



QPET: A Versatile and Portable Quantity-of-Interest-Preservation Framework for Error-Bounded Lossy Compression

Jinyang Liu*
University of Houston
Houston, TX, USA
jliu217@central.uh.edu

Pu Jiao*
University of Kentucky
Lexington, KY, USA
pujiao@uky.edu

Kai Zhao
Florida State University
Tallahassee, FL, USA
kzhao@cs.fsu.edu

Xin Liang†
University of Kentucky
Lexington, KY, USA
xliang@uky.edu

Sheng Di
Argonne National Laboratory
Lemont, IL, USA
sdi1@anl.gov

Franck Cappello
Argonne National Laboratory
Lemont, IL, USA
cappello@mcs.anl.gov

ABSTRACT

Error-bounded lossy compression has been widely adopted in many scientific domains because it can address the challenges in storing, transferring, and analyzing unprecedented amounts of scientific data. However, general error-bounded lossy compressors may fail to meet additional quality requirements for downstream analysis, a.k.a. Quantities of Interest (QoIs). This may lead to uncertainties and even misinterpretations in scientific discoveries, significantly limiting the use of lossy compression in practice. In this paper, we propose QPET, a novel, versatile, and portable framework for QoI-preserving error-bounded lossy compression, which overcomes the challenges of modeling diverse QoIs by leveraging numerical strategies. QPET features (1) high portability to multiple existing lossy compressors, (2) versatile preservation to most differentiable univariate and multivariate QoIs, and (3) significant compression improvements in QoI-preservation tasks. Experiments with six real-world datasets demonstrate that integrating QPET into state-of-the-art error-bounded lossy compressors can gain 2x to 10x compression speedups of existing QoI-preserving error-bounded lossy compression solutions, up to 1000% compression ratio improvements to general-purpose compressors, and up to 133% compression ratio improvements to existing QoI-integrated scientific compressors.

PVLDB Reference Format:

Jinyang Liu, Pu Jiao, Kai Zhao, Xin Liang, Sheng Di, and Franck Cappello. QPET: A Versatile and Portable Quantity-of-Interest-Preservation Framework for Error-Bounded Lossy Compression. PVLDB, 18(8): 2440 - 2453, 2025.
doi:10.14778/3742728.3742739

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/JLiu-1/QPET-Artifact>.

*Both authors contributed equally to this research.

†Corresponding author: Xin Liang, Department of Computer Science, University of Kentucky, Lexington, KY 40506.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 18, No. 8 ISSN 2150-8097.
doi:10.14778/3742728.3742739

1 INTRODUCTION

Today's high-performance computing facilities and high-resolution instruments produce vast amounts of scientific data that overwhelm the data storage and transmission systems. According to recent studies, climate simulations generate hundreds of TB of data every 16 seconds [14], and fusion applications may produce 1 EB of data in a single run with current- and next-generation exascale computing systems [4]. In the molecular dynamics (MD) domain, the EXAALT project produces trajectory data with a trillion time steps [3], which require hundreds of terabytes of disk space in a scientific data management system (e.g., HDF5 [1]). Those examples pose grand challenges in the underlying data management tasks, including I/O and data transfer, which necessitates effective data compression.

Data compression has been widely adopted in database and data management systems for reducing data volumes [19, 38, 41, 52] and accelerating queries [8, 10, 56, 57]. Nonetheless, challenges arise when applying existing compression techniques in the context of scientific data. On the one hand, while lossless compression techniques [2, 7, 11, 33, 38] can recover data to the exact precision, they suffer from limited compression ratios in floating-point scientific data. On the other hand, lossy compressors for natural media data and time-series databases, such as JPEGs [44, 51], H.26x series [9, 41, 52], and SummaryStore [5] cannot provide a quantifiable bound on the decompression errors, leading to uncertainties in the downstream data visualization and analytics. To this end, error-bounded lossy compression has been proposed as a viable way to reduce the data size while providing guaranteed error control. For instance, ModelarDB [19, 21] is a typical example of time-series databases that support error-bounded lossy compression.

The same trend is observed in scientific data management systems. Over the past years, several error-bounded lossy compressors [25, 29, 31, 32, 36, 45, 46] have been developed and adopted by multiple scientific domains, as they can effectively accelerate data management in scientific applications. These compressors feature strict error control based on user requirements, so domain scientists can control the impact of lossy compression according to specific needs. However, most compressors only provide error bounds on the raw data, even though the scientists care more about the errors in their derived quantities or downstream analyses. Such derived information, also known as Quantities of Interest (QoIs), is of utmost importance to the fidelity and integrity of scientific discoveries.

Typical QoIs include symbolic functions [20], derivatives and gradients [40], integrals [15], and topological features [27, 54, 55]. Failure to quantify the distortions in QoIs may lead to misinterpretation of the data and even falsified discoveries, significantly limiting the use of error-bounded lossy compressors in practice.

A few QoI-preserving compressors have already been proposed and applied in scientific applications to bridge the gap. For instance, MGARD [6] has enabled the support for bounded linear QoIs, which is further integrated into fusion applications to preserve critical derived quantities [16]; variations of SZ have been proposed to tackle the preservation of symbolic QoIs [20], contour tree [55], and critical points in vector fields [27, 54]. Despite those efforts, several significant challenges remain unaddressed in the context of QoI reservation. (1) Existing QoI-preserving lossy compressors are finely crafted for particular QoIs, which lose generalization and extensibility to more diverse QoI formats. (2) Most existing methods rely on custom point-wise error bounds to enable the preservation of QoIs, but effectively determining or storing these error bounds is non-trivial. (3) Existing approaches usually tightly couple the QoI-preserving mechanism and compression pipeline, which fail to leverage newly developed compression methods. To address those challenges, we propose QPET, a **novel, versatile, and portable framework for efficiently preserving symbolic QoIs**. In particular, we notice that various symbolic QoIs can be modeled in mathematical formats composed of differentiable functions. Moreover, the numerical approximations of the QoIs with a simplified format (e.g., Taylor expansion with second-order derivatives) can provide a unified, efficient, and well-performing algorithmic routine to determine the point-wise error bounds. Thus, we develop a highly optimized QoI-preserving framework that can easily adapt to diverse QoIs and scientific lossy compression pipelines. In summary, our contributions are three-fold:

- We propose QPET, a QoI-oriented Point-wise Error-bound auto-tuning framework to enable the preservation of diverse QoIs in scientific applications. In particular, QPET leverages numerical and probabilistic methods to derive sufficient error bounds on each data point, which easily adapts to most differentiable QoIs.
- QPET effectively determines the best-fit point-wise error bound for each data value in the QoI-preserving compression. It also jointly optimizes the error-bound storage overhead and the reduction in raw data by a dynamic global error bound selection.
- We integrate QPET into 3 state-of-the-art error-bounded lossy compressors. Experimental results demonstrate that QPET yields significant compression ratio (CR) and throughput gains in QoI-preserving compression tasks. Under the same QoI error threshold, QPET-integrated compressors achieve up to 1000% CR improvements over the general-purpose compressors and up to 133% CR improvements over existing QoI-preserving lossy compressors. It also achieves 2x to 10x compression speedups over existing QoI-preserving scientific lossy compression solutions.

We organize the rest of this paper as follows: Section 2 discusses the related work. In Section 3, we formulate and clarify our research motivation and target. In Section 4, we propose the high-level design of QPET. Section 5 demonstrates the detailed designs and algorithms in QPET. In Section 6, we present the evaluations of QPET. Section 7 concludes this work and discusses future plans.

2 RELATED WORK

Traditional compression techniques for database systems:

Various data compressors have been proposed for large databases in diverse data domains/formats. Generic lossless compressors such as GZIP [12], ZSTD [11], and LZ4 [37] are widely used to reduce storage requirements [19, 38, 41, 52] and accelerate queries [8, 10, 56, 57]. Several lossless compressors have also been proposed to deal specifically with the increasing amount of floating-point data. A typical example is ndZip [22, 23], which leverages Lorenzo prediction [18] and residual encoding to compress floating-point data with high efficiency. To handle different types of data, Gorilla [38] and Chimp [26] have been designed for time-series data, and Buff [34] has been tailored to process low-precision data. Despite their ability to recover the exact data, lossless compressors exhibit low compression ratios, severely limiting their use in practice.

To address the limited compression ratios in lossless compressors, lossy compression has been proposed as an alternative way to reduce data in database systems. Regarding lossy database compression methods, ModelarDB [19, 21] and summaryStore [5] are designed for time-series data, which leverage data-processing techniques including PMC-mean [24] and the linear Swing model [13].

Emerging lossy compression techniques for scientific data:

To process the high volumes of scientific data while preserving the data quality for accurate post hoc analytics, scientific data compressors require not only decent compression ratios but also strict error control, where error-bounded lossy compression is becoming a promising option. It allows users to specify an error bound as a parameter and ensures the element-wise difference between the original data and decompressed data is less than the error bound.

SZ3 [31, 58] is one of the most widely used error-bounded lossy compressors with a prediction-based design. It utilizes dynamic spline interpolations to approximate the data, followed by linear-scaling quantization [42] and lossless encoding to reduce data size with guaranteed error control. Later, QoZ/HPEZ [35, 36] is proposed to significantly improve the quality of SZ3 via more complex interpolation schemes and quality-oriented auto-tuning mechanisms. Another category of error-bounded lossy compressors relies on domain transform for decorrelation. For instance, ZFP [32] is a fast compressor using near-orthogonal block transform, and SPERR [25] utilizes more complex wavelet transforms to obtain better compression quality at the cost of lower compression throughput. While error-bounded lossy compressors have guaranteed data error control, they cannot provide a quantifiable error bound on the downstream QoIs, as the QoIs could be highly diverse and non-linear.

QoI-preserving lossy compression methods: Several solutions have been proposed to enable error control for various QoIs. MGARD [6] is one of the first compressors to enable QoI preservation, but it can only provide guaranteed error control for bounded-linear QoIs. In [30], cpSZ has been proposed to preserve the locations and types of critical points in piece-wise linear vector fields. In particular, it carefully derived sufficient point-wise error bounds based on the target QoI and leveraged it to guide the compression procedure for effective QoI preservation. This idea is further extended to cover the preservation of critical points in multilinear vector fields [27], critical points extracted by sign-of-determinant predicates [54], contour trees [55], and symbolic QoIs [20].

