	30
5	Implementing Bounded Sum on Computers	31
5.1	Overflow Attack on Integers with Modular Addition	33
5.2	Reordering Attack on Signed Integers with Saturation Addition	34
5.3	Reordering Attack on Floats	36
5.4	Rounding-Based Attack on Floats	38
5.5	Repeated Rounding Attack on Floats I	41
5.6	Repeated Rounding Attack on Floats II	41
5.7	Concrete Implementations of the Attacks on DP Libraries	42
5.7.1	Implementing Example Attack 5.5	43
5.7.2	Implementing Example Attacks 5.10 and 5.11	44
5.7.3	Implementing Example Attack 5.13	46

6	Proposed Solutions to Incorrect Sensitivities	48
6.1	Bounding n	49
6.1.1	Solution for Integers	49
6.1.2	Use Case for Floats	49
6.2	Modular Noise Addition	50
6.3	Split Summation for Saturation Arithmetic	54
6.4	Randomized Permutation for Saturating Integers	56
6.4.1	Sensitivities with Respect to Ordered Distance for Saturating Integers	57
6.4.2	Randomized Permutation is a Transformation between Distance Metrics	58
6.4.3	Sensitivities for Unordered Distance	59
6.5	Sensitivity from Accuracy	60
6.5.1	Shifting Bounds	65
6.6	Split Summation with Random Permutation	65
	References	73
A	Missing proofs from Section 2	77
B	Missing proofs from Section 4	80
C	Missing proofs form Section 5	82

1 Introduction

Differential privacy (DP) [13] has become the prevailing framework for protecting individual-level privacy when releasing statistics or training machine learning models on sensitive datasets. It has been the subject of a rich academic literature across many areas of research, and has seen major deployments by the US Census Bureau [34, 17, 2], Google [16, 4, 5, 6], Apple [46], Facebook/Meta [37, 23], Microsoft [10, 42], LinkedIn [43], and OhmConnect.¹ To facilitate the adoption of differential privacy, a number of researchers and organizations have released open-source software tools for differential privacy, starting with McSherry’s PINQ [36], and now including libraries and systems from companies like Google [21], Uber [29], IBM [26], and Facebook/Meta [50], and the open-source projects OpenMined² and OpenDP [18].

However, implementing differential privacy correctly is subtle and challenging. Previous works have identified and attempted to address vulnerabilities in implementations of differential privacy coming from side channels such as timing [22, 28] and the failure to faithfully emulate the noise infusion mechanisms needed for privacy when using floating-point arithmetic instead of idealized real arithmetic [38, 27, 28].

In this work, we identify a new and arguably more basic class of vulnerabilities in implementations of differential privacy. Specifically, they arise when computing basic statistics such as sums, thanks to discrepancies between the implemented arithmetic using finite data types (namely, ints or floats) and idealized arithmetic over the reals or integers. These discrepancies result in the *sensitivity* of the implemented statistics — how much one individual’s data can affect the result — to be much higher than the sensitivity we expect. Consequently, essentially all differential privacy libraries fail to introduce enough noise to hide individual-level information as required by differential privacy, and we show that this may be exploited in realistic attacks on differentially private query systems. In addition to presenting these vulnerabilities, we also provide a number of solutions, which modify or constrain the way in which the sum is implemented in order to recover the idealized or near-idealized bounds on sensitivity.

In this extended introduction, we give an overview of our results. Precise definitions, proofs, and more details about the attack experiments begin at Section 2.

1.1 Differential Privacy

Informally, a randomized algorithm \mathcal{M} is *differentially private* if for every two datasets u, u' that differ on one individual’s data, the probability distributions $\mathcal{M}(u)$ and $\mathcal{M}(u')$ are close to each other. Intuitively, this means that an adversary observing the output cannot learn much about any individual, since what the adversary sees is essentially the same as if that individual’s data was not used.

To make the definition of differential privacy precise we need to specify both what it means for two datasets u and u' to “differ on one individual’s data,” and how we measure the closeness of probability distributions $\mathcal{M}(u)$ and $\mathcal{M}(u')$. For the former, there are two common choices in the differential privacy literature. In one choice, we allow u' to differ from u by adding or removing any one record; this is called *unbounded differential privacy* and thus we denote this relation $u \simeq_{unbdd} u'$. Alternatively, we can allow u' to differ from u by *changing* any one record; this is called *bounded differential privacy* and thus we will write $u \simeq_{bdd} u'$. Note that if $u \simeq_{bdd} u'$,

¹ 

² 

then u and u' necessarily have the same number of records; thus, this is an appropriate definition when the size n of the dataset is known and public. When $u \simeq u'$ for whichever relation we are using (\simeq_{bdd} or \simeq_{unbdd}), we call u and u' *adjacent* with respect to \simeq .

For measuring the closeness of the probability distributions $\mathcal{M}(u)$ and $\mathcal{M}(u')$, we use the standard definition of (ε, δ) -differential privacy [12], requiring that:

$$\forall T \quad \Pr[\mathcal{M}(U) \in T] \leq e^\varepsilon \cdot \Pr[\mathcal{M}(u') \in T] + \delta, \quad (1)$$

where we quantify over all sets T of possible outputs. There are now a variety of other choices, like Concentrated DP [15, 8], Rényi DP [39], and f -DP [11]; our results are equally relevant to these forms of DP, but we stick with the basic (ε, δ) notion for simplicity. If \mathcal{M} satisfies (1) for all u, u' such that $u \simeq u'$, then we say that \mathcal{M} is (ε, δ) -DP with respect to \simeq . If $\delta = 0$, we say \mathcal{M} is ε -DP with respect to \simeq . Intuitively, ε measures *privacy loss* of the mechanism \mathcal{M} , whereas δ bounds the probability of failing to limit privacy loss to ε (so we typically take δ to be cryptographically small).

The fundamental building block of most differentially private algorithms is noise addition calibrated to the *sensitivity* of a function f that we wish to estimate.

Definition 1.1 (Sensitivity). Let f be a real-valued function on datasets, and \simeq a relation on datasets. The (*global*) *sensitivity* of f with respect to \simeq is defined to be:

$$\Delta_{\simeq} f = \sup_{u \simeq u'} |f(u) - f(u')|.$$

Theorem 1.2 (Laplace Mechanism [13]). *For every function f and relation \simeq on datasets, the mechanism*

$$\mathcal{M}(u) = f(u) + \text{Lap}\left(\frac{\Delta_{\simeq} f}{\varepsilon}\right),$$

is ε -DP, where $\text{Lap}(s)$ denotes a draw from a Laplace random variable with scale parameter s .

There are a number of other choices for the noise distribution, leading to the Discrete Laplace (a.k.a., Geometric) Mechanism [20], the Gaussian Mechanism [40], and the Discrete Gaussian Mechanism [9], where the latter two achieve (ε, δ) -DP with $\delta > 0$. The key point for us is that in all cases, the scale or standard deviation of the noise is supposed to grow linearly with the sensitivity $\Delta_{\simeq} f$, so it is crucial that the sensitivity is calculated correctly. If we incorrectly underestimate the sensitivity as $\Delta = (\Delta_{\simeq} f)/c$ for a large constant c , then we will only achieve $c\varepsilon$ -DP; that is, our privacy loss will be much larger than expected. Given that it is common to use privacy loss parameters like $\varepsilon = 1$ or $\varepsilon = .5$, a factor of 5 or 10 increase in the privacy-loss parameter can have dramatic effects on the privacy protection (since $e^5 > 148$, allowing a huge difference between the probability distributions $\mathcal{M}(u)$ and $\mathcal{M}(u')$).

The most widely used function f in DP noise addition mechanisms is the *Bounded Sum* function.

Definition 1.3 (Bounded Sum). For real numbers $L \leq U$, and a dataset v consisting of elements of the interval $[L, U]$, we define the *bounded sum function* on v to be:

$$BS_{L,U}(v) = \sum_{i=1}^{\text{len}(v)} v_i.$$

The restricting of $BS_{L,U}$ to datasets v of length n is denoted $BS_{L,U,n}$. When we do not constrain the data to lie in $[L, U]$, we omit the subscripts.

Typically, the bounds L and U are enforced on the dataset v via a record-by-record clamping operation. To avoid the extra notation of the clamp operation, throughout we will work with datasets that are assumed to already lie within the bounds.

It is well-known and straightforward to calculate the sensitivity of bounded sum.

Proposition 1.4. *$BS_{L,U}$ has sensitivity $\max\{|U|, |L|\}$ with respect to \simeq_{umbdd} , and $BS_{L,U,n}$ has sensitivity $U - L$ with respect to \simeq_{bdd} .*

Combining Proposition 1.4 and Theorem 1.2 (or analogues for other noise distributions), we obtain a differentially private algorithm for approximating bounded sums, which we will refer to as Noisy Bounded Sum. This algorithm is pervasive throughout both the differential privacy literature and software. Many more complex statistical analyses can be decomposed into bounded sums, and any machine learning algorithm that can be described in Kearns’ Statistical Query model [32] can be made DP using Noisy Bounded Sum. Indeed, the versatility of noisy sums was the basis of the SuLQ privacy framework [7], which was a precursor to the formal definition of differential privacy. Noisy Bounded Sum is also at the heart of differentially private deep learning [1], as each iteration of (stochastic) gradient descent amounts to approximating a sum (or average) of gradients. For this reason, every software package for differential privacy that we are aware of supports computing bounded sums via noise addition.

1.2 Previous Work

Despite their simplicity, Noisy Bounded Sum and related differentially private algorithms are surprisingly difficult to implement in a way that maintains the desired privacy guarantees.

In the early days of implementing differential privacy, several challenges were pointed out by Haeberlen, Pierce, and Narayan [22]. Their attacks concern not the bounded sum function itself, but dataset transformations that are applied before the bounded sum function. In a typical usage, bounded sum is not directly applied to a dataset $u = [u_1, \dots, u_n]$, but rather to the dataset

$$q(u) \stackrel{\text{def}}{=} [q(u_1), q(u_2), \dots, q(u_n)]$$

where q is a “microquery” mapping records to real numbers or to the interval $[L, U]$ (we can incorporate the clamping into q). If $u \simeq u'$, then $q(u) \simeq q(u')$, so we if we apply Noisy Bounded Sum (or any other differentially private algorithm) to the transformed dataset, we should still satisfy differential privacy with respect to the original dataset u . Haeberlen et al.’s attacks rely on a discrepancy between this mathematical abstraction and implementations. In code, q may not be a pure function, and may be able to access global state (allowing information to flow from the execution of $q(u_i)$ to the execution of $q(u_j)$ for $j > i$) or leak information to the analyst through side channels (such as timing). The authors of the main differentially private systems at the time, PINQ [36] and Airavat [44], were aware of and noted the possibility of such attacks, but the implemented software prototypes did not fully protect against them. As discussed in [36, 44, 22], remedies for attacks like these include using a domain-specific language for the microquery q to ensure that it is a pure function and enforcing constant-time execution for q .

At the other end of the DP pipeline (after Bounded Sum calculation), the seminal work of Mironov [38] demonstrated vulnerabilities coming from the noise addition step. Specifically, the Laplace distribution in Theorem 1.2 is a continuous distribution on the real numbers, but computers cannot manipulate arbitrary real numbers. Typical implementations approximate real

numbers using finite-precision floating-point numbers, and Mironov shows that these approximations can lead to complete failure of the differential privacy property. To remedy this, Mironov proposed the Snapping Mechanism, which adds a sufficiently coarse rounding and clamping after the Laplace mechanism to recover differential privacy with a slightly larger value of ϵ . Subsequent works [19, 3, 9] avoided floating-point arithmetic, advocating and studying the use of exact finite-precision arithmetic (e.g., using big integers) and using discrete noise distributions, such as the discrete Laplace distribution [20] and the discrete Gaussian distribution [9]. However, a recent paper by Jin et al. [28] shows that current implementations of the discrete distribution samplers are vulnerable to timing attacks (which leak information about the generated noise value and hence of the function value it was meant to obscure). They, as well as [25], also show that the floating-point implementations of the continuous Gaussian Mechanism are vulnerable to similar attacks as shown by Mironov [38] for the Laplace Mechanism.

Ilvento [27] studies the effect of floating-point approximations on another important DP building block, the Exponential Mechanism. On a dataset u , the exponential mechanism samples an outcome y from a finite set \mathcal{Y} of choices with probability $p_y(u)$ proportional to $\exp(\epsilon f_y(u)/(2 \max_y \Delta_{\simeq} f_y))$, where $f_y(u)$ is an arbitrary measure of the “quality” of outcome y for dataset u . She shows that the discrepancy between floating-point and real arithmetic can lead to incorrectly converting the quality scores $f_y(u)$ into the probabilities $p_y(u)$ and hence violate differential privacy. To remedy this, Ilvento proposes an exact implementation of the Exponential Mechanism using finite-precision base 2 arithmetic. Note that, like noise addition mechanisms, the Exponential Mechanism also is calibrated to the sensitivities $\Delta_{\simeq} f_y$ of the quality functions. Here too, if we underestimate the sensitivity by a factor of c , our privacy loss can be greater than intended by a factor of c , even if we perfectly implement the sampling or noise generation step.

Indeed, Mironov’s paper [38] also suggests that our sensitivity calculations may fail to translate from idealized, real-number arithmetic to implemented, floating-point arithmetic. He gives an example of two datasets $u \simeq_{bdd} u'$ of 64-bit floating-point numbers such that $|BS(u) - BS(u')| = 1$ but $|BS^*(u) - BS^*(u')| = 129$, where BS^* is the standard, iterative implementation of summation of floating-point numbers, illustrated in Figure 1:

```

1  def iterated(u):
2      total = 0
3      for elem in u:
4          total += elem
5      return total

```

Figure 1: Iterated Summation.

However, Mironov’s example does not immediately lead to an underestimation of sensitivity, because the datasets u and u' include values ranging from $L = 2^{-23}$ to $U = 2^{30}$, so the idealized sensitivity predicted by Proposition 1.4 is $U - L > 2^{30} \gg 129$.³ Mironov’s paper suggests that

³As pointed out in Mironov’s paper, this example does demonstrate an underestimation of sensitivity by a factor of 129 if we define $u \simeq u'$ to mean that u and u' agree on all but one coordinate and they differ by at most 1 on that coordinate. This notion of dataset adjacency is common in the DP literature when u and u' represent datasets in *histogram* format, where u_i is the number of individuals of type i (rather than individual i ’s data). However, in this case, the u_i ’s would be nonnegative integers (so this example would not be possible) and it would be strange to use a floating-point data type.

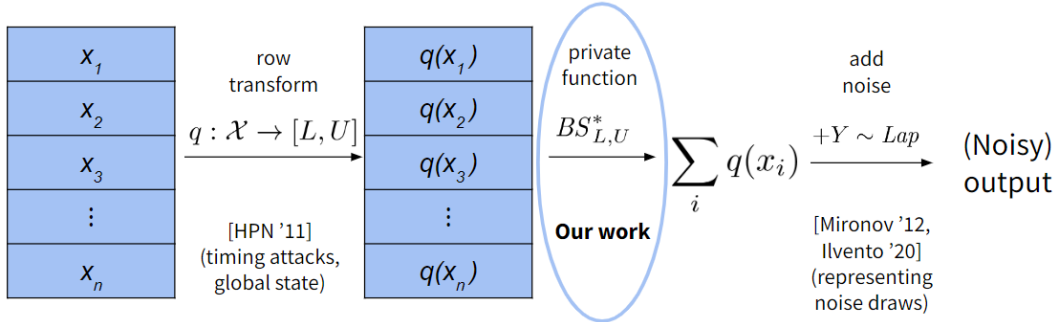


Figure 2: Illustration of the DP pipeline and the relationship of this paper to previous works uncovering vulnerabilities of DP implementations.

these potential sensitivity issues may be addressed by replacing Iterated Summation by a *Kahan Summation* [31], a different way of summing floating-point numbers that accumulates rounding errors more slowly. Unfortunately, this section of Mironov’s paper seems to have gone mostly unnoticed, and we are not aware of any prior work that has addressed the question of how to correctly bound and control sensitivity in implementations of differential privacy.

1.3 Underestimation of Sensitivity

In our work, we identify new privacy vulnerabilities arising from underestimation of the sensitivity of the Bounded Sum function when implemented using finite data types, including 32-bit and 64-bit integers and floats. Specifically, implementations of Bounded Sum often have sensitivity much larger than the idealized sensitivity given by Proposition 1.4, and consequently the privacy loss of DP mechanisms using Bounded Sum is much larger than specified. Our work covers vulnerabilities arising from the “middle step” of the DP pipeline—aggregation—sitting between the foci of previous work, which considered vulnerabilities in the preprocessing step before aggregation and the noise generation/sampling step after aggregation.

Throughout we will write T for a finite numerical data type. We think of $T \subseteq \mathbb{R}$, but adding two elements a, b of T can yield a number outside of T , so the *overflow mode* and/or *rounding mode* of T determine the result of the computation $a + b$. For $L, U \in T$ with $L \leq U$, we write $T_{[L,U]} = \{x \in T : L \leq x \leq U\}$. We will consider various implementations of the bounded sum function $BS_{L,U}^*$ and $BS_{L,U,n}^*$ on datasets consisting of elements of $T_{[L,U]}$. Except when otherwise stated, BS^* will use the standard Iterated Summation method from Figure 1. In such a case, the functions $BS_{L,U}^*$ and $BS_{L,U,n}^*$ and their sensitivities are completely determined by the choice of overflow and/or rounding modes.

The data types T we will consider in the paper are the following:

- k -bit unsigned integers, whose elements are $\{0, 1, \dots, 2^k - 1\}$. Standard choices are $k = 32$ and $k = 64$.
- k -bit signed integers, whose elements are $\{-2^{k-1}, -2^{k-1} + 1, \dots, 0, 1, \dots, 2^{k-1} - 1\}$. Standard choices are $k = 32$ and $k = 64$.
- (k, ℓ) -bit (normal) floats, which are represented in binary scientific notation as $(-1)^s \cdot$

$(1.M) \cdot 2^E$ with a k -bit mantissa M and an exponent $E \in [-(2^{\ell-1} - 2), 2^{\ell-1} - 1]$.⁴ In 32-bit floats (“singles”), we have $k = 23$ and $\ell = 8$; and in 64-bit floats (“doubles”) we have $k = 52$ and $\ell = 11$. In machine learning applications, it is common to use even lower-precision floating-point numbers for efficiency, such as $k = 7$ and $\ell = 8$ in Google’s `bfloat16` [47].

We find that for these data types, the sensitivity of bounded sum can be much larger than the idealized for several reasons, described in the paragraphs below.

Overflow. The default way of dealing with overflow in k -bit integers T (signed or unsigned) is *wraparound*, i.e., the result is computed modulo 2^k . It is immediately apparent how this phenomenon can lead to large sensitivity. If the idealized sum on one dataset u is equal to the largest element of T , call it $\max(T)$ and on an adjacent dataset u' equals $\max(T) + 1$, then modular arithmetic will yield results that differ by $2^k - 1$. If our parameter settings allow for us to construct two such datasets, then the implemented sensitivity of bounded sum will be $2^k - 1$, a completely useless bound.

In the case of bounded DP on datasets of size n , we can construct two such datasets u and u' if $n \cdot (U - L) \geq 2^k$. In practice, while n may be modest (e.g., $n = 2^{15}$) and much smaller than 2^k , the bounds U and L are typically user-specified parameters. So, for example, we can set $L = 0$ and $U = \lceil 2^k/n \rceil$, and we get a sensitivity that is roughly a factor of n larger than the idealized sensitivity.

In the case of unbounded DP, n is unconstrained, so we always get an implemented sensitivity of $2^k - 1$, provided that $U > L$. Specifically, there are datasets u and u' of size at most $n = \lceil 2^k/U \rceil$ such that $u \simeq_{\text{ubdd}} u'$ and $|BS_{L,U}^*(u) - BS_{L,U}^*(u')| = 2^k - 1$. Again, we get a blow-up in sensitivity of roughly a factor of n .

Overflow also arises with floating-point arithmetic (leading to results that are $\pm\infty$), but the semantics of arithmetic with $\pm\infty$ is not clearly defined in the IEEE standard, and we consider it better to avoid this possibility entirely (as discussed in our solutions below).

Rounding. When adding two floating-point numbers, the result may not be exactly representable as a float, but may lie strictly between two adjacent floats. Thus, the actual result is determined by the *rounding mode*. The standard, called *banker’s rounding*, rounds the result to the nearest float, breaking ties by rounding to the one whose mantissa has least-significant bit 0. By inspection, when we round a real number z to a (k, ℓ) -float $\text{BRound}(z)$, the *relative* effect is minimal. Specifically, $|\text{BRound}(z) - z| \leq |z|/2^{k+1}$.

This may have led to an incorrect impression that floating-point arithmetic is “close enough” to idealized real arithmetic to not cause a significant increase in sensitivity. Unfortunately, we find that this is not the case.

In the case of bounded DP, even a single rounding can cause a substantial blow-up in sensitivity, as we illustrate with the following example. Let Q be a float which is a power of 2, and

⁴Note that there are only $2^\ell - 2$ choices for E . The remaining choices are used to represent *subnormal* floats, as well as $\pm\infty$, and NaN. For simplicity in this introduction, we only consider normal floats.

let $n - 1$ be a power of 2 between 2 and 2^k . Then, we take

$$\begin{aligned} L &= \left(1 + \frac{n-1}{2^{k+1}}\right) \cdot Q, \\ U &= \left(1 + \frac{n+1}{2^{k+1}}\right) \cdot Q. \end{aligned}$$

Thus, L and U are both very close to Q , but their difference is only $U - L = Q/2^k$. Our datasets u and u' will both begin with $n - 1$ copies of L . The iterated sum of these first $n - 1$ elements can be shown to experience no rounding, and will result in a floating-point number with exponent $\log_2((n - 1)Q)$ (since $(n - 1)L < 2(n - 1)Q$). Thus, the space between adjacent floating-point numbers after this iterated sum is $(n - 1)Q/2^k$. Hence, when we add one more copy of L to obtain $BS^*(u)$, the result will lie exactly halfway between two adjacent floats (since L is an odd multiple of $(n - 1)Q/2^{k+1}$), and by banker's rounding, will get rounded down. On the other hand, when we add a copy of U to obtain $BS^*(u')$, the result will get rounded up. The effect of these two roundings gives us:

$$BS^*(u') - BS^*(u) = 2 \frac{n-1}{2^{k+1}} \cdot Q = (n-1) \cdot (U - L).$$

In contrast, the idealized sensitivity with respect to \simeq_{bdd} is $(U - L)$, so the sensitivity blows up by a factor of $(n - 1)$, which is dramatic even for very small datasets. We also show that this construction also applies to Kahan summation (contrary to Mironov's hope that it would salvage the sensitivity) and pairwise summation (another common method—the default in `numpy`—where the values are added in a binary tree).

Repeated Rounding. The above example does not give anything interesting for unbounded DP, since the idealized sensitivity is $\max\{|U|, |L|\}$, but $BS^*(u') - BS^*(u) \leq Q < U$.

However, we can obtain a sensitivity blow-up by exploiting the accumulated effect of *repeated rounding*. Consider a dataset whose first $(n/2)$ elements are U . After that, each rounding can have the effect of increasing the sum by $\Theta(nU/2^k)$, for a total rounding error of $\Theta(n^2U)/2^k$. We show how to exploit this phenomenon to construct two datasets $u \simeq_{unbdd} u'$ which differ on their middle element such that

$$|BS_{L,U}^*(u) - BS_{L,U}^*(u')| \geq \left(1 + \Omega\left(\frac{n^2}{2^k}\right)\right) \cdot U.$$

Thus, the sensitivity blows up by a factor of $\Theta(n^2/2^k)$. This allows us to exhibit a sensitivity blow-up with datasets of size $n = \Theta(2^{k/2})$, which are plausible even for 64-bit floats (where $k = 52$) and quite easy to obtain for 32-bit and lower-precision floats.

Reordering. Finally, we exhibit potential vulnerabilities based on ambiguity about whether datasets are ordered or unordered. In the DP theory literature, datasets are typically considered unordered (i.e., multisets) when using unbounded DP. (Indeed, adjacency is often described by requiring that the *symmetric difference* between the multisets u and u' has size at most 1, or by requiring that the histograms of u and u' have ℓ_1 distance at most 1.) On the other hand, when using bounded DP, it is common to denote datasets as ordered n -tuples. (Indeed, adjacency is often described by requiring that the *Hamming distance* between the n -tuples is at

most 1.) Most implementations of DP are not explicit about these choices, but the wording in the documentation suggests the same conventions (e.g., see Section 1.3 in [49]).

In theory, the distinction does not matter much, because most of the functions we compute (such as Bounded Sum) are symmetric functions and do not depend on the ordering.

However, when implementing these functions using finite data types, the ordering can matter a great deal. Indeed, we show that there are datasets u and u' where u' is a permutation of u (so define exactly the same multisets) but

$$|BS_{L,U}^*(u) - BS_{L,U}^*(u')| \geq \left(1 + \Omega\left(\frac{n^2}{2^k}\right)\right) \cdot U.$$

That is, we can obtain the same kind of blow-ups in sensitivity that we obtained due to rounding by instead just reordering the dataset. (Our rounding attacks preserved order, i.e., we obtained u' by changing one element of the ordered tuple u , or by inserting/deleting one element of u without changing the other elements.) One reason this may have been missed before is that floating-point arithmetic *is* commutative, e.g, $\text{BRound}(a + b) = \text{BRound}(b + a)$, so it may seem like order does not matter. However, for the sum to be invariant under arbitrary permutations also requires *associativity*, which fails for floating-point arithmetic (due to the rounding).

Modular integer arithmetic is associative, so this issue does not arise for k -bit (unsigned or signed) integers with wraparound. However, if instead we use *saturation* arithmetic, where values get clamped to the range $[\min(T), \max(T)]$ (this addresses the aforementioned sensitivity problems with wraparound), then addition of *signed* integers is no longer associative. Indeed, if $n \cdot \min\{U, |L|\} \geq 2^{k+1}$, we exhibit two datasets u, u' of length n such that u' is a permutation of u and

$$|BS_{L,U,n}^*(u) - BS_{L,U,n}^*(u')| \geq 2^k - 1.$$

That is, we get no improvement over the (trivial) sensitivity we had with modular arithmetic.

1.4 Attacks

We show that our attacks can be carried out on many existing implementations of differential privacy. In each attack, we set the privacy-loss parameter to $\epsilon \in \{0.5, 1\}$, set the parameters L , U , and possibly n as required by our sensitivity lower bounds, construct two appropriate adjacent datasets $u \simeq_{bdd} u'$ or $u \simeq_{unbdd} u'$, and run the supposedly ϵ -DP bounded sum mechanism \mathcal{M} on u and u' . We show that by applying a threshold test to the outputs, we can almost perfectly distinguish $\mathcal{M}(u)$ and $\mathcal{M}(u')$. For example, in one experiment, we succeed in correctly identifying whether the dataset is u or u' in all but 3 out of 20,000 runs of the mechanism, which we show would be astronomically unlikely for a mechanism that is truly 0.5-DP.

Specifically, we show that our overflow attack on integers works on IBM’s diffprivlib; our rounding attack on floats work on Google’s DP library, IBM’s diffprivlib, OpenDP / SmartNoise, and Chorus; our repeated rounding attacks on floats work on Google’s DP library, IBM’s diffprivlib, OpenDP / SmartNoise, Chorus, and PINQ; and our re-ordering attack works on Google’s DP library, IBM’s diffprivlib, OpenDP / SmartNoise, Chorus, and PINQ. (Chorus only offers sensitivities for the bounded DP setting, so slight adjustments were made for the repeated rounding attacks – specifically, a value of 0 was appended to dataset v .) With the exception of OpenDP, none of the libraries include a disclaimer indicating that the integer overflow or floating-point rounding behaviors we exploit could yield vulnerabilities.

Our attacks utilize contrived datasets u and u' that are unlikely occur “in the wild.” However, our overflow and rounding attacks can easily be applied to attack realistic datasets given access to a DP system that fields external queries with user-defined microqueries (a.k.a., mappers or row transforms). Specifically, consider a dataset $x = [x_1, \dots, x_n]$ where each record x_i contains a known/public user identifier uid_i for the i 'th individual, and a sensitive bit val_i . Suppose further that the dataset is sorted by the identifiers, i.e., $\text{uid}_1 \leq \text{uid}_2 \leq \dots \leq \text{uid}_n$. Then it is straightforward for an adversarial analyst to define a simple micro-query q such that

$$q(x) = [q(x_1), \dots, q(x_n)] = \begin{cases} u & \text{if } \text{val}_i = 1 \\ u' & \text{if } \text{val}_i = 0 \end{cases}$$

where u, u' are the datasets in our attacks, and they differ on the i 'th record. Thus, by seeing the result of the bounded-sum mechanism on $q(x)$, the adversary can extract the i 'th individual's bit val_i with very high probability. This can be generalized to not just attack a particular individual, but a constant fraction of the individuals in the dataset (e.g., any individual in the middle half of the dataset). Note that our use of microqueries here is different than in Haeberlen et al. [22]. Our microqueries are pure functions, not making use of any side channels or global state, and are simple enough to be implemented in virtually any DSL for microqueries. Indeed, see Figure 3 for a code snippet showing how our attack would look in SQL and Figure 4 for PINQ.

```

1 SELECT
2   FROM
3     WHERE uid < m
4     WHERE uid = m
5     WHERE uid % 2 = 1
6     WHERE val = pos_val
7   )
8   WHERE uid < 4 * U

```

Figure 3: Example SQL code for our repeated rounding attack. Here the upper clamping bound is U , the lower clamping bound is $L=-U$, uid denotes the user id attribute of an individual record, m is the middle user id (the one we wish to attack), val is the individual's sensitive bit, and pos_val and neg_val are particular floats in the interval $[L, U]$ from our attack.

```

1 var uid = new BigInteger(256, new BigInteger(
2   (1 > 1 ? ' ' : ' '))
3   .set(100, new BigInteger(100, 4 * (uid)), uid
4   (uid)))
5   .set(100, (uid < uid) ? U :
6     (uid = uid) ? U * uid :
7     (uid % 2 = 1) ? pos_val : neg_val)
8   .set(100.0, v > v);

```

Figure 4: Example PINQ code for our repeated rounding attack.

Since PINQ uses 64-bit floats, our repeated rounding attack requires quite a large dataset (e.g., $n = 2^{28}$), and we were not able to complete execution due to memory timeouts, but the

attack is possible in principle. Perhaps due to concerns about timing and other side-channel attacks (see Section 1.2), recent DP SQL systems such as Chorus do not support user-defined microqueries. However, it may be still be possible to carry out our basic rounding attack (not the repeated rounding one), since it can be implemented in such a way that microquery is simply the clamp function (applied automatically in most implementations of Noisy Bounded Sum), if we know that the dataset is sorted by the value we wish to attack (e.g., consider a dataset of salaries, and where the employee with second-highest salary wishes to find out the CEO’s salary). Finally, we also hypothesize that the attack can be carried out on DP Machine Learning systems that use DP-SGD, using a user-defined loss function $\ell(\theta, x_i)$, designed so that $\nabla_{\theta}\ell(\theta, x_i)$ equals our desired microquery $q(x_i)$. If we set parameters so that there is only one iteration, using the entire dataset as a batch, then the output parameter vector θ will exactly give us a Noisy Bounded Sum. Since these ML systems allow using low-precision floats, these attacks should be feasible even on very small datasets.

To carry out our reordering attacks on a DP query system, we would need a system that can be made to perform a data-dependent reordering. These attacks indicate that we need to be explicit and careful about how we treat ordering when we implement DP.

In any case, our attacks and experimental results reported in the appendix already demonstrate that essentially all of the implementations of DP fail to meet their promised DP guarantees.

Responsible disclosure: Immediately after the submission of this paper, we are sharing the paper with the maintainers of all of the DP software toolkits that we have attacked and informed them that we would wait 30 days to post the paper publicly, to give them time to implement any needed patches (like our solutions below). In response, all offered some indication that they are working to resolve these issues, and Google’s Bug Hunter program acknowledged our contribution to Google security on their Honorable Mentions.⁵ All of the authors of this paper are members of the OpenDP team and are involved with the development of the library. The findings of this paper occurred while the authors were writing mathematical proofs to accompany the algorithms that are part of the OpenDP library as part of our vetting process, upon which we realized the vulnerabilities described in this paper. We are updating the OpenDP library following the roadmap described in Section 1.6.⁶ We believe that our findings also illustrate the importance of vetting processes such as the one put in place for OpenDP.

Code: Code for our attacks and experiments is available in the Github repository at <https://github.com/xbqg846z/underestimate-sensitivity>.

1.5 Solutions

Dataset Metrics. Toward addressing the potential reordering vulnerability, we propose that when datasets $u = [u_1, \dots, u_n]$ are stored as ordered tables, we should define and distinguish four adjacency relations \simeq . Specifically, the bounded-DP relation \simeq_{bdd} should separate into the usual \simeq_{Ham} (“Hamming”), where $u \simeq_{Ham} u'$ means that there is at most one coordinate i such that $u_i \neq u'_i$, and \simeq_{CO} (“change-one”), where $u \simeq_{CO} u'$ means that we can convert the multiset of elements in u into the multiset of elements in u' by changing one element. Equivalently, $u \simeq_{CO} u'$ iff there is a permutation π such that $\pi(u) \simeq_{Ham} u'$. Similarly, the unbounded-DP relation \simeq_{unbdd} should split into an ordered version \simeq_{ID} (“insert-delete”) and an unordered version \simeq_{Sym} (“symmetric distance”). By being explicit about ordering in our

⁵ <https://bughunters.google.com/bug?id=846f72-98-484-9f7-86046511>

⁶ <https://github.com/xbqg846z/underestimate-sensitivity>

, <https://bughunters.google.com/bug?id=846f72-98-484-9f7-86046511>

adjacency relation, and in particular analyzing DP and sensitivity with respect to a specific relation, we can avoid the reordering vulnerabilities. In particular, a DP system using ordered adjacency relations should take care to disallow data-dependent reorderings, unless they can be shown to preserve ordered adjacency (or, more generally are stable with respect to Hamming distance or insert-delete distance).

Random permutations (RP). Given the above adjacency relations, it still remains to find implementations of Bounded Sum that achieve a desired sensitivity with respect to them. In general, it is easier to bound sensitivity with respect to the ordered relations \simeq_{Ham} and \simeq_{ID} . For example, we can show that iterated sum of signed integers with saturation arithmetic has the idealized sensitivities of $U - L$ and $\max\{U, |L|\}$ with respect to \simeq_{Ham} and \simeq_{ID} , respectively, whereas our reordering attacks show that the sensitivities can be as large as $2^k - 1$ with respect to \simeq_{CO} and \simeq_{Sym} .

Therefore, we give a general way to convert sensitivity bounds with respect to the ordered relations into the same bounds with respect to the unordered relations: randomly permute the dataset before applying the function. To formalize the effect of this transformation, we need to extend the definition of sensitivity to randomized functions. Following [45], we say that a randomized function f has *sensitivity at most Δ with respect to \simeq* if for all pairs of datasets u, u' such that $u \simeq u'$, there is a *coupling* of the random variables $f(u)$ and $f(u')$ such that $\Pr[|f(u) - f(u')| \leq \Delta] = 1$. Given this transformation, it suffices for us to obtain sensitivity bounds with respect to ordered metrics.

Checking or bounding parameters. Many of our attacks only lead to a large increase in sensitivity under certain parameter regimes, e.g., $n \cdot U \geq 2^k$. In the case of bounded DP, all of these parameters (n, U, L, k) are known in advance (before we touch the sensitive dataset), and we can prevent the problematic scenarios by either constraining the parameters or incorporating the dependence on those parameters into our sensitivity bounds. Indeed, we show that many of the sensitivity lower bounds discussed in Section 1.3 are tight by giving nearly matching upper bounds.

For the case of unbounded DP, n is not known, but we can achieve a similar effect by *truncating* the sum, namely implementing

$$BS_{L,U,n}^{**}(u) = BS_{L,U}^*(u_1, u_2, \dots, u_{\min\{n, \text{len}(u)\}}). \quad (2)$$

Note that this truncation operation behaves nicely with respect to the *ordered* adjacency relation \simeq_{ID} . That is, if u and u' differ only by a single insertion or deletion, then their truncations differ by at most a single insertion or deletion.

Changing the summation. Our matching upper bound on the sensitivity of Iterated Summation of floats with banker’s rounding is obtained via reduction to accuracy. Specifically,

$$\begin{aligned} |BS^*(u) - BS^*(u')| &\leq \\ &|BS(u) - BS(u')| + |BS^*(u) - BS(u)| + |BS^*(u') - BS(u')|. \end{aligned}$$

The first term is bounded by the idealized sensitivity, and the latter two terms can be bounded by using known numerical analysis results about the accuracy of iterated summation. Specifically,

a bound from Wilkinson [48] shows that

$$|BS^*(u) - BS(u)| = O\left(\frac{n}{2^k} \sum_{i=1}^n |u_i|\right) = O\left(\frac{n^2 \cdot \max\{|U|, |L|\}}{2^k}\right).$$

Thus, applying Truncation (Equation (2)) to Iterated Summation, we get a sensitivity bound of:

$$\Delta_{\simeq_{unbdd}} BS_{L,U,n}^* \leq \left(1 + O\left(\frac{n^2}{2^k}\right)\right) \cdot \max\{|U|, |L|\}. \quad (3)$$

Thus, when $n \ll 2^{k/2}$, we recover almost the idealized sensitivity bound of $\max\{|U|, |L|\}$. Higham [24] has proven that Pairwise Summation and Kahan Summation have better bounds on accuracy, allowing us to replace the $O(n^2)$ above with $O(n \log n)$ or $O(n)$, respectively, allowing us to closely approximate the idealized sensitivity whenever $n \ll 2^k$, which covers many practical scenarios even for 32-bit floats ($k = 23$).

Another change in summation that sometimes reduces sensitivity is what we call *split summation*, where we separately sum the positive numbers and the negative numbers in the dataset. Specifically, split summation provides another solution to the reordering attack on signed integers with saturation arithmetic; it recovers the idealized sensitivity with respect to unordered dataset adjacency relations (without needing to perform a random permutation).

Shifting bounds. For the case of bounded DP (known n), some of the sensitivity blow-ups (such as our basic rounding attack) come from having U (or $\max\{|U|, |L|\}$) rather than $U - L$ in the sensitivity bound, and U can be much larger than $U - L$. One way to address this is to subtract L from every element of the dataset, so that all elements lie in the interval $[L' = 0, U' = U - L]$, apply our solutions that have U' in the sensitivity bound, and then add $L \cdot n$ at the end as post-processing. That is, we convert a method BS^* with sensitivity depending on $\max\{|U|, |L|\}$ (such as Expression (3)) into a method BS^{**} with sensitivity depending on $U - L$ as follows:

$$BS_{L,U,n}^{**}(u) = BS_{L,U,n}^*(u_1 - L, u_2 - L, \dots, u_n - L) + L \cdot n,$$

where a noisy version of the $BS_{L,U,n}^*$ term is computed first and $L \cdot n$ is added as a post-processing step.

Modular sensitivity. One way to address the large sensitivity of Bounded Sum over k -bit integers with wraparound is to change our definition of sensitivity, measuring the distance between k -bit integers as if they are equally spaced points on a circle. That is, we replace $|f(u) - f(u')|$ in the definition of sensitivity with $\min\{f(u), 2^k - f(u)\}$; we call this the *modular sensitivity* of f . With this change, Bounded Sum recovers its idealized sensitivity. But this begs the question of whether functions with bounded modular sensitivity still be estimated in a DP manner. Fortunately, we show that the answer is *yes*: if we add *integer-valued* noise (such as in the Discrete Laplace [20] or Discrete Gaussian [8] Mechanisms) and also do the noise addition with wraparound, then the result achieves the same privacy parameters as if we had done everything exactly over the integers, with no wraparound. Indeed, by the fact that modular reduction is a ring homomorphism, we can analyze the output distribution as if we had done modular reduction only at the end, which amounts to post-processing.

	Input metric	Implemented sensitivity	Conditions
Idealized [Thm. 4.4]	Known n	$U - L$	
	Unknown n	U	
Modular addition [Thm. 5.4] (signed or unsigned ints)	Known n	$\max(T) - \min(T)$	$nU > \max(T)$
	Unknown n	$\max(T) - \min(T)$	
Saturation (signed ints) [Thm. 5.6]	Known n	$\max(T) - \min(T)$	$nU > \max(T)$ $U > 0$ and $L < 0$
	Unknown n	$\max(T) - \min(T)$	$U > 0$ and $L < 0$
Floats with k -bit mantissa	Known n [Thm. 5.12]	$\geq (n - 1) \cdot (U - L)$	$U \geq 2^k \cdot (U - L)$
	Unknown n [Thm. 5.19]	$(1 + \Theta(n^2/2^k)) \cdot U$	$L \leq -U \cdot n/2^{k+3}$

Table 1: Lower bounds obtained in our attacks for numerical type T and $-U \leq L \leq U$.

Several DP libraries already implement this solution, because it happens by default when everything is a k -bit integer. The reason we were able to attack IBM’s diffprivlib with the overflow attack is that, while it computes the Bounded Sum using integer arithmetic, the noise addition is done using floating-point arithmetic.

It is not clear whether there is an analogue of this solution for the use of sensitivity in the Exponential Mechanism. (See Section 1.2.)

Changing overflow mode. As mentioned above, another solution for the case of k -bit integers is to replace wraparound with saturation arithmetic, combining it with either the random permutation technique or split summation in order to handle unordered adjacency relations.

Changing rounding mode. For our rounding attacks against floating-point numbers, another solution (beyond bounding parameters) is to replace the default banker’s rounding with another standard rounding mode, namely round toward zero (RTZ). We show that this gives an implemented sensitivity of

$$\Delta_{\simeq_{unbdd}} BS_{L,U,n}^* \leq \left(\max \left\{ 1 + O \left(\frac{n}{2^k} \right), 2 \right\} \right) \cdot \max\{|U|, |L|\},$$

provided that (a) all elements of the datasets have the same sign (i.e., $L \geq 0$ or $U \leq 0$), and (b) we work with an ordered dataset adjacency relation. To handle the case of mixed signs, we can use the shifting technique (in case of bounded DP) or split summation (in the case of unbounded DP). To handle unordered dataset relations, we can apply the random permutation technique. Altogether these solutions maintain a sensitivity that is within a small constant factor of the idealized sensitivity in all cases.

1.6 A Roadmap for Implementing Our Solutions

In this section, we present a flowchart aimed at DP practitioners who wish to fix the vulnerabilities presented in Section 5. Table 1.5 presents several solutions and their associated sensitivities.

Solution name		$\frac{\textit{implemented sensitivity}}{\textit{idealized sensitivity}}$	Conditions
RP [Thm. 6.15] Ordered distance [Thm. 6.16] Checked multiplication [Thm. 6.1] Modular noise addition (<code>ints</code>) [Thm. 6.6]		1	
RP + Split summation + RTZ [Thms. 6.30, 6.31]		$\min\{(1 + O(n/2^k)), 2\}$	$\text{sign}(U) = \text{sign}(L)$
		$\min\{(2 + O(n/2^k)), 5\}$	$\text{sign}(U) \neq \text{sign}(L)$
Sensitivity from accuracy + Truncated summation [Thm. 6.23]	Iterative	$1 + \Theta(n^2/2^k)$	
	Pairwise	$1 + O(n \log(n)/2^k)$	$n < 2^k$
	Kahan	$1 + O(n/2^k)$	$n < 2^k$

Table 2: Upper bounds obtained in our solutions for numerical type T and $-U \leq L \leq U$. We remark that in the iterative, pairwise, and Kahan sensitivities, the 1 factor becomes a 2 in the case where n is unknown and $\text{sign}(U) \neq \text{sign}(L)$, as stated in Corollary 6.26.

While implementing some of the solutions would require developers to make significant changes to existing libraries, other solutions can be implemented with only a few alterations to current code.

k -bit integers. We recommend the following solutions:

1. Modular sensitivity (Section 6.2): In many programming languages, the default method for handling integer overflow is wraparound, which is equivalent to modular summation. Thus, many DP libraries luckily already implement the modular solution that we prove to be correct in Theorem 6.5. One point of caution is that both the summation and noise addition steps must occur in this modular fashion (modular summation coupled with non-modular noise addition caused the privacy violation found in IBM’s `diffprivlib` and described in Section 5.1 and demonstrated in Section 5.7.1). Wraparound is a standard feature of arithmetic on integer data types, so this solution should not create new unexpected behaviors for data analysts.

If library maintainers are uncomfortable with the possibility of wraparound or unable to offer modular noise addition, we encourage the following two solutions.

2. Checking parameters (Section 6.1.1): In the bounded DP setting, perform a check on the parameters ensuring that overflow cannot occur, e.g., check that $n \cdot U < 2^k$ in the setting of unsigned k -bit integers.
3. Split summation (Section 6.3): In the unbounded DP setting, we recommend switching to saturation arithmetic (where overflow is handled by clamping to the range $[\min(T), \max(T)]$) and applying split summation, where we separately sum the positive and negative numbers (Method 6.9). For L and U both non-negative or both non-positive, split summation is

equivalent to standard summation. If saturation arithmetic is not possible for the programming language being used, then we recommend applying a randomized permutation to the dataset (Method 6.14) or using an ordered notion of neighboring datasets and being careful about stability with respect to ordering in all dataset transformations.

(k, l) -bit floats. For starters, we recommend using the implemented sensitivity bounds that come from accuracy bounds (see Table 1.5 and Section 6.5) for whatever summation method is already implemented. The summation method could be changed from iterative to pairwise or Kahan for tighter sensitivities (Section 6.5). For bounded DP, using these sensitivities only requires also checking the parameters to ensure that overflow cannot occur (Section 6.1.2). We remark that this is a different check than the one required for integers above. To achieve unbounded DP, a random permutation (Method 6.14) and truncated summation (Method 6.24) should be used as well.

The above solutions should work very well for 64-bit floats, as they give an implemented sensitivity that is at most $1.5 \cdot U$ for datasets of size smaller than 67 million. If it is important to have sensitivity close to $U - L$, then the “shifting bounds” technique (Section 6.5.1) can be used as well. For data consisting of lower-precision floats, one attractive approach is to keep the values as low-precision floats (e.g., to keep the memory footprint small in machine learning pipelines) but accumulate the sum in a 64-bit float, which will allow the above solutions to apply. For huge datasets (e.g., more than 67 million records), a 128-bit accumulator could be used or the summation method could be switched to a method where the accuracy bounds grow more slowly with n (e.g., pairwise or Kahan summation). Otherwise, we recommend considering the solution described in Section 6.6, which requires changing the rounding mode from the standard banker’s rounding to round toward zero. Lastly, if practitioners wish to use fixed-point instead of floating-point arithmetic for small mantissa lengths, then bounded sums of floats could be performed by casting all values in the dataset to fixed-point numbers, using one of our integer summation solution, and then casting back to a float.

2 Preliminaries

2.1 Measuring Distances

Definition 2.1 (Dataset). For a data domain \mathcal{D} , a *dataset on \mathcal{D}* is a vector v of elements from \mathcal{D} , i.e., $v \in \text{Vec}(\mathcal{D}) := \bigcup_{n \geq 0} \mathcal{D}^n$. For $v \in \mathcal{D}^n$, we write $v = [v_1, \dots, v_n]$ for the elements of v and $\text{len}(v) = n$.

Note that datasets $v \in \text{Vec}(\mathcal{D})$ are *ordered*. When we are not concerned with the order of elements in v , we will often refer to a dataset’s *histogram*.

Definition 2.2 (Histogram). For a dataset $v \in \text{Vec}(\mathcal{D})$, the *histogram* of v is the function $h_v : \mathcal{D} \rightarrow \mathbb{N}$ defined as

$$h_v(z) = \#\{i : v_i = z\}.$$

Definition 2.3 (Permuting Datasets). Let S_n denote the set of permutations on n elements. For a permutation $\pi \in S_{\text{len}(v)}$, we write $\pi(v) = [v_{\pi(1)}, \dots, v_{\pi(n)}]$.

Lemma 2.4. For datasets $u, v \in \text{Vec}(\mathcal{D})$, $h_u = h_v$ if and only if $\text{len}(u) = \text{len}(v)$ and there exists a permutation $\pi \in S_{\text{len}(u)}$ such that $\pi(u) = v$.

Depending on which distance function we want to employ, there are different notions of *neighboring datasets*. In this paper we consider the following four distance functions between vectors/datasets: symmetric distance, Hamming distance, change-one distance, and insert-delete distance.

Definition 2.5 (Symmetric Distance). Let $u, v \in \text{Vec}(\mathcal{D})$. The *symmetric distance* between u and v , denoted $d_{Sym}(u, v)$, is

$$d_{Sym}(u, v) = \sum_{z \in \mathcal{D}} |h_u(z) - h_v(z)|,$$

where h_u and h_v are the histograms of u and v .

Equivalently, $d_{Sym}(u, v)$ is the size of the symmetric difference between u and v when viewed as multisets.

Observe that for every permutation $\pi \in S_{\text{len}(u)}$, $d_{Sym}(u, \pi(u)) = 0$, so d_{Sym} (and d_{CO} below) is only a pseudometric since unequal datasets can have distance 0.

Definition 2.6 (Change-One Distance). Let $u, v \in \mathcal{D}^n$. The *change-one distance* between u and v , denoted $d_{CO}(u, v)$, is

$$d_{CO}(u, v) = \sum_{\substack{z \in \mathcal{D} \text{ s.t.} \\ h_u(z) > h_v(z)}} (h_u(z) - h_v(z)) = \sum_{\substack{z \in \mathcal{D} \text{ s.t.} \\ h_v(z) > h_u(z)}} (h_v(z) - h_u(z)).$$

For $u, v \in \text{Vec}(\mathcal{D})$ with $\text{len}(u) \neq \text{len}(v)$, we define $d_{CO}(u, v) = \infty$.

Equivalently, $d_{CO}(u, v)$ is the minimum number of elements in u that need to be changed to produce v , when u and v are viewed as multisets.

Definition 2.7 (Insert-Delete Distance). For $u \in \mathcal{D}^m$, an *insertion* to u is an addition of an element z to some location within u , resulting in a vector $u' = [u_1, \dots, u_i, z, u_{i+1}, \dots, u_m] \in \mathcal{D}^n$. Likewise, a *deletion* from u is a removal of an element from some location within u , resulting in a vector $u' = [u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_m] \in \mathcal{D}^m$.

The *insert-delete distance* $d_{ID}(u, v)$ between $u \in \mathcal{D}^m, v \in \mathcal{D}^n$ is the minimum number of insertion operations and deletion operations needed to change u into some vector u' such that $u' = v$.

Definition 2.8 (Hamming Distance). Let $u, v \in \mathcal{D}^n$. The *Hamming distance* between vectors u, v is

$$d_{Ham}(u, v) = \#\{i : u_i \neq v_i\}.$$

For $u, v \in \text{Vec}(\mathcal{D})$ with $\text{len}(u) \neq \text{len}(v)$, we define $d_{Ham}(u, v) = \infty$.

Equipped with functions for measuring the distance between datasets, we can now define the notion of neighboring / adjacent datasets.

Definition 2.9 (Neighboring/Adjacent Datasets). We say that vectors $u, v \in \text{Vec}(\mathcal{D})$ are *neighbors* or *adjacent* with respect to dataset distance metric d whenever $d(u, v) \leq 1$. In this case, we write $u \simeq_d v$, or $u \simeq_{\text{name}} v$, where d corresponds to d_{name} (e.g., d_{Sym}, d_{Ham}, d_{ID}).

A key difference between these distance metrics, which we will use in our attacks, is the fact that d_{Sym} and d_{CO} are *unordered* metrics (meaning the distance between two vectors only depends on their histograms), whereas d_{Ham} and d_{ID} are *ordered* metrics (meaning the distance between two vectors is affected by the order of elements in these vectors). Another key difference is that d_{Ham} and d_{CO} can only be applied to vectors of the *same* length n . On the other hand, d_{Sym} and d_{ID} apply to vectors of potentially different length. Thus, we will use d_{CO} and d_{Ham} in privacy contexts where the dataset size n is known and public, and d_{Sym} and d_{ID} otherwise. Table 3 summarizes these differences.

	Unordered	Ordered
Unknown n	d_{Sym}	d_{ID}
Known n	d_{CO}	d_{Ham}

Table 3: Summary of the dataset metrics considered in this paper.

The relationships between these four metrics in terms of sensitivity are described in the following lemma, which can be summarized with the following diagram:

$$\begin{array}{ccc}
 d_{Sym} & \xrightarrow{\leq} & d_{ID} \\
 \downarrow = & & \downarrow \leq \\
 2 \cdot d_{CO} & \xrightarrow{\leq} & 2 \cdot d_{Ham}
 \end{array}$$

Lemma 2.10 (Relating Metrics). *These metrics are related as follows.*

1. For $u, v \in \text{Vec}(\mathcal{D})$,

$$d_{Sym}(u, v) = \min_{\pi \in S_{\text{len}(u)}} d_{ID}(\pi(u), v) \leq d_{ID}(u, v).$$

2. For $u, v \in \mathcal{D}^n$,

$$d_{CO}(u, v) = \min_{\pi \in S_n} d_{Ham}(\pi(u), v) \leq d_{Ham}(u, v).$$

3. For $u, v \in \mathcal{D}^n$,

$$d_{Sym}(u, v) = 2 \cdot d_{CO}(u, v).$$

4. For $u, v \in \mathcal{D}^n$,

$$d_{ID}(u, v) \leq 2 \cdot d_{Ham}(u, v).$$

The proof can be found in Appendix A.

2.2 Sensitivity of Functions

Now that we have defined the notion of neighboring datasets, we can define *sensitivity*, a term around which this paper revolves.

Definition 2.11 (Sensitivity). We define the (global) *sensitivity* of a function $f : \text{Vec}(\mathcal{D}) \rightarrow \mathbb{R}$ with respect to a metric d on $\text{Vec}(\mathcal{D})$ as

$$\Delta_d f = \sup_{\substack{u, v \in \text{Vec}(\mathcal{D}) \\ u \simeq_d v}} |f(u) - f(v)|.$$

The sensitivity of f captures how much the output of f can change if we change a single element of the input vector.

2.2.1 The Path Property

In this paper, we state our results in terms of neighboring datasets, as is customary in the DP literature. These results, though, readily generalize to datasets at arbitrary distance. This idea is captured by the *path property*.

Definition 2.12 (Path Property). A metric⁷ d is said to satisfy the *path property* if it fulfills the following two conditions:

1. For all datasets u, v , we have $d(u, v) \in \mathbb{Z}$.
2. If $d(u, v) = d$ then there exists a sequence $u = u^0, u^1, \dots, u^d = v$ such that $d(u^{i-1}, u^i) \leq 1$.

Theorem 2.13 (Applying the Path Property). *For every function $f : \text{Vec}(\mathcal{D}) \rightarrow \mathbb{R}$, every dataset metric d that satisfies the path property, and every $u, v \in \text{Vec}(\mathcal{D})$,*

$$|f(u) - f(v)| \leq \Delta_d f \cdot d(u, v).$$

Proof. The result follows directly from d applications of the triangle inequality to each of the pairs $d(u_{i-1}, u_i)$ given by the path property definition. \square

Lemma 2.14. *The dataset metrics d_{Sym}, d_{CO}, d_{Ham} , and d_{ID} all satisfy the path property.*

The proof can be found in Appendix A.

Due to Lemma 2.14, all of the theorems in this paper will be stated in terms of neighboring datasets, and the path property can be used to generalize these results to apply to any datasets at arbitrary distance d .

Lemma 2.15 (Convert Sensitivities). *We can convert between sensitivities in the following ways.*

1. For every function $f : \text{Vec}(\mathcal{D}) \rightarrow \mathbb{R}$, $\Delta_{ID} f \leq \Delta_{Sym} f$.
2. For every function $f : \mathcal{D}^n \rightarrow \mathbb{R}$, $\Delta_{Ham} f \leq \Delta_{CO} f$.
3. For every function $f : \mathcal{D}^n \rightarrow \mathbb{R}$, $\Delta_{CO} f \leq 2\Delta_{Sym} f$.
4. For every function $f : \mathcal{D}^n \rightarrow \mathbb{R}$, $\Delta_{Ham} f \leq 2\Delta_{ID} f$.

The proof can be found in Appendix A.

⁷We use the term *metric* throughout the paper, although some of our “metrics” (e.g., symmetric distance and change-one distance) allow distance 0 for unequal elements and should be considered *pseudometrics* instead.

2.3 Differential Privacy

We now recall the definition of (pure) differential privacy, also known as ϵ -DP. Although some of the libraries that we investigate also offer the more general (ϵ, δ) -DP, they all offer ϵ -DP, and we only consider the pure-DP versions of these implementations. Before proceeding with the definition of DP, we define a *mechanism*.

Definition 2.16 (Mechanism). A *mechanism* with input space $\text{Vec}(\mathcal{D})$ and output space \mathcal{Y} is a randomized algorithm $\mathcal{M} : \text{Vec}(\mathcal{D}) \rightarrow \mathcal{Y}$ that on input $u \in \text{Vec}(\mathcal{D})$ outputs a sample from the distribution $\mathcal{M}(u)$ over \mathcal{Y} .

Definition 2.17 (Pure Differential Privacy). For any $\epsilon \geq 0$, a mechanism $\mathcal{M} : \text{Vec}(\mathcal{D}) \rightarrow \mathcal{Y}$ is ϵ -differentially private with respect to metric d on $\text{Vec}(\mathcal{D})$ if for all subsets $S \subset \mathcal{Y}$ and for all adjacent datasets $u \simeq_d u'$,

$$\Pr[\mathcal{M}(u) \in S] \leq e^\epsilon \cdot \Pr[\mathcal{M}(u') \in S].$$

We note that differential privacy is “robust to post-processing”. Intuitively, this means that a data analyst – with any amount of additional knowledge – cannot take an ϵ -DP answer and make it less private.

Proposition 2.18 (DP is Robust to Post-Processing [14]). *Let $\mathcal{M} : \text{Vec}(\mathcal{D}) \rightarrow \mathcal{Y}$ be a randomized algorithm that is ϵ -DP. Let $f : \mathcal{Y} \rightarrow \mathcal{Z}$ be an arbitrary, randomized mapping. Then $f \circ \mathcal{M} : \text{Vec}(\mathcal{D}) \rightarrow \mathcal{Z}$ is ϵ -DP.*

Proof. See the proof of Proposition 2.1 in [14]. □

2.3.1 The Laplace Mechanism

A common way to fulfill DP is to have mechanisms that add noise to a true query response, with the noise scaled to the sensitivity of the underlying function. These mechanisms are called *additive mechanisms*. An elementary example is the Laplace mechanism, which was first shown to fulfill ϵ -DP in [13].

Definition 2.19 (The Laplace Mechanism [13]). Given a function $f : \text{Vec}(\mathcal{D}) \rightarrow \mathbb{R}$, the *Laplace mechanism for f with scale λ* is defined as

$$M(u) = f(u) + Y, \text{ where } Y \sim \text{Lap}(\lambda).$$

where $\text{Lap}(\lambda)$ denotes the zero-centered continuous probability distribution defined by the density function

$$g_\lambda(x) := \frac{1}{2\lambda} \exp\left(-\frac{|x|}{\lambda}\right).$$

Theorem 2.20. *The Laplace mechanism for f with scale λ is ϵ -DP with respect to dataset metric d if and only if $\lambda \geq \Delta_d f / \epsilon$.*

Proof. The forward direction is proven in [13]. For the other direction, suppose $\lambda < \Delta_d f / \epsilon$. Then there are $u \simeq_d u'$ such that $|f(u) - f(u')| > \lambda \epsilon$. Evaluating the PDFs for $M(u)$ and $M(u')$ at the point $f(u)$ then yields the fraction

$$\frac{\exp\left(-\frac{|f(u) - f(u)|}{\lambda}\right)}{\exp\left(-\frac{|f(u') - f(u)|}{\lambda}\right)} > \frac{1}{\exp\left(-\frac{\lambda \epsilon}{\lambda}\right)} = \exp(\epsilon).$$

Because the ratio of probability densities is not upper-bounded by e^ε , the Laplace mechanism with scale $\lambda < \Delta_{df}/\varepsilon$ does not offer ε -DP. \square

Therefore, even with a correctly implemented noise addition mechanism, incorrectly calibrated noise will imply that DP will not be fulfilled. In this paper, we demonstrate how the sensitivity is frequently underestimated, which means that the DP condition is not fulfilled even when the mechanism would fulfill DP if its noise were scaled correctly.

2.4 Integers

In this section, we review the facts about integer representation on computers that we require in our proofs.

2.4.1 Integer Values

Integers can be *signed* or *unsigned*: the signed integers can hold both positive and negative values, whereas the unsigned integers only hold non-negative values. Table 4 summarizes the data types that we will consider.

Definition 2.21 (Unsigned integers). A *k-bit unsigned integer* can take on any value in

$$[0, 2^k - 1] \cap \mathbb{Z}.$$

Definition 2.22 (Signed integers). A *k-bit signed integer* can take on any value in

$$[-2^{k-1}, 2^{k-1} - 1] \cap \mathbb{Z}.$$

2.4.2 Integer Arithmetic

In this paper, we are only concerned with the *addition* of integers. We use two varieties of integer addition: modular addition and saturating addition.

Many libraries and programming languages default to using modular addition when adding values of integer types. Before defining modular addition, we first note that the *k-bit unsigned integers* and *k-bit signed integers* each consist of unique representations of the congruence classes in $\mathbb{Z}/(2^k)\mathbb{Z}$.

Definition 2.23 (Modular Addition on the *k-bit Integers*). Let *T* be a (signed or unsigned) type for *k-bit integers*. For all values *x, y* of type *T*, we define $[x + y]_{2^k}$ as the unique value *z* such that $z \equiv x + y \pmod{2^k}$ and such that *z* is of type *T* (meaning $z \in [0, 2^k - 1] \cap \mathbb{Z}$ for unsigned *T*, and $z \in [-2^{k-1}, 2^{k-1} - 1] \cap \mathbb{Z}$ for signed *T*).

Saturation arithmetic is a strategy commonly used to prevent integer overflow from occurring. In our attacks and in our proposed solutions, we consider the effects of using saturating addition. We describe its specific behavior below.

Definition 2.24 (Saturation Addition). Let *T* be a (signed or unsigned) type for *k-bit integers*. Additionally, let $\max(T)$ represent the maximum representable integer of type *T*, and let $\min(T)$

represent the minimum representable integer of type T . Also, let $+$ represent addition in \mathbb{Z} , and let \boxplus represent saturation addition. For all values x, y of type T , we define

$$x \boxplus y = \begin{cases} \min(T) & \text{if } x + y < \min(T) \\ x + y & \text{if } \min(T) \leq x + y \leq \max(T) \\ \max(T) & \text{if } x + y > \max(T) \end{cases}$$

2.5 Floating-Point Representation

In this section, we review the facts about floating-point representation that we require in our proofs.

2.5.1 Floating-Point Values

According to the IEEE floating-point standard, floating-point numbers are represented with an exponent E and a mantissa m [35], in a base-2 analog of scientific notation. More precisely,

Definition 2.25 (Normal Floating-Point Number). A *normal* (k, ℓ) -bit floating-point number z is represented as

$$z = (-1)^s \cdot (1.M) \cdot 2^E,$$

where

- $s \in \{0, 1\}$ is used to represent the *sign* of z .
- $M \in \{0, 1\}^k$ is a k -bit string that represents the part of the *mantissa* to the right of the radix point. That is,

$$1.M = 1 + \sum_{i=1}^k M_i 2^{-i}.$$

- $E \in \mathbb{Z}$ represents the *exponent* of 2. When ℓ bits are allocated for representing E , then $E \in [-(2^{\ell-1} - 2), 2^{\ell-1} - 1] \cap \mathbb{Z}$.

Note that the number of possibilities for E is $2^\ell - 2$ rather than 2^ℓ . The remaining two choices for an ℓ -bit E are used to represent the following additional floating-point numbers.

The key characteristic of normal floating-point representation is that representable values are spaced *non-uniformly* throughout the real line [38]: for a (k, ℓ) -bit float, and for $m \in \mathbb{Z}$, successive intervals from 0 to ∞ of the form $[2^m, 2^{m+1})$ double in length (with the length always corresponding to a power of 2), and each interval contains exactly 2^k representable real values. Analogously, successive intervals from 0 to $-\infty$ are of the form $[-2^m, -2^{m+1})$ and double in length. This implies that, between each successive interval $[2^m, 2^{m+1})$ and $[2^{m+1}, 2^{m+2})$, and between $[-2^m, -2^{m+1})$ and $[-2^{m+1}, -2^{m+2})$, the spacing between adjacent representable floats also doubles. This phenomenon is depicted in Figure 5.

The definition of $ULP_{(k, \ell)}$ provided below enables discussion of the precision available when working with a normal (k, ℓ) -bit floating-point number z .

Definition 2.26 (ULP). We define $ULP_{(k, \ell)}(z)$ as the *unit in the last place* for a normal (k, ℓ) -bit floating-point number z ; i.e., the place-value of the least significant digit of a floating-point number z . That is, if $z = (-1)^s \cdot (1.M) \cdot 2^E$, then $ULP_{(k, \ell)}(z) = 2^{E-k}$.

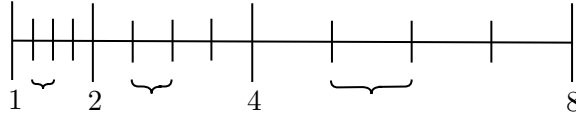


Figure 5: Depiction of the increased spacing between adjacent representable floats. For demonstration purposes, the mantissa is limited to a length of 2 bits. The curly braces illustrate how the distance between adjacent representable floats doubles across successive intervals of the form $[2^k, 2^{k+1})$.

We generalize this notion to apply to all $z \in \mathbb{R} \setminus \{0\}$. For z such that $|z| \in [2^m, 2^{m+1})$ for $m \in \mathbb{Z}$, we define $ULP_{(k,\ell)}(z) = 2^{m-k} = 2^{\lfloor \log_2 |z| \rfloor - k}$.

Lemma 2.27 shows how $ULP_{(k,\ell)}$ can be used to determine whether a value $z \in \mathbb{R}$ can be exactly represented as a normal (k, ℓ) -bit float.

Lemma 2.27. *A number $z \in \mathbb{R}$ can be represented exactly as a normal (k, ℓ) -bit float if and only if*

1. z is an integer multiple of $ULP_{(k,\ell)}(z) = 2^{\lfloor \log_2 |z| \rfloor - k}$, and
2. $\lfloor \log_2 |z| \rfloor \in [-(2^{\ell-1} - 2), 2^{\ell-1} - 1]$.

Proof. We prove each direction.

- (\Rightarrow) Let z be a normal (k, ℓ) -bit float. Then, z must be of the form $(1.j) \cdot 2^m = (1 + j \cdot 2^{-k}) \cdot 2^m$ for some $j \in \{0, 1\}^k$ and $m \in [-(2^{k-1} - 2), 2^{k-1} - 1]$. Equivalently, we must have $z = \text{sign}(z) \cdot (2^k + j) \cdot 2^{m-k}$. Therefore, because $(2^k + j)$ is an integer, z is an integer multiple of 2^{m-k} , and $\lfloor \log_2 |z| \rfloor = m \in [-(2^{k-1} - 2), 2^{k-1} - 1]$.
- (\Leftarrow) Let $m = \lfloor \log_2 |z| \rfloor$, and suppose z is an integer multiple of $ULP_{(k,\ell)}(z) = 2^{m-k}$. We can then write $|z| = 2^m + j \cdot 2^{m-k}$ for some $j \in [0, 2^k) \cap \mathbb{Z}$. Equivalently, $z = \text{sign}(z) \cdot (1 + j \cdot 2^{-k}) \cdot 2^m$.

The set of the bitstrings of the form $\{0, 1\}^k$ is the full set of possible values of j , so $(1 + j \cdot 2^{-k})$ corresponds directly with a k -bit mantissa $1.j$. We can write $z = \text{sign}(z) \cdot (1.j) \cdot 2^m$, for $m \in [-(2^{\ell-1} - 2), 2^{\ell-1} - 1] \cap \mathbb{Z}$ and $j \in \{0, 1\}^k$, so, by definition, z is representable as a normal (k, ℓ) -bit float.

This completes the proof. □

Definition 2.28 (NRF). Let $NRF(z)$ be to the *nearest (bigger) representable float* for a normal (k, ℓ) -bit floating-point number z . That is, $NRF(z)$ is the floating-point value z' such that $|z'| > |z|$ and such that $\forall z'' \neq z$ where $|z''| > |z|$, we have $|z' - z| \leq |z'' - z|$. (For the $z = 0$ case, we require that $ULP_{(k,\ell)}(z)$ is positive.)

Definition 2.29 (Subnormal Floats). A *subnormal (k, ℓ) -bit float* z is represented as

$$z = (-1)^s \cdot (0.M) \cdot 2^E,$$

where

- $s \in \{0, 1\}$ is used to represent the *sign* of z .
- $M \in \{0, 1\}^k$ is a k -bit string that represents the part of the *mantissa* to the right of the radix point. That is,

$$0.M = \sum_{i=1}^k M_i 2^{-i}.$$

- $E = -(2^{\ell-1} - 2)$. Note that E is a constant.