Most existing QoI-preserving lossy compressors suffer from low generalizability and adaptability. In particular, they feature specific error control mechanisms toward the target QoI, which is hard to generalize to new sets of QoIs. In addition, most existing works are tightly coupled to a specific compression pipeline, making it difficult to integrate with newly developed compression methods. In this work, we propose a portable module that can easily adapt to diverse QoIs and compression methods. This will significantly improve the efficiency and usability of QoI-preserving lossy compression.

3 BACKGROUND

In this section, we introduce the background for the proposed work. In particular, we define the target QoIs to preserve, followed by a formulation of the QoI-preserving lossy compression problem.

3.1 Quantities of Interest in Scientific Data

As previously discussed, scientific applications usually require the preservation of QoIs to ensure the integrity of scientific discoveries if lossy compression is used. QoIs represent any derived information that is computed from the raw data, including physical properties [16, 53], derivatives [40], and topological features [55].

In this work, we target the preservation of symbolic QoIs, which can be formulated as mathematical functions of the input data. This covers various QoIs in scientific domains, including kinetic energy in cosmology, density in fusion energy science, and vector magnitude in climatology. We list the categorizations of target QoIs that the proposed framework can preserve as follows:

Univariate QoIs: Any second-order differentiable QoI functions, particularly the elementary functions, which are the combination of the following basic functions:

- Polynomial functions such as x^2 , x^3 , and $x^2 + x + 1$;
- Exponential functions such as 2^x and e^x ;
- Logarithm functions such as $\log_2 x$ and $\ln x$;
- Generalize power functions such as $(x + c)^{-1}$ (i.e., $\frac{1}{x+c}$);

Composite elementary functions such as Sigmoid $\sigma(x) = \frac{1}{1+e^{-x}}$ and Hyperbolic Tangent $\tanh x = \frac{e^x - e^{-x}}{e^x + e^{-x}}$ can also be supported.

Multivariate QoIs: Any differentiable multivariate functions in the format of $F(x_1, x_2, \dots, x_n)$, in which $\{x_1, x_2, \dots, x_n\}$ are a set of data points, are also supported by QPET. The point set can be a data region or a vector data entry. Typical examples include:

- Regional weighted sum of univariate QoI ($\frac{1}{n} \sum x^2$);
- Velocity scalar of 3-D velocity vector ($\sqrt{v_x^2 + v_y^2 + v_z^2}$);
- Kinetic energy of 3-D velocity vector ($\frac{1}{2}m(v_x^2 + v_y^2 + v_z^2)$);

While several existing works managed to preserve certain subsets of the aforementioned QoIs [6, 20], they cannot be generalized to other QoIs easily because they heavily rely on the analytical form of the QoIs to make analysis.

3.2 QoI-preserving Lossy Compression

We mathematically define QoI-preserving lossy compression as providing guaranteed error control in both the downstream QoIs and the raw data, which is similar to existing works [6, 20, 27].

Given an input data $X = \{x_i\}$, and a pair of compressor and decompressor (Cmp, Dec) that process X to compressed data $C =$

Cmp(X) and decompressed data $D = \{d_i\} = \text{Dec}(C)$, for a QoI function Q (which maps X to another set of values $Q(X)$), a data error bound ϵ , and a QoI error threshold τ , in QoI-preserving error-bounded lossy compression, we require the errors of data and QoI are both bounded, i.e., $\|X - D\|_\infty \leq \epsilon$ and $\|Q(X) - Q(D)\|_\infty \leq \tau$.

Our research would like to optimize Cmp and Dec to maximize the compression ratio. Because most QoIs are non-linear, setting a uniform error bound on each data point will result in divergent QoI errors. To address this issue, existing approaches [20, 27, 30] apply different point-wise error bounds on each data point to achieve effective QoI preservation. Nonetheless, these methods focused more on deriving the error bounds and overlooked their storage overhead, leading to suboptimal overall compression ratios. To optimize this strategy, selecting the best-fit point-wise error bound (composing a set of $\{\epsilon_i\}$) for each data point such that $|x_i - x'_i| \leq \epsilon_i$ to maximize compression ratio, we need to jointly optimize the compressed data size ($|\text{Cmp}_{\{\epsilon_i\}}(X)|$) and the compressed error-bound size of $\{\epsilon_i\}$ themselves ($|\text{Cmp}(\{\epsilon_i\})|$). This optimization problem is shown in Eq. 1 and will be specified in Section 5.

$$\begin{aligned} \{\epsilon_i\} = \arg \max_{\{\epsilon_i\}} & \frac{|X|}{|\text{Cmp}_{\{\epsilon_i\}}(X)| + |\text{Cmp}(\{\epsilon_i\})|} \\ \text{s.t. } & \|X - D\|_\infty \leq \epsilon \text{ and } \|Q(X) - Q(D)\|_\infty \leq \tau \end{aligned} \quad (1)$$

4 QPET DESIGN OVERVIEW

In this section, we provide an overview of the proposed QPET (QoI-oriented Point-wise Error-bound Tuning) framework, describing its modules, underlying algorithm, and its integration into several existing error-bounded lossy compression pipelines.

4.1 QPET composition and algorithm

In Figure 1, we demonstrate the composition of the QPET framework that provides QoI-preservation functionality, presenting how it integrates and interacts with general modular scientific error-bounded lossy compression pipelines. Compared to existing QoI-preserving work [20], QPET introduces several brand-new or improved strategies to cover a wide range of QoI constraints and deliver high compression ratios. The components of QPET are:

- **QoI Analyzer:** This symbolic module creates interfaces for efficient evaluations of the QoI function and its derivatives.
- **Point-wise Error Bound Estimator:** To constrain the QoI error, this module computes an estimated best-fit compression error bound separately for each input data point.
- **Global Error Bound Auto-tuner:** This module auto-tunes the best-fit global data error bound for the compression.
- **Error Bound Compressor:** This module compresses and stores the point-wise error bounds (will be skipped if the base compression pipeline cannot support point-wise error bounds) with improvements from [20, 30].
- **QoI Validator:** This module verifies the QoI constraints during online data prediction-quantization processes and corrects out-of-constraint data points by lossless storage.

Leveraging our proposed QPET framework, Algorithm 1 presents the process of QoI-preserving lossy compression. In Line 1, the QoI Analyzer analyzes input QoI and creates numerical interfaces for

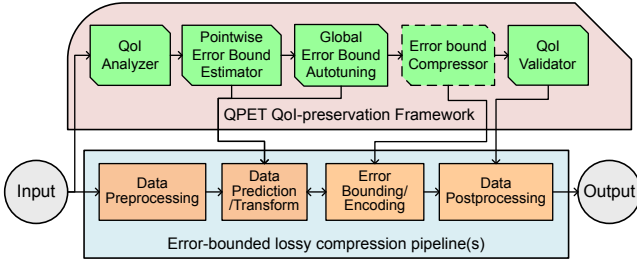


Figure 1: QPET framework and integration

Algorithm 1 QoI ERROR-CONSTRAINED ERROR-BOUNDED LOSSY COMPRESSION WITH QPET

Input: input data $X = \{x_i\}$ of size n , initial global error bound ϵ , QoI function Q , QoI error threshold t
Output: Compressed data C

```

1: QoIAnalysis( $Q$ ) /*Pre-analysis of QoI constraint*/
2:  $\{\epsilon_i\} \leftarrow \{\epsilon\}$  /*Initialize point-wise error bounds*/
3: if isUnivariate( $Q$ ) then /* $F$  is univariate*/
4:   for  $i = 1 \rightarrow n$  do
5:      $\epsilon_i \leftarrow \text{GetPointwiseEB\_UniVar}(x_i, Q, t, \epsilon)$  /*Compute point-wise data error bounds by Algorithm 2*/
6:   end for
7: else /* $Q$  is multivariate*/
8:   for  $i = 1 \rightarrow n_r$  do /* $n_r$  is number of data regions*/
9:      $R_i \leftarrow \text{GetDataRegion}(i)$ 
10:     $E_i \leftarrow \text{GetPointwiseEB\_MulVar}(R_i, Q, t, \epsilon)$  /*Compute point-wise data error bounds  $E_i$  for data region  $R_i$  by Algorithm 3*/
11:   end for
12: end if
13:  $\epsilon_g \leftarrow \text{TuneGlobalEB}(\{\epsilon_i\})$  /*Auto-tune global error bound by Algorithm 4*/
14: for  $i = 1 \rightarrow n$  do
15:    $\epsilon_i \leftarrow \min(\epsilon_i, \epsilon_g)$  /*Update point-wise data errors bound with new global error bound*/
16: end for
17:  $X', C \leftarrow \text{Compress}(X, \{\epsilon_i\}, \epsilon_g)$  /*Compress input data with tuned error bounds.  $C$  and  $X'$  are the compressed data and decompressed data*/
18:  $X_o \leftarrow \text{QoIValidation}(X, X', Q, t)$  /*Validate QoI errors and collect out-of-constraint data*/
19:  $C \leftarrow C + \text{Lossless\_Compress}(X_o)$  /*Losslessly compress  $X_o$  and append the compressed data to  $C$ */
20: return  $C$ 

```

QoI/derivatives evaluations; In Line 5 and Line 10, the point-wise data error bounds are computed (detailed in Section 5.1 and 5.2); In Line 13, the global data error bound is auto-tuned (detailed in Section 5.3); In Line 17, the point-wise data error bounds are compressed and stored; In Line 18, the QoI validator module checks the decompressed QoI errors. The outlier data points that bring out-of-bound QoI errors will be losslessly encoded and appended to the compressed data (detailed in Section 5.4).

4.2 Integrating QPET into error-bounded compression pipelines

Because the design of QPET is only based on error-bound estimation/tuning and post-hoc data validation/correction, it does not rely on certain data compression techniques and implementation details. As such, QPET has high portability and adaptability to many existing error-bounded lossy compression pipelines. In this work, we have integrated QPET into three state-of-the-art scientific error-bounded lossy compressors with two archetypes: SZ3 [31, 58], HPEZ [36], and SPERR [25]. SZ3 and HPEZ are interpolation-prediction-based error-bounded lossy compressors, and SPERR is a

Table 1: Mathematical notations for QoI-preserving

Symbol	Description
$X(\{x_i\})$	Input Data(set)
x_i	Input data point
$X'(\{x'_i\})$	Decompressed Data(set)
x'_i	Decompressed data point
ϵ	General data error bound
ϵ_g	Global data error bound
$\epsilon_i(\{\epsilon_i\})$	Pointwise data error bound (set)
Q	General QoI function
f	Univariate QoI function
F	Multivariate QoI function
τ	General QoI error threshold
t	Univariate QoI error threshold
T	Multivariate QoI error threshold

wavelet-transform-based error-bounded lossy compressor. In the integration of QPET to each data compressor, the error-bound estimation and tuning modules (according to the QoI function and error threshold) are added before the data compression process, and the data validation/correction module serves as a data post-processing step in the compression pipeline.

5 QPET DESIGN DETAILS

In this section, we present the design details of QPET introduced in Section 4. First, we discuss the core of QPET: providing optimized point-wise error bounds for the QoI-preserving lossy compression. Then, we explain how to auto-tune the global data error bound according to the computed point-wise error bounds to optimize the compression ratio. Last, we specify how to correct the outlier data values to strictly respect the QoI error threshold.

5.1 Computing Point-wise Error Bounds with Univariate QoI

Regarding the point-wise error-bound computation, we first visit the univariate QoI cases and then extend it to multivariate QoI functions. We formulate the optimization of the error-bounded lossy compression with univariate QoI as follows. For each data point x_i in the input data, given a global error bound ϵ_g , a QoI function f , and a QoI error threshold t , we need to find the optimized point-wise error bound ϵ_i of x_i from the following problem:

$$\begin{aligned}
 \epsilon_i &= \max \epsilon \\
 \text{s.t. } \epsilon &\leq \epsilon_g, \\
 \text{and } |x'_i - x_i| &\leq \epsilon \implies |f(x'_i) - f(x_i)| \leq t
 \end{aligned} \tag{2}$$

Eq. 2 guarantees that compressing x_i with error bound ϵ_i can bound both the global data and QoI errors with required thresholds. Meanwhile, it is maximized to optimize the data compression ratio. We can optimize this compression task if we can solve ϵ_i effectively from Eq. 2. When the QoI function f has a simple format, such as linear, quadratic, or simply logarithmic, Eq. 2 may also have analytical solutions [20]. However, analytical solutions have limitations: their existence and computational efficiency depend highly on the QoIs. Each analytical solution is customized for a separate QoI, and the computation method cannot be extended to generalized QoIs.

To this end, our work proposes a numerical solution that provides a generalized and efficient approximation method for ϵ_i with a wide range of QoI formats. Specifically, as long as f is second-order differentiable (any elementary function is infinitely differentiable), when computing ϵ_i for data point x_i , we will use its second-order Taylor expansion $f_2(x'_i) = f(x_i) + f'(x_i)(x'_i - x_i) + \frac{f''(x_i)}{2}(x'_i - x_i)^2$ to estimate the QoI value at the decompressed data x'_i . In particular, this Taylor expansion with the third-order remainder is:

$$f(x'_i) = f_2(x'_i) + \frac{f'''(\xi)}{3!}(x'_i - x_i)^3, \quad (3)$$

where ξ is a specific value in the interval $[x_i, x'_i]$ (or $[x'_i, x_i]$ if $x'_i < x_i$). Given the error bound ϵ_i which ensures $x'_i \in [x_i - \epsilon_i, x_i + \epsilon_i]$ and assuming $f'''(x)$ is bounded in $[x_i - \epsilon_i, x_i + \epsilon_i]$ by $|f'''(x)| \leq M$, the remainder $\frac{f'''(\xi)}{3!}(x - x_0)^3$ would be smaller than $\frac{M^3}{6}\epsilon_i^3$, i.e. $|f_2(x'_i) - f(x_i)| \leq \frac{M^3}{6}\epsilon_i^3$. With f_2 , we can efficiently compute a good estimation for the best-fit ϵ_i by the following theorem:

THEOREM 5.1. *Given a second-order differentiable QoI function f , a global error bound ϵ_g , and a QoI error threshold t , for each data point x_i , if $f''(x_i) \neq 0$, an estimation for the best-fit data error bound ϵ_{x_i} to fit its QoI error threshold will be $\min\left(\epsilon_g, \frac{\sqrt{|a|^2 + 2|b|t - |a|}}{|b|}\right)$, in which $a = f'(x_i)$ and $b = f''(x_i)$.*

PROOF. Using $f_2(x) = x_i + a(x - x_i) + \frac{b}{2}(x - x_i)^2$ as an approximation of f , we first find the maximum ϵ so that $|x'_i - x_i| \leq \epsilon \implies |f(x'_i) - f(x_i)| \approx |f_2(x'_i) - f_2(x_i)| \leq |a(x'_i - x_i) + \frac{b}{2}(x'_i - x_i)^2| \leq t$. This is equivalent to maximizing ϵ which satisfies that $[-\epsilon, \epsilon] \subseteq A$, in which A is the solution set of $|ax + \frac{b}{2}x^2| \leq t$.

By solving $|ax + \frac{b}{2}x^2| \leq t$ (for simplicity we omit the detailed process), we can acquire that $\max \epsilon = \frac{\sqrt{|a|^2 + 2|b|t - |a|}}{|b|}$. Considering the global ϵ_g , finally $\epsilon_{x_0} = \min\left(\epsilon_g, \frac{\sqrt{|a|^2 + 2|b|t - |a|}}{|b|}\right)$. \square

Theorem 5.1 is valid when $b = f''(x_0) \neq 0$. Under $f''(x_0) = 0$, we have the following theorem as a supplement of Theorem 5.1, whose proof is intuitively trivial:

THEOREM 5.2. *Given a second-order differentiable QoI function f , a global error bound ϵ_g , and a QoI error threshold t , for a data point x_0 that $f''(x_0) = 0$, an estimation for the best-fit data error bound of x_0 will be $\min\left(\epsilon_g, \frac{t}{|f'(x_0)|}\right)$ if $f'(x_0) \neq 0$, or ϵ_g when $f'(x_0) = 0$.*

Theorem 5.1 and Theorem 5.2 cover all the cases for point-wise error-bound estimation. In Figure 2, based on different pointwise QoIs and error tolerances, we plot the point-wise data error bounds for different data values from both analytical solutions and our proposed estimation method. For various function formats (Figure 2 (a-c)), QPET acquires accurate error bound estimations compared to the analytical solutions, even when $|f'''(x)|$ has large values (see Figure 2 (c) where $f'''(x) = -1000 \cos 10x$). In Figure 2 (d), we show an example of $f(x) = x^3$, which has close-to-zero first- and second-order derivatives when x is close to zero. While QPET may produce over-estimated error bounds in this case, the over-estimated error bounds will be cropped by a pre-given or auto-tuned global error bound (to be detailed in Section 5.3) and thus do no harm to the compression process (notice that in Figure 2 we set ϵ_g as $+\infty$ to

Algorithm 2 POINT-WISE ERROR BOUND ESTIMATION WITH UNIVARIATE QOI

Input: input data $X = \{x_i\}$ of size n , global error bound ϵ_g , QoI function f , point-wise QoI error threshold t

Output: Approximated optimized point-wise data error bounds $\{\epsilon_i\}$

```

1:  $d_1(x) \leftarrow f'(x)$  /*Get the efficient process for the first-order derivative*/
2:  $d_2(x) \leftarrow f''(x)$  /*Get the efficient process for the second-order derivative*/
3: for  $i = 1 \rightarrow n$  do
4:    $a \leftarrow d_1(x_i)$ 
5:    $b \leftarrow d_2(x_i)$ 
6:   if  $b \neq 0$  then
7:      $\epsilon_i \leftarrow \frac{\sqrt{|a|^2 + 2|b|t - |a|}}{|b|}$  /*Compute  $\epsilon_i$  with both derivatives*/
8:   else if  $a \neq 0$  then
9:      $\epsilon_i \leftarrow \frac{t}{|a|}$  /*Compute  $\epsilon_i$  with only first-order derivative*/
10:  else
11:     $\epsilon_i \leftarrow \epsilon_g$  /*set  $\epsilon_i$  as  $\epsilon_g$ */
12:  end if
13:   $\epsilon_i \leftarrow \min(\epsilon_g, \epsilon_i)$ 
14: end for
15: return  $\{\epsilon_i\}$ 

```

focus on the estimation quality from QoI error bounds). As long as QPET presents good estimations for small error bounds (it does), it will be optimized for the compression ratio.

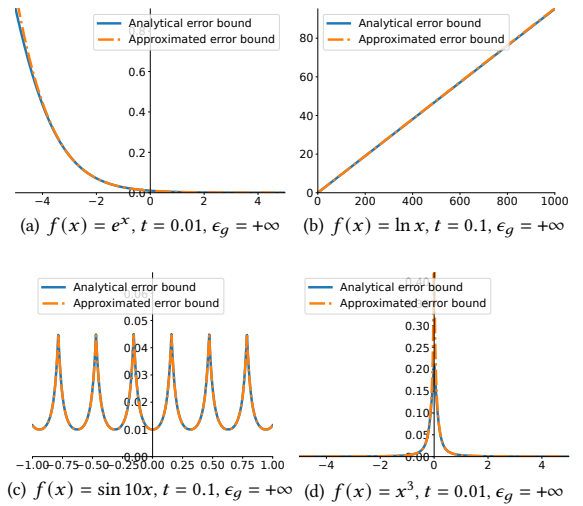


Figure 2: Analytical and estimation point-wise data error bounds of different QoI function $f(x)$ and error threshold t .

5.2 Computing Point-wise Error Bounds with Multivariate QoI

For multivariate QoIs, we actually deal with QoI functions in the format of $F(x_1, \dots, x_n)$. $\{x_i\}$ is a set of data points from the input, which can be a data block, scalar components of a vector, and so on. Given $F(x_1, \dots, x_n)$, and a tolerance T for which we requires that $|F(x'_1, \dots, x'_n) - F(x_1, \dots, x_n)| \leq t$ ($\{x'_i\}$ are decompressed values of $\{x_i\}$), we will determine a set of error bounds $\{\epsilon_i\}$ so that

$$\forall i, |x'_i - x_i| \leq \epsilon_i \implies |F(x'_1, \dots, x'_n) - F(x_1, \dots, x_n)| \leq T.$$

In other words, compressing each x_i with an error bound ϵ_i will lead the QoI error to be constrained. In the next, we will first discuss optimizing $\{\epsilon_i\}$ when $F(x_1, \dots, x_n)$ is the linear combination of univariate terms, i.e., the variables are separated. Then, we will show how to extend the methodology to generalized cases.