Definition 2.30 ($\pm\text{inf}$). Let n_{max} and n_{min} be the largest and smallest (k, ℓ) -bit normal floating-point numbers. Then, for arithmetic and rounding purposes, $\text{inf} = n_{max} + \text{ULP}_{(k,\ell)}(n_{max})$ and $-\text{inf} = n_{min} - \text{ULP}_{(k,\ell)}(n_{min})$. This means, for example, that any real number $x \geq n_{max} + \text{ULP}_{(k,\ell)}(n_{max})$ is rounded to inf . We assume the behavior that, for all $z \neq \pm\text{inf}$, $\text{inf} + z = \text{inf}$ and $-\text{inf} + z = -\text{inf}$.

Definition 2.31 ((k, ℓ) -bit Float). The set of (k, ℓ) -bit floats is the union of the set of all normal (k, ℓ) -bit floats, all subnormal (k, ℓ) -bit floats, and $\pm\text{inf}$. Any element of this set is termed a (k, ℓ) -bit float.

In the libraries of DP functions that we investigate, floating-point numbers come in two main varieties: single-precision floating-point numbers, which use 32 bits of memory; and double-precision floating-point numbers, which use 64 bits of memory.

Definition 2.32 (32-Bit and 64-Bit Floats). We define the two types of floating-point numbers we encounter.

1. *Single-precision floating-point* numbers are (23, 8)-bit floating-point numbers. We denote the set of single-precision floats with the symbol \mathbb{S} .
2. *Double-precision floating-point* numbers are (52, 11)-bit floating-point numbers. We denote the set of double-precision floats with the symbol \mathbb{D} .

As described in Section 2.4, integers can be *signed* or *unsigned*: the signed integers can hold both positive and negative values, whereas the unsigned integers only hold non-negative values. Table 4 summarizes the data types that we will consider.

Integers (Sec. 2.4)	32-bit signed	64-bit signed
	32-bit unsigned	64-bit unsigned
Floats (Sec. 2.5)	32-bit	64-bit

Table 4: Summary of the data types considered in this paper.

2.5.2 Rounding Modes

Because floating-point representations on computers have a limited number of bits, sometimes rounding is necessary. There are many different rounding modes, but the IEEE 754 standard [35] for floating-point arithmetic prescribes using *banker’s rounding* as the default rounding mode. Thus in this paper (unless stated otherwise) we will always assume the use of banker’s rounding.

Definition 2.33 (Banker’s Rounding). For $x \in \mathbb{R}$, we define $BRound(x)$ to be the floating-point value z such that, for all floating-point values $z' \neq z$, we have $|z - x| \leq |z' - x|$. In the event of a tie (meaning x is equally close to some z and some $z' \neq z$), x is rounded to the value $\in \{z, z'\}$ whose mantissa ends in an even bit.

There is one more rounding mode that we consider in this paper: *round toward zero*.

Definition 2.34 (Round Toward Zero). For $x \in \mathbb{R}$, we define $RTZ(x)$ to be the floating-point value z such that $|z| < |x|$ and for all floating-point values with $|v'| < |x|$, we have $|v - x| \leq |v' - x|$.

Lastly, we remark that *fixed-point arithmetic* constitutes an alternative way of representing real numbers on computers, which is simpler than floating-point arithmetic. A fixed-point representation of a fractional number is essentially treated as an integer, which is implicitly multiplied by a fixed scaling factor. Thus, a fixed-point number has a fixed number of digits after the decimal point, whereas a floating-point number allows for a varying number of digits after the decimal point.

2.5.3 Floating-Point Arithmetic

In this paper, we are only concerned with the *addition* of floating-point numbers. We use two varieties of floating-point addition: addition using banker’s rounding, and addition using round toward zero.

Before providing these definitions, though, we provide a note about the behavior of all floating-point addition.

Note 2.35 (Floating-Point Arithmetic Saturates). All floating-point addition is *saturating* (see Definition 2.24), with the exception that, as described in Section 2.5.2, the results of computations are floating-point values (which means that rounding occurs if necessary). That is, if the sum $x + y$ would be larger than $+\text{inf}$ (respectively, smaller than $-\text{inf}$), then the result returned is $+\text{inf}$ (respectively, $-\text{inf}$).

Definition 2.36 (Addition With banker’s Rounding). We define $x \oplus y = BRound(x + y)$, where $BRound(z)$ returns $z \in \mathbb{R}$ banker’s rounded to a floating-point value. (See Definition 2.33 for the definition of banker’s rounding.)

Definition 2.37 (Addition With Round Toward Zero). We use $RTZ(x + y)$ to denote the result of adding two floating-point values x and y using round toward zero, where $RTZ(z)$ returns $z \in \mathbb{R}$ “rounded toward zero” to a floating-point value. (See Definition 2.34 for the definition of round toward zero.)

Lemma 2.38. For (k, ℓ) -bit floats x, y , $x + y = BRound(x + y) = RTZ(x + y)$ if and only if $x + y$ is a (k, ℓ) -bit float.

Proof. This follows immediately from Definitions 2.36 and 2.37. □

3 An Overview of DP Libraries

Google DP Library. Google’s DP Library [49] was first released in September 2019 and remains under active development. It already supports many fundamental functions, such as

count, sum, mean, variance, the Laplace mechanism, and the Gaussian mechanism, among others.

The repository (<https://github.com/google/differential-privacy>) contains an end-to-end differential privacy framework built on top of Apache Beam, and three DP building block libraries, in C++, Go, and Java, which implement basic noise addition primitives and differentially private aggregations. It also contains a differential privacy accounting library, which is used for tracking privacy budget, and a command line interface for running differentially private SQL queries with ZetaSQL. The GoogleDP library contains documentation concerning sampling algorithms for the Laplace and Gaussian distributions⁸ which they claim circumvent problems with naïve floating-point implementations. These sampling mechanisms essentially follow the snapping mechanism proposed by Mironov [38].

Specifically, this document says, “A key property that we rely on is that the IEEE floating-point standard guarantees that the results of basic arithmetic operations are the same as if the computation was performed exactly and then rounded to the closest floating-point number.” Although this is true for a single operation, this is *not* true when several basic arithmetic operations are *iteratively* applied, as it is the case of bounded sum. Issues related to accumulated rounding and non-associativity of floating-point addition can arise when more than a single addition is performed (as happens when computing the bounded sum of a dataset with more than two elements). We demonstrate the severity of these issues in Section 5.

OpenMined PyDP. A related library is OpenMined’s PyDP (<https://github.com/OpenMined/PyDP>), which consists of a Python wrapper for Google’s DP library. Similarly, they support the fundamental functions of bounded mean, bounded sum, max, min, and median, among others. For now, they only use Laplace noise.

SmartNoise/OpenDP. SmartNoise is a toolkit for differential privacy on tabular data built through a collaboration between Harvard and Microsoft as part of the OpenDP project. It was originally driven by the SmartNoise-Core library of differentially private algorithms, which has since been deprecated and replaced by the OpenDP library (<https://opendp.org/>). The OpenDP library has a stringent framework for expressing and validating privacy properties [18]. The library is built in Rust but includes several Python bindings.

OpenDP releases are built by assembling a number of constituent transformations and measurements using operations such as “chaining” and “composition”. As the other libraries, it supports fundamental functions such as count, clamp, mean, bounded sum, and resize, and it also supports other types of noise such as Gaussian noise. Also like the other libraries, it maintains a privacy budget tracker.

IBM’s diffprivlib. IBM’s differentially private library [26] is implemented in Python and is more geared towards applications of DP in machine learning (<https://github.com/IBM/differential-privacy-library>). The library includes a host of mechanisms alongside a number of applications to machine learning and other data analytics tasks.

In addition to implementing fundamental functions such as bounded sum, the library implements DP mechanisms such as the Exponential mechanism, the Geometric mechanism, and the Bingham mechanism. Due to floating-point imprecision, functions in the library sample from the Laplace distribution using the method described in Holohan and Braghin [25]. The library also implements tools for computing DP histograms, training naive DP Bayes classifiers,

⁸   .

and performing linear regressions. The library is organized into three main modules: mechanisms, models, and tools. Like the other libraries described in this section, it uses a so-called budget accountant to track the privacy budget and calculate total privacy loss using advanced composition techniques.

Opacus. Facebook’s DP library [1] is called Opacus and is heavily geared towards machine learning applications of differential privacy, and so in particular they implement the well-known differentially-private Stochastic Gradient Descent (DP-SGD) function [1]. Their library (<https://github.com/pytorch/opacus>) explicitly enables training PyTorch models with differential privacy.

Opacus allows building image classifiers with DP, training a differentially private LSTM model for name classification, and building a text classifier with differential privacy on BERT, among others. It also maintains a privacy tracking with an accountant. A main difference between Opacus and the rest of the libraries is that it is based on Rényi differential privacy [39].

Chorus. Chorus [30] consists of a query analysis and rewriting framework to enforce differential privacy for general-purpose SQL queries. The library (<https://github.com/uber-archive/sql-differential-privacy>) was deployed at Uber for its internal analytics tasks as part of the company’s efforts to comply with the General Data Protection Regulation (GDPR). Chorus supports integration with any standard SQL database and is focused on the large-scale deployment of DP methods, designed to process datasets consisting of billions of rows. At Uber, Chorus processed more than 10,000 queries per day, and they evaluated 18,774 real-world queries with a database of 300 million rows. Beyond implementing fundamental functions such as sum and the Laplace mechanism, they also implement more complex functions such as the matrix mechanism.

While these are the modern main libraries that we will use to demonstrate our attacks, other relevant DP libraries and systems include PINQ [36], Airavat [44], Fuzz [22], and Ektelo [51].

4 Bounded Sum

4.1 Bounded Sum is a Building Block in DP Libraries

In this paper, we focus on the bounded sum function, an elementary function which adds the elements of a (numerical) dataset after clipping all values in the dataset to a bounded range $[L, U]$ and then returns the resulting sum. This function is important for two reasons: first, it is simple, and thus allows us to clearly convey the general problems that arise from the use of finite precision arithmetic. Second, it is the building block of many important and complex DP functions, such as functions used for computing averages and performing stochastic gradient descent (SGD). For this reason, bounded sum is a common function in DP libraries, and it appears in Google’s DP library, SmartNoise / OpenDP, IBM’s `diffprivlib`, Chorus, and Facebook’s Opacus, to name a few. For these reasons, bounded sum is often a function that is chosen as a motivating example when presenting DP libraries, as was done for Chorus [30] and Airavat [44].

Whenever the bounded sum function is implemented, the code implementation must specify the sensitivity of the function. We will now review the sensitivity of this function over \mathbb{R} (i.e., its *idealized* sensitivity) and point to evidence that this is the sensitivity used by DP libraries. In Section 5, we will show that this idealized sensitivity is not always exhibited by the

implementation of the bounded sum algorithm, and we will show how mismatches between the idealized sensitivity can cause blatant privacy violations.

4.2 The Idealized Bounded Sum Function

We begin by formally defining the bounded sum function.

Definition 4.1 (Clamped Domain). For a numeric domain or data type T , and $L \leq U \in T$ we define

$$T_{[L,U]} = \{x \in \mathbb{T} \mid L \leq x \leq U\}.$$

Definition 4.2 (Bounded Sum). For a numeric domain or data type T , and $L \leq U \in T$ we define the *bounded sum* function $BS_{L,U} : \text{Vec}(T_{[L,U]}) \rightarrow T$ as

$$BS_{L,U}(v) := \sum_{i=1}^{\text{len}(v)} v_i.$$

We write $BS_{L,U,n}$ to denote the restriction of $BS_{L,U}$ to $T_{[L,U]}^n$.

Remark 4.3. To differentiate between the ideal and real-world realizations of the bounded sum function, we will later use BS when referring to the bounded sum function over \mathbb{R} (i.e., with infinite precision) and BS^* when referring to the computer implementations of it.

4.3 Sensitivity of the Idealized Bounded Sum

We now recall the sensitivity for the bounded sum function $BS_{L,U}$ over \mathbb{Z} and \mathbb{R} . Because the sensitivity depends on whether or not n is known, we will prove the sensitivity for each case.

Theorem 4.4 (Idealized Sensitivities of $BS_{L,U}$ and $BS_{L,U,n}$). *The sensitivities of the bounded sum function are the following.*

1. (Unknown n .) $\Delta_{Sym} BS_{L,U} = \Delta_{ID} BS_{L,U} = \max\{|L|, U\}$.
2. (Known n .) $\Delta_{CO} BS_{L,U,n} = \Delta_{Ham} BS_{L,U,n} = U - L$.

The proof can be found in Appendix B.

4.3.1 Uses of the Idealized Sensitivity in Libraries of DP Functions

Noise is scaled according to the idealized sensitivity of the bounded sum function in following libraries:

- Unknown n (Theorem 4.4, Part 1): Google’s sum function,⁹ SmartNoise’s sum function,¹⁰ Opacus,¹¹ Chorus [30, §3], and Airavat [44, §4.1].

⁹Exact link: [https://github.com/google/differential-privacy/blob/master/ops/sum.py#L46](#)

¹⁰Exact link: [https://github.com/SmartNoise/smartnoise-dp/blob/master/ops/sum.py#L26](#)

¹¹Exact link: [https://github.com/Opacus/differential-privacy/blob/master/ops/sum.py#L16](#)

- Known n (Theorem 4.4, Part 2): IBM `diffprivlib`'s sum and mean,¹² Google's mean,¹³ and SmartNoise's sized sum function.¹⁴

Note 4.5. SmartNoise-Core is deprecated, and, though the functions have been ported to OpenDP, they are now kept behind the `contrib` and `floating-point` flags with the disclaimer that no guarantees are made about whether these functions offer differential privacy; the other libraries listed above and mentioned in Section 3 do not appear to have such a disclaimer.

As far as we are aware, none of these libraries concretely specify the metric used for defining neighboring datasets, although we infer that all of them implicitly use an unordered distance metric.¹⁵ If the sensitivity used by a library matches the sensitivity with respect to a known n distance, we infer that the library defines neighboring datasets in terms of known n distances, and likewise for unknown n distances.

However, as we develop below in Section 5, computers work with finite precision arithmetic, and thus implementations of the bounded sum function no longer fulfill the idealized sensitivity of Theorem 4.4. We will now analyze the mismatches between the idealized and the implemented sensitivities of the bounded sum function, and show that they can lead to blatant violations of privacy. Thus, whenever the libraries above work with floating-point types, they should use a different sensitivity bound than is used in their code.

5 Implementing Bounded Sum on Computers

Iterative summation. While the proofs of idealized sensitivity presented in Section 2.5 are over the reals \mathbb{R} , computers cannot work with the reals and typically use finite-precision integers or the floating point numbers, as defined in Section 2.4 and Section 2.5, respectively. Note that, as occurs in many of the libraries discussed in Section 3, the sum is calculated *iteratively*, meaning that the sum is accumulated in a single variable, and that each term in the vector is added to this variable in order of occurrence in the vector. As before, we assume that all elements of the dataset u are numerical elements of the same type. Libraries thus implement a bounded sum function $BS_{L,U}^*$ in the following iterative way (written in Python-style pseudocode):

Definition 5.1. We let $BS_{L,U}^* : T_{[L,U]} \rightarrow T$ denote the iterative bounded sum function on type T as follows:

```

1 def bounded_sum(u):
2     total = 0
3     for element in u:
4         total += element
5     return total

```

¹²Exact link for sum: <https://diffprivlib.org/docs/sum/#sum>
<https://diffprivlib.org/docs/sum/#sum>
<https://diffprivlib.org/docs/sum/#sum>

Exact link for mean: <https://diffprivlib.org/docs/mean/#mean>

¹³Exact link: <https://diffprivlib.org/docs/mean/#mean>
<https://diffprivlib.org/docs/mean/#mean>

¹⁴Exact link: <https://diffprivlib.org/docs/sum/#sum>
<https://diffprivlib.org/docs/sum/#sum>

¹⁵E.g., see the documentation <https://diffprivlib.org/docs/sum/#sum> for the GoogleDP library.

We remark that in Section 6.5 we consider other orderings than the iterative one when performing the summation of the elements u_i in a database, such as *pairwise summation* (also known as *cascade summation*) *Kahan summation*. We introduce both summation methods in this subsection.

Definition 5.2 (Pairwise Summation [24]). Given n values u_i , for $i = 1, \dots, n$, the *pairwise summation* method consists of adding the u_i as follows:

```

1 def pairwise(u):
2     if n < 2:
3         return u[0]
4     m = int(n/2)
5     return pairwise(u[0:m]) +
6         pairwise(u[m:n])

```

That is, we add the u_i terms as follows:

$$v_i = u_{2i-1} + u_{2i}, \quad i = 1 : \lfloor n/2 \rfloor$$

($y_{\lfloor (n+1)/2 \rfloor} = u_n$ if n is odd). This pairwise summation process is repeated recursively on the v_i , $i = 1 : \lfloor (n+1)/2 \rfloor$. The final sum is obtained in $\lceil \log_2 n \rceil$ steps.

For example, to add 4 values u_i , we would compute $((u_1 + u_2) + (u_3 + u_4))$. Pairwise summation is the default summation mode in NumPy.

Another summation method is that of *Kahan's summation*, which is also known as *compensated summation*.

Definition 5.3 (Kahan Summation [33]). Given n values u_i , for $i = 1, \dots, n$, *Kahan's summation* method consists of adding the u_i values as follows:

```

1 def kahan(u):
2     s = 0.0
3     c = 0.0
4     for i in u:
5         y = i - c
6         t = s + y
7         c = (t - s) - y
8         s = t
9     return s

```

We will return to the question of the summation function in Section 6.5. Meanwhile, for the rest of Section 5 we will focus on detailing the attacks. We will present two different styles of attacks which are based on two different principles: those which take advantage of non-associativity, and those which take advantage of rounding errors.

Additionally, while we do not run attacks on all DP libraries, similar issues to the ones we demonstrate throughout this section can occur in these libraries. In Section 4.3 we pointed out where exactly in the code of the different DP libraries one can find the implementation of the sensitivity of the bounded sum function, which in all cases corresponds to the idealized sensitivity (Theorem 4.4). As we have shown throughout this section, the idealized sensitivity does not hold in practice, and hence the same vulnerabilities that we exploit in this section also hold for the rest of the libraries listed in Section 4.3.

5.1 Overflow Attack on Integers with Modular Addition

We prove that the sensitivities $\Delta_{ID}BS_{L,U}^*$ and $\Delta_{Ham}BS_{L,U,n}^*$ on the integers with modular addition (defined in Definition 2.23) are in fact much larger than the idealized sensitivities $\Delta_{ID}BS_{L,U}$ and $\Delta_{Ham}BS_{L,U,n}$ presented in Section 4.3.

Theorem 5.4 ($\Delta_{ID}BS_{L,U}^*$ and $\Delta_{Ham}BS_{L,U,n}^*$ on the Integers with Modular Addition). *Let $BS_{L,U}^* : Vec(T) \rightarrow T$ be iterative bounded sum with modular addition on the type T of k -bit integers. For all $U \geq 1$ and $L = 0$, the sensitivity $\Delta_{ID}BS_{L,U}^* = 2^k - 1 \geq \Delta_{ID}BS_{L,U} = \max\{|L|, U\}$. In fact, for all n satisfying the condition*

$$n = \left\lceil \frac{\max(T)}{U} \right\rceil + 1,$$

there are datasets $u, v \in T_{[L,U]}^n$ such that $d_{Ham}(u, v) = 1$ and $|BS_{L,U,n}^*(u) - BS_{L,U,n}^*(v)| = 2^k - 1$, so $\Delta_{Ham}BS_{L,U,n}^* = 2^k - 1$.

Example Attack 5.5 (64-bit Unsigned Integers). Let $L = 0$ and $U = 2^{47}$. Theorem 5.4 gives us datasets $u, v \in T_{[L,U]}^n$ of size

$$n = \left\lceil \frac{\max(T)}{U} \right\rceil + 1 = 2^{17} + 1$$

where

$$|BS_{L,U,n}^*(u) - BS_{L,U,n}^*(v)| = 2^{64} - 1.$$

The idealized sensitivities presented in Section 4.3, however, suggest a maximum difference in sums of $\Delta_{Ham}BS_{L,U} = U - L = 2^{47}$. However, u and v actually experience a difference in sums of $2^{64} - 1$, more than a factor of 2^{16} larger than these idealized sensitivities. This means, then, that a DP mechanism that claims to offer ε -DP here but calibrates its random distribution to the idealized sensitivity would instead offer no better than $2^{16}\varepsilon$ -DP.

Proof of Theorem 5.4. Let $M = \max(T) - (n - 2) \cdot U$. We see that $M \in [L, U]$. Consider the datasets

$$u = [U_1, \dots, U_{n-2}, M], v = [U_1, \dots, U_{n-2}, M, 1].$$

By the definition of modular addition in Definition 2.23, we see that

$$BS_{L,U}^*(u) = U_1 + \dots + U_{n-2} + M = \max(T),$$

while

$$BS_{L,U}^*(v) = (\max(T) + 1) \bmod 2^k$$

such that $BS_{L,U}^*(v) \in [\min(T), \max(T)]$. This means, then, that $BS_{L,U}^*(v) = \min(T)$. Therefore,

$$|BS_{L,U}^*(u) - BS_{L,U}^*(v)| = \max(T) - \min(T) = 2^k - 1,$$

so, because $d_{ID}(u, v) = 1$, $\Delta_{ID}BS_{L,U}^* = 2^k - 1 \geq \max\{|L|, U\} = \Delta_{ID}BS_{L,U}$.

Now, consider $u' = [U_1, \dots, U_{n-2}, M, 0]$.

$$BS_{L,U}^*(u') = U_1 + \dots + U_{n-2} + M = \max(T),$$

so, because $d_{Ham}(u', v) = 1$, $\Delta_{Ham}BS_{L,U,n}^* = 2^k - 1 \geq (U - L) = \Delta_{Ham}BS_{L,U,n}$. □

5.2 Reordering Attack on Signed Integers with Saturation Addition

We prove that the sensitivity $\Delta_{Sym} BS_{L,U}^*$ on the signed integers with saturation addition (defined in Definition 2.24) is in fact much larger than the idealized sensitivity $\Delta_{Sym} BS_{L,U}$ presented in Section 4.3.

Theorem 5.6 ($\Delta_{Sym} BS_{L,U}^*$ With Saturation Addition on the Integers). *Let $BS_{L,U}^* : Vec(T) \rightarrow T$ be iterative bounded sum with saturation addition on the type T of k -bit signed integers. For $U > 0$ and $L < 0$, the sensitivity is $\Delta_{Sym} BS_{L,U}^* = 2^k - 1 > \Delta_{Sym} BS_{L,U}$.*

In fact, for all n satisfying the condition

$$n \geq \left\lceil \frac{\max(T) - \min(T)}{|L|} \right\rceil + \left\lceil \frac{\max(T) - \min(T)}{U} \right\rceil,$$

there are datasets $u, v \in T_{[L,U]}^n$ with $d_{CO}(u, v) = 0$ such that $|BS_{L,U}^(u) - BS_{L,U}^*(v)| = 2^k - 1$.*

Example Attack 5.7 (32-bit Signed Integers). Let $L = -2^{14} = -16,384$, $U = 2^{15} = 32,768$, and $k = 32$. Theorem 5.6 gives us datasets $u, v \in T_{[L,U]}^n$ of size

$$n = \left\lceil \frac{\max(T) - \min(T)}{|L|} \right\rceil + \left\lceil \frac{\max(T) - \min(T)}{U} \right\rceil$$

where

$$|BS_{L,U,n}^*(u) - BS_{L,U,n}^*(v)| = 2^{32} - 1,$$

despite the fact that $d_{Sym}(u, v) = 0$ would suggest that the difference in sums should be 0.

Moreover, the idealized sensitivities presented in Section 4.3 suggest a maximum difference in sums of $\Delta_{Sym} BS_{L,U} = \max\{|L|, |U|\} = U = 2^{15}$ and $\Delta_{CO} BS_{L,U,n} = (U - L) = 2^{15} + 2^{14}$. However, u and v actually experience a difference in sums of $2^{32} - 1$, more than a factor of 2^{16} larger than these idealized sensitivities. This means, then, that a DP mechanism that claims to offer ϵ -DP here but calibrates its random distribution to the idealized sensitivity would instead offer no better than $2^{16}\epsilon$ -DP.

The proof of Theorem 5.6 relies on the fact that saturation addition is not associative, i.e., it is not always the case that $(a \boxplus b) \boxplus c = a \boxplus (b \boxplus c)$. For example, consider $a = \max(T)$, $b = \max(T)$, $c = -1$. Then $(a \boxplus b) \boxplus c = \max(T) \boxplus -1 = \max(T) - 1$, while $a \boxplus (b \boxplus c) = \max(T) \boxplus (\max(T) - 1) = \max(T)$. The operation \boxplus is commutative, so if it were associative we could freely reorder the summation without changing the result, which would imply that, for all u, v such that $d_{Sym}(u, v) = 0$, $BS_{L,U}^*(u) = BS_{L,U}^*(v)$. However, \boxplus is not associative, and we explore the impacts of that below.

We first state a property of saturation arithmetic.

Property 5.8. *For any starting summation value z_0 of type T and any sequence of values z_1, \dots, z_n of type T , if $z_i \geq 0$ for all $i > 0$,*

$$z_0 \boxplus \dots \boxplus z_n = \min\{z_0 + \dots + z_n, \max(T)\}.$$

Likewise, if $z_i \leq 0$ for all $i > 0$, then

$$z_0 \boxplus \dots \boxplus z_n = \max\{z_0 + \dots + z_n, \min(T)\}.$$

Proof of Property 5.8. We first prove the case where, for all $i > 0$, $z_i \geq 0$. We use a proof by induction. By the definition of saturation addition in Definition 2.24, we see that the base case holds: $z_0 \boxplus z_1 = \min\{z_0 + z_1, \max(T)\}$. Assume that $z_0 \boxplus \dots \boxplus z_{k-1} = \min\{z_0 + \dots + z_{k-1}, \max(T)\}$. Then, $z_0 \boxplus \dots \boxplus z_k = \min\{z_0 + \dots + z_{k-1}, \max(T)\} \boxplus z_k = \min\{z_0 + \dots + z_k, \max(T)\}$. Therefore, the claim holds.

A symmetric argument can be used to prove the claim associated with the case where $z_i \leq 0$ for all $i > 0$. \square

Note that Property 5.8 means that, for $L, U \geq 0$ or $L, U \leq 0$, the idealized sensitivity presented in Theorem 4.4 is an upper bound on the sensitivity $\Delta_{Sym} BS_{L,U}^*$ and $\Delta_{CO} BS_{L,U,n}^*$, and thus the condition $U > 0$ and $L < 0$ in Theorem 5.6 is essential. We now proceed with the proof of Theorem 5.6.

Proof of Theorem 5.6. Let $b = \lceil \frac{\max(T) - \min(T)}{|L|} \rceil$ and $c = \lceil \frac{\max(T) - \min(T)}{U} \rceil$, and consider

$$u = [L_1, \dots, L_b, U_1, \dots, U_c], v = [U_1, \dots, U_c, L_1, \dots, L_b],$$

where, for all i , $L_i = L$ and $U_i = U$. We note that $d_{Sym}(u, v) = 0$, so $u \simeq_{Sym} v$. We now evaluate and compare the results of $BS_{L,U}^*(u)$ and $BS_{L,U}^*(v)$.

We first show that $BS_{L,U}^*(u) = \max(T)$.

We begin by showing that $BS_{L,U}^*[L_1, \dots, L_b] = \min(T)$. By Property 5.8, we know that $L_1 \boxplus \dots \boxplus L_b = \max\{L_1 + \dots + L_b, \min(T)\}$. We see that $L_1 + \dots + L_b = b \cdot L \leq \frac{\max(T) - \min(T)}{|L|} \cdot L = -\max(T) + \min(T)$. Therefore,

$$BS_{L,U}^*[L_1, \dots, L_b] = \max\{-\max(T) + \min(T), \min(T)\} = \min(T).$$

To complete the evaluation of $BS_{L,U}^*(u)$, we now consider the addition of the U_i terms to the intermediate sum $BS_{L,U}^*[L_1, \dots, L_b] = \min(T)$. By Property 5.8, we know that $\min(T) \boxplus U_1 \boxplus \dots \boxplus U_c = \min\{\min(T) + U_1 + \dots + U_c, \max(T)\}$. We see that $U_1 + \dots + U_c = c \cdot U \geq \frac{\max(T) - \min(T)}{U} \cdot U = \max(T) - \min(T)$. Therefore, by Property 5.8, $BS_{L,U}^*(u) = \min\{\max(T) - \min(T) + \min(T), \max(T)\} = \max(T)$.

We next show that $BS_{L,U}^*(v) = \min(T)$. By an argument similar to the one used above, $BS_{L,U}^*[U_1, \dots, U_c] = \max(T)$; then, when we consider the addition of the L_i terms, we find that $BS_{L,U}^*(v) = \min(T)$.

We see that $|BS_{L,U}^*(u) - BS_{L,U}^*(v)| = \max(T) - \min(T)$. Because $u \simeq_{Sym} v$, we have $\Delta_{Sym} BS_{L,U}^* \geq \max(T) - \min(T) = 2^k - 1$. Additionally, by Definition 2.22, $U \leq 2^{k-1} - 1$ and $|L| \leq 2^{k-1}$, so we necessarily have $\Delta_{Sym} BS_{L,U}^* = 2^k - 1 > \max\{|L|, U\} = \Delta_{Sym} BS_{L,U}$.

The results described above will hold for any pair of datasets

$$u' = [L_1, \dots, L_{b'}, U_1, \dots, U_{c'}], v' = [U_1, \dots, U_{c'}, L_1, \dots, L_{b'}]$$

with $b' \geq b$ and $c' \geq c$. We note that $d_{CO}(u', v') = 0$. Therefore, for all

$$n = b' + c' \geq \left\lceil \frac{\max(T) - \min(T)}{|L|} \right\rceil + \left\lceil \frac{\max(T) - \min(T)}{U} \right\rceil,$$

$\Delta_{CO} BS_{L,U,n}^* = 2^k - 1$. By Definition 2.22, $U \leq 2^{k-1} - 1$ and $|L| \leq 2^{k-1}$, so we necessarily have $\Delta_{CO} BS_{L,U,n}^* = 2^k - 1 \geq U - L = \Delta_{CO} BS_{L,U,n}$. \square

5.3 Reordering Attack on Floats

We now provide a generalized counterexample showing that the idealized sensitivities $\Delta_{Sym}BS_{L,U}$ and $\Delta_{CO}BS_{L,U,n}$ in Section 4.3 do not apply to the floating-point numbers. Similarly to how the example in Section 5.2 leverages the non-associativity of saturating addition on integers, this counterexample leverages the non-associativity of addition on floats.

Because the sensitivities $\Delta_{Sym}BS_{L,U}^*$ and $\Delta_{CO}BS_{L,U,n}^*$ on the floats are neither easy to calculate nor useful, we only prove lower bounds on the sensitivities to show that $\Delta_{Sym}BS_{L,U}^* \neq \Delta_{Sym}BS_{L,U}$ and $\Delta_{CO}BS_{L,U,n}^* \neq \Delta_{CO}BS_{L,U,n}$. We prove useful upper bounds on sensitivities for alternative implementations of floating-point summation in Section 6.6.

Theorem 5.9. *Let $BS_{L,U}^* : Vec(T) \rightarrow T$ be iterative bounded sum on the type T of (k, ℓ) -bit floats. Let $j, a, d \in \mathbb{Z}$ satisfy $-(2^{\ell-1} - 2) \leq j \leq 2^{\ell-1} - 3 - k$, $a \geq 0$, $d > 0$, and $a + d \leq k + 1$. For $L = 2^j$, $U = 2^{j+d}$, and $n = 2^{(k+1)-a} + 2^{(k+1)-d}$, there are datasets $u, v \in T_{L,U}^n$ with $d_{Sym}(u, v) = d_{CO}(u, v) = 0$ such that*

$$|BS_{L,U}^*(u) - BS_{L,U}^*(v)| \geq 2^{(k+1)-(a+d)} \cdot U.$$

In particular,

$$\frac{\Delta_{Sym}BS_{L,U}^*}{\Delta_{Sym}BS_{L,U}} \geq 2^{(k+1)-(a+d)}$$

and

$$\frac{\Delta_{CO}BS_{L,U,n}^*}{\Delta_{CO}BS_{L,U,n}} \geq 2^{(k+1)-(a+d)} \cdot \left(\frac{U}{U-L} \right).$$

To maximize the blow-up in sensitivity, we should minimize $a + d$, namely by setting $a = 0$ and $d = 1$, which yields a blow-up of 2^k , albeit on datasets of size $n = 3 \cdot 2^k$. To obtain smaller datasets, we should maximize both a and d , subject to the constraint $a + d \leq k + 1$. For example, setting $a = d = k/2$, we obtain a blow-up by a factor of 2 (or a factor of $2 \cdot (\frac{U}{U-L})$ in the Δ_{CO} case) on datasets of size $n = 4 \cdot 2^{k/2}$.

Example Attack 5.10 (32-bit floats). Here, $k = 23$ and $\ell = 8$. We set $a = 0$, $d = 1$, and $j = 0$. We note that these settings satisfy the conditions of Theorem 5.9, which lead to $L = 1$, $U = 2$, and $n = 2^{24} + 2^{23}$. This dataset construction results in a difference in sums of $2^{24-1} \cdot U = 2^{24}$. This is a factor of 2^{24} larger than the sensitivity $\Delta_{CO}BS_{L,U,n} = (U - L) = 1$, and it is a factor of 2^{23} larger than the sensitivity $\Delta_{Sym}BS_{L,U} = \max\{|L|, U\} = 2$.

This means, then, that a DP mechanism that claims to offer ε -DP but calibrates its random distribution to the idealized sensitivity would instead offer no better than $2^{24}\varepsilon$ -DP when dealing with the change-one distance notion of neighboring datasets, and no better than $2^{23}\varepsilon$ -DP when dealing with the symmetric distance notion of neighboring datasets.

Example Attack 5.11 (64-bit floats). Here, $k = 52$ and $\ell = 11$. We set $a = 26$, $d = 26$, and $j = -26$. We note that these settings satisfy the conditions of Theorem 5.9, which leads to $L = 2^{-26}$, $U = 1$, and $n = 2^{28}$. This dataset construction results in a difference in sums of $2 \cdot U = 2$. This is more than a factor of 2 larger than the sensitivity $\Delta_{CO}BS_{L,U,n} = (U - L) = 1 - 2^{-26} < 1$, and it is a factor of 2 larger than the sensitivity $\Delta_{Sym}BS_{L,U} = \max\{|L|, U\} = 2$.