5.2.1 *Multivariate QoI in variable-separated format.* We first demonstrate our strategy for multivariate QoI functions in the format of

$$F(x_1, \dots, x_n) = C + \sum_i \alpha_i f(x_i) \quad (4)$$

in which C and α_i are constant coefficients, and f can be regarded as a point-wise QoI function. Many common multivariate QoIs [20] are in this format, such as region sum, the (weighted) average of a data block, or the difference-based Laplacian operator [40]. Therefore, it is critical that we present a good solution for this variable-separated format. Suppose we have a set of optimized point-wise tolerances $\{t_i\}$ from the following optimization problem:

$$\begin{aligned} & \text{Optimize } \{t_i\} \\ & \text{s.t. } \forall i, |f(x'_i) - f(x_i)| \leq t_i \\ & \implies |F(x'_1, \dots, x'_n) - F(x_1, \dots, x_n)| \leq T \end{aligned} \quad (5)$$

Then we can optimize the point-wise data error bound $\{\epsilon_i\}$ for set $\{x_i\}$ from those $\{t_i\}$ with Algorithm 2, as in this case we are just bounding point-wise QoI f error within t_i . Therefore, by solving Eq. 5 we can find a way to optimize $\{\epsilon_i\}$. Regarding the optimization target of Eq. 5, intuitively, we can maximize $\sum_i t_i$. However, this optimization target sometimes leads to skewed distributions of t_i and ϵ_i values over different i . Preliminary works [20] and our experiments showed that point-wise error bounds with skewed distribution would cause the compression ratio to decrease significantly. Therefore, we select to maximize $\min t_i$, equivalent to ensure $t_i = t$ for each i . According to our preliminary experiments, this scheme is better than other intuitive strategies, such as aligning $\alpha_i t_i$. Moreover, if we denote D_i as $f(x'_i) - f(x_i)$, we will notice that $F(x'_1, \dots, x'_n) - F(x_1, \dots, x_n) = \sum_i \alpha_i D_i$. To this end, we can convert Eq. 5 to a more concise format:

$$\begin{aligned} & \text{Optimize } t \\ & \text{s.t. } \forall i, |D_i| \leq t \implies \left| \sum_i \alpha_i D_i \right| \leq T \end{aligned} \quad (6)$$

Next, we will show two solutions of Eq. 6, which contribute collaboratively to optimizing t , $\{\epsilon_i\}$, and the compression ratio. First, we have a solution of t that deterministically matches Eq. 6:

THEOREM 5.3. *If $t = \frac{T}{\sum_i |\alpha_i|}$, then $|D_i| \leq t \implies |\sum_i \alpha_i D_i| \leq T$.*

PROOF. $|\sum_i \alpha_i D_i| \leq \sum_i |\alpha_i| |D_i| \leq t \sum_i |\alpha_i| = T$. \square

Unfortunately, Theorem 5.3 sometimes presents over-conservative values for t . Fortunately, we have another strategy for determining t . If we regard the $\{D_i\}$ as random variables with the following assumptions: 1) $\{D_i\}$ are independent; 2) D_i has symmetric distribution regarding zero (so $E(D_i) = 0$). Since $|D_i| \leq t$, according to [49, 50], $\{D_i\}$ are sub-Gaussian, and the following concentration inequality holds for each $T > 0$ [17, 49, 50]:

$$P(|\sum_i \alpha_i D_i| \geq T) \leq 2e^{-\frac{T^2}{2 \sum_i \alpha_i^2 \sigma_i^2}} \quad (7)$$

σ_i^2 is called the *variance proxy* of D_i . For finite-bounded $|D_i| \leq t$, we have $\sigma_i \leq t$ [49, 50]. QPET uses a single estimation $\sigma_0 = \frac{t}{c}$ for all D_i , in which $c \leq 1$ and σ_0 represents the estimated standard variance of D_i . For example, $c = \sqrt{3}$ and $\sigma_0 = \frac{t}{\sqrt{3}}$ correspond to the

Table 2: Point-wise QoI ($f(x)$) error thresholds from different multivariate QoI function F and variable numbers ($c = 2$, $\beta = 0.9999$). Error threshold of $F(X)$ is T .

QoI	Variable Num. (n)	t from Theo. 5.3	t from Theo. 5.4
$f(x) + f(y) + f(z)$	3	$0.33T$	$0.26T$
$\frac{1}{n} \sum f(x)$	2^3	T	$1.27T$
	4^3	T	$3.6T$

uniform distribution over $[-t, t]$. With σ_0 and Eq. 7, we can set up another estimation of an optimized t for Eq. 6:

THEOREM 5.4. *A multivariate QoI function $F(x_1, \dots, x_n) = C + \sum_i \alpha_i f(x_i)$ and the corresponding QoI error threshold T are given. On $\{x_1, \dots, x_n\}$, if the point-wise QoI error D_i is σ_i -Sub-Gaussian in which $\sigma_i \leq \frac{\max D_i}{c}$, for each confidence level of β , taking point-wise QoI error threshold $t = \max |f(x'_i) - f(x_i)| = cT \sqrt{\frac{1}{2 \sum_i \alpha_i^2 \ln \frac{2}{1-\beta}}}$, the global QoI error threshold can be guaranteed in a confidence level of β , i.e.: $P(|F(x'_1, \dots, x'_n) - F(x_1, \dots, x_n)| \leq T) \geq \beta$.*

PROOF. Using $\sigma_i = \frac{t}{c}$, for $\{x_1, \dots, x_n\}$, from Eq. 7 we have $P(|F(x'_1, \dots, x'_n) - F(x_1, \dots, x_n)| \geq T) = P(|\sum_i \alpha_i D_i| \geq T) \leq 2e^{-\frac{T^2}{2 \sum_i \alpha_i^2 \sigma_i^2}} = 2e^{-\ln \frac{2}{1-\beta}} = 1 - \beta$. Therefore, $P(|F(x'_1, \dots, x'_n) - F(x_1, \dots, x_n)| \leq T) = 1 - P(|F(x'_1, \dots, x'_n) - F(x_1, \dots, x_n)| \geq T) \geq \beta$. \square

Because Theorem 5.3 and 5.4 fits in different cases when optimizing t in Eq. 6. In practice, we select the larger one between those two estimations to optimize the compression ratio. For a small variable number n , Theorem 5.3 can often provide better estimations and vice versa. Table 2 shows a few examples of the point-wise QoI ($f(x)$) error tolerances deduced from Theorem 5.3 and 5.4. As the number of variables increases, Theorem 5.4 provides more aggressive estimations.

Admittedly, Theorem 5.4 needs certain preliminary conditions of $\{D_i\}$, which are not always true. The variance of compression errors is dependent on diverse factors, such as data characteristics and input error bound, and $\{D_i\}$ are not theoretically independent of each other (such as in interpolation-based data compression). Nevertheless, existing research [35, 36, 58] showed that, in practical cases, distributions of $\{D_i\}$ are not severely distorted from the assumptions, exhibiting distributions that are close to uniform or Gaussian, and the autocorrelation of decompression errors drop rapidly as the compression accuracy increases [35]. Moreover, to handle the unclear characteristics of $\{D_i\}$, Theorem 5.4 can still acquire adequate QoI error threshold estimations by selecting a conservative value of c . To optimize the compression ratio, Theorem 5.4 does not bring a strict guarantee for bounding all QoI errors. In Section 5.4, we will discuss how QPET strictly guarantees the QoI error threshold and can optimize the compression ratio through more aggressive data error-bound settings.

5.2.2 *Multivariate QoI in non-variable-separated format.* Beyond the variable-separated format described in Eq. 4, regarding non-variable-separated multivariate QoI function $F(x_1, \dots, x_n)$, to separate the variables so that techniques described in Section 5.2.1 can be applied, we use its approximation for computing point-wise error bounds, leveraging the first-order differentials of F (using

Algorithm 3 POINT-WISE ERROR BOUND ESTIMATION WITH MULTI-VARIATE QoI CONSTRAINT

Input: input data $X = \{x_i\}$, Data regions $\mathcal{R} = \{R_i\}$ ($R_i = \{r_{ij}, id(r_{ij})\}$ is the global index of r_{ij} so $r_{ij} = x_{id(r_{ij})}$), global error bound ϵ_g , multivariate QoI function F , QoI error threshold T , estimation parameter c , and confidence level β .
Output: Point-wise data error bounds $\{\epsilon_i\}$ For multivariate QoI error constraint.

```

1: for  $i = 1 \rightarrow |X|$  do
2:    $\epsilon_i \leftarrow \epsilon_g$  /*Initialize all  $\epsilon_i$  with global error bound*/
3: end for
4: for  $R_i \in \mathcal{R}$  do /*Iterate through regions*/
5:   if  $F(R_i) = C + \sum_j a_j g(r_{ij})$  then /*Linear Format*/
6:      $f \leftarrow g$ 
7:      $\{\alpha_j\} \leftarrow \{a_j\}$ 
8:   else /*Non-linear*/
9:      $f \leftarrow f(x) = x$ 
10:     $\{\alpha_j\} \leftarrow \left\{ \frac{\partial F}{\partial x_j}(r_{ij}) \right\}$ 
11:   end if
12:    $t_1 \leftarrow \frac{T}{\sum_i |\alpha_i|}$  /*Theorem 5.3*/
13:    $t_2 \leftarrow cT \sqrt{\frac{1}{2 \sum_j \alpha_j^2 \ln \frac{2}{1-\beta}}}$  /*Theorem 5.4*/
14:    $t \leftarrow \max(t_1, t_2)$ 
15:   for  $r_{ij} \in R_i$  do /*Iterate through data value in regions*/
16:      $\epsilon \leftarrow \text{Compute\_eb\_univar}(r_{ij}, t, f, \epsilon)$  /*By Algorithm 2*/
17:      $\epsilon_{id(r_{ij})} \leftarrow \min(\epsilon_{id(r_{ij})}, \epsilon)$  /*update  $\epsilon_{id(r_{ij})}$ */
18:   end for
19: end for
20: return  $\{\epsilon_i\}$ 

```

higher-order differentials will bring cross partial derivatives, failing in separate the variables). Specifically, we have:

$$F(x'_1, \dots, x'_n) \approx F(x_1, \dots, x_n) + \sum_i \frac{\partial F}{\partial x_i}(x_i)(x'_i - x_i) \quad (8)$$

Eq. 8 falls into the format of Eq. 4, with $C = F(x_1, \dots, x_n)$, $\alpha_i = \frac{\partial F}{\partial x_i}(x_i)$, and $f(x) = x - x_i$ (equivalent to $f(x) = x$). Therefore, we can use the Eq. 8 and Theorem 5.3/5.4 to determine the point-wise error bounds from F . It is worth noticing that, in this case, for different $\{x_i\}$, we will have different t from Theorem 5.3 and 5.4 as there are different α_i values from the partial derivatives.

In Algorithm 3, we combine the computation of point-wise data error bounds from multivariate QoI in either linear or non-linear format. When dealing with the linear format, the computation can often get simplified. For example, for linear multivariate QoIs, the t in Line 14 can be shared over data regions.

5.3 Optimizing Global Error Bound

After QPET acquires the point-wise error bound for each data point x_i , gathering up a set of error bounds $\{\epsilon_i\}$, QPET auto-tunes a new global error bound ϵ_g , cropping all larger point-wise error bounds to ϵ_g in the compression. It is for two purposes: 1) For compressors that support point-wise accuracy (like SZ3/HPEZ), QPET will store the quantized point-wise error bound, and cropping most error bounds will significantly reduce their storage, which improves the overall compression ratio in turn; 2) For compressors that do not support point-wise accuracy (like SPERR), QPET needs an optimized global accuracy setting as the compression configuration.

To this end, we proposed a two-step dynamic global error-bound auto-tuning strategy in QPET. After acquiring the set of point-wise error bounds $\{\epsilon_i\}$, in the first step, a set of quantiles (20%, 10%, 5%, 2%, 1%, 0.5%, 0.25%) is drawn from it as global error bound

Algorithm 4 GLOBAL ERROR BOUND AUTO-TUNING

Input: point-wise data error bounds $\{\epsilon_i\}$ with size n , a dropping threshold coefficient $c_0 \leq 1$

Output: Auto-tuned global error bound ϵ_g

```

1:  $E_c \leftarrow \{\}$ 
2: for  $q \in \{0.2, 0.1, 0.05, 0.02, 0.01, 0.005, 0.0025\}$  do
3:    $k \leftarrow \lfloor qn \rfloor$ 
4:    $\epsilon \leftarrow \text{KthSmallest}(\{\epsilon_i\}, k)$  /* $\epsilon$  is the  $k^{\text{th}}$ -smallest of  $\{\epsilon_i\}$ */
5:    $E_c.\text{insert}((\epsilon, k))$ 
6: end for
7:  $\epsilon_0, k_0 \leftarrow \text{cmpTest\_BestCR}(E_c)$  /* $\epsilon_g$  achieves the best compression ratio in compression tests among all error bounds in  $E_c$ , save its current value in  $\epsilon_0$ */
8:  $\epsilon_g \leftarrow \epsilon_0$ 
9:  $E_p \leftarrow \text{partialSort}(\{\epsilon_i\}, k_0)$  /* $E_p$  is the sorted  $k$  smallest values in  $\{\epsilon_i\}$ */
10:  $k \leftarrow k_0 - 1$  /*Start checking smaller error bounds*/
11: while  $k \geq 0$  /*we regard minimum as 0th-smallest*/ do
12:    $\epsilon \leftarrow E_p[k]$ 
13:   if  $\epsilon \geq (c_0 + \frac{k}{k_0}(1 - c_0))\epsilon_0$  /* $\epsilon$  is above a slope*/ then
14:      $\epsilon_g \leftarrow \epsilon$  /*Update  $\epsilon_g$ */
15:   else
16:     break /*Stop decreasing  $\epsilon_g$ */
17:   end if
18:    $k \leftarrow k - 1$ 
19: end while
20: return  $\epsilon_g$ 

```

candidates, and then QPET identify the best from them in a sample-and-tests scheme described in [35]. In the second step, QPET further reduces ϵ_g over the distribution of $\{\epsilon_i\}$ until a steep decrease is detected. According to our preliminary experiments, saving fewer error-bound values is often more beneficial in improving the compression ratio than applying a slightly larger global error bound.

The details of the described strategy are featured in Algorithm 4. In Lines 1 to 7, QPET optimizes the global error bound by tests on sampled data. In Line 13, QPET checks whether the current error bound ϵ is still over a gentle linear slope between the initial ϵ_0 and a termination value ϵ_t (e.g., $\epsilon_t = 0.95\epsilon_0$), and stop the decreasing of ϵ_g if it will beneath that slope. The KthSmallest() operations in Algorithm 4 can be implemented by fast-selection and only need to be performed on reduced ranges of $\{\epsilon_i\}$. Therefore, all the operations in Algorithm 4 can be efficiently done with $O(n + qn \log qn)$ time complexity in average, in which q is the quantile of ϵ_0 . In practice, QPET skips the second part of tuning (from Line 8) when q is larger than a threshold (0.005) to guarantee efficiency, bringing negligible overhead to the overall compression pipeline.

5.4 QoI Error Validation and Correction

As previously described in Section 3.2, the design purpose of QPET is to optimize the compression ratio when strictly constraining the QoI error threshold. To fulfill this target, the point-wise and global error-bound calculations demonstrated in Section 5.1, Section 5.2, and Section 5.3 do not conservatively make the QoI value strictly-preserved on each data element, but require a subsequent process to correct and store the out-of-bound data points (i.e., outliers). By selecting proper error-bound values, QPET can significantly reduce the average storage cost of each compressed data value, meanwhile introducing negligible overheads for correcting outliers as they are only a tiny portion ($< 1\%$) of the whole input data. Specifically, QPET detects and corrects the outliers in the following scheme: Given the decompressed data (which are naturally acquired after the compression process), QPET computes the QoI values

on it and compares them to the original QoI values. For out-of-bound univariate QoI errors, QPET losslessly encodes and stores all the corresponding original data values. For each out-of-bound multivariate QoI error, there are several related data values. To minimize the number of data values that are losslessly stored, QPET stores those values one by one and updates the QoI result after each value update until the QoI error gets bounded.

6 EVALUATIONS

We evaluate QPET using six real-world datasets and diverse QoIs. In particular, we integrate QPET into three top-performing compressors (SZ3 [31], HPEZ [36], and SPERR [25]) and compare them with diverse baselines in terms of QoI preservation.

6.1 Experimental Setup

6.1.1 Experimental environment and datasets. We perform the evaluations on 6 real-world scientific datasets from diverse domains (details in Table 3). Experiments are operated on Purdue Anvil computing cluster [39] (each node is equipped with two 64-core AMD EPYC 7763 CPUs and 512GB DDR4 memory).

Table 3: Information of the datasets in experiments

App.	# fields	Dimensions	Total Size	Domain
Miranda	7	256×384×384	1GB	Turbulence
Hurricane	13	100×500×500	1.2GB	Weather
RTM	11	449×449×235	2.0GB	Seismic Wave
NYX	6	512×512×512	3.1GB	Cosmology
SEGSalt	3	1008×1008×352	4.0GB	Geology
SCALE-LetKF	12	98×1200×1200	6.3GB	Climate

6.1.2 Baselines. Besides the QPET-integrated compressors, we included several existing solutions for QoI-preserving scientific lossy compression in the evaluations, which are in 2 categories: (1) **Parameter-search-based solutions:** general-purpose compressors (SZ3/HPEZ/SPERR) cannot directly bound specified QoI errors, so we apply parameter-search methods on top of them to figure out the best-fit (yielding highest compression ratio) data error bound for preserving the QoI. The parameter-search methods, including binary-search-based and FraZ [48], are from OptZConfig [47], the state-of-the-art scientific lossy compression parameter-search toolkit. When applying, we slightly revised those methods to get them better adapted and accelerated in the QoI-preserving tasks. Those baselines are named SZ3/HPEZ/SPERR-OptZ-R (R is short for revised), and they represent the best (fastest) results from both parameter-search methods on every base compressor. The parameter-search process often requires multiple iterations of data compression and QoI validation to fit the QoI error bound, so they are typically quite slow. (2) **Direct QoI-preserving solutions:** [20] provided the SZ3-based QoI-preserving compressor QoI-SZ3, and we further ported its QoI-preserving features to HPEZ, creating QoI-HPEZ. Moreover, we evaluated MGARD-QoI [6]. Those QoI-preserving compressors have limited support for diverse QoI formats. QoI-SZ3/HPEZ only supports square and logarithm (also their block-averages), and MGARD-QoI only supports linear QoIs. So, we only tested them on the QoIs supported. Other existing QoI-integrated compressors are designed for different tasks (e.g., cpSZ [28] only works for critical points in vector field data), so they are not included in our evaluation baselines.

6.1.3 QoI functions, experimental configurations, and evaluation metrics. Table 4 shows the QoI functions in the evaluation tasks. Among them, there are three different categories: point-wise, regional, and vector. They have diverse mathematical formats, and for many among them (such as $\tanh x$, $\frac{1}{n} \sum x^3$, and vector QoIs), QPET is the first framework that supports compression with preservation of those QoIs. The selection of QoI functions in our evaluation is based on existing investigations and analysis [20, 40, 53] of QoIs in practical scientific data analysis tasks, such as physical transform (kinetic energy as velocity's square), and clustering ($\sqrt{x^2 + y^2 + z^2}$ is the distance from origin when x , y , and z are coordinates).

Table 4: QoI functions in the evaluation

QoI type	QoI function	QoI type	QoI function
Pointwise	x^2	Regional	x (average)
	x^3		x^2 (average)
	$\log_2 x$		x^3 (average)
	$\sin 10x$	Vector	$x^2 + y^2 + z^2$
	$\tanh x$		$\sqrt{x^2 + y^2 + z^2}$

For the OptZ-R parameter search, we set an early termination that triggers when a maximum QoI error between 90% and 100% of the required threshold is found because it presents a near-optimal compression ratio with reasonable search time. Regarding the compression configurations, we apply the default optimization level and compression-ratio-preferred mode for HPEZ. Regarding QPET parameters in Algorithms 3, we set $c = 3$, $\beta = 0.999$ for SPERR, $c = 2$, $\beta = 0.999$ for HPEZ, and $c = 2$, $\beta = 0.99999$ for SZ3. On SZ3 and HPEZ, the autocorrelation of decompression errors are relatively high when the error bound is large [35], so we linearly dynamically decrease c when τ increases over 10^{-3} , eventually to 1.0 when τ becomes 10^{-2} .