This means, then, that a DP mechanism that claims to offer ε -DP but calibrates its random distribution to the idealized sensitivity would instead offer no better than 2ε -DP when dealing with the change-one distance notion of neighboring datasets, and no better than 2ε -DP when dealing with the symmetric distance notion of neighboring datasets.

Proof of Theorem 5.9. Let $b = 2^{(k+1)-a}$ and $c = 2^{(k+1)-d}$, and consider the datasets

$$u = [L_1, \dots, L_b, U_1, \dots, U_c], v = [U_1, \dots, U_c, L_1, \dots, L_b]$$

where, for all i , $L_i = L$ and $U_i = U$. Note that $d_{Sym}(u, v) = d_{CO}(u, v) = 0$.

We first show that the first $b - 1$ terms of u can be added exactly; that is, we show that $BS_{L,U}^*(u) = BS_{L,U}(u) = b \cdot L + c \cdot U$.

We begin by showing that

$$BS_{L,U}^*[L_1, \dots, L_b] = BS_{L,U}[L_1, \dots, L_b] = b \cdot L.$$

We take the approach of showing that all intermediate sums computed in the calculation of $BS_{L,U}^*[L_1, \dots, L_b]$ can be represented exactly as (k, ℓ) -bit floats, which (with the exception of $b \cdot L$) is done by first showing that all of these intermediate sums are multiples of their $ULP_{(k,\ell)}$ and then applying Lemma 2.27.

We observe that $\{L, \dots, L \cdot (b - 1)\}$ is the set of all intermediate sums that result from calculating $BS_{L,U}[L_1, \dots, L_{b-1}]$ and note that $b \cdot L = 2^{j+k+1-a}$. For all $x \in \{L, \dots, L \cdot (b - 1)\}$,

$$\begin{aligned} ULP_{(k,\ell)}(x) &\leq ULP_{(k,\ell)}(L \cdot (b - 1)) \\ &< ULP_{(k,\ell)}(2^{j+k+1-a}) \\ &= 2^{j+1-a} \\ &\leq 2^{j+1}. \end{aligned}$$

All such x are multiples of $L = 2^j$, so they are necessarily multiples of $ULP_{(k,\ell)}(x)$. By Lemma 2.27, all of these values can be represented exactly as (k, ℓ) -bit floats. We also note that $b \cdot L = 2^{j+k+1-a}$ can be represented exactly as a (k, ℓ) -bit float (by our constraint on j), so by Lemma 2.38, $BS_{L,U}^*[L_1, \dots, L_b] = BS_{L,U}[L_1, \dots, L_b] = b \cdot L$.

We now complete the evaluation of $BS_{L,U,n}^*(u)$. We use Lemma 2.27 to show that all of the intermediate sums of $BS_{L,U}(u) = b \cdot L + c \cdot U$ are exactly representable as (k, ℓ) -bit floats, which then shows that $BS_{L,U}^*(u) = BS_{L,U}(u) = b \cdot L + c \cdot U$.

We consider the addition of the U_i terms to the intermediate sum $b \cdot L$. We note that $b \cdot L = 2^{j+k+1-a}$ is an integer multiple of $U = 2^{j+d}$ since $a + d \leq k + 1$. Thus each intermediate sum $x \in \{b \cdot L, b \cdot L + U, \dots, b \cdot L + c \cdot U\}$ is a multiple of U . Moreover,

$$\begin{aligned} ULP_{(k,\ell)}(x) &\leq ULP_{(k,\ell)}(b \cdot L + c \cdot U) \\ &= ULP_{(k,\ell)}(2^{(k+1)+j-a} + 2^{(k+1)+j}) \\ &\leq 2^{j+1}, \end{aligned}$$

unless $a = 0$ and $x = b \cdot L + c \cdot U$. Thus $U = 2^{j+d}$ is a multiple of $ULP_{(k,\ell)}(x)$ and hence x is representable as a normal (k, ℓ) -bit float. For the case $a = 0$ and $x = b \cdot L + c \cdot U$, we note that $b \cdot L + c \cdot U = 2^{j+k+2}$, which is a representable (k, ℓ) -bit float by the condition on j .

We next evaluate $BS_{L,U,n}^*(v)$. We use Lemma 2.27 to show that $BS_{L,U}^*[U_1, \dots, U_c] = BS_{L,U}[U_1, \dots, U_c] = c \cdot U$.

We observe that $\{U, \dots, c \cdot U\}$ is the set of intermediate sums that result from calculating $BS_{L,U}[U_1, \dots, U_c]$ and note that $c \cdot U = 2^{k+1+j}$. For all $x \in \{U, \dots, c \cdot U\}$, $ULP_{(k,\ell)}(x) = 2^{j+1-m}$ for some integer $m \geq 0$. All $x \in \{U, \dots, c \cdot U\}$ are multiples of $U = 2^{j+d}$ with $d > 0$, so they

are necessarily multiples of $ULP_{(k,\ell)}(x) \leq ULP_{(k,\ell)}(c \cdot U) = 2^{j+1}$. By Lemma 2.27, all of these values can be represented exactly as (k, ℓ) -bit floats, so by Lemma 2.38, $BS_{L,U}^*[U_1, \dots, U_c] = BS_{L,U}[U_1, \dots, U_c] = c \cdot U$.

We now complete the evaluation of $BS_{L,U}^*(v)$. We use Lemma 2.27 to show that the intermediate sums in the calculation of $BS_{L,U}(v)$ are not exactly representable as (k, ℓ) -bit floats, and we show how to calculate the result $BS_{L,U}^*(v)$.

We note that $c \cdot U + L_1 = 2^{k+1+j} + 2^j \in [2^{k+1+j}, 2^{k+2+j})$, so $ULP_{(k,\ell)}(c \cdot U + L_1) = 2^{j+1}$. However, $c \cdot U + L_1 = 2^{k+1+j} + 2^j$ is not a multiple of 2^{j+1} so by Lemma 2.27, $c \cdot U + L_1$ is not exactly representable as a (k, ℓ) -bit float. $c \cdot U + L_1$ lies exactly halfway between the consecutive (k, ℓ) -bit floats 2^{k+1+j} and $2^{k+1+j} + 2^{j+1}$. Because the mantissa associated with 2^{k+1+j} ends in an even bit, by banker's rounding (Definition 2.33), $c \cdot U \oplus L_1 = 2^{k+1+j}$. For all i , then, the summation $c \cdot U \oplus L_i = 2^{(k+1)+j} = c \cdot U$. Therefore, $BS_{L,U}^*(v) = c \cdot U$.

Thus,

$$|BS_{L,U}^*(u) - BS_{L,U}^*(v)| = b \cdot L = 2^{(k+1)-a} \cdot 2^j = 2^{(k+1)-(a+d)} \cdot U,$$

so $\Delta_{Sym} BS_{L,U}^* \geq 2^{(k+1)-(a+d)} \cdot U$ and $\Delta_{CO} BS_{L,U,n}^* \geq 2^{(k+1)-(a+d)} \cdot U$.

We compare this result to the idealized sensitivities presented in Theorem 4.4: $\Delta_{Sym} BS_{L,U} = \max\{|L|, U\} = U$ and $\Delta_{CO} BS_{L,U,n} = U - L$. We see that

$$\frac{\Delta_{Sym} BS_{L,U}^*}{\Delta_{Sym} BS_{L,U}} \geq 2^{(k+1)-(a+d)}$$

and

$$\frac{\Delta_{CO} BS_{L,U,n}^*}{\Delta_{CO} BS_{L,U,n}} \geq 2^{(k+1)-(a+d)} \cdot \left(\frac{U}{U - L} \right).$$

□

5.4 Rounding-Based Attack on Floats

In the proof of Theorem 5.9, $d_{Sym}(u, v) = 0$, but $d_{Ham}(u, v)$ is large. Now we show that the sensitivity can be unexpectedly large even when neighboring datasets are defined in terms of an ordered distance metric such as d_{Ham} or d_{ID} .

Theorem 5.12 ($\Delta_{Ham} BS_{L,U,n}^* \neq \Delta_{Ham} BS_{L,U,n}$). *Let $BS_{L,U,n}^* : T^n \rightarrow T$ be iterative bounded sum on the type T of (k, ℓ) -bit floats. Let $j, k \in \mathbb{Z}$ satisfy $1 < j < (k+1)/2$ and $-2^{(\ell-1)} + 2 \leq m \leq 2^{\ell-1} - 1 - j$. For $n = 2^j + 1$, $L = (1 + 2^{-(k+1)+j}) \cdot 2^m$, and $U = L + 2^{m-k}$, there are datasets $u, v \in T_{[L,U]}^n$ with $d_{Ham}(u, v) = 1$ such that*

$$|BS_{L,U,n}^*(u) - BS_{L,U,n}^*(v)| \geq 2^{j+m-k} = 2^j \cdot (U - L).$$

In particular,

$$\frac{\Delta_{Ham} BS_{L,U,n}^*}{\Delta_{Ham} BS_{L,U,n}} \geq (n - 1).$$

Example Attack 5.13 (64-bit floats). Here, $k = 52$ and $\ell = 11$. We set $j = 4$ and $m = -1$. We note that these settings satisfy Theorem 5.12, which lead to $L = (1 + 2^{-49})/2$, $U = (1 + 2^{-49} + 2^{-53})/2$, and $n = 2^4 + 1 = 17$. The result is a difference in sums of $2^{4-1-52} = 2^{-49}$. This is a factor of 2^4 larger than $\Delta_{Ham} BS_{L,U,n} = (U - L) = 2^{-53}$.

This means, then, that a DP mechanism that claims to offer ε -DP but calibrates its random distribution to the idealized sensitivity would instead offer no better than 16ε -DP in the known n case. More generally, no better than $2^j\varepsilon$ -DP is offered for $1 < j < 26.5$.

Proof of Theorem 5.12. Let $b = 2^j$ and $m = 0$, and consider the datasets

$$u = [L_1, \dots, L_b, U], v = [L_1, \dots, L_b, L_{b+1}],$$

where, for all i , $L_i = L$. Note that $d_{Ham}(u, v) = 1$.

We begin by showing that

$$BS_{L,U}^*[L_1, \dots, L_b] = BS_{L,U}[L_1, \dots, L_b] = b \cdot L.$$

We take the approach of showing that all intermediate sums computed in the calculation of $BS_{L,U}^*[L_1, \dots, L_b]$ can be represented exactly as (k, ℓ) -bit floats, which is done by first showing that all of these intermediate sums are multiples of their $ULP_{(k,\ell)}$ and then applying Lemma 2.27.

We observe that $\{L, \dots, L \cdot (b-1)\}$ is the set of intermediate sums that result from calculating $BS_{L,U}[L_1, \dots, L_{b-1}]$. Let \mathcal{L} denote this set $\{L, \dots, (b-1) \cdot L\}$. We note that

$$(b-1) \cdot L = b \cdot (L-1) - L + b < b = 2^j,$$

where the inequality holds because $b \cdot (L-1) = 2^{2j-(k+1)} \leq 1 < L$. Therefore, $ULP_{(k,\ell)}((b-1) \cdot L) = 2^{j-k-1}$. For all $x \in \mathcal{L}$, then,

$$ULP_{(k,\ell)}(x) \leq ULP_{(k,\ell)}(L \cdot (b-1)) = 2^{-(k+1)+j}.$$

Next, we show that for all $x \in \mathcal{L}$, $L = 1 + 2^{-(k+1)+j}$ is a multiple of $ULP_{(k,\ell)}(x)$. We note that $ULP_{(k,\ell)}(x) = 2^{-(k+1)+j}$ divides $2^{-(k+1)+j}$ and, because we require $j < \frac{k+1}{2}$, we note that $ULP_{(k,\ell)}(x) = 2^{j-k-1-a}$ also divides $1 = 2^0$. Therefore, L is a multiple of $ULP_{(k,\ell)}(x)$ for all $x \in \mathcal{L}$.

All $x \in \mathcal{L}$ are multiples of L , so they are necessarily multiples of $ULP_{(k,\ell)}(x)$. By Lemma 2.27, all $x \in \mathcal{L}$ can be represented exactly as (k, ℓ) -bit floats. We also note that $b \cdot L = 2^j + 2^{-(k+1)+2j} = (1 + 2^{-(k+1)+j}) \cdot 2^j$ can be represented exactly as a (k, ℓ) -bit float, so by Lemma 2.38, $BS_{L,U}^*[L_1, \dots, L_b] = BS_{L,U}[L_1, \dots, L_b] = b \cdot L$.

We now evaluate $BS_{L,U,n}^*(u)$. We use Lemma 2.27 to show that $BS_{L,U,n}(u) = b \cdot L + U$ is not exactly representable as a (k, ℓ) -bit float, and we show how to calculate the result $BS_{L,U,n}^*(u)$.

As shown above, the sum of the first b terms is

$$BS_{L,U}^*[L_1, \dots, L_b] = b \cdot L.$$

Thus,

$$BS_{L,U}^*(u) = b \cdot L \oplus U.$$

We note that

$$b \cdot L + U = 2^j + 2^{-(k+1)+2j} + 1 + 2^{-(k+1)+j} + 2^{-k} \in [2^j, 2^{j+1}),$$

so $ULP_{(k,\ell)}(b \cdot L + U) = 2^{j-k}$. However, $b \cdot L + U$ is not an integer multiple of 2^{j-k} , so by Lemma 2.27, $b \cdot L + U$ is not exactly representable as a (k, ℓ) -bit float. $b \cdot L + U$ lies between

the consecutive (k, ℓ) -bit floats $2^{j-k}(2^k + 2^{j-1} + 2^{k-j})$ and $2^{j-k}(2^k + 2^{j-1} + 2^{k-j} + 1)$. The sum $b \cdot L + U$ is closer to the larger value, $2^{j-k}(2^k + 2^{j-1} + 2^{k-j} + 1)$, so by banker's rounding as defined in Definition 2.33, the result is $BS_{L,U,n}^*(u) = 2^{j-k}(2^k + 2^{j-1} + 2^{k-j} + 1)$.

We next evaluate $BS_{L,U,n}^*(v)$. We again use Lemma 2.27 to show that $BS_{L,U,n}(v) = b \cdot L + L$ is not exactly representable as a (k, ℓ) -bit float, and we show how to calculate the result $BS_{L,U,n}^*(v)$.

As shown above, the sum of the first b terms is

$$BS_{L,U}^*[L_1, \dots, L_b] = b \cdot L.$$

Thus, $BS_{L,U}^*(v) = b \cdot L \oplus L$. We note that

$$b \cdot L + L = 2^j + 2^{-(k+1)+2j} + 1 + 2^{-(k+1)+j} \in [2^j, 2^{j+1}),$$

so $ULP_{(k,\ell)}(b \cdot L + L) = 2^{j-k}$. However, $b \cdot L + L$ is not a multiple of 2^{j-k} (it is a multiple of 2^{j-k-1}) so by Lemma 2.27, $b \cdot L + L$ is not exactly representable as a (k, ℓ) -bit float. $b \cdot L + L$ lies exactly between the consecutive (k, ℓ) -bit floats $2^{j-k}(2^k + 2^{j-1} + 2^{k-j})$ and $2^{j-k}(2^k + 2^{j-1} + 2^{k-j} + 1)$. Because the mantissa associated with the smaller value $2^{j-k}(2^k + 2^{j-1} + 2^{k-j})$ ends in an even bit, by banker's rounding, $b \cdot L \oplus L = 2^{j-k}(2^k + 2^{j-1} + 2^{k-j})$, so $BS_{L,U,n}^*(u) = 2^{j-k}(2^k + 2^{j-1} + 2^{k-j})$.

We see that $|BS^*(u) - BS^*(v)| = 2^{j-k}$.

From the bounds $-2^{(\ell-1)} + 2 \leq m \leq 2^{\ell-1} - 1 - j$, we see that m only affects the exponent of floating point values in the proof and does not affect whether values are representable as (k, ℓ) -bit normal floats or the direction of rounding. All values in the proof above, then, can be multiplied by 2^m , so, for all m such that $-2^{\ell-1} + 2 \leq m \leq 2^{\ell-1} - 1 - j$, $\Delta_{Ham} BS_{L,U,n}^* \geq 2^{j+m-k}$.

We observe that $U - L = 2^{m-k}$, $\Delta_{Ham} BS_{L,U,n} = U - L = 2^{m-k}$, and $n - 1 = 2^j$. Therefore,

$$\frac{\Delta_{Ham} BS_{L,U,n}^*}{\Delta_{Ham} BS_{L,U,n}} \geq (n - 1).$$

□

Definition 5.14 (Compensated Summation). A *compensated summation* algorithm, denoted $CS : \text{Vec}(T) \rightarrow T$, attempts to track the error that is accumulated in the computation of an iterative summation, and uses this error-tracking to adjust its final answer accordingly. One limitation of a compensated summation algorithm that we assume here is that its output type must be the same as its input type. An idealized version of a compensated summation algorithm, then, computes the exact sum, and then rounds or casts this value to some output of type T .

Remark 5.15 (Resistance to Compensated Summation). The same increased sensitivity seen in the proof of Proposition 5.12 also applies when Kahan summation [24] (or any type of compensated summation as defined in Definition 5.14) is used. This is because the inaccuracy in the final answer occurs due to an inability to represent the idealized sum, rather than due to an accumulation of rounding error throughout the summation. (Rounding only occurs in the addition of the final term.) We see, then that the idealized sensitivity of $\Delta_{CO} BS_{L,U,n} = (U - L)$ and $\Delta_{Ham} BS_{L,U,n} = (U - L)$ is fundamentally unsalvageable under the assumption that the data type of elements in the input vector is the same as the data type of the output of $BS_{L,U,n}^*$.

Remark 5.16 (Resistance to Pairwise Summation). The same increased sensitivity seen in the proof of Proposition 5.12 also applies when pairwise summation is used. This is because the dataset u is of length $2^k + 1$. As a result, the final term is only added once the sum of the rest of the dataset has been calculated, so the same analysis presented in the proof of Proposition 5.12 applies when calculating the sum by way of pairwise summation.

5.5 Repeated Rounding Attack on Floats I

The counterexample described in Section 5.4, which shows that we can have $\Delta_{Ham}BS_{L,U,n}^* \gg \Delta_{Ham}BS_{L,U,n}$, only takes advantage of rounding that happens at the very last summation step, when we add the last elements u_n and v_n in the datasets u and v . Since these last elements are bounded in absolute value by U , the rounding changes each result by at most $\pm U$, so $|BS_{L,U}^*(u) - BS_{L,U}^*(v)| \leq 2U$. The reason for which we get a large blow-up in sensitivity is that $\Delta_{Ham}BS_{L,U,n} = (U - L)$, which can be much smaller than U . To get a large blow-up relative to $\Delta_{ID}BS_{L,U} = U$, we need to take advantage of rounding throughout the summation. Here, we show that when u and v differ in their first element, rounding errors can accumulate throughout the summation, yielding $\Delta_{ID}BS_{L,U}^* \gg \Delta_{ID}BS_{L,U}$.

Theorem 5.17. *Let $BS_{L,U}^* : \text{Vec}(T_{[L,U]}) \rightarrow T$ be iterative bounded sum on the type T of (k, ℓ) -bit floats. Let $j, m \in \mathbb{Z}$ satisfy $0 < j \leq k$ and $-(2^{\ell-1} - 2) \leq m \leq 2^{\ell-1} - 2 - j - k$.*

Then for $L = 0$, $U = 2^{k+m}$, and $n = j \cdot 2^k + 2$, there are datasets $u, v \in \text{Vec}(T)$ where $\text{len}(u) = n$ and $\text{len}(v) = n - 1$ such that $d_{ID}(u, v) = 1$ and

$$|BS_{L,U}^*(u) - BS_{L,U}^*(v)| \geq 2^{k+j+m}.$$

In particular,

$$\frac{\Delta_{ID}BS_{L,U}^*}{\Delta_{ID}BS_{L,U}} \geq 2^j.$$

Example Attack 5.18 (32-bit floats). Here, $k = 23$ and $\ell = 8$. We set $j = 15$ and $m = 0$. We note that these settings satisfy the conditions of Theorem 5.17, which lead to $L = 0$, $U = 2^{23}$, and $\text{len}(u) = 15 \cdot 2^{23} + 2$ and $\text{len}(v) = 15 \cdot 2^{23} + 1$. This dataset construction results in a difference in sums of $2^{15} \cdot U = 2^{38}$. This is a factor of 2^{15} larger than the sensitivity $\Delta_{ID}BS_{L,U} = \max\{|L|, U\} = 2^{23}$.

This means, then, that a DP mechanism that claims to offer ε -DP but calibrates its random distribution to the idealized sensitivity would instead offer no better than $2^{15}\varepsilon$ -DP when dealing with the insert-delete distance notion of neighboring datasets.

The proof of Theorem 5.17 can be found in Appendix C.

5.6 Repeated Rounding Attack on Floats II

We now show that we can have an even larger blow-up in sensitivity relative to $d_{ID}BS_{L,U} = U$ as a function of the dataset length n when we consider datasets consisting of both positive and negative values.

Theorem 5.19. *Let $BS_{L,U}^* : \text{Vec}(T_{[L,U]}) \rightarrow T$ be iterative bounded sum on the type T of (k, ℓ) -bit floats. Additionally, let $a, j \in \mathbb{Z}$ satisfy $2 \leq j < k$, $-(2^{\ell-1} - 2) \leq a$, $-(2^{\ell-1} - 2) + 1 + k \leq j + a \leq 2^{\ell-1} - 1$ (these conditions ensure that values are representable as (k, ℓ) -bit floats), $n = 2^j$, $m = n/2$. Then, for*

$$U = 2^a, L = -\left(\frac{U \cdot m}{2^k}\right) \cdot \left(\frac{1}{2} - \frac{1}{2^k}\right),$$

there are datasets $u, v \in \text{Vec}(T_{[L,U]})$ with $d_{ID}(u, v) = 1$ such that

$$|BS_{L,U}^*(u) - BS_{L,U}^*(v)| \geq \frac{n^2}{2^{k+3}} \cdot U + U.$$

In particular,

$$\frac{\Delta_{ID}BS_{L,U}^*}{\Delta_{ID}BS_{L,U}} = \frac{n^2}{2^{k+3}} + 1.$$

Example Attack 5.20 (64-bit Floats). Here, $k = 52$ and $\ell = 11$. We set $j = 30$ and $a = 0$. We note that these settings satisfy the conditions of Theorem 5.19, which lead to $L = 2^{-23} \cdot -(\frac{1}{2} + 2^{-52})$, $U = 1$, and $\text{len}(u) = 2^{30}$ and $\text{len}(v) = 2^{30} - 1$. This dataset construction results in a difference in sums of $2^5 + U = 33$. This is a factor of 33 larger than the sensitivity $\Delta_{ID}BS_{L,U} = \max\{|L|, U\} = 1$.

This means, then, that a DP mechanism that claims to offer ε -DP but calibrates its random distribution to the idealized sensitivity would instead offer no better than 33ε -DP when dealing with the insert-delete distance notion of neighboring datasets.

The proof of Theorem 5.19 can be found in Appendix C.

5.7 Concrete Implementations of the Attacks on DP Libraries

In this section, we show how malicious data analysts can exploit the fact that, as shown in Section 4.3.1, all libraries with a (supposedly) DP bounded sum function scale its noise according to the idealized sensitivity $\Delta_dBS_{L,U}$ rather than to the implemented sensitivity $\Delta_dBS_{L,U}^*$.

As shown in the previous parts of Section 5, there are many cases for which $\Delta_dBS_{L,U} < \Delta_dBS_{L,U}^*$. By Theorem 2.20, the fact that, for a user-specified ε , the scale parameter λ is set to $\lambda = \Delta_dBS_{L,U}/\varepsilon < \Delta_dBS_{L,U}^*/\varepsilon$ means that ε -DP is not fulfilled by functions which claim to offer ε -DP.

The fact that ε -DP is not fulfilled indicates that there could be opportunities for malicious data analysts to get more precise information about datasets for a given value of ε than should be allowed. In this section, we demonstrate how a malicious data analyst can take advantage of these errors and use the outputs of allegedly DP functions to perform membership-inference attacks and reconstruction attacks. Our attacks use the Python wrappers for OpenDP / SmartNoise and for IBM’s `diffprivlib`; and we use OpenMined’s PyDP Python wrapper for our interactions with Google’s DP libraries.

Definition 5.21 ($\mathcal{M}_\varepsilon^*$). For all of the examples below, let $\mathcal{M}_\varepsilon^* : \text{Vec}(\mathcal{D}) \rightarrow \mathbb{R}$ represent the library’s attempt to implement the function being used in the attack (either a mean function or a summation function) such that it fulfills ε -DP. In all of the libraries we explore, this is done using a discrete version of the Laplace mechanism, often implemented using the snapping mechanism described in [38].

In our attacks, we perform the following process: First, we create (sometimes slightly modified) versions of the datasets in the example attacks from the previous parts of Section 5, and we execute supposedly ε -DP functions on these datasets. We then apply post-processing to the result which, by Proposition 2.18, will not make the result any less differentially private. We show, though, that the results we achieve after post-processing would be very unlikely to occur for the specified value of ε . Finally, we show how the fact that ε -DP is not fulfilled can be used to recover information about the datasets that would have been protected by an ε -DP function.

5.7.1 Implementing Example Attack 5.5

Only IBM’s diffprivlib was susceptible to Example Attack 5.5. This is because, although . Many libraries – perhaps intentionally, perhaps naively – compute the bounded sum function with modular arithmetic (i.e., overflow is allowed) and then also add noise such that overflow is allowed, a solution which we prove to be correct in Section 6.2. IBM’s diffprivlib, on the other hand, appears to compute the bounded sum with modular arithmetic and then add noise using saturation addition.

Library-Mediated Attack 5.22 (IBM’s diffprivlib; Signed 32-bit Integers). We create datasets as described in Section 5.1, and we set $U = 2^{24}$, $L = 0$, and $n = 2^7$. Let $\mathcal{M}_\varepsilon^*$ be the bounded sum function as implemented by IBM’s diffprivlib for U and L .

We run $\mathcal{M}_\varepsilon^*(u)$ and $\mathcal{M}_\varepsilon^*(v)$, where $\mathcal{M}_\varepsilon^*$ is the mechanism used by the library for computing $BS_{L,U}^*$ with DP and where $\varepsilon = 0.5$. We define the post-processing function $f : \mathbb{R} \rightarrow \{0, 1\}$ as follows: $f(x) = 1$ if and only if $x > 0$; this value was chosen since it lies between the maximum and minimum values.

Running each of $f(\mathcal{M}_\varepsilon^*(u))$ and $f(\mathcal{M}_\varepsilon^*(v))$ 10,000 times gives the following results.

	$f(\mathcal{M}_\varepsilon^*(u))$	$f(\mathcal{M}_\varepsilon^*(v))$
0	0	10,000
1	10,000	0

Figure 6: “0.5-DP” results from IBM’s diffprivlib (with post-processing).

By the fact that $u \simeq_{ID} v$, the definition of 0.5-DP, and the fact that DP is robust to post-processing per Proposition 2.18, we would expect to have

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 0]}{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 0]} \leq e^{0.5} < 1.65$$

and

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]}{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]} \leq e^{0.5} < 1.65.$$

Let $p_u = \Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]$ and $p_v = \Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]$, and suppose that $\frac{p_u}{p_v} \leq e^\varepsilon$ – that is, suppose the queries did satisfy ε -DP. Then,

$$\begin{aligned} (p_u)^{10,000} \cdot (1 - p_v)^{10,000} &\leq (e^\varepsilon p_v \cdot (1 - p_v))^{10,000} \\ &\leq \left(\frac{e^\varepsilon}{4}\right)^{10,000}. \end{aligned}$$

The results in Fig. 6, then, would occur with probability $\leq 2^{-12,000}$ if $\mathcal{M}_\varepsilon^*$ were truly an ε -DP function for $\varepsilon = 0.5$. These results, then, support the calculation in Example Attack 5.5 that $\mathcal{M}_\varepsilon^*$ does not offer ε -DP for $\varepsilon = 0.5$.

5.7.2 Implementing Example Attacks 5.10 and 5.11

We show how a malicious data analyst can use the results shown in Section 5.3 to perform membership-inference attacks on datasets. We work with the “unknown n ” summation functions provided in Section 4.3.1 (where “unknown n ” means that the function is designed to offer DP under the d_{Sym} notion of neighboring datasets).

Attack Idea 5.23. We create the datasets as described in Example Attack 5.10 if we are working with 32-bit floats or as described in 5.11 if we are working with 64-bit floats, with the exception that we set $v = [v_1, \dots, v_{\text{len}(v)-1}]$ (i.e., the last element is removed).

We now describe a setting in which the example attacks can be used to perform membership-inference attacks.

Suppose we know that we either have dataset u' which contains individual i and has the property $h_{u'} = h_u$, or dataset v' which does not contain individual i and has the property $h_{v'} = h_v$. (Note that the presence or absence of individual i is why $d_{Sym}(u', v') = 1$.)

Additionally, suppose the library of functions we are working with provides us with the bounded sum function used in Example Attacks 5.10 and 5.11, and suppose that the DP library also provides us with the ability to reorder datasets based on some condition. Since this implementation of bounded sum indicates that the creator of the library thinks floating point addition is associative, the creator should have no issues with offering this functionality.

Then, we can conditionally reorder the dataset based on the absence or presence of i . If i is in the dataset (meaning we have u'), some permuted dataset $\pi(u') = u$ is returned where $\pi \in S_{\text{len}(u)}$ and where $\pi(u')$ has the property that all rows with value L appear before any row with value U . Likewise, if i is not in the dataset (meaning we have v'), some permuted dataset $\sigma(v') = v$ is returned where $\sigma \in S_{\text{len}(v)}$ and where $\sigma(v')$ has the property that all rows with value U appear before any row with value L .

Finally, we perform an “ ε -DP” bounded sum for some ε on the reordered dataset, which is either $\pi(u') = u$ or $\sigma(v') = v$. Using the ideas and evidence described in the attacks below, we will be able to identify whether we are working with u' or v' with higher probability than ε should allow – and, correspondingly, we will be able to determine whether i is in the dataset.

(Although this attack still relies on having specific values in the dataset, a library that provides access to a “join” function like the one offered in [49] and a “resize” function like the one offered in [41] could enable a malicious data analyst to manipulate a dataset such that this attack could work in less specific settings.)

Library-Mediated Attack 5.24 (SmartNoise; 32-bit Floats). We create the datasets as described in Example Attack 5.10, with the exception that we set $v = [v_1, \dots, v_{\text{len}(v)-1}]$ (i.e., the last element is removed). The result, then, is that $d_{Sym}(u, v) = 1$, so $u \simeq_{Sym} v$. From Example Attack 5.10, we expect to have $|BS_{L,U}^*(u) - BS_{L,U}^*(v)| = b \cdot L = 2^{24}$, despite the fact that $\Delta_{Sym} BS_{L,U} = \max\{|L|, U\} = 2$.

We run $\mathcal{M}_\varepsilon^*(u)$ and $\mathcal{M}_\varepsilon^*(v)$, where $\mathcal{M}_\varepsilon^*$ is the mechanism used by the library for computing $BS_{L,U}^*$ with DP and where $\varepsilon = 0.5$. We define the post-processing function $f : \mathbb{R} \rightarrow \{0, 1\}$ as follows: $f(x) = 1$ if and only if $x \geq 25,165,824$ (this value was chosen since it lies halfway between $BS_{L,U}^*(u)$ and $BS_{L,U}^*(v)$).

Running each of $f(\mathcal{M}_\varepsilon^*(u))$ and $f(\mathcal{M}_\varepsilon^*(v))$ 100 times gives the following results. (We only run these computations 100 times each due to the time needed to compute these summations of large datasets.)

	$f(\mathcal{M}_\varepsilon^*(u))$	$f(\mathcal{M}_\varepsilon^*(v))$
0	0	100
1	100	0

Figure 7: “0.5-DP” results from SmartNoise (with post-processing).

By the fact that $u \simeq_{Sym} v$, the definition of 0.5-DP, and the fact that DP is robust to post-processing per Proposition 2.18, we would expect to have

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 0]}{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 0]} \leq e^{0.5} \leq 1.65$$

and

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]}{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]} \leq e^{0.5} \leq 1.65.$$

Let $p_u = \Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]$ and $p_v = \Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]$, and suppose that $\frac{p_u}{p_v} \leq e^\varepsilon$ – that is, suppose the queries did satisfy ε -DP. Then,

$$\begin{aligned} (p_u)^{100} \cdot (1 - p_v)^{100} &\leq (e^\varepsilon p_v \cdot (1 - p_v))^{100} \\ &\leq \left(\frac{e^\varepsilon}{4}\right)^{100}. \end{aligned}$$

The results in Fig. 7, then, would occur with probability $\leq 2^{-127}$ if $\mathcal{M}_\varepsilon^*$ were truly an ε -DP function for $\varepsilon = 0.5$.

Moreover, for these 100 rounds, $\min(\mathcal{M}_\varepsilon^*(u)) > 33,554,400$ and $\max(\mathcal{M}_\varepsilon^*(v)) < 16,777,230$. These results support the calculation in Example Attack 5.10 that we do not have ε -DP for $\varepsilon = 0.5$ but instead have at least $2^{23}\varepsilon$ -DP.

Library-Mediated Attack 5.25 (Google DP (OpenMined wrapper); 64-bit Floats). We create the datasets as described in Example Attack 5.11, with the exception that we set $v = [v_1, \dots, v_{\text{len}(v)-1}]$ (i.e., the last element is removed). The result, then, is that $d_{Sym}(u, v) = 1$, so $u \simeq_{Sym} v$. From Example Attack 5.11, we expect to have $|BS_{L,U}^*(u) - BS_{L,U}^*(v)| = 3$, despite the fact that $\Delta_{Sym} BS_{L,U} = \max\{|L|, U\} = 1$.

We run $\mathcal{M}_\varepsilon^*(u)$ and $\mathcal{M}_\varepsilon^*(v)$, where $\mathcal{M}_\varepsilon^*$ is the mechanism used by the library for computing $BS_{L,U}^*$ with DP and where $\varepsilon = 1.0$. We define the post-processing function $f : \mathbb{R} \rightarrow \{0, 1\}$ as follows: $f(x) = 1$ if and only if $x \geq 134,217,729.5$ (this value was chosen since it lies exactly between $BS_{L,U}^*(u)$ and $BS_{L,U}^*(v)$).

Running each of $f(\mathcal{M}_\varepsilon^*(u))$ and $f(\mathcal{M}_\varepsilon^*(v))$ 10 times gives the following results. (We only run these computations 10 times each due to the time needed to compute these summations of large datasets.)

	$f(\mathcal{M}_\varepsilon^*(u))$	$f(\mathcal{M}_\varepsilon^*(v))$
0	1	10
1	9	0

Figure 8: “1-DP” results from Google’s DP library (with post-processing).