In evaluating the compression performance, the following widely adopted metrics [20, 43] are used: (1) Compression and decompression speeds (throughputs). (2) Compression ratio $CR = \frac{|X|}{|C|}$, which is the input data size $|X|$ divided by the data size $|C|$; (3) Bit rate $BR = \frac{|C| * 8 * \text{sizeof}(x)}{|X|}$, which is the number of bits in compressed data to store each value in the input. (3) Maximum data error and QoI error between the input and output;




Figure 3: Compression and decompression speed for $Q(x) = x^2$ and $\tau = 1e-3$.




Figure 4: Compression and decompression speed for $Q(x) = x^3$ and $\tau = 1e-3$.

6.2 Evaluation Results

Here, we propose our evaluation results for QPET, which prove its advantages and versatility. Due to page limits, we selected the most representative results to be presented in this paper.

6.2.1 Outperforming speeds. First, we measure and profile the execution speeds of baselines and QPET-integrated compressors to examine how QPET has accelerated the QoI-preserving error-bounded lossy compression. From the compression and decompression speeds on the Anvil cluster with the QoI functions $Q(x) = x^2$ and $Q(x) = x^3$ and the relative QoI error threshold $\tau = 1e-3$, reported in Figure 3 to 4 (no results from QoI-SZ3/HPEZ for $Q(x) = x^3$ because their current implementation don't support this QoI), we can conclude as follows: compared to existing parameter-search-based baselines and QoI-preserving compressors, QPET-integrated compressors have significantly improved the compression throughputs (4x ~ 6x of *-OptZ-R and 2x ~ 2.5x of QoI-SZ3/HPEZ), with similar decompression throughput with the fast general-purpose compressors. Parameter search on general-purpose compressors (*-OptZ-R) need to perform repeated executions for acquiring a certain QoI error threshold, and in our evaluations, this task may cost 4 to 10 or even more compression executions and QoI validations. Moreover, QPET brings quite limited time overheads in the decompression process (because neither error-bound computation nor QoI validation is needed), and sometimes it even improves the decompression speed over baselines due to the reduced accuracy required in the compression to ensure the same QoI error threshold. Last, Compared to the existing QoI-SZ3/HPEZ, for the same use cases, QPET gains 100% ~ 150% compression speed improvement and up to 400% decompression speed improvement, which is highly attributed to the efficient error-bound selection and tuning process proposed by QPET, and also the revised implementation in storing and recovering the point-wise error bounds.

6.2.2 Representative showcases. In Table 5, on diverse data domains, we present 15 QoI-preserving error-bounded lossy compression showcases for the 10 QoI functions in Table 4, leveraging QPET and all available baselines. We apply different error-bound targets and report both compression ratio and sequential throughputs.

From those showcases, key findings are: (1) Matching the discussion in Section 6.2.1, QPET-integrated compressors deliver better compression throughputs than all baselines. The QPET-integrated compressor achieves 2x to 10x compression speedups over the parameter-search-based solutions. Moreover, SZ3/HPEZ-QPET has 2x to 3x compression speeds and 3x to 6x decompression speeds over QoI-SZ3/HPEZ. (2) With better throughputs, QPET-integrated compressors bring no worse (and mostly improved) compression ratios over the baselines. On $q(x) = x^2$, under error bounds of 10^{-4} , QPET gained similar compression ratios with baselines in much shorter time costs. On preserving $q(x) = x^3$ of RTM-3500 and Hurricane-Cloud data (they both feature small absolute values less than 0.002), QPET-integrated compressors achieved 30% - 700% compression ratio improvements over the best parameter-search solution. On NYX-Velocity_x data ranging from zero to tens of millions, general-purpose compressors failed to preserve the QoI $\tanh x$ well due to the required point-wise data accuracy extremely diverging among small and large values. In contrast, QPET handled this case excellently. On multivariate QoIs, QPET significantly improved the compression ratio over baselines, as benefited from Theorem 5.4. Compared with existing QoI-preserving compressors on basic QoIs, QPET-integrated compressors still achieve noticeable compression ratio gains. On preserving $q(x) = \log_2 x$ on Scale-LetKF-T data, SPERR-QPET achieved 100% CR improvement over QoI-HPEZ under $\epsilon = \tau = 10^{-4}$. On preserving $q(X) = \frac{1}{n} \sum x_i^2$ on SegSalt-Pressure-3000 data, SPERR-QPET achieved 133% CR improvement over QoI-HPEZ under $\tau = 10^{-3}$. General-purpose compressors usually over-preserve data quality to preserve the QoI error threshold because they can neither apply point-wise data accuracy nor auto-correct outliers. This explains why non-QoI-integrated compressors have limited compression ratios on QoI-preserving compression tasks, even after parameter searches on optimizing the compression ratio. Meanwhile. Existing QoI-preserving compressors (like QoI-SZ3/HPEZ and MGARD-QoI) bring significant computational overheads for the QoI-preserving compression tasks. With lightweight error-bound tuning and outlier correction mechanisms, QPET successfully avoided the abovementioned limitations, proving its effectiveness and usability in preserving the QoI.

6.2.3 Improved Rate-distortions. To comprehensively understand the compression ratio improvement brought by QPET in the QoI-preserving error-bounded data compression, the compressors mentioned above are evaluated on a large variety of QoI functions and error thresholds. Then, we plot the bit rates and maximum QoI errors (relative, absolute value divided by value range), showing that QPET-integrated compressors have substantially improved the compression ratio when achieving the same maximum QoI errors as others. Due to the page limit, we will show the most representative results among the excellent outcomes we acquired.

Regarding different point-wise QoI functions, Figure 5 presents the bit rate-maximum QoI error plots we acquired on 6 data domains. In this figure and all other similar ones in the paper, the results from the same base compressors are plotted in the curves of the same color, with QPET-integrated compressors solid and baselines dashed. Moreover, the results from QoI-SZ3/HPEZ appear if and only if the corresponding QoI is supported by them.

Table 5: Showcases of QoI-preserving errorlossy compression. For blocked QoIs, $n = 4^3$ (i.e., average on 4x4x4 blocks). ϵ : data error bound. τ : QoI error threshold. #It: Number of iterations. CR: Compression ratio. T_c : Compression throughput (in MB/s). T_d : Decompression throughput (in MB/s). All experiments are performed on Purdue Anvil, and all error bounds are relative (absolute value/value range). N/A indicates that the baseline is not available in the case.