By the fact that $u \simeq_{sym} v$, the definition of 1-DP, and the fact that DP is robust to post-processing per Proposition 2.18, we would expect to have

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 0]}{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 0]} \leq e \text{ and } \frac{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]}{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]} \leq e.$$

Let $p_u = \Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]$ and $p_v = \Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]$, and suppose that $\frac{p_u}{p_v} \leq e^\varepsilon$ – that is, suppose the queries did satisfy ε -DP. Maximizing

$$(10(p_u)^9(1 - p_u) + (p_u)^{10}) \cdot (1 - p_v)^{10}$$

over $p_u, p_v \in [0, 1]$ such that $p_u \leq e^\varepsilon \cdot p_v$ and $(1 - p_v) \leq e^\varepsilon \cdot p_u$ for $\varepsilon = 1$ yields a maximum of < 0.015 .

The results in Fig. 7, then, would occur with probability < 0.015 if $\mathcal{M}_\varepsilon^*$ were truly an ε -DP function for $\varepsilon = 1$.

These results support the calculation in Example Attack 5.11 that we do not have ε -DP for $\varepsilon = 1$ but instead have at least 3ε -DP.

5.7.3 Implementing Example Attack 5.13

We next show how a malicious data analyst can use the results shown in Section 5.3 to recover sensitive bits in a dataset. We work with the “known n ” summation functions provided in Section 4.3.1 (where “known n ” means that the function is designed to offer DP under the d_{CO} notion of neighboring datasets). The Google DP library does not have a “known n ” bounded sum, but it does have a “known n ” bounded mean, so we use this.

Attack Idea 5.26. Suppose that our library has a clamp function as described in Definition 4.2 (all of the libraries that we work with offer this function) and that we can cast between data types.

In this attack, we essentially create the datasets described in Example Attack 5.13, but we do not need to start with these datasets directly. Instead, we can start with datasets of binary values, cast these binary values to floats, and apply clamping function to these datasets to effectively get the datasets described in Example Attack 5.13.

Suppose we have a dataset with a field containing binary values (for example, a dataset with a “had_covid” field, where the specified column has a value of 1 if and only if the individual has had COVID-19), and suppose we know 2^j individuals in the dataset with “had_covid” value 0. We begin by selecting these 2^j individuals. Then, suppose we want to learn the “had_covid” value of the individual located at index i in the dataset. We select this individual as well, and, in our selected dataset of $2^j + 1$ items, we cast the values in the “had_covid” column to 64-bit floats. Then, we carefully set U and L as in Example Attack 5.13 and clamp our values. This has the effect of setting 0 to L and 1 to U .

Finally, we perform an “ ε -DP” bounded sum (or bounded mean) on this dataset for some ε . Using the ideas and evidence in the attacks below, we will be able to correctly recover the specified individual’s “had_covid” value with much higher probability than the specified value of ε should allow. We perform one example using the “known n ” bounded sum and one example using the “known n ” bounded mean.

Library-Mediated Attack 5.27 (IBM `diffprivlib`: Bounded Sum, 64-bit Floats). We create the datasets as described in Example Attack 5.13, except we set $j = 5$ and set L, U correspondingly. Additionally, in the construction of our datasets, we set all L_i in the datasets to 0 instead of to L ; and for all U in the datasets, we instead put 1. The result is two datasets u, v of binary values, with $\text{len}(u) = \text{len}(v) = 33$ and $d_{\text{Ham}}(u, v) = 1$.

We next clamp these datasets to values between L and U . Since $0 \leq L < U \leq 1$, in clamped datasets u and v we have $0 \mapsto L$ and $1 \mapsto U$.

Finally, we run $\mathcal{M}_\varepsilon^*(u)$ and $\mathcal{M}_\varepsilon^*(v)$, where $\mathcal{M}_\varepsilon^*$ is the mechanism used by the library for computing $BS_{L,U}^*$ with ε -DP and where $\varepsilon = 0.5$. We define the post-processing function $f : \mathbb{R} \rightarrow \{0, 1\}$ as follows: $f(x) = 1$ if and only if $x > \text{BRound}(33 \cdot L)$.

Running each of $f(\mathcal{M}_\varepsilon^*(u))$ and $f(\mathcal{M}_\varepsilon^*(v))$ 10,000 times gives the following results.

	$f(\mathcal{M}_\varepsilon^*(u))$	$f(\mathcal{M}_\varepsilon^*(v))$
0	1	9998
1	9999	2

Figure 9: “0.5-DP” results from IBM’s `diffprivlib` (with post-processing).

By the fact that $u \simeq_{\text{Ham}} v$, the definition of 0.5-DP, and the fact that DP is robust to post-processing per Proposition 2.18, we would expect to have

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 0]}{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 0]} \leq e^{0.5} < 1.65$$

and

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]}{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]} \leq e^{0.5} < 1.65.$$

Let $p_u = \Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]$ and $p_v = \Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]$, and suppose that $\frac{p_u}{p_v} \leq e^\varepsilon$ – that is, suppose the queries did satisfy ε -DP. Maximizing

$$\begin{aligned} & (p_u^{10,000} + 10,000p_u^{9999}(1 - p_u)) \\ & \cdot ((1 - p_v)^{10,000} + 10,000(1 - p_v)^{9999}p_v) \\ & + \binom{10,000}{2}(1 - p_v)^{9998}p_v^2 \end{aligned}$$

over $p_u, p_v \in [0, 1]$ such that $p_u \leq e^\varepsilon \cdot p_v$ and $(1 - p_v) \leq e^\varepsilon \cdot p_u$ for $\varepsilon = 0.5$ yields a maximum of $< 2^{-100}$. The results in Fig. 9, then, would occur with probability $< 2^{-100}$ if $\mathcal{M}_\varepsilon^*$ were truly an ε -DP function for $\varepsilon = 0.5$.

These results, then, support the calculation in Example Attack 5.13 that $\mathcal{M}_\varepsilon^*$ does not offer ε -DP for $\varepsilon = 0.5$.

Library-Mediated Attack 5.28 (Google DP (OpenMined wrappings): Bounded Mean, 64-bit Floats). We perform the same attack as described in Library-Mediated Attack 5.27, except we set $j = 7$ and set L, U correspondingly. The result is two datasets u, v of binary values, with $\text{len}(u) = \text{len}(v) = 33$ and $d_{\text{Ham}}(u, v) = 1$.

Finally, we run $\mathcal{M}_\varepsilon^*(u)$ and $\mathcal{M}_\varepsilon^*(v)$, where $\mathcal{M}_\varepsilon^*$ is the mechanism used by the library for computing the bounded mean (with values bounded by L and U) with ε -DP and where $\varepsilon = 0.5$. We define the post-processing function $f : \mathbb{R} \rightarrow \{0, 1\}$ as follows: $f(x) = 1$ if and only if $x > L$. Running each of $f(\mathcal{M}_\varepsilon^*(u))$ and $f(\mathcal{M}_\varepsilon^*(v))$ 10,000 times gives the following results.

	$f(\mathcal{M}_\varepsilon^*(u))$	$f(\mathcal{M}_\varepsilon^*(v))$
0	0	10,000
1	10,000	0

Figure 10: “0.5-DP” results from Google’s DP library (with post-processing).

By the fact that $u \simeq_{Ham} v$, the definition of 0.5-DP, and the fact that DP is robust to post-processing per Proposition 2.18, we would expect to have

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 0]}{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 0]} \leq e^{0.5} < 1.65$$

and

$$\frac{\Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]}{\Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]} \leq e^{0.5} < 1.65.$$

Let $p_u = \Pr[f(\mathcal{M}_\varepsilon^*(u)) = 1]$ and $p_v = \Pr[f(\mathcal{M}_\varepsilon^*(v)) = 1]$, and suppose that $\frac{p_u}{p_v} \leq e^\varepsilon$ – that is, suppose the queries did satisfy ε -DP. Then,

$$\begin{aligned} (p_u)^{10,000} \cdot (1 - p_v)^{10,000} &\leq (e^\varepsilon p_v \cdot (1 - p_v))^{10,000} \\ &\leq \left(\frac{e^\varepsilon}{4}\right)^{10,000}. \end{aligned}$$

The results in Fig. 10, then, would occur with probability $\leq 2^{-12,000}$ if $\mathcal{M}_\varepsilon^*$ were truly an ε -DP function for $\varepsilon = 0.5$. These results, then, support the calculation in Example Attack 5.13 that $\mathcal{M}_\varepsilon^*$ does not offer ε -DP for $\varepsilon = 0.5$.

6 Proposed Solutions to Incorrect Sensitivities

In this section, we propose several solutions in which we modify the implementation of bounded sum in order to recover a sensitivity that equals or is close to its idealized sensitivity. Each of the five methods is fitting in different types of situations and has different advantages and disadvantages, as we will explain.

We can broadly classify our solutions into two different types of fixes: the ones that modify the “standard” iterative sum function and that ones that check for the parameters before the call of the sum function.

We remark that all theorems in this section hold for normal floats (Definition 2.25). We believe that the results also hold for subnormal floats (Definition 2.29), and we will formalize this in a future version of this paper.

6.1 Bounding n

This strategy can be used as a stand-alone solution for the integers, and ideas from this strategy are important when implementing a solution over floats.

6.1.1 Solution for Integers

A natural first solution for integer data types is to directly avoid any overflow. We can achieve this by checking for overflow before calling the bounded sum function. This method applies to integer data types. Its main benefit is that it allows for the use of the idealized sensitivity (as proven in Theorem 4.4). However, this method only works when the length of the input dataset is known.

The direct check for overflow method, which we call *check multiplication* for conciseness, is defined as follows.

The check multiplication method only works when $\text{len}(u)$ is known to the data analyst (or when a DP count is performed to estimate $\text{len}(u)$, or when an upper bound on $\text{len}(u)$ is known), and it is overprotective in the sense that the potential overflow may only occur in a worst-case scenario. This method does not immediately work for floating-point arithmetic because blow-up can occur even when $U \cdot n < \max(T)$ and $L \cdot n > \min(T)$ due to rounding effect. The examples in Sections 5.4, 5.5, and 5.6 illustrate this effect.

Theorem 6.1 (Sensitivity of Bounded Sum With Checked Multiplication). *Let $BS_{L,U,n}^* : T_{[L,U]}^n \rightarrow T$ be the iterative bounded sum function on the type T of k -bit integers such that $U \cdot n \leq \max(T)$, $L \cdot n \geq \min(T)$. Then,*

$$\Delta_{CO} BS_{L,U,n}^* = \Delta_{Ham} BS_{L,U,n}^* = U - L.$$

Proof. Let $u, v \in T_{[L,U]}^n$ be two datasets such that $u \simeq_{CO} v$. If $n \cdot U \leq \max(T)$ and $n \cdot L \geq \min(T)$ then we claim that $BS^*(u) = BS(u)$ and $BS^*(v) = BS(v)$, per Definition 2.24. Thus,

$$|BS^*(u) - BS^*(v)| = |BS(u) - BS(v)| \leq U - L,$$

so

$$\Delta_{CO} BS_{L,U,n}^* = \Delta_{Ham} BS_{L,U,n}^* = U - L.$$

□

6.1.2 Use Case for Floats

When working with floats, it is also necessary to ensure that overflow cannot occur – that is, we must check that a computation cannot result in an answer of $\pm\text{inf}$. Per Definition 2.30, this is necessary for all the solutions we present for floating-point values. As a result of this behavior, a dataset that results in a sum of $\pm\text{inf}$ will always return an answer of $\pm\text{inf}$, even after noise addition; this deterministic behavior will cause privacy violations.

We now describe a generic check for overflow. Let $BS_{L,U,n}$ be the idealized bounded sum, $BS_{L,U,n}^*$ be an implementation of the bounded sum function, and $acc \in \mathbb{R}$ (where acc stands for accuracy) be some bound such that $|BS_{L,U,n}(v) - BS_{L,U,n}^*(v)| \leq acc$ for all $v \in T_{[L,U]}^n$. Additionally, let $Round_{\infty} : \mathbb{R} \rightarrow \mathbb{R}$ return the smallest float greater than or equal to the input; and let $Round_{-\infty} : \mathbb{R} \rightarrow \mathbb{R}$ return the largest float less than or equal to the input. We remark

that in the IEEE 754 standard, $Round_{-\infty}$ is referred to as “round toward $+\infty$, and $Round_{\infty}$ is referred to as “round toward $-\infty$ [35].

1. Compute $L' = Round_{-\infty}(L \cdot n)$ and $U' = Round_{\infty}(U \cdot n)$ using floating-point arithmetic. (This can be computed by performing floating-point multiplication with banker’s rounding and then going to the next float toward $\pm\text{inf}$ as appropriate.)
2. Ensure that $Round_{-\infty}(L' - acc) \neq -\text{inf}$ and that $Round_{\infty}(U' + acc) \neq \text{inf}$. (This can be computed by performing floating-point addition with banker’s rounding and then going to the next float toward $\pm\text{inf}$ as appropriate.)
3. If these checks pass, overflow cannot occur, and the computation can proceed; otherwise, reject computations using these parameter values n, L, U .

Remark 6.2. While this check for overflow works for any value of acc such that $|BS_{L,U,n}(v) - BS_{L,U,n}^*(v)| \leq acc$ for all $v \in T_{[L,U]}^n$, the accuracy bounds presented in Theorem 6.22 can be used to compute acc .

6.2 Modular Noise Addition

As a next step, we would like a solution that works for integer types when the length of the input dataset is not known, and which is not as conservative as the check multiplication method, given that check multiplication might not be applicable even if no overflow would actually occur in the computation of the bounded sum.

For this reason, we introduce *modular bounded sum*, which computes a bounded sum using modular arithmetic. We remark that modular bounded sum applies to both *signed* and *unsigned* integers, since we can treat the signed integers as a re-naming of the unsigned integers. This idea is described more formally in Section 2.4.2.

This method allows for overflow, which has the effect that two adjacent datasets can produce answers that, in terms of the absolute value of the difference between outputs, are as far from each other as possible; e.g., 0 and $2^k - 1$ for k -bit unsigned integers. However, when doing modular arithmetic, we should consider these two values to have a distance of only 1, captured by the following definition:

Definition 6.3 (Modular Distance). Let $x, y \in \mathbb{Z}/m\mathbb{Z}$ for some $m \in \mathbb{Z}$. Then, the *modular distance* between x and y is defined as

$$d_{Mod}^m(x, y) = \min\{x - y, y - x\}$$

where we use $\{0, \dots, m-1\}$ to represent the elements of $\mathbb{Z}/m\mathbb{Z}$ and, for the purposes of evaluating “min”, we specify that the elements are ordered as $0 < 1 < \dots < m - 1$.

Because we are working with modular distance, we measure sensitivity using *modular sensitivity*, as defined below.

Definition 6.4 (Modular Sensitivity). We define the *modular sensitivity* of a function $f : \text{Vec}(\mathcal{D}) \rightarrow \mathbb{Z}/m\mathbb{Z}$, with respect to a metric d on $\text{Vec}(\mathcal{D})$, as

$$\Delta_d^m f = \sup_{\substack{u, v \in \text{Vec}(\mathcal{D}) \\ u \simeq_d v}} d_{Mod}^m(f(u), f(v)).$$

We recall from the Library-Mediated Attack 5.22 that modular bounded sum does not fulfill DP when used with non-modular noise.

If noise were added to these bounded sum results in a non-modular fashion, u and u' would be very unlikely to result in similar answers, so a data analyst could readily tell the difference between a sum resulting from u and a sum resulting from u' , even though $u \simeq_d u'$. For this reason, we define *modular noise addition* and prove that modular noise addition can fulfill DP.

Theorem 6.5 (Modular Noise Addition Offers ε -DP). *Let Z be an integer-valued random variable and d a metric on $\text{Vec}(D)$. Suppose that for every function $f : \text{Vec}(D) \rightarrow \mathbb{Z}$ of sensitivity at most c with respect to d , the mechanism $\mathcal{M}(n) = f(n) + Z$ is ε -DP. Then for every function $f : \text{Vec}(D) \rightarrow \mathbb{Z}/n\mathbb{Z}$ of modular sensitivity at most c with respect to d , the mechanism $\mathcal{M}_{\text{mod}}(n) = f(n) + Z \bmod m$ is ε -DP.*

Proof. Let $\Delta_d f \leq c$. This means that, for $u \simeq_d u'$, we have $|f(u) - f(u')| \leq c$. Therefore, for all $m \in \mathbb{Z}$, there exists $k \in \mathbb{Z}$ such that $f(u) - f(u') \leq b + km$, where $b \in [-c, c] \cap \mathbb{Z}$. This is equivalent to $|f(u) - (f(u') + km)| \leq c$.

Because $|f(u) - (f(u') + km)| \leq c$, by the theorem statement $f(u) + Z$ and $f(u') + km + Z$ are ε -close. By post-processing, then, $f(u) + Z \bmod m$ and $f(u') + km + Z \equiv f(u') + Z \bmod m$ are ε -close. Therefore, adding Z modulo m provides ε -DP for $\mathbb{Z}/m\mathbb{Z}$ -valued functions with modular sensitivity $\leq c$ with respect to d on input datasets. \square

We now proceed to the proof that modular bounded sum results in a useful sensitivity.

Theorem 6.6 (Modular Sensitivity of Modular Bounded Sum). *Let T be a type of (signed or unsigned) k -bit integers and let $L, U \in T$. Let $BS_{L,U}^* : \text{Vec}(T_{[L,U]}) \rightarrow T$ be the iterative bounded sum function on k -bit integers of type T implemented with modular addition, and similarly for $BS_{L,U,n}^* : T_{[L,U]}^n \rightarrow T$. Then, where we let $m = 2^k$,*

1. (Unknown n .)

$$\begin{aligned} \Delta_{\text{Sym}}^m BS_{L,U}^* &= \Delta_{\text{ID}}^m BS_{L,U}^* = \\ &= \min \left\{ \left\lfloor \frac{m}{2} \right\rfloor, \max\{|L|, U\} \right\} \leq \max\{|L|, U\}. \end{aligned}$$

2. (Known n .)

$$\Delta_{\text{CO}}^m BS_{L,U,n}^* = \Delta_{\text{Ham}}^m BS_{L,U,n}^* = \min \left\{ \left\lfloor \frac{m}{2} \right\rfloor, U - L \right\} \leq U - L.$$

Note 6.7 (Implement with caution, unknown n). When working with k -bit signed integers, we must ensure that computing the sensitivity $|L|$ does not cause overflow. This is most readily done by disallowing $L = -2^k$ or by storing the sensitivity as an unsigned k -bit integer. When working with the 32-bit integers, for example, in the case where $L = -2^{31}$ we would have $|L|$ overflow to -2^{31} , which would lead to an underestimate when computing $\max\{|L|, U\}$. An easy method is to disallow $L = -2^{31}$; or, more generally, to disallow $L = \min(T)$ when working with the signed integers of type T .

Note 6.8 (Implement with caution, known n). When working with k -bit signed integers, the computation of the sensitivity $(U - L)$ should be stored as the corresponding unsigned value (i.e., the value in $[0, 2^k - 1] \cap \mathbb{Z}$, rather than a value in $[-2^{k-1}, 2^{k-1} - 1] \cap \mathbb{Z}$) to avoid issues of overflow. For example, when working with the 32-bit integers, setting $U = 0$ and $L = -2^{31}$ would result in a sensitivity of 2^{31} , which would overflow to -2^{31} in the signed integers and result in a negative sensitivity.

Proof of Theorem 6.6. We prove each part below. Let $m = 2^k$. Because addition over $\mathbb{Z}/m\mathbb{Z}$ is associative, we can essentially use the proof of Theorem 4.4, with slight modifications.

(1) Let $u, u' \in \text{Vec}(T_{[L,U]})$ be two datasets such that $u \simeq_{Sym} u'$. From the formal definition of a histogram in Definition 2.2, we observe that, for all $v \in \text{Vec}(T_{[L,U]})$,

$$\begin{aligned} BS_{L,U}^*(v) &= \left(\sum_{i=1}^{\text{len}(v)} v_i \right) \bmod m \\ &= \left(\sum_{z \in T} h_v(z) \cdot z \right) \bmod m \end{aligned}$$

Because $u \simeq_{Sym} u'$, we know that there is at most one value z^* such that $|h_u(z^*) - h_{u'}(z^*)| = 1$, and that, for all $z \neq z^*$, we have $|h_u(z) - h_{u'}(z)| = 0$.

We now calculate $d_{Mod}^m(BS_{L,U}^*(u), BS_{L,U}^*(u'))$. Recall that $d_{Mod}^m(x, y) = \min\{x - y, y - x\}$. We only evaluate the left term within the “min” expression and then compute the other side of the expression by a symmetric argument. Assume, without loss of generality, that $h_u(z^*) \geq h_{u'}(z^*)$. We can then write the following expressions.

$$\begin{aligned} &BS_{L,U}^*(u) - BS_{L,U}^*(u') \\ &= \left(\sum_{z \in T} h_u(z) \cdot z - \sum_{z \in T} h_{u'}(z) \cdot z \right) \bmod m \\ &= \left(\sum_{z \in T \setminus z^*} h_u(z) \cdot z - \sum_{z \in T \setminus z^*} h_{u'}(z) \cdot z \right) + \\ &\quad + (h_u(z^*) \cdot z^* - h_{u'}(z^*) \cdot z^*) \bmod m \\ &= z^* \cdot (h_u(z^*) - h_{u'}(z^*)) \bmod m \\ &\in \{0, \dots, z^*\}. \end{aligned}$$

By a symmetric argument, $BS_{L,U}^*(u') - BS_{L,U}^*(u) \in \{0, \dots, m - z^*\}$. Modular sensitivity is defined as the maximum possible modular distance between the sums of neighboring datasets. We observe that setting

$$z^* = \min\{\lfloor \frac{m}{2} \rfloor, \max\{|L|, U\}\}$$

will maximize the expression $\min\{z^*, -z^*\}$. By the definition of sensitivity, then,

$$\Delta_{Sym}^m BS_{L,U}^* \leq \min\{\lfloor \frac{m}{2} \rfloor, \max\{|L|, U\}\} \leq \max\{|L|, U\}.$$

By Theorem 2.15, we also have

$$\Delta_{ID}^m BS_{L,U}^* \leq \min\{\lfloor \frac{m}{2} \rfloor, \max\{|L|, U\}\} \leq \max\{|L|, U\}.$$

For the lower bound, consider the datasets $u = [0]$, $u' = [0, \min\{\lfloor \frac{m}{2} \rfloor, \max\{|L|, U\}\}]$. We note that $u \simeq_{ID} u'$. We also note that

$$d_{Mod}^m(BS_{L,U}^*(u), BS_{L,U}^*(u')) = \min\{\lfloor \frac{m}{2} \rfloor, \max\{|L|, U\}\}.$$

This means, then, that $\Delta_{ID}^m BS_{L,U}^* \geq \min\{\lfloor \frac{m}{2} \rfloor, \max\{|L|, U\}\}$. By the contrapositive of Theorem 2.15, then, $\Delta_{Sym}^m BS_{L,U}^* \geq \max\{|L|, U\}$.

Combining these upper and lower bounds on the idealized sensitivity tells us, then, that

$$\Delta_{Sym}^m BS_{L,U}^* = \Delta_{ID}^m BS_{L,U}^* = \min\{\lfloor \frac{m}{2} \rfloor, \max\{|L|, U\}\}.$$

(2) Let $u, u' \in \mathbb{R}^n$ be two datasets such that $u \simeq_{CO} u'$. By Lemma 2.10, there is a permutation $\pi \in S_n$ such that $\pi(u) \simeq_{Ham} u'$. This means there is at most one index i^* such that $\pi(u)_{i^*} \neq u'_{i^*}$. We now calculate $d_{Mod}^m(BS_{L,U}^*(u), BS_{L,U}^*(u'))$. Recall that $d_{Mod}^m(x, y) = \min\{x - y, y - x\}$. We only evaluate the left term within the “min” expression and then compute the other side of the expression by a symmetric argument. We can then write the following expressions.

$$\begin{aligned} BS_{L,U,n}^*(u) - BS_{L,U,n}^*(u') &= BS_{L,U,n}^*(\pi(u)) - BS_{L,U,n}^*(u') \\ &= \left(\sum_{i=1}^n \pi(u)_i - \sum_{i=1}^n u'_i \right) \bmod m \\ &= (\pi(u)_{i^*} - u'_{i^*}) \bmod m \\ &\in \{0, \dots, U - L\}, \end{aligned}$$

where we represent $U - L \in \{0, \dots, m - 1\}$. The final line follows from the fact that all values are clamped to the interval $[L, U]$, so the largest difference in sums arises when, without loss of generality, $\pi(u)_{i^*} = U$ and $u'_{i^*} = L$. By similar logic, $BS_{L,U,n}^*(u') - BS_{L,U,n}^*(u) \in \{0, \dots, m - (U - L)\}$. Thus

$$\begin{aligned} d_{Mod}^m(BS_{L,U}^*(u'), BS_{L,U}^*(u)) &\leq \min\{U - L, m - (U - L)\} = \\ &= \min\left\{\left\lfloor \frac{m}{2} \right\rfloor, U - L\right\} \leq U - L. \end{aligned}$$

For the lower bound, consider the datasets $u = [z]$ and $u' = [z']$, where $z, z' \in [L, U] \cap \mathbb{Z}$ such that $z - z' = \min\{\lfloor \frac{m}{2} \rfloor, U - L\}$. We note that $u \simeq_{Ham} u'$ and $d_{Mod}^m(BS_{L,U,n}^*(u), BS_{L,U,n}^*(u')) = \min\{\lfloor \frac{m}{2} \rfloor, U - L\}$. This means, then, that $\Delta_{ID}^m BS_{L,U}^* \geq \min\{\lfloor \frac{m}{2} \rfloor, U - L\}$. By the contrapositive of Theorem 2.15, it follows that $\Delta_{CO}^m BS_{L,U,n}^* \geq \min\{\lfloor \frac{m}{2} \rfloor, U - L\}$. Combining these upper and lower bounds on the idealized sensitivity tells us, then, that

$$\Delta_{CO}^m BS_{L,U,n}^* = \Delta_{Ham}^m BS_{L,U,n}^* = \min\left\{\left\lfloor \frac{m}{2} \right\rfloor, U - L\right\}.$$

This concludes the proof. □

Note that Theorem 6.6 relies on modular arithmetic, so it can only be applied to integer types. The primary disadvantage of this summation method is that computing a DP bounded sum on a dataset with a very large true sum may output a very small sum (or vice versa) due to the modular reduction. An advantage of this method is that it does *not* require loosening the sensitivity relation, and thus Theorem 6.6 achieves the same sensitivities as the corresponding idealized sensitivities in Section 4.3.

Another benefit of the modular addition method is that it preserves associativity. This means that applying bounded sum with modular addition to different orderings of an unordered dataset will always yield the same output. As we show in Sections 5.2 and 5.3, some implementations of bounded sum do not necessarily preserve associativity, in which case applying a function to different orderings of an unordered dataset can yield different outputs.

A final benefit of modular sum is that many libraries perform modular arithmetic on the integers by default. For example, overflow in both Python and C++ occurs in the modular fashion described in this section. Therefore, a naive implementation of bounded sum on the integers may actually be implementing the modular bounded sum and modular noise addition method described here. Interestingly, this means that libraries that ignore issues of overflow often fulfill the privacy guarantees; on the other hand, libraries that try to prevent overflow by using a strategy like saturation addition may encounter the vulnerabilities described in Section 5.2.

Moreover, as demonstrated in Library-Mediated Attack 5.22, it is essential that the added noise is also of integer type and that noise addition is done with modular arithmetic. Lastly, we remark that Theorem 6.5 (modular noise addition offers ϵ -DP) only applies to additive mechanisms, and hence this modular approach does not necessarily work for other uses of sensitivity, such as in the exponential mechanism.

6.3 Split Summation for Saturation Arithmetic

Before we turn to methods that work for floats, we present two methods that work for integers with saturation arithmetic even when the length of the input dataset is not known, as an alternative to the modular arithmetic method presented in Section 6.2. In this section we present Method 6.9, which we call *split summation for saturating integers*, and in the next section we present Method 6.14, called *randomized permutation* (RP). We will show that the implemented sensitivity for both methods matches the idealized sensitivity for bounded sum stated in Theorem 4.4.

As we showed in Section 5.2, integers with saturating arithmetic do *not* satisfy associativity. However, as noted in Property 5.8, we observe that non-associativity only arises from the mixing of positive and negative integers. This motivates the introduction of the *split summation method*, which consists of adding the negative terms and non-negative terms separately. We remark that this is the first of our methods which changes the summation function: instead of adding the elements in the order in which they appear in the dataset, we add them in a different pre-defined order.

Method 6.9 (Split Summation). Let U denote the upper bound, L the lower bound, and u the input dataset with numeric elements of type T . We define the *split summation* function as follows:

```

1  def BS*(u)
2      P = 0
3      N = 0
4      for i in u
5          if i < 0:
6              N = N + i
7          else
8              P = P + i
9      return P + N

```

Theorem 6.10 (Sensitivity of bounded sum with split summation). *Let $BS_{L,U}^* : \text{Vec}(T_{[L,U]}) \rightarrow T$ be the bounded sum function with split summation (Method 6.9) on the k -bit integers of type $T_{[L,U]}$, and similarly for $BS_{L,U,n}^* : T_{[L,U]}^n \rightarrow T$. Then,*

1. (Unknown n .) $\Delta_{\text{Sym}} BS_{L,U}^* = \Delta_{\text{ID}} BS_{L,U}^* = \max\{|L|, U\}$.
2. (Known n .) $\Delta_{\text{CO}} BS_{L,U,n}^* = \Delta_{\text{Ham}} BS_{L,U,n}^* = U - L$.

Proof. We prove each part below. Let $m = 2^k$.

1. Let $u, u' \in \text{Vec}(T)$ such that $u \simeq_{\text{Sym}} u'$. Additionally, for computing $BS_{L,U}^*(u)$, let P, N be the final values of P, N in Method 6.9. Likewise, for computing $BS_{L,U}^*(u')$, let P', N' be the final values of P, N in Method 6.9.

Note that saturation addition is associative. Because $u \simeq_{\text{Sym}} u'$, we know that there is at most one value z^* such that $|h_u(z^*) - h_{u'}(z^*)| = 1$, and that, for all $z \neq z^*$, we have $|h_u(z) - h_{u'}(z)| = 0$.

Consider the case where $z^* \geq 0$. By the logic in the proof of Theorem 4.4, then, $|N - N'| = 0$ and $|P - P'| \leq U$.

Now, consider the case where $z^* < 0$. By the logic in the proof of Theorem 4.4, then, $|N - N'| \leq |L|$ and $|P - P'| = 0$.

By the definition of sensitivity, then, $\Delta_{\text{Sym}} BS_{L,U}^* \leq \max\{|L|, U\}$. By Theorem 2.15, we also have $\Delta_{\text{ID}} BS_{L,U}^* \leq \max\{|L|, U\}$.

By the logic used in the proof of Theorem 4.4, we see that this is also a lower bound on the sensitivity, which implies that $\Delta_{\text{Sym}} BS_{L,U}^* = \Delta_{\text{ID}} BS_{L,U}^* = \max\{|L|, U\}$.

2. Let $u, u' \in T^n$ such that $u \simeq_{\text{CO}} u'$. Additionally, for computing $BS_{L,U,n}^*(u)$, let P, N be the final values of P, N in Method 6.9. Likewise, for computing $BS_{L,U,n}^*(u')$, let P', N' be the final values of P, N in Method 6.9.

Recall that saturation addition is associative. By Lemma 2.10, there is a permutation $\pi \in S_n$ such that $\pi(u) \simeq_{\text{Ham}} u'$. This means there is at most one index i^* such that $\pi(u)_{i^*} \neq u'_{i^*}$.

Let $\pi(u)^P$ represent the vector of values in $\pi(u) \geq 0$; likewise for $\pi(u)^N, u'^P, u'^N$. We consider the following cases and apply the logic of Theorem 4.4 due to associativity.

- $\pi(u)_{i^*} \geq 0$ and $u'_{i^*} \geq 0$. Here, $\pi(u)^P \simeq_{\text{Ham}} u'^P$, so $|BS_{L,U,n}^*(\pi(u)^P) - BS_{L,U,n}^*(u'^P)| \leq U - \max\{0, L\}$. Also, $\pi(u)^N = u'^N$, so $|BS_{L,U,n}^*(\pi(u)^N) - BS_{L,U,n}^*(u'^N)| \leq U - \max\{0, L\}$.

- $\pi(u)_{i^*} \geq 0$ and $u'_{i^*} < 0$. Here, $\pi(u)^P \simeq_{Sym} u'^P$, so $|BS_{L,U,n}^*(\pi(u)^P) - BS_{L,U}^*(u'^P)| \leq U$. Likewise, $\pi(u)^N \simeq_{Sym} u'^N$, so $|BS_{L,U,n}^*(\pi(u)^N) - BS_{L,U}^*(u'^N)| \leq L$. Overall, $|BS_{L,U,n}^*(u) - BS_{L,U,n}^*(u')| \leq U - L$.
- $\pi(u)_{i^*} < 0$ and $u'_{i^*} \geq 0$. Here, $\pi(u)^P \simeq_{Sym} u'^P$, so $|BS_{L,U,n}^*(\pi(u)^P) - BS_{L,U}^*(u'^P)| \leq U$. Likewise, $\pi(u)^N \simeq_{Sym} u'^N$, so $|BS_{L,U,n}^*(\pi(u)^N) - BS_{L,U}^*(u'^N)| \leq L$. Overall, $|BS_{L,U,n}^*(u) - BS_{L,U,n}^*(u')| \leq U - L$.
- $\pi(u)_{i^*} < 0$ and $u'_{i^*} < 0$. Here, $\pi(u)^N \simeq_{Ham} u'^N$, so $|BS_{L,U,n}^*(\pi(u)^N) - BS_{L,U,n}^*(u'^N)| \leq \min\{0, U\} - L$. Also, $\pi(u)^P = u'^P$, so $|BS_{L,U,n}^*(u) - BS_{L,U,n}^*(u')| \leq \min\{0, U\} - L$.

By the definition of sensitivity, then, $\Delta_{CO} BS_{L,U,n}^* \leq U - L$. By Theorem 2.15, we have $\Delta_{Ham} BS_{L,U,n}^* \leq U - L$.

By the logic used in the proof of Theorem 4.4, we see that this is also a lower bound on the sensitivity, which implies that $\Delta_{CO} BS_{L,U,n}^* = \Delta_{Ham} BS_{L,U,n}^* = U - L$.

□

Note that Method 6.9 would also work with modular addition instead of saturation arithmetic. This is because, when working with integers, issues with integer arithmetic only arise due to non-associativity. However, modular addition is associative, so the same procedure would hold.

6.4 Randomized Permutation for Saturating Integers

We now present the second alternative method for integers with saturating arithmetic, called *randomized permutation (RP) for saturating integers*. Like the split summation method, RP matches the sensitivity of bounded sum, but when combined with the split summation method in Method 6.9, will lead to one of our methods for float types (Section 6.6). The RP method randomly permutes the dataset before carrying out naive iterative summation. A potential accuracy advantage of RP over split summation is that since the positive and negative values are spread out, it is less likely that saturation will occur and cause a deviation from the true answer. On the other hand, the randomization has the downside that different runs may produce different results, but that is less of a concern given that we will be adding noise for privacy at the end anyway.

Definition 6.11 (Randomized Permutation (RP) Function). We define Π as a randomized function that, on input some dataset $u \in \mathcal{D}^n$ returns $\pi(u)$, where $\pi \in S_n$ is chosen uniformly at random from S_n , and $\pi(n)$ is as in Definition 2.3.

Until now, the notion of sensitivity that we have employed (Definition 2.11) was for deterministic functions. To define sensitivity for randomized functions, we follow [45] and require *couplings* between output distributions of neighboring datasets, and we then work with a definition of sensitivity based on these couplings. The idea of using couplings for this purpose was introduced in [45].

Definition 6.12 (Coupling). Let r_1, r_2 be two random variables defined over the probability spaces \mathcal{R}_1 and \mathcal{R}_2 , respectively. A *coupling* of r_1 and r_2 is a joint variable $(\tilde{r}_1, \tilde{r}_2)$ taking values in the product space $\mathcal{R}_1 \times \mathcal{R}_2$ such that \tilde{r}_1 has the same marginal distribution as r_1 and \tilde{r}_2 has the same marginal distribution as r_2 .

Definition 6.13 (Sensitivity of Randomized Functions). A randomized function $f : \text{Vec}(\mathcal{D}) \rightarrow \mathbb{R}$ has *sensitivity* Δf if for all neighboring datasets $u \simeq_d u'$ there exists a coupling $(\tilde{s}_u, \tilde{s}_{u'})$ of the random variables $s_u = f(u)$ and $s_{u'} = f(u')$ such that with probability 1, $d(\tilde{s}_u, \tilde{s}_{u'}) \leq \Delta f$.

Having defined the necessary terms, we now state our new summation method based on randomized permutations.

Method 6.14 (Randomized Permutation). Let $u \in \text{Vec}(T)$ for some type T . We define the *random permutation* of u as follows:

```

1 def random(u):
2     return perm(u) // pr 6.11

```

Theorem 6.15 (Sensitivity of Bounded Sum with RP-based Summation). *Let T be the type of k -bit signed integer with saturation arithmetic. Let $BS_{L,U}^* : \text{Vec}(T_{[L,U]}) \rightarrow T$ be the bounded sum function on a dataset to which RP has been applied (Method 6.14), and similarly for $BS_{L,U,n}^* : T_{[L,U]}^n \rightarrow T$. Then,*

1. (Unknown n .) $\Delta_{Sym} BS_{L,U}^* = \max\{|L|, U\}$.
2. (Known n .) $\Delta_{CO} BS_{L,U,n}^* = (U - L)$.

We first show that the idealized sensitivity is preserved if an *ordered* distance metric is used. Then, we show that the randomized permutation allows datasets at some unordered distance c to be treated as if they are at ordered distance c .

6.4.1 Sensitivities with Respect to Ordered Distance for Saturating Integers

We first show that the implemented sensitivity with respect to ordered distance metrics is equal to the idealized sensitivity.

Theorem 6.16 (Ordered Distance Sensitivities). *Let T be the type of k -bit integers with saturation arithmetic, let $BS_{L,U}^{*I} : \text{Vec}(T) \rightarrow T$ be the iterative bounded sum function from Definition 5.1, and similarly for $BS_{L,U,n}^* : T^n \rightarrow T$. Then,*

1. (Unknown n .) $\Delta_{ID} BS_{L,U}^* = \max\{|L|, U\}$.
2. (Known n .) $\Delta_{Ham} BS_{L,U,n}^* = (U - L)$.

Before we begin with the proof of Theorem 6.16, we provide two lemmas.

Lemma 6.17. *Let T be the type of k -bit integers. Then, for all $y, z, z' \in T$,*

$$|(y \boxplus z) - (y \boxplus z')| \leq |z - z'|.$$

Proof. Consider the cases where saturation does not occur (i.e., where $\min(T) \leq y+z \leq \max(T)$, and $\min(T) \leq y+z' \leq \max(T)$). Then $|(y \boxplus z) - (y \boxplus z')| = |(y+z) - (y+z')| = |z - z'|$.

Now, consider the case where exactly one of the sums saturates, say $y \boxplus z$. Then, $|y \boxplus z| < |y+z|$, and thus

$$|(y \boxplus z) - (y \boxplus z')| \leq |(y+z) - (y \boxplus z')| \leq |(y+z) - (y+z')| = |z - z'|.$$

Finally, if both $y \boxplus z$ and $y \boxplus z'$ saturate, then they must both equal $\max(T)$ (if $y > 0$) or both equal $\min(T)$ (if $y < 0$), so $|(y \boxplus z) - (y \boxplus z')| = 0 \leq |z - z'|$. □

Corollary 6.18. *Let T be the type of k -bit integers. Then, for all $y, z, z' \in T$,*

$$|(z \boxplus y) - (z' \boxplus y)| \leq |z - z'|.$$

Proof. Saturation addition is commutative. Therefore, by Lemma 6.17, we see that $|(z \boxplus y) - (z' \boxplus y)| \leq |z - z'|$. \square

Proof of Theorem 6.16. We begin with the proof of (2).

(2) Let $u \simeq_{Ham} u'$ be two datasets of length n with integer values of type T . Because $u \simeq_{Ham} u'$, there is at most one value $i \in [n]$ such that $u_i \neq u'_i$.

We now consider what happens at index i . For all $j < i$, we have $u_j = u'_j$, so we see that

$$BS_{L,U,i-1}^*[u_1, \dots, u_{i-1}] = BS_{L,U,i-1}^*[u'_1, \dots, u'_{i-1}].$$

By Lemma 6.17, we see that

$$|BS_{L,U,i}^*[u_1, \dots, u_i] - BS_{L,U,i}^*[u'_1, \dots, u'_i]| \leq |u_i - u'_i| \leq U - L. \quad (4)$$

Note that, for $j > i$, we have $u_j = u'_j$. By Corollary 6.18 and by induction, the sums will get no further apart from each other than they were in the expression above. Therefore, $|BS_{L,U,n}^*(s) - BS_{L,U,n}^*(s')| \leq (U - L)$. By the logic used in the proof of Theorem 4.4, we see that this is also a lower bound on the sensitivity, so

$$\Delta_{Ham} BS_{L,U,n}^* = (U - L).$$

(1) Let $u \simeq_{ID} u'$ be two datasets with integer values of type T . There is at most one index i^* such that, without loss of generality, $u_i = u'_i$ for all $i < i^*$, and $u_j = u'_{j+1}$ for all $j > i^*$.

Consider the case where there is no such i^* . Then $u = u'$, so $BS_{L,U}^*(u) = BS_{L,U}^*(u')$ and $|BS_{L,U}^*(u) - BS_{L,U}^*(u')| = 0$.

Now, consider the case where there is such an i^* . We note that, for all $y \in T$, we have $y \boxplus 0 = y$. Therefore, inserting a value $u_{i^*} = 0$ will not affect the final sums. It will also provide the effect that, for all $i \neq i^*$, $u_i = u'_i$.

We can now borrow the logic from Expression (4) to see that

$$|BS_{L,U}^*[u_1, \dots, u_i] - BS_{L,U}^*[u'_1, \dots, u'_i]| \leq |u_{i^*} - u'_{i^*}| \leq \max\{|L|, U\}. \quad (5)$$

By Corollary 6.18 and by induction, the sums will get no further apart from each other than they were in the expression above. Therefore, $|BS_{L,U}^*(u) - BS_{L,U}^*(u')| \leq \max\{|L|, U\}$. By the logic used in the proof of Theorem 4.4, we see that this is also a lower bound on the sensitivity, so

$$\Delta_{ID} BS_{L,U}^* = \max\{|L|, U\}.$$

\square

6.4.2 Randomized Permutation is a Transformation between Distance Metrics

Having shown that useful sensitivities can be established for ordered distance metrics, we show that randomized permutation provides a transformation from unordered distance metrics to ordered distance metrics. Using randomized permutation in combination with the theorems in Section 6.4.1 will then allow us to establish useful sensitivities for unordered distance metrics.

Theorem 6.19 (Coupling, Unknown n). *Let $u \in \mathcal{D}^m$ and $u' \in \mathcal{D}^n$. Additionally, let $\Pi : \mathcal{D}^m \rightarrow \mathcal{D}^m, \Pi' : \mathcal{D}^n \rightarrow \mathcal{D}^n$ be randomized permutations of datasets u and u' , respectively. For all datasets u, u' such that $u \simeq_{Sym} u'$, there exists a coupling $(\tilde{s}_u, \tilde{s}_{u'})$ of the datasets $s_u = \Pi(x)$ and $s_{u'} = \Pi'(u')$ such that $\tilde{s}_u \simeq_{ID} \tilde{s}_{u'}$.*

Proof. Without loss of generality, let $\text{len}(u') = \text{len}(u) + 1$. (In the $u \simeq_{Sym} u'$ case, we could have $\text{len}(u) = \text{len}(u')$, but then we would have a trivial coupling where each ordering $\Pi(u)$ is matched with the same ordering of elements $\Pi'(u')$, so we only consider the $|\text{len}(u) = \text{len}(u') + 1$ case.) We can then think of u' as being u , but with one “insertion”.

We now describe a way to enumerate u and u' so that equivalent terms are matched. Let $\mathcal{J}[i]$ be the i^{th} term in a sequence \mathcal{J} . Because $d_{Sym}(u, u') = 1$, we know that u and u' are identical from an unordered perspective, except u' has one additional row. This means, then, that there is some sequence \mathcal{J} of terms $\{1, \dots, m+1\}$ such that, for all $i \in \{1, \dots, m\}$, we have $u_i = u'_{\mathcal{J}[i]}$. As stated, there will be an unmatched u' term since $n = m + 1$; we call this term $u'_{\mathcal{J}[m+1]}$.

We now use this enumeration scheme to show that, for every permutation $\Pi(u)$, there is a permutation $\Pi'(u')$ such that $d_{ID}(\Pi(u), \Pi'(u')) = 1$. We first construct an intermediate ordering u'' . Let $\Pi(u)_k$ denote the k^{th} value in the vector $\Pi(u)$. Suppose that $\Pi(u)_k = u_i$. We pair this with $u_{\mathcal{J}[i]}$. We proceed in this way to construct a dataset u'' where $d_{ID}(\Pi(u), u'') = 0$.

Note that u'' does not include the leftover term $u'_{\mathcal{J}[m+1]}$. This means, then, that there are $(m+1)$ permutations $\Pi'(u')$ such that $d_{ID}(\Pi(u), \Pi'(u')) = 1$. This construction applies to every $\Pi(u)$. Therefore, every $\Pi(u)$ has $m+1$ permutations $\Pi'(u')$ such that $d_{ID}(\Pi(u), \Pi'(u')) = 1$.

Recall that, without loss of generality, we are assuming that $m \geq n$. This means that u has $m!$ permutations, while u' has $n!$ permutations. Note that our method pairs each permutation of u with $(|u| + 1)$ permutations of u' , and the mapping from permutations of u to permutations of u' is injective, with its inverse surjective. Because RP samples over permutations uniformly at random, we have shown that there is a coupling of randomness $(\tilde{s}_u, \tilde{s}_{u'})$ between $s_u = \Pi(u)$ and $s_{u'} = \Pi'(u')$ such that $\tilde{s}_u \simeq_{ID} \tilde{s}_{u'}$. □

Corollary 6.20 (Coupling, Known n). *Let $\Pi : \mathcal{D}^n \rightarrow \mathcal{D}^n$ be a randomized permutation of datasets of length n . For all datasets $u, u' \in \mathcal{D}^n$ such that $u \simeq_{CO} u'$, there exists a coupling $(\tilde{s}_u, \tilde{s}_{u'})$ of the datasets $s_u = \Pi(u)$ and $s_{u'} = \Pi(u')$ such that $\tilde{s}_u \simeq_{Ham} \tilde{s}_{u'}$.*

Proof. Suppose $u \simeq_{CO} u'$. By Theorem 2.10, this means that $d_{Sym}(u, u') = 2$. Using the path property of d_{Sym} (Lemma 2.14) and by Theorem 6.19, then, there is a coupling of randomness such that $d_{ID}(\Pi(u), \Pi(u')) \leq 2$ [45]. By Theorem 2.10, then, $d_{Ham}(\Pi(u), \Pi(u')) \leq 1$, so there is a coupling of randomness $(\tilde{s}_u, \tilde{s}_{u'})$ between $s_u = \Pi(u)$ and $s_{u'} = \Pi(u')$ such that $\tilde{s}_u \simeq_{Ham} \tilde{s}_{u'}$. □

We remark that the proofs of Theorem 6.19 and Corollary 6.20 are *constructive*, and so not only do we show that the coupling *exists*, but we also show how to find it efficiently.

6.4.3 Sensitivities for Unordered Distance

Now that we have proven the coupling between unordered distance metrics and ordered distance metrics, we can finally prove the sensitivity of bounded sum with RP-based summation; i.e., Theorem 6.15.

Proof of Theorem 6.15. We now prove each part of Theorem 6.15, which shows that the implemented sensitivities for Method 6.14 match the idealized sensitivities of bounded sum, with respect to unordered distances between datasets.

1. Method 6.14 first applies RP to the input dataset, and Theorem 6.19 tells us that RP transforms datasets that are neighboring with respect to d_{Sym} to a coupling of datasets that are neighboring with respect to d_{ID} . We can now work with datasets that are neighboring with respect to d_{ID} . Theorem 6.16 tells us that $\Delta_{ID}BS_{L,U}^* \leq \max\{|L|, U\}$, so we must also have $\Delta_{Sym}BS_{L,U}^* \leq \max\{|L|, U\}$. By the logic used in the proof of Theorem 4.4, we see that this is also a lower bound on the sensitivity, so

$$\Delta_{Sym}BS_{L,U}^* = \max\{|L|, U\},$$

completing the proof.

2. Method 6.14 first applies RP to the input dataset, and Corollary 6.20 tells us that RP transforms datasets that are neighboring with respect to d_{CO} to a coupling of datasets that are neighboring with respect to d_{Ham} . We can now work with datasets that are neighboring with respect to d_{Ham} . Theorem 6.16 tells us that $\Delta_{Ham}BS_{L,U,n}^* \leq (U - L)$, so we must also have $\Delta_{CO}BS_{L,U,n}^* \leq (U - L)$. By the logic used in the proof of Theorem 4.4, we see that this is also a lower bound on the sensitivity, so

$$\Delta_{CO}BS_{L,U,n}^* = U - L,$$

completing the proof.

□

6.5 Sensitivity from Accuracy

In this section, we show a relationship between an upper bound on a function's implemented sensitivity, and the function's idealized sensitivity and accuracy. The sensitivities that we derive apply directly to an unaltered implementation of iterative summation. The accuracy of a summation function is dependent on the number of terms being summed (meaning the sensitivities we derive are dependent on the number of terms being summed), so we also introduce a method called *truncated summation* that allows these sensitivities to apply to the unbounded DP setting in which the number of terms being summed is not necessarily known to the data analyst.

Because these upper bounds on implemented sensitivity depend on a function's accuracy, we also present methods from the numerical analysis literature for improving the accuracy of a summation function, and we compute the resulting sensitivities for these methods.

Lemma 6.21 (Relating Sensitivity and Accuracy).

$$\begin{aligned} \max_{u \simeq u'} |BS^*(u) - BS^*(u')| &\leq \\ &\max_{u \simeq u'} (|BS(u) - BS(u')| + |BS^*(u) - BS(u)| + |BS^*(u') - BS(u')|). \end{aligned}$$

Proof. This follows immediately from two applications of the triangle inequality. □

The term $|BS(u) - BS(u')|$ is bounded by the idealized sensitivity. The terms $|BS^*(u) - BS(u)|$ and $|BS^*(u') - BS(u')|$ can be bounded by the accuracy of the function BS^* as compared to its idealized counterpart BS . We turn to Higham, Wilkinson's, and Kahan's work on the accuracy of the sum function [24, 48, 31] for derivations of the accuracies of various summation strategies, including iterative summation.

Theorem 6.22 (Accuracy of Summation Algorithms). *Let $BS_{L,U,n}^* : T_{[L,U]}^n \rightarrow T$ be an implementation of the bounded sum function on the normal (k, ℓ) -bit floats of type T , and let $BS_{L,U,n} : \mathbb{R}_{[L,U]}^n \rightarrow \mathbb{R}$ compute the idealized bounded sum.*

1. If BS^* corresponds to iterative summation (Definition 5.1), then for all $u \in T_{[L,U]}^n$

$$|BS_{L,U,n}^*(u) - BS_{L,U,n}(u)| \leq \frac{n^2}{2^{k+1}} \cdot \max\{|L|, U\}.$$

2. If BS^* corresponds to pairwise summation (Definition 5.2) and $n < 2^k$, then for all $u \in T_{[L,U]}^n$

$$|BS_{L,U,n}^*(u) - BS_{L,U,n}(u)| \leq O\left(\frac{n \log n}{2^k}\right) \cdot \max\{|L|, U\}.$$

3. If BS^* corresponds to Kahan summation (Definition 5.3) and $n < 2^k$, then for all $u \in T_{[L,U]}^n$

$$|BS_{L,U,n}^*(u) - BS_{L,U,n}(u)| \leq O\left(\frac{n}{2^k}\right) \cdot \max\{|L|, U\}.$$

Proof. We show each bound for the corresponding sum function separately. For all of the proofs below, let $U \geq |L|$; a symmetric result follows for $U < |L|$, which allows us to get the bounds stated in the theorem.

1. We first show that

$$|BS^*(u) - BS(u)| \leq C_{n,k} \cdot \sum_{i=1}^n |u_i| \leq C_{n,k} \cdot nU,$$

where

$$C_{n,k} = \left(1 + \frac{1}{2^{k+1}}\right)^{n-1} - 1 \approx \frac{n-1}{2^{k+1}}$$

for $n < 2^k$. We proceed by induction on n . Hence assume that the proposition holds for $n-1$, and consider the addition step $BS^*(u) = BS^*(u_1, \dots, u_{n-1}) \oplus u_n$. Then,

$$\begin{aligned} |BS^*(u) - BS(u)| &\leq |BS^*(u) - (BS^*(u_1, \dots, u_{n-1}) + u_n)| + \\ &|BS^*(u_1, \dots, u_{n-1} + u_n) - BS(u_1, \dots, u_{n-1}) + u_n|. \end{aligned}$$

Let $s^* = BS^*(u_1, \dots, u_{n-1})$, $s = BS(u_1, \dots, u_{n-1})$. Then,

$$\begin{aligned}
& |BS^*(u) - BS(u)| \\
&= |BRound(s^* + u_n) - (s^* + u_n)| + |s^* - s| \\
&\leq \frac{|s^* + u_n|}{2^{k+1}} + |s^* - s| \\
&\leq \frac{|s^*| + |u_n|}{2^{k+1}} + \left(1 + \frac{1}{2^{k+1}}\right) \cdot |s^* - s| \\
&\leq \frac{\sum_{i=1}^n |u_i|}{2^{k+1}} + \left(1 + \frac{1}{2^{k+1}}\right) \cdot |s^* - s| \\
&\leq \left(\frac{1}{2^{k+1}} + \left(1 + \frac{1}{2^{k+1}}\right)C_{n-1}\right) \cdot \sum_{i=1}^n |u_i|.
\end{aligned}$$

Since

$$C_n = \frac{1}{2^{k+1}} + \left(1 + \frac{1}{2^{k+1}}\right)C_{n-1},$$

the induction step concludes. We remark that a similar accuracy bound is derived in [24] and [31].

2. We employ the accuracy bound shown in [24]. Let $t = 1/2^{k+1}$ (sometimes referred to as “unit roundoff” or “machine epsilon”). Then, [24, Eq. 3.6] shows that

$$|BS^*(u) - BS(u)| \leq \frac{\log_2 nt}{1 - \log_2 nt} \sum_{i=1}^n |v_i|.$$

In our setting, we have an upper bound on $\sum_{i=1}^n |v_i|$, and hence we can write

$$|BS^*(u) - BS(u)| \leq \frac{\log_2 nt}{1 - \log_2 nt} \cdot nU.$$

3. We employ the accuracy bound shown in [24]. Let $t = 1/2^{k+1}$. Then, [24, Eq. 3.11] shows that

$$|BS^*(u) - BS(u)| \leq (2t + O(nt^2)) \sum_{i=1}^n |v_i|$$

In our setting, we have an upper bound on $\sum_{i=1}^n |v_i|$ (namely, nU), and hence we can write

$$|BS^*(u) - BS(u)| \leq (2t + O(nt^2)) \cdot nU.$$

□

Theorem 6.23 (Sensitivity upper bounds). *Let $BS_{L,U,n}^* : T_{[L,U]}^n \rightarrow T$ be a bounded sum function on the normal (k, ℓ) -bit floats of type T . Additionally, let $-U \leq L < U$. Then,*

1. If BS^* corresponds to iterative summation (Definition 5.1), then

$$\Delta_{CO} BS_{L,U,n}^* \leq (U - L) + \left(\frac{n^2}{2^k}\right) \cdot \max\{|L|, U\}.$$

2. If BS^* corresponds to pairwise summation (Definition 5.2) and $n < 2^k$, then

$$\Delta_{CO}BS_{L,U,n}^* \leq (U - L) + O\left(\frac{n \log n}{2^k}\right) \cdot \max\{|L|, U\}.$$

3. If BS^* corresponds to Kahan summation (Definition 5.3) and $n < 2^k$, then

$$\Delta_{CO}BS_{L,U,n}^* \leq (U - L) + O\left(\frac{n}{2^k}\right) \cdot \max\{|L|, U\}.$$

Proof. The proof of each part follows from Lemma 6.21 and the corresponding part of Theorem 6.22. \square

In the cases where n is unknown, we can still make use of the upper bounds on the sensitivity that we showed in Theorem 6.23 by introducing a truncation method. To that end, we introduce the definition of the *truncated bounded sum* method.

Method 6.24 (Truncated Bounded Sum). Let u be a dataset of integers of type $T_{[L,U]}$, U an upper bound on the largest number in the dataset, L a lower bound on the smallest number in the dataset, n_{\max} a point at which to truncate the summation, and $BS_{L,U}^*$ any summation method. Then, the *truncated bounded sum* $BS_{L,U,n_{\max}}^{**} : \text{Vec}(T_{[L,U]}) \rightarrow T$ returns

- $BS_{L,U}^*(u)$ if $\text{len}(u) \leq n_{\max}$.
- $BS_{L,U}^*[u_1, \dots, u_{n_{\max}}]$ if $\text{len}(u) > n_{\max}$.

Theorem 6.25 (Idealized Sensitivity of Truncated Bounded Sum). *When the summation method $BS_{L,U,n_{\max}}^{**} : \text{Vec}(T_{[L,U]}) \rightarrow T$ corresponds to the idealized bounded sum $BS_{L,U}$ (Definition 4.2), the truncated bounded sum $BS_{L,U,n_{\max}}^{**} : \text{Vec}(T_{[L,U]}) \rightarrow T$ as defined in Method 6.24 has sensitivity $\Delta_{ID}BS_{L,U,n_{\max}}^{**} = \max\{|L|, U, U - L\}$.*

Proof. Consider all pairs of neighboring datasets $u \simeq_{ID} u'$. There are two possible cases: (1) Truncation does not occur for either dataset. (2) Truncation occurs for at least one dataset.

1. In case 1, we consider the standard bounded sum over a subset of all pairs of datasets $u \simeq_{ID} u'$, so the idealized sensitivity can be bounded by $\max\{|L|, U\}$ due to Theorem 4.4.
2. In case 2, we consider the standard bounded sum over all pairs of datasets $u \simeq_{Ham} u'$, where $u, u' \in T_{[L,U]}^{n_{\max}}$, so the idealized sensitivity can be bounded by $U - L$ due to Theorem 4.4.

Therefore, $\Delta_{ID}BS_{L,U,n_{\max}}^{**} \leq \max\{|L|, U, U - L\}$.

We next prove that $\Delta_{ID}BS_{L,U,n_{\max}}^{**} \geq \max\{|L|, U, U - L\}$. Consider the datasets $u = [U_1, \dots, U_{n_{\max}}, L]$ and $u' = [U_1, \dots, U_{n_{\max}-1}, L]$. We see that $u \simeq_{ID} u'$. We also see that $BS_{L,U,n_{\max}}^{**} = U - L$, and $BS_{L,U,n_{\max}+1}^{**} = \max\{|L|, U\}$.

Therefore, $\Delta_{ID}BS_{L,U,n_{\max}}^{**} = \max\{|L|, U, U - L\}$. \square

Corollary 6.26. *Let $BS_{L,U,n_{\max}}^{**} : \text{Vec}(T_{[L,U]}) \rightarrow T$ be the bounded sum function implemented with Method 6.24 on normal (k, ℓ) floats. Then,*

1. If Method 6.24 is implemented with iterative summation, then

$$\begin{aligned}\Delta_{ID}BS_{L,U,n_{\max}}^{**} &\leq \left(1 + \frac{n_{\max}^2}{2^k}\right) \cdot \max\{|L|, U, U - L\} \\ &\leq \left(2 + \frac{n_{\max}^2}{2^k}\right) \cdot \max\{|L|, U\}.\end{aligned}$$

2. If Method 6.24 is implemented with pairwise summation, then

$$\begin{aligned}\Delta_{ID}BS_{L,U,n_{\max}}^{**} &\leq \left(1 + O\left(\frac{n_{\max} \log n}{2^k}\right)\right) \cdot \max\{|L|, U, U - L\} \\ &\leq \left(2 + O\left(\frac{n_{\max} \log n}{2^k}\right)\right) \cdot \max\{|L|, U\}.\end{aligned}$$

3. If Method 6.24 is implemented with Kahan summation, then

$$\begin{aligned}\Delta_{ID}BS_{L,U,n_{\max}}^* &\leq \left(1 + O\left(\frac{n_{\max}}{2^k}\right)\right) \cdot \max\{|L|, U, U - L\} \\ &\leq \left(2 + O\left(\frac{n_{\max}}{2^k}\right)\right) \cdot \max\{|L|, U\}.\end{aligned}$$

Remark 6.27. We note that, where $U, L \geq 0$ or $U, L \leq 0$, we have $U - L \leq \max\{|L|, U\}$, which gives us the following improvements on the sensitivity bounds provided in Corollary 6.26.

1. If Method 6.24 is implemented with iterative summation, then

$$\Delta_{ID}BS_{L,U,n_{\max}}^{**} \leq \left(1 + \frac{n_{\max}^2}{2^k}\right) \cdot \max\{|L|, U\}.$$

2. If Method 6.24 is implemented with pairwise summation, then

$$\Delta_{ID}BS_{L,U,n_{\max}}^{**} \leq \left(1 + O\left(\frac{n_{\max} \log n}{2^k}\right)\right) \cdot \max\{|L|, U\}.$$

3. If Method 6.24 is implemented with Kahan summation, then

$$\Delta_{ID}BS_{L,U,n_{\max}}^* \leq \left(1 + O\left(\frac{n_{\max}}{2^k}\right)\right) \cdot \max\{|L|, U\}.$$

Note that, although Corollary 6.26 is stated in terms of ordered sensitivities, pre-processing datasets with RP (Method 6.14) will yield the same sensitivity bounds with respect to unordered distance metrics. Additionally, we can also employ a truncation strategy in the known n case if we want to establish an upper bound on the sensitivity of the implemented bounded sum function.

Data analysts should carefully choose n_{\max} . Because the implemented sensitivity grows with the length of the dataset, a larger n_{\max} means that the result will be noisier. However, if $n_{\max} \ll n$, causing many values of the dataset to be truncated, then the statistic may not be meaningful to the data analyst. Often, setting n_{\max} to an unrealistically large value can ensure that the sum is computed over the entire dataset while offering a minimal blow-up in the implemented sensitivity; for example, when working with 64-bit floats, the accuracy term will be less than 1 times the idealized sensitivity for $n_{\max} < 67$ million (meaning that the majority of the implemented sensitivity will be attributable to the idealized sensitivity). In some cases, a noisy count can be helpful for choosing n_{\max} .

6.5.1 Shifting Bounds

We note that, when working in the bounded DP (known n) setting with $|U|, |L| \gg U - L$, the sensitivity provided in Theorem 6.23 may be much larger than the idealized sensitivity $\Delta_{Ham}BS_{L,U,n} = \Delta_{CO}BS_{L,U,n} = U - L$. (To reduce notation, the following discussion considers the case where $U, L > 0$; a symmetric method can be used when $U, L < 0$.) We can achieve a sensitivity much closer to the idealized sensitivity by subtracting L from every element of the dataset so that all elements lie in the interval $[L' = 0, U' = U - L]$. We can then apply our solutions from Section 6.5, using 0 and U' in the sensitivity bound, and adding $L \cdot n$ at the end as post-processing. That is, we convert a method BS^* from Section 6.5 into a method BS^{**} with sensitivity depending on $U - L$ as follows:

$$BS_{L,U,n}^{**}(u) = BS_{L,U,n}^*(u_1 - L, u_2 - L, \dots, u_n - L) + L \cdot n,$$

where a noisy version of the $BS_{L,U,n}^*$ term is computed first and $L \cdot n$ is added as a post-processing step.

6.6 Split Summation with Random Permutation

In this section, we present our method for a bounded sum function on floats that recovers sensitivity within a constant factor of the idealized sensitivity. While the attack presented in Section 5.6 shows that the application of RP alone does not solve issues related to floating point arithmetic, we will show that we can achieve a good sensitivity for bounded sum by combining RP with our previous method of split summation and by changing the rounding mode from banker’s rounding to round toward 0 (see Definition 2.34).

Round toward 0 allows us to achieve tighter bounds due to the property that, in many cases, the rounding never leads to an increase in the magnitude of the sum. We use this to argue that the sums of adjacent datasets rarely diverge. On the other hand – as we saw in the counterexamples in Sections 5.4 and 5.5 – banker’s rounding can make the magnitude of the sum increase much more quickly or much less quickly than expected, causing similar datasets to sometimes produce surprisingly dissimilar sums.

Method 6.28 (Split Summation on Floats). Let U denote the upper bound, L the lower bound, and u the input dataset with elements of type T with round toward zero. We define `fsplit_sum` as follows:

```

1  def fsplit(u)
2      P = 0
3      N = 0
4      for i in s
5          if i < 0:
6              N = N + i
7          else:
8              P = P + i
9      return P ⊕ N

```

There are two important features to note in the algorithm above.

1. The iterative sums N and P are computed using round toward zero as defined in Definition 2.34.

2. The returned sum $P \oplus N$ is computed using standard banker's rounding as defined in Definition 2.33. (This is the only use of banker's rounding in the algorithm.)

Remark 6.29 (Avoiding Overflow). Let \max be the largest float $< \text{inf}$, and let \min be the smallest float $> -\text{inf}$. To prevent overflow from occurring (which is necessary, as described in Section 6.1.2), we do the following. If $P = \text{inf}$, we set $P = \max$, and if $N = -\text{inf}$, we set $N = \min$. Issues related to non-associativity will not arise because we are using split summation and, per Section 6.3, issues of non-associativity do not affect split summation.

We now state a theorem about the sensitivity of the function described in Method 6.28.

Theorem 6.30 (Δ_{Sym} and Δ_{CO} of Method 6.28). *Let $BS_{L,U}^* : \text{Vec}(T) \rightarrow T$ be the bounded sum function with split summation and round toward zero (Method 6.28) on the normal (k, ℓ) -bit floats of type T , and similarly for $BS_{L,U,n}^* : T^n \rightarrow T$. Define $f_{L,U} = BS_{L,U}^* \circ RP$ and $f_{L,U,n} = BS_{L,U,n}^* \circ RP$. Then,*

1. (Unknown n .) Let $b_U = \lfloor \log_2 U \rfloor$, $b_L = \lfloor \log_2 |L| \rfloor$. Then,

$$\begin{aligned} \Delta_{Sym} f_{L,U} &\leq \max\{2^{b_U} + U, 2^{b_L} + |L|\} + \max\{2^{b_U}, 2^{b_L}\} \\ &\leq 3 \cdot \max\{U, |L|\}. \end{aligned}$$

2. (Known n .) Let $c_U = \lfloor \log_2 U \rfloor$, $b_U = \min\{c_U, \lfloor \log_2(nU) \rfloor - k\}$, $c_L = \lfloor \log_2 |L| \rfloor$, $b_L = \min\{c_L, \lfloor \log_2(|nL|) - k \rfloor\}$.

Then,

$$\begin{aligned} \Delta_{CO} f_{L,U,n} &\leq 2^{b_U} + U + 2^{b_L} + |L| \\ &\quad + \max\{\min\{2^{b_U}, 2^{c_U+1}\}, \min\{2^{b_L}, 2^{c_L+1}\}\} \\ &\leq 5 \cdot \max\{U, |L|\}. \end{aligned}$$

Recall, from Theorem 6.19 for unknown dataset length and Corollary 6.20 for known dataset length, that RP allows us to reason about the ordered distance between datasets, even when we are only provided with a guarantee about the unordered distance between the datasets. For this reason, we state a theorem about the sensitivities of $BS_{L,U}^*$ and $BS_{L,U,n}^*$ with respect to ordered distance metrics, since sensitivities with respect to ordered distance metrics will be able to be converted to statements about sensitivities with respect to unordered distances.

Theorem 6.31 (Ordered Sensitivity of Split Summation with RP for Floats). *Let $BS_{L,U}^* : \text{Vec}(T) \rightarrow T$ be the bounded sum function with split summation and round toward zero (Method 6.28) on the normal (k, ℓ) -bit floats of type T , and similarly for $BS_{L,U,n}^* : T^n \rightarrow T$. Then,*

1. (Unknown n .) Let $b_U = \lfloor \log_2 U \rfloor$, $b_L = \lfloor \log_2 |L| \rfloor$. Then,

$$\begin{aligned} \Delta_{ID} BS_{L,U}^* &\leq \max\{2^{b_U} + U, 2^{b_L} + |L|\} + \max\{2^{b_U}, 2^{b_L}\} \\ &\leq 3 \cdot \max\{U, |L|\}. \end{aligned}$$

2. (Known n .) Let $c_U = \lfloor \log_2 U \rfloor$, $b_U = \min\{c_U, \lfloor \log_2(nU) \rfloor - k\}$, $c_L = \lfloor \log_2 |L| \rfloor$, $b_L = \min\{c_L, \lfloor \log_2(|nL|) \rfloor - k\}$. Then,

$$\begin{aligned} \Delta_{\text{Ham}} BS_{L,U,n}^* &\leq 2^{b_U} + U + 2^{b_L} + |L| + \max\{\min\{2^{b_U}, 2^{c_U+1}\}, \\ &\quad \min\{2^{b_L}, 2^{c_L+1}\}\} \\ &\leq 5 \cdot \max\{U, |L|\}. \end{aligned}$$

Before proceeding with the proof of Theorem 6.31, we state and prove some lemmas about the behavior of floating point arithmetic. The motivation for these lemmas is to determine the situations in which rounding can cause the iterative sums of adjacent datasets to diverge. Specifically, we prove that this divergence-causing rounding only occurs when the intermediate sum moves from a floating point interval of the form $[2^k, 2^{k+1})$ for $k \in \mathbb{Z}$ to a floating point interval of the form $[2^\ell, 2^{\ell+1})$ where $k \neq \ell \in \mathbb{Z}$.