QoI			$Q(x) = x^2$				$Q(x) = x^3$				$Q(x) = \log_2 x$				$Q(x) = \sin 10x$				$Q(x) = \tanh x$			
Data field			SegSalt-Pressure2000				RTM-3500				NYX-Baryon Density				Miranda-Pressure				Scale-LetKF-RH			
ϵ	τ	Compressor	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d
10^{-1}	10^{-2}	SZ3-OptZ-R	9	283	21	693	9	167	19	695	23	14.1	3	226	12	101	7	379	12	9.88	8	182
		HPEZ-OptZ-R	9	448	14	413	6	231	20	428	3	14.7	21	182	12	151	6	369	12	8.78	8	214
		SPERR-OptZ-R	6	413	10	144	9	237	10	220	23	19.4	2	71	12	132	3	118	12	11.4	4	60
		QoI-SZ3		458	36	88		N/A	N/A	N/A		21.9	19	80		N/A	N/A	N/A		N/A	N/A	N/A
		QoI-HPEZ		687	30	70		N/A	N/A	N/A		24	17	71		N/A	N/A	N/A		N/A	N/A	N/A
		SZ3-QPET		689	102	513		619	119	630		28.1	46	376		110	27	453		74.4	77	641
		HPEZ-QPET		954	75	398		1466	92	413		27.9	44	336		154	26	364		196	71	503
		SPERR-QPET		849	49	135		728	63	260		35.6	22	67		134	21	120		11.9	22	53
10^{-2}	10^{-3}	SZ3-OptZ-R	9	76.5	19	505	6	28	25	444	3	5.9	22	152	12	32.9	7	320	12	5.4	7	144
		HPEZ-OptZ-R	9	98.8	14	381	6	30.6	18	322	3	6.2	21	183	12	43.9	6	320	12	4.7	7	154
		SPERR-OptZ-R	10	108	5	126	10	39.6	7	136	15	6.1	2	38	12	52	3	99	12	5.8	3	39
		QoI-SZ3		169	35	86		N/A	N/A	N/A		7.4	24	112		N/A	N/A	N/A		N/A	N/A	N/A
		QoI-HPEZ		206	30	70		N/A	N/A	N/A		8	20	123		N/A	N/A	N/A		N/A	N/A	N/A
		SZ3-QPET		181	76	332		244	88	505		8.8	33	126		34.4	25	288		18.1	57	350
		HPEZ-QPET		196	64	276		320	78	405		9.2	33	128		44.5	24	253		23.4	52	297
		SPERR-QPET		224	47	128		210	59	219		9.1	18	40		52.5	20	100		5.9	17	35
QoI			$Q(x) = x^2$				$Q(x) = x^3$				$Q(x) = \log_2 x$				$Q(x) = \sin 10x$				$Q(x) = \tanh x$			
Data field			Miranda-Density				Hurricane-Cloud				Scale-LetKF-T				Miranda-Viscosity				NYX-Velocity_x			
ϵ	τ	Compressor	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d
10^{-3}	10^{-3}	SZ3-OptZ-R	4	87.3	37	379	8	49.3	23	596	3	73.9	27	727	8	47.4	10	343	17	1.2	5	223
		HPEZ-OptZ-R	4	112.7	33	407	4	53	40	502	3	84.3	21	460	8	58.9	9	411	8	1.2	11	235
		SPERR-OptZ-R	4	115	12	116	5	26.4	10	98	3	142	15	154	8	63.4	5	104	7	1.3	3	15
		QoI-SZ3		80.2	39	84		N/A	N/A	N/A		67.3	30	85		N/A	N/A	N/A		N/A	N/A	N/A
		QoI-HPEZ		90.9	36	73		N/A	N/A	N/A		95.8	28	77		N/A	N/A	N/A		N/A	N/A	N/A
		SZ3-QPET		96.4	120	450		69.3	97	353		88.6	106	391		51.2	31	400		86.5	72	420
		HPEZ-QPET		121	108	390		62.1	82	329		105	46	329		63.4	29	327		89.2	56	291
		SPERR-QPET		118	44	116		36.3	44	114		158	33	131		65.2	22	106		72	32	81
10^{-4}	10^{-4}	SZ3-OptZ-R	4	31.1	35	317	5	29.5	35	502	3	16	21	202	8	21.5	10	288	3	1.2	30	274
		HPEZ-OptZ-R	4	36.1	31	343	4	28.2	34	433	3	15.6	20	279	8	25.5	9	294	3	1.2	25	195
		SPERR-OptZ-R	4	49.2	11	97	2	14.3	21	73	3	28.1	12	96	8	31.5	4	86	3	1.2	6	14
		QoI-SZ3		29.5	39	80		N/A	N/A	N/A		13.2	29	75		N/A	N/A	N/A		N/A	N/A	N/A
		QoI-HPEZ		33.4	35	71		N/A	N/A	N/A		15.1	28	70		N/A	N/A	N/A		N/A	N/A	N/A
		SZ3-QPET		32.4	116	369		36.6	98	329		16.6	45	192		22.2	28	267		15.3	58	202
		HPEZ-QPET		37	106	327		28.2	70	258		16.7	44	216		26.4	27	229		16.1	54	183
		SPERR-QPET		50.2	39	98		17.3	34	79		30	26	85		32.1	20	87		14.1	23	51
QoI			$Q(X) = \frac{1}{n} \sum x_i$				$Q(X) = \frac{1}{n} \sum x_i^2$				$Q(X) = \frac{1}{n} \sum x_i^3$				$Q(x, y, z) = x^2 + y^2 + z^2$				$Q(x, y, z) = \sqrt{x^2 + y^2 + z^2}$			
Data field			RTM-3200				SegSalt-Pressure3000				Scale-LetKF-V				Hurricane-UVW				Miranda-VXYZ			
ϵ	τ	Compressor	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d	#It	CR	T_c	T_d
10^{-2}	10^{-3}	SZ3-OptZ-R	7	76.1	20	617	6	66.7	24	576	14	37	11	605	10	13	11	270	8	122	14	393
		HPEZ-OptZ-R	7	99.2	14	397	7	115	15	388	9	66.8	13	434	10	14	8	260	8	179	13	411
		SPERR-OptZ-R	5	155.6	16	205	6	126	9	129	7	94	9	142	8	20.4	7	48	9	167	5	121
		MGARD-QoI		15.7	3	4		N/A	N/A	N/A		N/A	N/A	N/A		N/A	N/A	N/A		N/A	N/A	N/A
		QoI-SZ3		53.1	32	84		104	31	85		N/A	N/A	N/A		N/A	N/A	N/A		N/A	N/A	N/A
		QoI-HPEZ		55.2	25	68		109	27	70		N/A	N/A	N/A		N/A	N/A	N/A		N/A	N/A	N/A
		SZ3-QPET		93	97	430		156	80	455		329	71	339		26.6	46	172		163	29	497
		HPEZ-QPET		140	75	361		234	70	369		300	59	288		26.8	43	152		223	28	389
10^{-3}	10^{-4}	SPERR-QPET		238	65	213		254	44	131		351	41	148		39	41	57		202	22	62
		SZ3-OptZ-R	7	15.2	18	341	7	19.4	19	353	10	14.3	10	190	4	6.5	17	132	8	35.8	13	325
		HPEZ-OptZ-R	6	15.2	16	256	8	28.1	13	314	8	15.5	13	286	4	5.9	17	172	8	45.5	12	363
		SPERR-OptZ-R	6	24.4	10	106	9	30.3	5	91	13	22.7	4	87	8	8.3	5	42	9	58.6	5	102
		MGARD-QoI		5.9	3	4		N/A	N/A	N/A		N/A	N/A	N/A		N/A	N/A	N/A		N/A	N/A	N/A
		QoI-SZ3		13.6	31	75		34.5	31	82		N/A	N/A	N/A		N/A	N/A	N/A		N/A	N/A	N/A
		QoI-HPEZ		15.1	28	62		50	27	65		N/A	N/A	N/A		N/A	N/A	N/A		N/A	N/A	N/A
		SZ3-QPET		19	87	301		49.4	59	265		56.5	64	304		8.7	46	130		43.6	29	419
10^{-3}	10^{-4}	HPEZ-QPET		20.8	71	251		67.2	54	236		82.7	57	269		9.1	43	124		54.3	28	343
		SPERR-QPET		43.4	50	128		73.6	41	116		83.3	38	121		13	30	32		67.4	21	53

In Figure 5, QPET-integrated compressors have shown the optimized compression ratio (bit rate) in every test case and achieve substantial improvements in most. For example, on the Scale-LetKF dataset (Figure 5 (b)), SPERR-QPET has around 100%/50% compression ratio improvements over the best baseline (QoI-SZ3/HPEZ) when the error threshold for the square of data ($Q(x) = x^2$) is set as $1e-2/1e-3$ (relative, and all the followings are same). On other QoIs for which existing QoI-preserving data compressors fail to support, such as x^3 and $\tanh x$ (Figure 5 (c, d, f)), QPET makes great use of the point-wise data accuracy to preserve QoIs, significantly outperforming parameter-search-based baselines (*-OptZ-R) by up to over 1000% compression ratio improvement (on NYX dataset

with $Q(x) = x^3$) and higher throughputs. Regarding preserving $Q(x) = \sin 10x$, which is close to a linear function in several local data regions, point-wise data accuracy brings a limited contribution to the compression, so the compression ratios of QPET-integrated compressors are not outperforming parameter-search-based baselines. Nevertheless, QPET-integrated compressors can acquire the same compression ratios at much lower time costs, which is still a critical improvement for high-performance data processing tasks.

Multivariate QoI functions contain two sub-categories: QoI on data blocks within one data field and QoI vector data among multiple fields. In Figure 6, for better comparison with baselines (particularly QoI-SZ3/HPEZ and MGARD-QoI), we perform evaluations

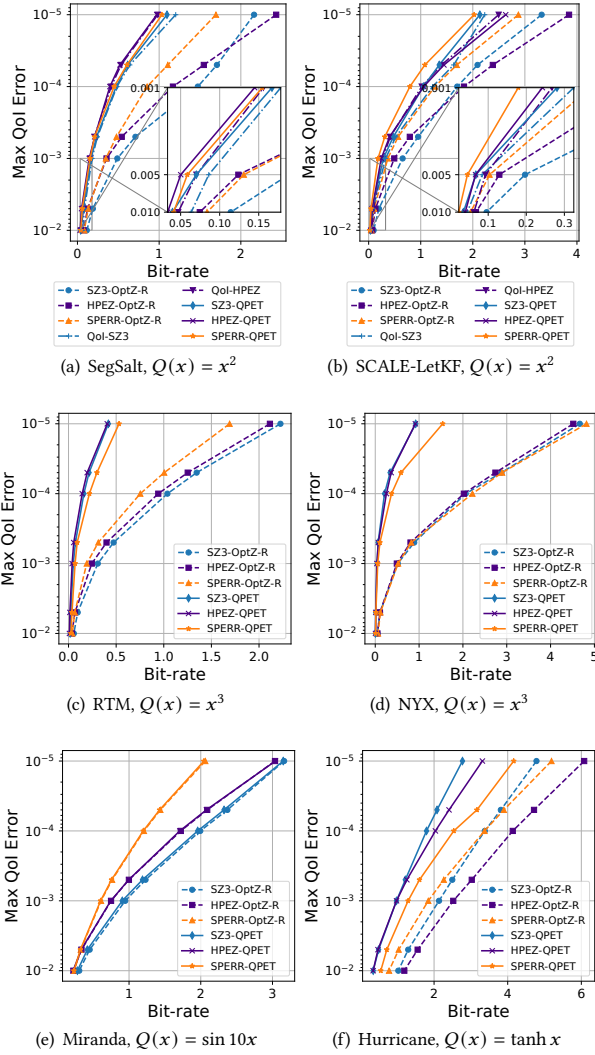


Figure 5: Bit rate and max QoI error plots (point-wise QoIs).

of QPET and the baselines on the average of x (data value itself), x^2 (square of the data value), and x^3 (cubic of the data value) on partitioned fixed-size data blocks. On Figure 6 (a) and (b), we can find that MGARD-QoI and QoI-SZ3 present unstable and suboptimal compression ratios, and QoI-HPEZ cannot outperform HPEZ-QPET and SPERR-QPET with an obvious gap between them. For $Q(X) = \frac{1}{n_b} \sum x^2$ and $Q(X) = \frac{1}{n_b} \sum x^3$, the 3 QPET-integrated compressors are the top-3 of all compressors, outperforming both parameter-search-based baselines and QoI-SZ3/HPEZ. On preserving the $Q(X) = \frac{1}{n_b} \sum x^2$ when compressing the RTM dataset (Figure 6 (c)), SPERR-QPET gains an aggregated compression ratio of ≈ 215 (bit rate of ≈ 0.15) under the QoI error tolerance of $1e-4$, which is 108% better than the best baseline SPERR-OptZ-R (CR ≈ 103). On preserving the $Q(X) = \frac{1}{n_b} \sum x^3$ when compressing the Hurricane dataset (Figure 6 (e)), HPEZ-QPET gains aggregated compression ratios of $\approx 121/45.4$ under the QoI error tolerance of $1e-3/1e-4$, which is 200%/220% of HPEZ-OptZ-R compression ratios ($\approx 60.9/20.6$).