Lemma 6.32 (Distance Between Sums – Part 1). *Let s, s', v be (k, ℓ) -bit floats such that $\text{sign}(s) = \text{sign}(s') = \text{sign}(v)$, $|s'| \geq |s|$, and $\text{ULP}_{(k,\ell)}(s+v) = \text{ULP}_{(k,\ell)}(s)$. Then it follows that $|\text{RTZ}(s'+v) - \text{RTZ}(s+v)| \leq |s - s'|$.*

Proof. We present the proof for the case that s, s', v are all positive; the case in which they are all negative is similar. We first consider $\text{RTZ}(s+v)$.

$$\begin{aligned} \text{RTZ}(s+v) &= \left\lfloor \frac{s+v}{\text{ULP}_{(k,\ell)}(s+v)} \right\rfloor \cdot \text{ULP}_{(k,\ell)}(s+v) \\ &= \left\lfloor \frac{s+v}{\text{ULP}_{(k,\ell)}(s)} \right\rfloor \cdot \text{ULP}_{(k,\ell)}(s) \\ &= s + \left\lfloor \frac{v}{\text{ULP}_{(k,\ell)}(s)} \right\rfloor \cdot \text{ULP}_{(k,\ell)}(s), \end{aligned}$$

with the final equality following from the fact that s is an integer multiple of $\text{ULP}_{(k,\ell)}(s)$.

We now consider $\text{RTZ}(s'+v)$.

$$\begin{aligned} \text{RTZ}(s'+v) &= \left\lfloor \frac{s'+v}{\text{ULP}_{(k,\ell)}(s'+v)} \right\rfloor \cdot \text{ULP}_{(k,\ell)}(s'+v) \\ &\leq s' + \left\lfloor \frac{v}{\text{ULP}_{(k,\ell)}(s'+v)} \right\rfloor \cdot \text{ULP}_{(k,\ell)}(s'+v) \\ &\leq s' + \left\lfloor \frac{v}{\text{ULP}_{(k,\ell)}(s)} \right\rfloor \cdot \text{ULP}_{(k,\ell)}(s). \end{aligned}$$

Therefore, we see that $|\text{RTZ}(s'+v) - \text{RTZ}(s+v)| \leq |s - s'|$. \square

Observe that Lemma 6.32 is not affected by $\text{ULP}_{(k,\ell)}(\text{RTZ}(s'+v))$. This is because the spacing between NRFs increases in each interval, so adding v to s' will not have a larger effect RTZ on the sum than adding v to s , where $|s'| > |s|$. Note that this does not hold if banker's rounding were used instead of round toward zero.

Lemma 6.33 (Distance Between Sums – Part 2). *If two (k, ℓ) -bit floats $s, s' \geq 0$ are such that $s < s'$ (with $s' - s = d$), and we add a term $v > 0$ to both terms such that $\text{ULP}_{(k,\ell)}(s+v) = 2 \cdot \text{ULP}_{(k,\ell)}(s)$, it follows that $\text{RTZ}(s'+v) - \text{RTZ}(s+v) \leq d + \text{ULP}_{(k,\ell)}(s)$.*

Proof. Let $s \in [2^{k+m}, 2^{k+m+1})$ for some $m \in \mathbb{Z}$, so $ULP_{(k,\ell)}(s) = \frac{2^{k+m}}{2^{k+m}} = 2^m$. Also, suppose that $v \in [p \cdot 2^m, p \cdot 2^{m+1})$ for some $p \in \mathbb{Z}$, and that $2^{k+m+1} - s = q \cdot 2^m$ for some $q \in \mathbb{Z}$. Additionally, suppose that $s' \geq 2^{k+m+1}$.

We now note that $RTZ(s+v) = s + q \cdot 2^m + \lfloor \frac{p-q}{2} \rfloor \cdot 2^{m+1}$. We proceed with a proof by cases.

1. Suppose that p is even, but that $(p - q)$ is odd. This means, then, that computing $RTZ(s' + v)$ will yield $RTZ(s' + v) \leq s' + \frac{p}{2} \cdot 2^{m+1} = s' + p \cdot 2^m$. On the other hand, computing $RTZ(s + v)$ will result in a value of $RTZ(s + v) = s + q \cdot 2^m + \frac{p-q-1}{2} \cdot 2^{m+1} = s + q \cdot 2^m + (p - q - 1) \cdot 2^m = s + (p - 1) \cdot 2^m$.

In this case, then, $RTZ(s' + v) - RTZ(s + v) \leq d + 2^k = d + ULP_{(k,\ell)}(s)$.

2. Now, suppose that p is even and $(p - q)$ is even. This means, then, that computing $RTZ(s' + v)$ will yield $RTZ(s' + v) \leq s' + \frac{p}{2} \cdot 2^{m+1} = s' + p \cdot 2^m$. On the other hand, computing $RTZ(s + v)$ will result in a value of $RTZ(s + v) = s + q \cdot 2^m + \frac{p-q}{2} \cdot 2^{m+1} = s + p \cdot 2^m$.

In this case, then, $RTZ(s' + v) - RTZ(s + v) \leq d$.

3. Finally, suppose that p is odd. This means, then, that computing $RTZ(s' + v)$ will yield $RTZ(s' + v) \leq \frac{p-1}{2} \cdot 2^{m+1} = s' + (p - 1) \cdot 2^m$. On the other hand, computing $RTZ(s + v)$ will result in a value of $RTZ(s + v) = s + q \cdot 2^m + \lfloor \frac{p-q}{2} \rfloor \cdot 2^{m+1} \geq s + q \cdot 2^m + (p - q - 1) \cdot 2^m = s + (p - 1) \cdot 2^m$.

In this case, then, $RTZ(s' + v) - RTZ(s + v) \leq d$.

Note that, in the cases above, we only considered $s' \geq 2^{k+m+1}$. We now consider the case where $s, s' \in [2^{k+m}, 2^{k+m+1})$. Let q' be the corresponding “ q value” for s' . The worst case arises when $(p - q)$ is odd and $(p - q')$ is even. We then see the behavior provided in case (1), where p is even but $(p - q)$ is odd, so the bounds on the difference as provided in the theorem statement still hold.

Therefore, $RTZ(s' + v) - RTZ(s + v) \leq d + 2^m = d + ULP_{(k,\ell)}(s)$. □

Corollary 6.34. *If two (k, ℓ) -bit floats $s, s' \geq 0$ are such that $s < s'$ (with $s' - s = d$), and we add a term $v > 0$ to both terms such that $ULP_{(k,\ell)}(RTZ(s + v)) = 2^m \cdot ULP_{(k,\ell)}(s)$ for $m \in \mathbb{Z}$, it follows that $RTZ(s' + v) - RTZ(s + v) \leq d + (2^m \cdot ULP_{(k,\ell)}(s) - ULP_{(k,\ell)}(s))$.*

Before proceeding to the next lemma, we introduce a new term.

Definition 6.35 (Boundary). We refer to a value of the form 2^m for $m \in \mathbb{Z}$ as a *boundary*. This is because, for values $s \in [2^m, 2^{m+1})$, the distance between NRFs is twice that of the distance between NRFs for values $s \in [2^{m-1}, 2^m)$. Whenever we have a floating point value $s \in [2^m, 2^{m+1})$, and we add a value v where $RTZ(s + v) \in (2^{m+j}, 2^{m+j+1})$ for $j \in \mathbb{Z}^+$, we say that the sum *crossed a boundary*.

Lemma 6.36 (Maximum Boundary Crossing Given U). *Let U be the upper bound for the bounded sum function over the (k, ℓ) -bit floats, and let $b \in \mathbb{Z}$ be such that $U \in [2^b, 2^{b+1})$. Then, in the worst case, the final crossed boundary corresponds to the boundary 2^{k+b} , between the intervals $[2^{k+b-1}, 2^{k+b})$ and $[2^{k+b}, 2^{k+b+1})$.*

Proof. Under round toward zero, the maximum sum that a dataset can reach with $U \in [2^b, 2^{b+1})$ is $2^{b+1} \cdot 2^k$. Therefore, the final boundary that we cross is no greater than the boundary between the intervals $[2^{k+b-1}, 2^{k+b})$ and $[2^{k+b}, 2^{k+b+1})$. Note that, because the maximum sum is 2^{k+b+1} , we have not yet *crossed* into the interval $(2^{k+b+1}, 2^{k+b+2})$. \square

Lemma 6.37 (Maximum Boundary Crossing Given U, n). *Let U be the upper bound for the bounded sum function over the (k, ℓ) -bit floats, let n be the length of the dataset, and let $m \in \mathbb{Z}$ be such that $U \cdot n \in [2^m, 2^{m+1})$. Then, in the worst case, the final crossed boundary corresponds to the boundary 2^m , between the intervals $[2^{m-1}, 2^m)$ and $[2^m, 2^{m+1})$.*

Proof. Since we are in the round-toward-zero rounding mode, the maximum sum of a length n dataset with upper bound U is $n \cdot U$. Because we define $m \in \mathbb{Z}$ be such that $U \cdot n \in [2^m, 2^{m+1})$ and $U \cdot n \in [2^m, 2^{m+1})$, we see that 2^m is the largest boundary. \square

Having developed lemmas about the behavior of floating point arithmetic under round toward 0, we are now ready to reason about the sensitivity of the bounded sum function described in Method 6.28.

We remark that the intervals stated in Lemma 6.36 are a worst-case scenario, and thus are not necessarily reached by the sum value.

Lemma 6.38 (Sensitivity $\Delta_{Ham} BS_{0,U,n}^*$). *Suppose we have two length n datasets $u \simeq_{Ham} u'$ of (k, ℓ) -bit floats, with values clipped within the interval $[0, U]$ with $U \in [2^c, 2^{c+1})$, where the maximum possible sum (in this case, $\min\{n \cdot U, 2^{k+c}\}$) is $\in [2^{k+b}, 2^{k+b+1})$. Then, the maximum difference in sums will be*

$$\Delta_{Ham} BS_{0,U,n}^* \leq \sum_{i=-\infty}^{k+b-1} \frac{2^i}{2^k} + U = 2^b + U \leq 2U,$$

where $2^b \leq \min\{U, \frac{n \cdot U}{2^k}\}$.

Proof. Let $u \simeq_{Ham} u'$ be two length n datasets of (k, ℓ) -bit floats. In the worst case, this means that u and u' are identical, except they differ at some value $u_i \neq u'_i$. In the worst case, without loss of generality, $u_i = 0$ and $u'_i = U$.

We note that we have $BS_{0,U}^*([u'_1, \dots, u'_i]) - BS_{0,U}^*([u_1, \dots, u_i]) \leq U$.

By Lemma 6.32, if for all $j \geq i$ and some $m \in \mathbb{Z}$ we have

$$BS_{0,U}^*([u_1, \dots, u_{j-1}]), BS_{0,U}^*([u_1, \dots, u_j]) \in [2^m, 2^{m+1}),$$

then we will continue to have

$$BS_{0,U}^*([u'_1, \dots, u'_j]) - BS_{0,U}^*([u_1, \dots, u_j]) \leq U.$$

However, we may have $BS_{0,U}^*([u_1, \dots, u_{j-1}]) \in [2^m, 2^{m+1})$ but $BS_{0,U}^*([u_1, \dots, u_j]) \notin [2^m, 2^{m+1})$. In this case, we apply Lemma 6.33. This tells us that the rounding that causes the difference in sums to grow only occurs when adding a value to the smaller sum (in this case, adding u_j) causes the sum to cross from an interval of values in $[2^m, 2^{m+1})$ to an interval of values in $[2^p, 2^{p+1})$ for some $p \in \mathbb{Z} > m$. By Lemma 6.32, the difference in sums does *not* grow when the smaller sum does *not* cross from one interval into another interval. Therefore, we only experience growth

in the difference of sums (recall that the “true” difference is $BS_{0,U}(u') - BS_{0,U}(u) = U$, so we expect this much of a difference in sums) when these interval crossings occur.

Suppose we are crossing from a value r with $ULP_{(k,\ell)}(r) = 2^q$ to a value r' with $ULP_{(k,\ell)}(r') = 2^{q+1}$. By Lemma 6.33, we experience growth in the difference in sums of at most 2^q .

By Lemma 6.36, the maximum round-up occurs when crossing from an interval with a $ULP_{(k,\ell)}$ of 2^{b-2} to an interval with a $ULP_{(k,\ell)}$ of 2^{b-1} . Therefore, the greatest growth in the difference in sums that we can experience is 2^{b-1} . We note that the crossing immediately prior to this can result in a growth in the difference in sums of 2^{b-2} . Therefore, we can bound the overall growth in difference in sums by $2^{b-1} + 2^{b-2} + 2^{b-3} + \dots = \sum_{m=-\infty}^{b-1} 2^m = \sum_{m=-\infty}^{k+b-1} \frac{2^m}{2^k}$. Because m ranges from $-\infty$ to $(k+b-1)$, the growth in the difference in sums is bounded by $2 \cdot 2^{b-1} = 2^b$, which is turn bounded by U given that, by assumption, $U < 2^{b+1}$. (Note that the bound provided in Corollary 6.34 would have a smaller effect than the roundings at every boundary crossing that we have here, so we only consider the case described in Lemma 6.33 when evaluating worst-case behavior.)

Therefore, the overall difference in sums for two datasets with difference in sums $|BS_{L,U}(u) - BS_{L,U}(u')| \leq U$ is at most the bound provided in the lemma statement,

$$|BS_{0,U,n}^*(u) - BS_{0,U,n}^*(u')| \leq \sum_{i=-\infty}^{k+b-1} \frac{2^i}{2^k} + U = 2^b + U \leq 2U.$$

□

Corollary 6.39 (Sensitivity $\Delta_{ID} BS_{0,U}^*$). *Suppose we have two datasets $u \simeq_{ID} u'$ of (k, ℓ) -bit floats, with values clipped to the interval $[0, U]$ and with $U \in [2^b, 2^{b+1})$, meaning that the maximum possible sum is $\in [2^{k+b}, 2^{k+b+1})$. Then, the maximum difference in sums will be*

$$\Delta_{ID} BS_{0,U}^* \leq \sum_{i=-\infty}^{k+b-1} \frac{2^i}{2^k} + U = 2^b + U \leq 2U.$$

Proof. Let $u \simeq_{ID} u'$ be two datasets of (k, ℓ) -bit floats. Without loss of generality, let $\text{len}(u) < \text{len}(u')$. Let u'' be the dataset that results from inserting a 0 immediately before the first location u_i where we have $u_i \neq u'_i$. (Because 0's have no effect on a floating point summation, this insertion will not affect the difference in sums.)

We now note that $u'' \simeq_{Ham} u'$, so we can apply Lemma 6.38. Because n is not known, we can pretend that we have $n = \infty$, which causes the “min” in the lemma statement to evaluate to 2^{k+b} . (We could not actually have $n = \infty$ since computers are finite, but this helps evaluate the worst case value for the “min” term.) Using this lemma and by the definition of sensitivity, we have

$$|BS_{0,U,n}^*(u'') - BS_{0,U,n}^*(u')| \leq \sum_{i=-\infty}^{k+b-1} \frac{2^i}{2^k} + d = 2^b + d \leq 2U.$$

Because of how we constructed u'' , we equivalently have

$$|BS_{0,U}^*(u) - BS_{0,U}^*(u')| \leq \sum_{i=-\infty}^{k+b-1} \frac{2^i}{2^k} + d = 2^b + d \leq 2U.$$

We know that we have $u \simeq_{ID} u'$, so by the definition of sensitivity, we have

$$\Delta_{ID} BS_{0,U}^* \leq \sum_{i=-\infty}^{k+b-1} \frac{2^i}{2^k} + d = 2^b + d \leq 2U.$$

□

Proof of Theorem 6.31. Recall that we are working with the bounded sum function as described in Method 6.28. Additionally, note that we prove (1) $\Delta_{ID} BS_{L,U}^*$ and (2) $\Delta_{Ham} BS_{L,U,n}^*$.

1. Let $u \simeq_{Ham} u'$ be two datasets of length n with (k, ℓ) -bit floating point values of type T . In the worst case, this means that u and u' are identical, except they differ at some value $u_i \neq u'_i$.

By Corollary 6.39, where $U \in [2^b, 2^{b+1})$, the calculation of P has a sensitivity of at most $2^b + U$. By generalizing Corollary 6.39 to work with values in $[L, 0]$ instead, it follows by symmetry that, where $L \in (-2^{b'+1}, -2^{b'})$, the calculation of N has a sensitivity of at most $2^{b'} + |L|$.

The last step of Method 6.28 consists of adding P and N using banker's rounding. We need to account for the rounding that can occur in this step. Let M_P be the maximum value of P , and let M_N be the minimum value (i.e., maximum magnitude) of N . We see, then, that $M_P = 2^b \cdot 2^{k+1}$, and that $M_N = -2^{b'} \cdot 2^{k+1}$.

Suppose that $M_P \geq |M_N|$. In the worst case, for all values $v < M_P$, we have $ULP_{(k,\ell)}(v) \leq 2^b$, so the worst-case round-down that we can get from combining P and N (using banker's rounding) is $\frac{2^b}{2}$. Likewise, the worst-case round-up that we can get from combining P and N (using banker's rounding) is $\frac{2^b}{2}$. Therefore, the worst-case overall rounding is 2^b . Likewise, for the $M_P \leq |M_N|$ case, the worst-case rounding is $2^{b'}$.

We recall that the sensitivity for computing N is $\leq 2^{b'} + |L|$; and that the sensitivity for P is $\leq 2^b + U$. Therefore, the overall sensitivity initially appears to be

$$\Delta_{ID} BS_{L,U}^* \leq 2^b + U + 2^{b'} + |L| + \max\{2^b, 2^{b'}\}.$$

We can improve on this upper bound on the sensitivity, though, with the observation that we have $d_{ID}(s, s') = 1$. A single insertion or a single deletion will only affect one of M_P or M_N (whereas the sensitivity above assumes that both M_P and M_N would be affected). In the worst case, the single insertion or deletion would affect the value that causes maximum rounding. Therefore, we get

$$\Delta_{ID} BS_{L,U}^* \leq \max\{2^b + U, 2^{b'} + |L|\} + \max\{2^b, 2^{b'}\}.$$

For $U \geq |L|$, then, the overall sensitivity is $\leq 3U$, and for $U < |L|$, the overall sensitivity is $\leq 3 \cdot |L|$. The overall sensitivity, then, can be upper-bounded by $3 \cdot \max\{U, |L|\}$.

2. Let $u \simeq_{ID} u'$ be two datasets with (k, ℓ) -bit floating-point values of type T . Without loss of generality, let $\text{len}(u) < \text{len}(u')$.

By Lemma 6.38, where $U \in [2^c, 2^{c+1})$, and where $\min\{n \cdot U, 2^{k+c}\} \in [2^{k+b}, 2^{k+b+1})$, the calculation of P has a sensitivity of at most $2^b + U$. By generalizing Lemma 6.38 to work

with values in $[L, 0]$ instead, it follows by symmetry that, where $L \in (-2^{c'+1}, -2^{c'})$, and where $\min\{|n \cdot L|, 2^{k+c'}\} \in (-2^{k+b'+1}, -2^{k+b'})$, the calculation of N has a sensitivity of at most $2^{b'} + |L|$.

The last step of Method 6.28 consists of adding P and N using banker's rounding. We need to account for the rounding that can occur in this step.

If $n \cdot U \geq 2^{k+c}$, then – by the reasoning in (1) – the maximum rounding effect is 2^b when $M_P \geq |M_L|$. Likewise, if $|n \cdot L| \geq 2^{k+c'}$, then – by the reasoning in (1) – the maximum rounding effect is $2^{b'}$ when $M_P < |M_L|$.

We now consider the case where $n \cdot U < 2^{k+c}$, and we also consider the case where $|n \cdot L| \geq 2^{k+c'}$.

Let M_P be the maximum value of P , and let M_N be the minimum value (i.e., maximum magnitude) of N . We see, then, that $M_P \leq 2^{c+1} \cdot 2^k$, and that $|M_N| \leq | -2^{c'+1} \cdot 2^k |$.

Suppose that $M_P \geq |M_N|$. In the worst case, $ULP_{(k,\ell)}(M_P) = 2^{c+1}$, so the worst-case round-down that we can get from combining P and N (using banker's rounding) is $\frac{2^{c+1}}{2}$. Likewise, the worst-case round-up that we can get from combining P and N (using banker's rounding) is $\frac{2^{c+1}}{2}$. Therefore, the worst-case overall rounding is 2^{c+1} . Likewise, for the $M_P \leq |M_N|$ case, the worst-case rounding is $2^{c'+1}$.

We recall that the sensitivity for computing N is $\leq 2^{b'} + |L|$; and that the sensitivity for P is $\leq 2^b + U$. Therefore, the overall sensitivity is

$$\begin{aligned} \Delta_{Ham} BS_{L,U}^* & \\ & \leq 2^b + U + 2^{b'} + |L| + \max\{\min\{2^b, 2^{c+1}\}, \min\{2^{b'}, 2^{c'+1}\}\}. \end{aligned}$$

For $U \geq |L|$, then, the overall sensitivity is $\leq 5U$, and for $U < |L|$, the overall sensitivity is $\leq 5L$. The overall sensitivity, then, can be upper-bounded by $5 \cdot \max\{U, |L|\}$.

□

Having proven sensitivities with respect to ordered distance metrics, we now prove sensitivities with respect to unordered distance metrics.

Proof of Theorem 6.30. Recall that we are working with the bounded sum function as described in the theorem statement. Additionally, note that we prove (1) $\Delta_{Sym} BS_{L,U}^*$ and (2) $\Delta_{CO} BS_{L,U,n}^*$.

1. We first apply RP to the input dataset, and Theorem 6.19 tells us that RP transforms datasets that are neighboring with respect to d_{Sym} to a coupling of datasets that are neighboring with respect to d_{ID} . We can now work with datasets that are neighboring with respect to d_{ID} . Theorem 6.31 tells us that

$$\Delta_{ID} BS_{L,U}^* \leq \max\{2^b + U, 2^{b'} + |L|\} + \max\{2^b, 2^{b'}\},$$

so we also have

$$\Delta_{Sym} BS_{L,U}^* \leq \max\{2^b + U, 2^{b'} + |L|\} + \max\{2^b, 2^{b'}\},$$

completing the proof.

2. We first apply RP to the input dataset, and Corollary 6.20 tells us that RP transforms datasets that are neighboring with respect to d_{CO} to a coupling of datasets that are neighboring with respect to d_{Ham} . We can now work with datasets that are neighboring with respect to d_{Ham} . Theorem 6.31 tells us that

$$\Delta_{Ham}BS_{L,U}^* \leq 2^b + U + 2^{b'} + |L| + \max\{\min\{2^b, 2^{c+1}\}, \min\{2^{b'}, 2^{c'+1}\}\},$$

so we also have

$$\Delta_{CO}BS_{L,U}^* \leq 2^b + U + 2^{b'} + |L| + \max\{\min\{2^b, 2^{c+1}\}, \min\{2^{b'}, 2^{c'+1}\}\},$$

completing the proof. □

References

- [1] Martín Abadi et al. “Deep Learning with Differential Privacy”. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’16. Vienna, Austria: ACM, 2016, pp. 308–318. ISBN: 978-1-4503-4139-4. DOI: [10.1145/2976749.2978318](https://doi.org/10.1145/2976749.2978318). URL: <http://doi.acm.org/10.1145/2976749.2978318>.
- [2] John M. Abowd. “The U.S. Census Bureau Adopts Differential Privacy”. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD ’18. London, United Kingdom: ACM, 2018, pp. 2867–2867. ISBN: 978-1-4503-5552-0. DOI: [10.1145/3219819.3226070](https://doi.org/10.1145/3219819.3226070). URL: <https://digitalcommons.ilr.cornell.edu/ldi/49/>.
- [3] Victor Balcer and Salil Vadhan. “Differential Privacy on Finite Computers”. In: *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Ed. by Anna R. Karlin. Vol. 94. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, 43:1–43:21. ISBN: 978-3-95977-060-6. DOI: [10.4230/LIPIcs.ITCS.2018.43](https://doi.org/10.4230/LIPIcs.ITCS.2018.43). URL: <http://drops.dagstuhl.de/opus/volltexte/2018/8353>.
- [4] Aleix Bassolas et al. “Hierarchical organization of urban mobility and its connection with city livability”. In: *Nature communications* 10.1 (2019), pp. 1–10.
- [5] Shailesh Bavadekar et al. “Google COVID-19 Search Trends Symptoms Dataset: Anonymization Process Description (version 1.0)”. In: *arXiv preprint arXiv:2009.01265* (2020).
- [6] Shailesh Bavadekar et al. “Google COVID-19 Vaccination Search Insights: Anonymization Process Description”. In: *arXiv preprint arXiv:2107.01179* (2021).
- [7] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. “Practical privacy: the SuLQ framework”. In: *Proceedings of the 24th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. ACM. 2005, pp. 128–138.
- [8] Mark Bun and Thomas Steinke. “Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds”. In: *CoRR* abs/1605.02065 (2016). URL: <http://arxiv.org/abs/1605.02065>.

- [9] Clément L. Canonne, Gautam Kamath, and Thomas Steinke. “The Discrete Gaussian for Differential Privacy”. In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*. Ed. by Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin. 2020. URL: <https://proceedings.neurips.cc/paper/2020/hash/b53b3a3d6ab90ce0268229151c9bde11-Abstract.html>.
- [10] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. “Collecting telemetry data privately”. In: *Advances in Neural Information Processing Systems* 30 (2017).
- [11] Jinshuo Dong, Aaron Roth, and Weijie J Su. “Gaussian differential privacy”. In: *arXiv preprint arXiv:1905.02383* (2019).
- [12] Cynthia Dwork, Krishnamurthy Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. “Our data, ourselves: privacy via distributed noise generation”. In: *Advances in cryptology—EUROCRYPT 2006*. Vol. 4004. Lecture Notes in Comput. Sci. Springer, Berlin, 2006, pp. 486–503. DOI: [10.1007/11761679_29](https://doi.org/10.1007/11761679_29). URL: http://dx.doi.org/10.1007/11761679_29.
- [13] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. “Calibrating noise to sensitivity in private data analysis”. In: *Journal of Privacy and Confidentiality* 7.3 (2016). To appear. Preliminary version in *Proc. TCC ‘06*.
- [14] Cynthia Dwork and Aaron Roth. “The Algorithmic Foundations of Differential Privacy”. In: *Found. Trends Theor. Comput. Sci.* 9.3-4 (2014), pp. 211–407. DOI: [10.1561/0400000042](https://doi.org/10.1561/0400000042). URL: <https://doi.org/10.1561/0400000042>.
- [15] Cynthia Dwork and Guy N. Rothblum. “Concentrated Differential Privacy”. In: *CoRR* abs/1603.01887 (2016). arXiv: [1603.01887](https://arxiv.org/abs/1603.01887). URL: <http://arxiv.org/abs/1603.01887>.
- [16] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. “RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response”. In: *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’14. Scottsdale, Arizona, USA: ACM, 2014, pp. 1054–1067. ISBN: 978-1-4503-2957-6. DOI: [10.1145/2660267.2660348](https://doi.org/10.1145/2660267.2660348). URL: <http://doi.acm.org/10.1145/2660267.2660348>.
- [17] Andrew David Foote, Ashwin Machanavajjhala, and Kevin McKinney. “Releasing Earnings Distributions using Differential Privacy: Disclosure Avoidance System For Post-Secondary Employment Outcomes (PSEO)”. In: *Journal of Privacy and Confidentiality* 9.2 (2019).
- [18] Marco Gaboardi, Michael Hay, and Salil Vadhan. “A programming framework for OpenDP”. In: *Manuscript* (2020).
- [19] Ivan Gazeau, Dale Miller, and Catuscia Palamidessi. “Preserving differential privacy under finite-precision semantics”. In: *Theoretical Computer Science* 655 (2016), pp. 92–108.
- [20] A. Ghosh, T. Roughgarden, and M. Sundararajan. “Universally Utility-maximizing Privacy Mechanisms”. In: *SIAM Journal on Computing* 41.6 (2012), pp. 1673–1693. DOI: [10.1137/09076828X](https://doi.org/10.1137/09076828X). eprint: <https://doi.org/10.1137/09076828X>. URL: <https://doi.org/10.1137/09076828X>.
- [21] *Google’s differential privacy libraries*. GitHub repository. URL: <https://github.com/google/differential-privacy>.

- [22] Andreas Haeberlen, Benjamin C. Pierce, and Arjun Narayan. “Differential Privacy Under Fire”. In: *20th USENIX Security Symposium, San Francisco, CA, USA, August 8-12, 2011, Proceedings*. USENIX Association, 2011. URL: http://static.usenix.org/events/sec11/tech/full%5C_papers/Haeberlen.pdf.
- [23] A Herdağdelen, A Dow, S Bogdan, M Payman, and A Pompe. “Protecting privacy in Facebook mobility data during the COVID-19 response”. In: *Facebook Research* (2020).
- [24] Nicholas J. Higham. “The Accuracy of Floating Point Summation”. In: *SIAM J. Sci. Comput.* 14.4 (1993), pp. 783–799. DOI: [10.1137/0914050](https://doi.org/10.1137/0914050). URL: <https://doi.org/10.1137/0914050>.
- [25] Naoise Holohan and Stefano Braghin. “Secure Random Sampling in Differential Privacy”. In: *Computer Security - ESORICS 2021 - 26th European Symposium on Research in Computer Security, Darmstadt, Germany, October 4-8, 2021, Proceedings, Part II*. Ed. by Elisa Bertino, Haya Shulman, and Michael Waidner. Vol. 12973. Lecture Notes in Computer Science. Springer, 2021, pp. 523–542. DOI: [10.1007/978-3-030-88428-4_26](https://doi.org/10.1007/978-3-030-88428-4_26). URL: https://doi.org/10.1007/978-3-030-88428-4_26.
- [26] Naoise Holohan, Stefano Braghin, Pól Mac Aonghusa, and Killian Levacher. “Diffprivlib: The IBM Differential Privacy Library”. In: *CoRR* abs/1907.02444 (2019). arXiv: [1907.02444](https://arxiv.org/abs/1907.02444). URL: <http://arxiv.org/abs/1907.02444>.
- [27] Christina Ilvento. “Implementing the Exponential Mechanism with Base-2 Differential Privacy”. In: *CCS ’20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*. Ed. by Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna. ACM, 2020, pp. 717–742. DOI: [10.1145/3372297.3417269](https://doi.org/10.1145/3372297.3417269). URL: <https://doi.org/10.1145/3372297.3417269>.
- [28] Jiankai Jin, Eleanor McMurtry, Benjamin IP Rubinstein, and Olga Ohrimenko. “Are We There Yet? Timing and Floating-Point Attacks on Differential Privacy Systems”. In: *arXiv preprint arXiv:2112.05307* (2021).
- [29] Noah Johnson, Joseph P. Near, and Dawn Song. “Towards Practical Differential Privacy for SQL Queries”. In: *Proc. VLDB Endow.* 11.5 (Jan. 2018), pp. 526–539. ISSN: 2150-8097. DOI: [10.1145/3187009.3177733](https://doi.org/10.1145/3187009.3177733). URL: <https://doi.org/10.1145/3187009.3177733>.
- [30] Noah M. Johnson, Joseph P. Near, Joseph M. Hellerstein, and Dawn Song. “Chorus: a Programming Framework for Building Scalable Differential Privacy Mechanisms”. In: *IEEE European Symposium on Security and Privacy, EuroS&P 2020, Genoa, Italy, September 7-11, 2020*. IEEE, 2020, pp. 535–551. DOI: [10.1109/EuroSP48549.2020.00041](https://doi.org/10.1109/EuroSP48549.2020.00041). URL: <https://doi.org/10.1109/EuroSP48549.2020.00041>.
- [31] William Kahan. “Pracniques: further remarks on reducing truncation errors”. In: *Communications of the ACM* 8.1 (1965), p. 40.
- [32] Michael J. Kearns. “Efficient Noise-Tolerant Learning from Statistical Queries”. In: *Journal of the Association for Computing Machinery* 45.6 (1998), pp. 983–1006. DOI: [10.1145/293347.293351](https://doi.org/10.1145/293347.293351). URL: <http://doi.acm.org/10.1145/293347.293351>.
- [33] Andreas Klein. “A Generalized Kahan-Babuska-Summation-Algorithm”. In: *Computing* 76.3-4 (2006), pp. 279–293. DOI: [10.1007/s00607-005-0139-x](https://doi.org/10.1007/s00607-005-0139-x). URL: <https://doi.org/10.1007/s00607-005-0139-x>.

- [34] Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber. “Privacy: Theory Meets Practice on the Map”. In: *Proceedings of the 2008 IEEE 24th International Conference on Data Engineering*. ICDE ’08. Washington, DC, USA: IEEE Computer Society, 2008, pp. 277–286. ISBN: 978-1-4244-1836-7. DOI: [10.1109/ICDE.2008.4497436](https://doi.org/10.1109/ICDE.2008.4497436). URL: <https://doi.org/10.1109/ICDE.2008.4497436>.
- [35] Peter W. Markstein. “The New IEEE-754 Standard for Floating Point Arithmetic”. In: *Numerical Validation in Current Hardware Architectures, 6.1. - 11.1.2008*. Ed. by Annie A. M. Cuyt, Walter Krämer, Wolfram Luther, and Peter W. Markstein. Vol. 08021. Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2008. URL: <http://drops.dagstuhl.de/opus/volltexte/2008/1448>.
- [36] Frank McSherry. “Privacy Integrated Queries: An Extensible Platform for Privacy-preserving Data Analysis”. In: *Communications of the ACM* 53.9 (Sept. 2010), pp. 89–97. ISSN: 0001-0782. DOI: [10.1145/1810891.1810916](https://doi.org/10.1145/1810891.1810916). URL: <http://doi.acm.org/10.1145/1810891.1810916>.
- [37] Solomon Messing et al. *Facebook Privacy-Protected Full URLs Data Set*. Version DRAFT VERSION. 2020. DOI: [10.7910/DVN/TDOAPG](https://doi.org/10.7910/DVN/TDOAPG). URL: <https://doi.org/10.7910/DVN/TDOAPG>.
- [38] Ilya Mironov. “On Significance of the Least Significant Bits for Differential Privacy”. In: *Proceedings of the 2012 ACM Conference on Computer and Communications Security*. CCS ’12. Raleigh, North Carolina, USA: ACM, 2012, pp. 650–661. ISBN: 978-1-4503-1651-4. DOI: [10.1145/2382196.2382264](https://doi.org/10.1145/2382196.2382264). URL: <http://doi.acm.org/10.1145/2382196.2382264>.
- [39] Ilya Mironov. “Rényi Differential Privacy”. In: *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*. IEEE Computer Society, 2017, pp. 263–275. ISBN: 978-1-5386-3217-8. DOI: [10.1109/CSF.2017.11](https://doi.org/10.1109/CSF.2017.11). URL: <https://doi.org/10.1109/CSF.2017.11>.
- [40] Ilya Mironov, Kunal Talwar, and Li Zhang. “Rényi Differential Privacy of the Sampled Gaussian Mechanism”. In: *arXiv preprint arXiv:1908.10530* (2019).
- [41] *OpenDP*. GitHub repository. URL: <https://github.com/opensdp/opensdp>.
- [42] Mayana Pereira et al. “US Broadband Coverage Data Set: A Differentially Private Data Release”. In: *arXiv preprint arXiv:2103.14035* (2021).
- [43] Ryan Rogers et al. “A Members First Approach to Enabling LinkedIn’s Labor Market Insights at Scale”. In: *arXiv preprint arXiv:2010.13981* (2020).
- [44] Indrajit Roy, Srinath T. V. Setty, Ann Kilzer, Vitaly Shmatikov, and Emmett Witchel. “Airavat: Security and Privacy for MapReduce”. In: *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*. NSDI’10. San Jose, California: USENIX Association, 2010, pp. 20–20. URL: <http://dl.acm.org/citation.cfm?id=1855711.1855731>.
- [45] Jayshree Sarathy and Salil Vadhan. “Analyzing the Differentially Private Theil-Sen Estimator for Simple Linear Regression.” In: *Theory and Practice of Differential Privacy*. 2021.

- [46] Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and XiaoFeng Wang. “Privacy Loss in Apple’s Implementation of Differential Privacy on MacOS 10.12”. In: *CoRR* abs/1709.02753 (2017). arXiv: [1709.02753](https://arxiv.org/abs/1709.02753). URL: <http://arxiv.org/abs/1709.02753>.
- [47] Shibo Wang and Pankaj Kanwar. *BFloat16: The secret to high performance on Cloud TPUs*. en. Aug. 2019. URL: <https://cloud.google.com/blog/products/ai-machine-learning/bfloat16-the-secret-to-high-performance-on-cloud-tpus/>.
- [48] James H Wilkinson. “Error analysis of floating-point computation”. In: *Numerische Mathematik* 2.1 (1960), pp. 319–340.
- [49] Royce J. Wilson et al. “Differentially Private SQL with Bounded User Contribution”. In: *CoRR* abs/1909.01917 (2019). arXiv: [1909.01917](https://arxiv.org/abs/1909.01917). URL: <http://arxiv.org/abs/1909.01917>.
- [50] Ashkan Yousefpour et al. “Opacus: User-Friendly Differential Privacy Library in PyTorch”. In: *CoRR* abs/2109.12298 (2021). arXiv: [2109.12298](https://arxiv.org/abs/2109.12298). URL: <https://arxiv.org/abs/2109.12298>.
- [51] Dan Zhang et al. “EKTELO: A Framework for Defining Differentially-Private Computations”. In: *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*. Ed. by Gautam Das, Christopher M. Jermaine, and Philip A. Bernstein. ACM, 2018, pp. 115–130. DOI: [10.1145/3183713.3196921](https://doi.org/10.1145/3183713.3196921). URL: <https://doi.org/10.1145/3183713.3196921>.

A Missing proofs from Section 2

Lemma A.1 (Relating Metrics, Lemma 2.10). *These metrics are related as follows.*

1. For $u, v \in \text{Vec}(\mathcal{D})$,

$$d_{Sym}(u, v) = \min_{\pi \in S_{\text{len}(u)}} d_{ID}(\pi(u), v) \leq d_{ID}(u, v).$$

2. For $u, v \in \mathcal{D}^n$,

$$d_{CO}(u, v) = \min_{\pi \in S_n} d_{Ham}(\pi(u), v) \leq d_{Ham}(u, v).$$

3. For $u, v \in \mathcal{D}^n$,

$$d_{Sym}(u, v) = 2 \cdot d_{CO}(u, v).$$

4. For $u, v \in \mathcal{D}^n$,

$$d_{ID}(u, v) \leq 2 \cdot d_{Ham}(u, v).$$

Proof. We prove each part below.

1. Let $u, v \in \text{Vec}(\mathcal{D})$.

We first show that there exists a permutation $\pi \in S_{\text{len}(u)}$ such that $d_{ID}(\pi(u), v) \leq d_{Sym}(u, v)$.

We begin by constructing v' such that $d_{ID}(v, v') \leq d_{Sym}(u, v)$ and $h_{v'} = h_u$. We show that we can use $d_{Sym}(u, v)$ insertions and deletions on v to get v' such that $h_{v'} = h_u$. We create

v' from v as follows: (1) for all $z \in \mathcal{D}$ such that $h_v(z) > h_u(z)$, we delete $h_v(z) - h_u(z)$ copies of z from v ; (2) for all $z \in \mathcal{D}$ such that $h_v(z) < h_u(z)$, we insert $h_v(z) - h_u(z)$ copies of z to v . By Definition 2.7, then,

$$d_{ID}(v', v) \leq d_{Sym}(u, v). \quad (6)$$

The resulting vector v' also has the property that, for all $z \in \mathcal{D}$, $h_{v'}(z) = h_u(z)$, so $h_u = h_{v'}$. By Lemma 2.4, $h_u = h_{v'}$ implies that there exists a permutation $\pi \in S_{\text{len}(u)}$ such that $\pi(u) = v'$. Substituting this into Expression (6), we get

$$d_{ID}(\pi(u), v) \leq d_{Sym}(u, v). \quad (7)$$

We next show that, for all permutations $\pi \in S_{\text{len}(u)}$, we have $d_{Sym}(u, v) \leq d_{ID}(\pi(u), v)$, which means that $d_{Sym}(u, v) \leq \min_{\pi \in S_{\text{len}(u)}} d_{ID}(\pi(u), v)$. This is true because, for all permutations $\pi \in S_{\text{len}(u)}$, $d_{Sym}(u, v) = d_{Sym}(\pi(u), v) \leq d_{ID}(\pi(u), v)$, with the equality following from the fact that $h_{\pi(u)} = h_u$.

2. Let $u, v \in \mathcal{D}^n$.

We first show that there is a permutation $\pi \in S_n$ such that $d_{CO}(u, v) = d_{Ham}(\pi(u), v)$.

We begin by constructing v' such that $d_{Ham}(v, v') \leq d_{CO}(u, v)$ and $h_{v'} = h_u$. We create v' from v as follows: (1) for all $z \in \mathcal{D}$ such that $h_v(z) > h_u(z)$, replace $h_v(z) - h_u(z)$ copies of z with some value $\diamond \notin \mathcal{D}$; (2) for all $z \in \mathcal{D}$ such that $h_v(z) < h_u(z)$, replace $h_u(z) - h_v(z)$ \diamond values with z . By Definition 2.8, then,

$$d_{Ham}(v', v) \leq d_{CO}(u, v). \quad (8)$$

The resulting vector v' also has the property that, for all $z \in \mathcal{D}$, $h_{v'}(z) = h_u(z)$, so $h_u = h_{v'}$. By Lemma 2.4, $h_u = h_{v'}$ implies that there exists a permutation $\pi \in S_n$ such that $\pi(u) = v'$. Substituting this into Expression (8), we get

$$d_{Ham}(\pi(u), v) \leq d_{CO}(u, v). \quad (9)$$

We next show that, for all permutations $\pi \in S_n$, we have $d_{CO}(u, v) \leq d_{Ham}(\pi(u), v)$, which means that $d_{CO}(u, v) \leq \min_{\pi \in S_n} d_{Ham}(\pi(u), v)$. This is true because, for every permutation $\pi \in S_n$,

$$d_{CO}(u, v) = d_{CO}(\pi(u), v) \leq d_{Ham}(\pi(u), v),$$

with the equality following from the fact that $h_{\pi(u)} = h_u$.

3. Let $u, v \in \mathcal{D}^n$. Then,

$$\begin{aligned} d_{Sym}(u, v) &= \sum_{z \in \mathcal{D}} |h_u(z) - h_v(z)| \\ &= \sum_{\substack{z \in \mathcal{D} \text{ s.t.} \\ h_u(z) > h_v(z)}} (h_u(z) - h_v(z)) \\ &\quad + \sum_{\substack{z \in \mathcal{D} \text{ s.t.} \\ h_v(z) > h_u(z)}} (h_v(z) - h_u(z)) \\ &= d_{CO}(u, v) + d_{CO}(u, v) \\ &= 2 \cdot d_{CO}(u, v). \end{aligned}$$

4. Let $u, v \in \mathcal{D}^n$, and let $d_{Ham}(u, v) = c$. Let \mathcal{I} be the set of all indices i^* such that $u_{i^*} \neq v_{i^*}$. We note that, to make it so that $u_{i^*} = v_{i^*}$, we can perform 1 insertion and 1 deletion to change u_{i^*} to v_{i^*} . We note that the cardinality of \mathcal{I} is c . Therefore, a total of (at most) $2c$ insertions and deletions need to be performed to change u into u' such that $u' = v$. By the definition of d_{ID} , then, $d_{ID}(u, v) \leq 2c = 2 \cdot d_{Ham}(u, v)$.

□

Lemma A.2 (Lemma 2.14). *The dataset metrics d_{Sym}, d_{CO}, d_{Ham} , and d_{ID} all satisfy the path property.*

Proof. By Definitions 2.5, 2.6, 2.7, 2.8 and the fact that the image of the histogram function is \mathbb{N} it is clear that all of d_{Sym}, d_{CO}, d_{Ham} , and d_{ID} satisfy Condition 1 of the Path Property.

For Condition 2, we show constructively how to build the sequence u^0, u^1, \dots, u^d for each of the metrics. Let u, v be datasets such that $d(u, v) = d$ for some $d \in \mathbb{Z}$.

1. For d_{Sym} , let set S be the symmetric difference of u and v . Let \mathcal{I}_u be the set of indices k such that $u_k \in S$, and similarly for \mathcal{J}_v . Wlog, assume that $|\mathcal{I}_u| \leq |\mathcal{J}_v|$ (otherwise, swap u and v in what follows). To construct the $u^{(i)}$, we iteratively apply the following two steps. To go from $u^{(i)}$ to $u^{(i+2)}$, first pick one of the indices $k \in \mathcal{I}_u$ and delete it from $u^{(i)}$, which yields vector $u^{(i+1)}$. Next, insert into $u^{(i+1)}$ element v_j such that $j \in \mathcal{J}_v$, never repeating either index k or j in the process. After $2|\mathcal{I}_u|$ steps, to continue constructing the $u^{(i)}$ iteratively we insert elements v_j for all the remaining $j \in \mathcal{J}_v$ one at a time (without repeating any index j). Hence, after $|\mathcal{J}_v| - |\mathcal{I}_u|$ more steps, $u^{(i)} = v$.
2. For d_{CO} , to construct each $u^{(i)}$ we iteratively change an element of u which is different from v (when viewed as multisets). It follows directly from the definition of d_{CO} that this procedure requires $d_{CO}(u, v) = d$ steps.
3. For d_{Ham} , we apply the same procedure, except that u and v are now viewed as ordered vectors.
4. For d_{ID} , we apply the same procedure as in the case of d_{CO} , where we do one insert or one delete operation at a time. It then follows directly from the definition of d_{ID} that this procedure requires $d_{ID}(u, v) = d$ steps.

□

Lemma A.3 (Convert Sensitivities, Lemma 2.15). *We can convert between sensitivities in the following ways.*

1. For every function $f : \text{Vec}(\mathcal{D}) \rightarrow \mathbb{R}$, $\Delta_{ID}f \leq \Delta_{Sym}f$.
2. For every function $f : \mathcal{D}^n \rightarrow \mathbb{R}$, $\Delta_{Ham}f \leq \Delta_{CO}f$.
3. For every function $f : \mathcal{D}^n \rightarrow \mathbb{R}$, $\Delta_{CO}f \leq 2\Delta_{Sym}f$.
4. For every function $f : \mathcal{D}^n \rightarrow \mathbb{R}$, $\Delta_{Ham}f \leq 2\Delta_{ID}f$.

Proof. We use Lemma 2.10 to prove each part. Let each f below be defined as described in its respective part above.

1. By Lemma 2.10, we see that, for all $u, v \in \text{Vec}(\mathcal{D})$, we have $d_{Sym}(u, v) \leq d_{ID}(u, v)$. Therefore, for all $u \simeq_{ID} u'$, we have $u \simeq_{Sym} u'$. By the definition of sensitivity in Definition 2.11, this means that $\Delta_{ID}f = \max_{u \simeq_{ID} u'} |f(u) - f(u')| \leq \max_{u \simeq_{Sym} u'} |f(u) - f(u')| = \Delta_{Sym}f$.
2. By Lemma 2.10, we see that, for all $u, v \in \mathcal{D}^n$, we have $d_{CO}(u, v) \leq d_{Ham}(u, v)$. Therefore, for all $u \simeq_{Ham} u'$, we have $u \simeq_{CO} u'$. By the definition of sensitivity in Definition 2.11, this means that $\Delta_{Ham}f = \max_{u \simeq_{Ham} u'} |f(u) - f(u')| \leq \max_{u \simeq_{CO} u'} |f(u) - f(u')| = \Delta_{CO}f$.
3. We have:

$$\begin{aligned}
\Delta_{CO}f &= \max_{u \simeq_{CO} u'} |f(u) - f(u')| \\
&\leq \max_{u \simeq_{CO} u'} \Delta_{Sym}f \cdot d_{Sym}(u, u') \quad (\text{Theorem 2.13}) \\
&= \max_{u \simeq_{CO} u'} \Delta_{Sym}f \cdot 2 \cdot d_{CO}(u, u') \quad (\text{Lemma 2.10}) \\
&= 2\Delta_{Sym}f,
\end{aligned}$$

so $\Delta_{CO}f \leq 2\Delta_{Sym}f$.

4. We have:

$$\begin{aligned}
\Delta_{Ham}f &= \max_{u \simeq_{Ham} u'} |f(u) - f(u')| \\
&\leq \max_{u \simeq_{Ham} u'} \Delta_{ID}f \cdot d_{ID}(u, u') \quad (\text{Theorem 2.13}) \\
&\leq \max_{u \simeq_{Ham} u'} \Delta_{ID}f \cdot 2 \cdot d_{Ham}(u, u') \quad (\text{Lemma 2.10}) \\
&= 2\Delta_{ID}f,
\end{aligned}$$

so $\Delta_{Ham}f \leq 2\Delta_{ID}f$.

□

B Missing proofs from Section 4

Theorem B.1 (Idealized sensitivities of $BS_{L,U}$ and $BS_{L,U,n}$, Theorem 4.4). *The sensitivities of the bounded sum function are the following.*

1. (Unknown n .) $\Delta_{Sym}BS_{L,U} = \Delta_{ID}BS_{L,U} = \max\{|L|, U\}$.
2. (Known n .) $\Delta_{CO}BS_{L,U,n} = \Delta_{Ham}BS_{L,U,n} = U - L$.

Proof. Each part is proven below.

1. Let $u, u' \in \text{Vec}(\mathbb{R}_{[L,U]})$ be two datasets such that $u \simeq_{Sym} u'$. From the formal definition of a histogram in Definition 2.2, we observe that, for all $v \in \text{Vec}(\mathbb{R})$, $BS_{L,U}(v) = \sum_{i=1}^{\text{len}(v)} v_i = \sum_{z \in \mathbb{R}} h_v(z) \cdot z$, where the last sum is well defined because $h_v(z) \neq 0$ for only finitely many values of z . Because $u \simeq_{Sym} u'$, we know that there is at most one value z^* such that $|h_u(z^*) - h_{u'}(z^*)| = 1$, and that, for all $z \neq z^*$, we have $|h_u(z) - h_{u'}(z)| = 0$.

We can then write the following expressions.

$$\begin{aligned}
& |BS_{L,U}(u) - BS_{L,U}(u')| \\
&= \left| \sum_{z \in \mathbb{R}} h_u(z) \cdot z - \sum_{z \in \mathbb{R}} h_{u'}(z) \cdot z \right| \\
&= \left| \sum_{z \in \mathbb{R} \setminus z^*} h_u(z) \cdot z - \sum_{z \in \mathbb{R} \setminus z^*} h_{u'}(z) \cdot z \right| \\
&\quad + |h_u(z^*) \cdot z^* - h_{u'}(z^*) \cdot z^*| \\
&= 0 + |z^* \cdot (h_u(z^*) - h_{u'}(z^*))| \\
&\leq |z^*| \\
&\leq \max\{|L|, U\},
\end{aligned}$$

with the final inequality following from the fact that all values are clamped to the interval $[L, U]$, so the largest difference in sums arises when $z^* = \max\{|L|, U\}$.

By the definition of sensitivity, then, $\Delta_{Sym} BS_{L,U} \leq \max\{|L|, U\}$. By Theorem 2.15, we also have $\Delta_{ID} BS_{L,U} \leq \max\{|L|, U\}$.

For the lower bound, consider the datasets $u = [0]$, $u' = [0, \max\{|L|, U\}]$. We note that $u \simeq_{ID} u'$. We also note that $|BS_{L,U}(u) - BS_{L,U}(u')| = \max\{|L|, U\}$. This means, then, that $\Delta_{ID} BS_{L,U} \geq \max\{|L|, U\}$. By the contrapositive of Theorem 2.15, then, $\Delta_{Sym} BS_{L,U} \geq \max\{|L|, U\}$.

Combining these upper and lower bounds on the idealized sensitivity tells us, then, that

$$\Delta_{Sym} BS_{L,U} = \Delta_{ID} BS_{L,U} = \max\{|L|, U\}.$$

2. Let $u, u' \in \mathbb{R}^n$ be two datasets such that $u \simeq_{CO} u'$. By Lemma 2.10, there is a permutation $\pi \in S_n$ such that $\pi(u) \simeq_{Ham} u'$. This means there is at most one index i^* such that $\pi(u)_{i^*} \neq u'_{i^*}$. We can then write the following expressions.

$$\begin{aligned}
& |BS_{L,U,n}^*(u) - BS_{L,U,n}^*(u')| \\
&= |BS_{L,U,n}^*(\pi(u)) - BS_{L,U,n}^*(u')| \\
&= \left| \sum_{i=1}^n (\pi(u)_i) - \sum_{i=1}^n (u'_i) \right| \\
&= |(\pi(u)_{i^*}) - (u'_{i^*})| \\
&\leq U - L.
\end{aligned}$$

with the final inequality following from the fact that all values are clamped to the interval $[L, U]$, so the largest difference in sums arises when, without loss of generality, $\pi(u)_{i^*} = U$ and $u'_{i^*} = L$.

By the definition of sensitivity, then, $\Delta_{CO} BS_{L,U,n} \leq U - L$. By Theorem 2.15, we have $\Delta_{Ham} BS_{L,U,n} \leq U - L$.

For the lower bound, consider the datasets $u = [L]$ and $u' = [U]$. Then, $u \simeq_{Ham} u'$ and $|BS_{L,U,n}(u) - BS_{L,U,n}(u')| = U - L$. By the contrapositive of Theorem 2.15, it follows that $\Delta_{CO} BS_{L,U,n} \geq U - L$.

Combining these upper and lower bounds on the idealized sensitivity, we then conclude that

$$\Delta_{Ham} BS_{L,U,n} = \Delta_{CO} BS_{L,U,n} = U - L.$$

□

C Missing proofs form Section 5

Theorem C.1 (Theorem 5.17). *Let $BS_{L,U}^* : \text{Vec}(T_{[L,U]}) \rightarrow T$ be iterative bounded sum on the type T of (k, ℓ) -bit floats. Let $j, m \in \mathbb{Z}$ satisfy $0 < j \leq k$ and $-(2^{\ell-1} - 2) \leq m \leq 2^{\ell-1} - 2 - j - k$. Then for $L = 0$, $U = 2^{k+m}$, and $n = j \cdot 2^k + 2$, there are datasets $u, v \in \text{Vec}(T)$ where $\text{len}(u) = n$ and $\text{len}(v) = n - 1$ such that $d_{ID}(u, v) = 1$ and*

$$|BS_{L,U}^*(u) - BS_{L,U}^*(v)| \geq 2^{k+j+m}.$$

In particular,

$$\frac{\Delta_{ID} BS_{L,U}^*}{\Delta_{ID} BS_{L,U}} \geq 2^j.$$

Proof. Let $f(x)$ be defined as $x \mapsto (2^x + 2^{x-k}) \cdot 2^m$, and consider the datasets

$$\begin{aligned} u = & [2^{k+m}, 2^{k+m}, f(0)_1, \dots, f(0)_{2^k}, \\ & f(1)_{2^{k+1}}, \dots, f(1)_{2 \cdot 2^k}, \dots, \\ & f(j-1)_{(j-1) \cdot 2^{k+1}}, \dots, f(j-1)_{j \cdot 2^k}], \end{aligned}$$

and v such that $v_i = u_{i+1}$ for all i (so the first term of u is not in v). Note that $d_{ID}(u, v) = 1$. Additionally, we consider $m = 0$ until the conclusion of the proof.

We first evaluate $BS_{L,U}^*(u)$. We begin by seeing that the sum of the first two terms can be computed exactly: $2^k \oplus 2^k = 2^{k+1}$.

We next use Lemma 2.27 to show that the sum of the first three terms cannot be computed exactly: $2^{k+1} + f(0) = 2^{k+1} + 1 + 2^{-k} \in [2^{k+1}, 2^{k+2})$, so $ULP_{(k,\ell)}(2^{k+1} + f(0)) = 2$. However, $2^{k+1} + f(0)$ is not a multiple of 2, so by Lemma 2.27, $2^{k+1} + f(0)$ cannot be represented exactly as a (k, ℓ) -bit float. This means that we must use banker's rounding (described in Definition 2.33) to determine the value of $2^{k+1} \oplus f(0)$. The value $2^{k+1} + f(0)$ is between the adjacent (k, ℓ) -bit floats $2^{k+1} + 0$ and $2^{k+1} + 2$. It is closer to $2^{k+1} + 2$, so we have $2^{k+1} \oplus f(0) = 2^{k+1} + 2$.

Having determined the result of adding $f(0)_1$ to the intermediate sum 2^{k+1} , we now determine the result of adding the remaining $f(0)$ values. By reasoning similar to the logic used above, we find that for every addition of $f(0)_i$, adding $f(0) = 1 + 2^{-k}$ has the effect of adding 2. Therefore, $2^{k+1} \oplus f(0)_1 \oplus \dots \oplus f(0)_{2^k} = 2^{k+2}$.

We use Lemma 2.27 to show that none of the remaining additions are exact, and we show the rules of banker's rounding affect the computation of $BS_{L,U}^*(u)$. By Lemma 2.27, because $ULP_{(k,\ell)}(2^{k+2}) = 4$ and $2^{k+2} + f(1)$ is not a multiple of 4, we must use banker's rounding (described in Definition 2.33) to determine the value of $2^{k+2} \oplus f(1)$. Reasoning similar to the logic above shows that adding each of the $f(1) = 2 + 2^{-k+1}$ terms has the effect of adding 4, so $2^{k+2} \oplus f(1)_{k+1} \oplus \dots \oplus f(1)_{2 \cdot 2^k} = 2^{k+3}$. Such logic follows for all of the floating point additions, resulting in an overall sum of $BS_{L,U}^*(u) = 2^{k+j+1}$.

We now evaluate $BS_{L,U}^*(v)$. The general structure of the proof is similar, but the resulting sums are different due to different effects of banker's rounding.

We next use Lemma 2.27 to show that the sum of the first three terms cannot be computed exactly: $2^k + f(0) = 2^k + 1 + 2^{-k} \in [2^k, 2^{k+1})$, so the $ULP_{(k,\ell)}(2^k + f(0)) = 2^0$. However, $2^{k+1} + f(0)$ is not a multiple of 1, so by Lemma 2.27, $2^k + f(0)$ cannot be represented exactly as a (k, ℓ) -bit float. This means that we must use banker's rounding (described in Definition 2.33) to determine the value of $2^k \oplus f(0)$. The value $2^k + f(0)$ is between the adjacent (k, ℓ) -bit floats $2^k + 0$ and $2^k + 1$. It is closer to $2^k + 1$, so we have $2^k \oplus f(0) = 2^k + 1$.

Having determined the result of adding $f(0)_1$ to the intermediate sum 2^k , we now determine the result of adding the remaining $f(0)$ values. By reasoning similar to the logic used above, we find that for every addition of $f(0)_i$, adding $f(0) = 1 + 2^{-k}$ has the effect of adding 1 (note in the calculation of $BS_{L,U}^*(u)$ above that adding $f(0) = 1 + 2^{-k}$ has the effect of adding 2 – this is a critical difference). Therefore, $2^k \oplus f(0)_1 \oplus \dots \oplus f(0)_{2^k} = 2^{k+1}$.

We now use Lemma 2.27 to show that none of the remaining additions are exact, and we show the rules of banker's rounding affect the computation of $BS_{L,U}^*(v)$. By Lemma 2.27, because $ULP_{(k,\ell)}(2^{k+1}) = 2$ and $2^{k+1} + f(1)$ is not a multiple of 2, we must use banker's rounding (described in Definition 2.33) to determine the value of $2^{k+1} \oplus f(1)$. Reasoning similar to the logic above shows that adding each of the $f(1) = 2 + 2^{-k+1}$ terms has the effect of adding 2, so $2^{k+1} \oplus f(1)_{k+1} \oplus \dots \oplus f(1)_{2 \cdot 2^k} = 2^{k+2}$. Such logic follows for all of the floating point additions, resulting in an overall sum of $BS_{L,U}^*(u) = 2^{k+j}$.

We see that $|BS^*(u) - BS^*(v)| = 2^{k+j}$.

From the bounds $-(2^{\ell-1}-2) \leq m \leq 2^{\ell-1}-2-j-k$, we see that m only affects the exponent of floating point values in the proof and does not affect whether values are representable as (k, ℓ) -bit normal floats or the direction of rounding. All values in the proof above, then, can be multiplied by 2^m , so, for all m such that $-(2^{\ell-1}-2) \leq m \leq 2^{\ell-1}-2-j-k$, $\Delta_{ID}BS_{L,U}^* \geq 2^{k+j+m}$.

We observe that $U = 2^{k+m}$ and $\Delta_{ID}BS_{L,U} = \max\{|L|, U\} = U$. Therefore,

$$\frac{\Delta_{ID}BS_{L,U}^*}{\Delta_{ID}BS_{L,U}} \geq 2^j.$$

□

Theorem C.2 (Theorem 5.19). *Let $BS_{L,U}^* : \text{Vec}(T_{[L,U]}) \rightarrow T$ be iterative bounded sum on the type T of (k, ℓ) -bit floats. Additionally, let $a, j \in \mathbb{Z}$ satisfy $2 \leq j < k$, $-(2^{\ell-1}-2) \leq a$, $-(2^{\ell-1}-2) + 1 + k \leq j + a \leq 2^{\ell-1}-1$ (these conditions ensure that values are representable as (k, ℓ) -bit floats), $n = 2^j$, $m = n/2$. Then, for*

$$U = 2^a, L = -\left(\frac{U \cdot m}{2^k}\right) \cdot \left(\frac{1}{2} - \frac{1}{2^k}\right),$$

there are datasets $u, v \in \text{Vec}(T_{[L,U]})$ with $d_{ID}(u, v) = 1$ such that

$$|BS_{L,U}^*(u) - BS_{L,U}^*(v)| \geq \frac{n^2}{2^{k+3}} \cdot U + U.$$

In particular,

$$\frac{\Delta_{ID}BS_{L,U}^*}{\Delta_{ID}BS_{L,U}} = \frac{n^2}{2^{k+3}} + 1.$$

Proof. Let

$$x = \left(\frac{U \cdot m}{2^k} \right) \cdot \left(\frac{1}{2} + \frac{1}{2^k} \right),$$

and consider the datasets

$$u = [U_1, \dots, U_{m-1}, U_m, x_1, L_2, x_3, L_4, \dots, x_{m-1}, L_m]$$

and

$$v = [U_1, \dots, U_{m-1}, x_1, L_2, x_3, L_4, \dots, x_{m-1}, L_m].$$

where, for all i , $L_i = L$ and $U_i = U$. Note that u and v are equivalent with the exception that v contains $m - 1$ copies of U rather than m copies of U , so $d_{ID}(u, v) = 1$.

We first show that $BS_{L,U}^*[U_1, \dots, U_{m-1}] = (m - 1) \cdot U$ and $BS_{L,U}^*[U_1, \dots, U_m] = m \cdot U$. We take the approach of showing that all intermediate sums computed in the calculation of $BS_{L,U}^*[U_1, \dots, U_m]$ can be represented exactly as (k, ℓ) -bit floats, which is done by first showing that all of these intermediate sums are multiples of their $ULP_{(k,\ell)}$ and then applying Lemma 2.27.

We observe that $\{U, \dots, m \cdot U\}$ is the set of intermediate sums that result from calculating $BS_{L,U}[U_1, \dots, U_m]$. Let \mathcal{U} denote this set $\{U, \dots, m \cdot U\}$. We note that $m \cdot U = 2^{a+j-1}$, so $ULP_{(k,\ell)}(m \cdot U) = 2^{a+j-1-k}$. For all $x \in \mathcal{U}$, then,

$$ULP_{(k,\ell)}(x) \leq ULP_{(k,\ell)}(m \cdot U) = 2^{a+j-1-k}.$$

Note that $j < k$, so $2^{a+j-1-k}$ necessarily divides $U = 2^a$. All such x are multiples of $U = 2^a$, so they are necessarily multiples of $ULP_{(k,\ell)}(x)$. By Lemma 2.27, all $x \in \mathcal{U}$ can be represented exactly as (k, ℓ) -bit floats, so by Lemma 2.38, $BS_{L,U}^*[U_1, \dots, U_{m-1}] = (m - 1) \cdot U$ and $BS_{L,U}^*[U_1, \dots, U_m] = m \cdot U$.

We now evaluate $BS_{L,U,n}^*(u)$. We use Lemma 2.27 to show that $BS_{L,U,n}(u) = m \cdot U + x$ is not exactly representable as a (k, ℓ) -bit float, and we show how to calculate the result $BS_{L,U,n}^*(u)$.

As shown above, the sum of the first m terms is

$$BS_{L,U}^*[U_1, \dots, U_m] = m \cdot U.$$

Thus, $BS_{L,U}^*[U_1, \dots, U_m, x] = m \cdot U \oplus x$. We note that

$$m \cdot U + x \in [m \cdot U, m \cdot U + ULP_{(k,\ell)}(m \cdot U)].$$

Because $m \cdot U + x$ is closer to $m \cdot U + ULP_{(k,\ell)}(m \cdot U)$, $m \cdot U \oplus x = m \cdot U + ULP_{(k,\ell)}(m \cdot U)$. We now consider the addition of L , $m \cdot U + ULP_{(k,\ell)}(m \cdot U) \oplus L$. We see that

$$m \cdot U + ULP_{(k,\ell)}(m \cdot U) + L \in [m \cdot U, m \cdot U + ULP_{(k,\ell)}(m \cdot U)],$$

and $m \cdot U + ULP_{(k,\ell)}(m \cdot U) + L$ is closer to $m \cdot U + ULP_{(k,\ell)}(m \cdot U)$, so $m \cdot U + ULP_{(k,\ell)}(m \cdot U) \oplus L = m \cdot U + ULP_{(k,\ell)}(m \cdot U) \oplus L$. Similar logic applies for all x, L , in which adding L has the effect of adding 0 and adding x has the effect of adding $ULP_{(k,\ell)}(m \cdot U)$ for a total sum of $BS_{L,U}^*(u) = m \cdot U + \frac{m^2 \cdot U}{2^{k+1}}$.

We now evaluate $BS_{L,U,n}^*(v)$. As shown above, the sum of the first $m - 1$ terms corresponds to $BS_{L,U}^*[U_1, \dots, U_{m-1}] = (m - 1) \cdot U$. Thus, $BS_{L,U}^*[U_1, \dots, U_{m-1}, x] = (m - 1) \cdot U \oplus x$. We note that

$$(m - 1) \cdot U + x \in [(m - 1) \cdot U, (m - 1) \cdot U + ULP_{(k,\ell)}((m - 1) \cdot U)].$$

Because $(m-1) \cdot U + x$ is closer to $(m-1) \cdot U + ULP_{(k,\ell)}((m-1) \cdot U)$, $(m-1) \cdot U \oplus x = (m-1) \cdot U + ULP_{(k,\ell)}((m-1) \cdot U)$. We now consider the addition of L , $(m-1) \cdot U + ULP_{(k,\ell)}((m-1) \cdot U) \oplus L$. We see that

$$\begin{aligned} & (m-1) \cdot U + ULP_{(k,\ell)}((m-1) \cdot U) + L \\ & \in [(m-1) \cdot U, (m-1) \cdot U + ULP_{(k,\ell)}((m-1) \cdot U)], \end{aligned}$$

and $(m-1) \cdot U + ULP_{(k,\ell)}((m-1) \cdot U) + L$ is closer to $(m-1) \cdot U$, so $(m-1) \cdot U + ULP_{(k,\ell)}((m-1) \cdot U) \oplus L = (m-1) \cdot U$. Similar logic applies for all x, L , in which adding x and L in succession has the effect of adding 0 for a total sum of $BS_{L,U}^*(v) = (m-1) \cdot U$.

Therefore,

$$|BS_{L,U}^*(u) - BS_{L,U}^*(v)| = \frac{m^2 \cdot U}{2^{k+1}} + U = \frac{n^2 \cdot U}{2^{k+3}} + U,$$

so, because $d_{ID}(u, v) = 1$, $d_{ID}BS_{L,U}^* = \frac{m^2 \cdot U}{2^{k+1}} + U$. We observe that $\Delta_{ID}BS_{L,U} = \max\{|L|, U\} = U$, so

$$\frac{\Delta_{ID}BS_{L,U}^*}{\Delta_{ID}BS_{L,U}} = \frac{n^2}{2^{k+3}} + 1.$$

□