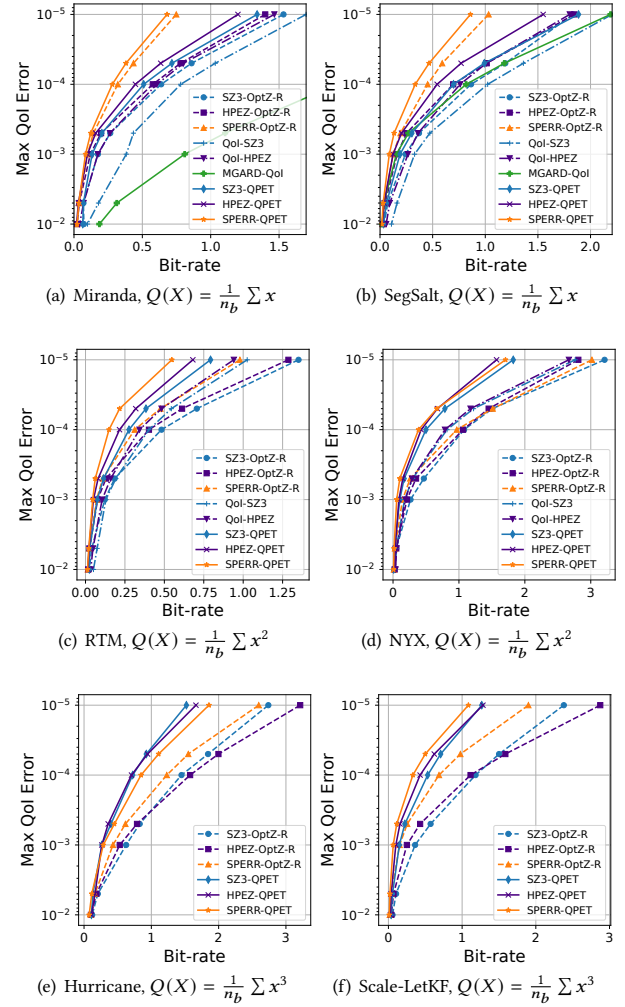


Figure 6: Bit rate and max QoI error plots of compressors on Regional QoIs ($n_b = 4^3$, i.e., average on $4 \times 4 \times 4$ blocks).

The improvement of compression ratios in those test cases is highly attributed to Theorem 5.4 because of the large number of variables (same as the block size 64) in the QoIs so that QPET can apply much higher point-wise QoI ($x/x^2/x^3$) error thresholds on each data point compared to the overall error threshold.

In Figure 7, we also evaluate several vector-style multivariate QoIs. On those vector QoIs, QPET also provides satisfactory compression ratio improvements under all test cases. On the Hurricane dataset, with $Q(x) = \sqrt{u^2 + v^2 + w^2}$ and $\tau = 1e-3$, SPERR-QPET improves the compression ratio of SPERR-OptZ-R by 25%. On the NYX dataset, with $Q(x) = u^2 + v^2 + w^2$ and $\tau = 1e-3/1e-4$, HPEZ-QPET improves the CR of the baselines by over 140%/80%.

6.2.4 Ablation study. At the end of our evaluation, we commit an ablation study of QPET design components, separately validating how different strategies described in Algorithm 2, 3, and 4 contribute to the performance of QPET.

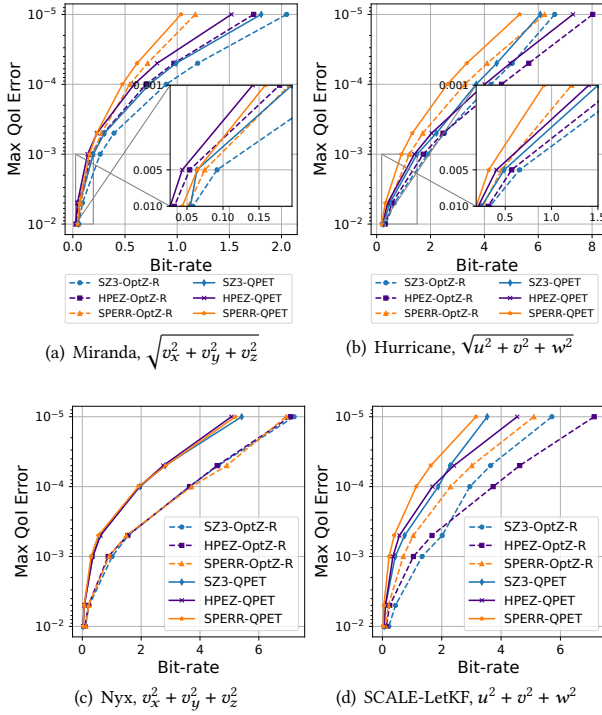


Figure 7: Bit rate and max QoI error plots (vector QoIs).




Figure 8: Ablation study on QPET Components.

On bounding the error of $Q(x) = x^3$, Figure 8 (a) and (b) compare QPET to its derivation QPET-A, which uses the analytical solution of inequality $|(\epsilon_i + x_i)^3 - x_i^3| \leq \tau$ instead of Algorithm 2 to determine

the point-wise error bound ϵ_0 for x_0 . On both of the datasets, QPET achieves the same compression ratio as QPET-A, meaning that Algorithm 2 performed well to replace analytic solutions. This fact indicates that the mechanism of QPET has worked well in determining the compression error bounds.

Next, we investigate the impact of Theorem 5.4 and the parameter c by evaluating SPERR-QPET on the SCALE-LetKF with the QoI of the average of x^2 on $4 \times 4 \times 4$ data blocks with different configurations. Under $c = 0$, SPERR-QPET deactivates Theorem 5.4, and larger c brings large error bound estimations from Theorem 5.4. Figure 8 (c) illustrates the corresponding results. Higher c does improve compression by proposing higher and still trustworthy estimations of point-wise error bounds. As the improvement of compression ratios becomes minor when c arrives at 3.0, we selected this value as the default in SPERR-QPET.

Last, in Figure 8 (d), we verify the stability and quality of QPET's global error-bound tuning mechanism (Algorithm 4). To this end, with $Q(x) = x^2$, we compared SZ3/HPEZ-QPET to QoI-SZ3/HPEZ. On the Miranda dataset, SZ3-QPET resolves the instability of QoI-SZ3 (which presents performance degradation on large error bounds due to suboptimal error bound selection), and HPEZ-QPET provides better compression ratios (5% ~ 20% across different QoI error thresholds) than QoI-HPEZ. According to these results and more similar ones we acquired, we can state that Algorithm 4 achieves state-of-the-art for global error-bound tuning.

7 CONCLUSION

To meet the significant requirements of Quantity-of-Interest (QoI) preserving in error-bounded lossy compression and address the critical limitations of existing solutions, we propose QPET, a versatile and portable framework for error-bounded lossy compression that enables the preservation of a broad range of QoI while significantly improving the performance and quality of existing state-of-the-arts. We integrate QPET into 3 representative error-bounded lossy compressors of different archetypes. Experimental results demonstrate that QPET is faster than all existing QoI-preserving error-bounded lossy compression solutions, leading to 2x to 10x compression speedups over the baselines of both parameter-search methods and direct QoI-preserving methods. QPET also delivers up to 1000% compression ratio improvements over existing QoI-preserving error-bounded lossy compression solutions. In future research, we will optimize QPET in the following three aspects: (1) Enable the support for more QoI formats in practical use cases; (2) Optimize the strategies for determining the point-wise and global error bounds in QoI-preserving; (3) Enhance the throughput to further reduce the computational cost for QoI preservation;

8 ACKNOWLEDGMENTS

This research was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research (ASCR), under contracts DE-AC02-06CH11357. This work was also supported by the National Science Foundation (Grant Nos. 2104023, 2311875, 2344717, OAC-2330367, OAC-2313122, OAC-2311756). We also acknowledge the computing resources provided by Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) – Purdue Anvil.

REFERENCES

- [1] [n.d.]. HDF5. <http://www.hdfgroup.org/HDF5>. Last Accessed: 2025-05-25.
- [2] [n.d.]. nvCOMP. <https://github.com/NVIDIA/nvcomp>. Last Accessed: 2025-05-25.
- [3] 2020. EXAALT: Molecular Dynamics at the Exascale. <https://www.exascaleproject.org/wp-content/uploads/2019/10/EXAALT.pdf>. Online, Last Accessed: 2025-05-25.
- [4] 2021. Team at Princeton Plasma Physics Laboratory employs DOE supercomputers to understand heat-load width requirements of future ITER device. <https://www.olcf.ornl.gov/2021/02/18/scientists-use-supercomputers-to-study-reliable-fusion-reactor-design-operation>. Online, Last Accessed: 2025-05-25.
- [5] Nitin Agrawal and Ashish Vulimiri. 2017. Low-latency analytics on colossal data streams with summarystore. In *Proceedings of the 26th Symposium on Operating Systems Principles*. 647–664.
- [6] Mark Ainsworth, Ozan Tugluk, Ben Whitney, and Scott Klasky. 2019. Multilevel techniques for compression and reduction of scientific data-quantitative control of accuracy in derived quantities. *SIAM Journal on Scientific Computing* 41, 4 (2019), A2146–A2171.
- [7] Jyrki Alakuijala, Andrea Farruggia, Paolo Ferragina, Eugene Kliuchnikov, Robert Obyrk, Zoltan Szabadka, and Lode Vandevenne. 2018. Brotli: A general-purpose data compressor. *ACM Transactions on Information Systems (TOIS)* 37, 1 (2018), 1–30.
- [8] Andrei Arion, Angela Bonifati, Ioana Manolescu, and Andrea Pugliese. 2007. XQueC: A query-conscious compressed XML database. *ACM Transactions on Internet Technology (TOIT)* 7, 2 (2007), 10–es.
- [9] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. 2021. Overview of the versatile video coding (VVC) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology* 31, 10 (2021), 3736–3764.
- [10] Zhiyuan Chen, Johannes Gehrke, and Flip Korn. 2001. Query optimization in compressed database systems. In *Proceedings of the 2001 ACM SIGMOD international conference on Management of data*. 271–282.
- [11] Yann Collet. 2015. Zstandard – Real-time data compression algorithm. <http://facebook.github.io/zstd/> (2015).
- [12] L Peter Deutsch. 1996. GZIP file format specification version 4.3.
- [13] Hazem Elmeleegy, Ahmed K. Elmagarmid, Emmanuel Cecchet, Walid G. Aref, and Willy Zwaenepoel. 2009. Online Piece-Wise Linear Approximation of Numerical Streams with Precision Guarantees. *Proc. VLDB Endow.* 2, 1 (Aug. 2009), 145–156.
- [14] Ian Foster, Mark Ainsworth, Bryce Allen, Julie Bessac, Franck Cappello, Jong Youl Choi, Emil Constantinescu, Philip E Davis, Sheng Di, Wendy Di, Hanqi Guo, Scott Klasky, Kerstin Kleese Van Dam, Tahsin Kurc, Qing Liu, Abid Malik, Kshiti Mehta, Klaus Mueller, Todd Munson, George Ostouchov, Manish Parashar, Tom Peterka, Line Pouchard, Dingwen Tao, Ozan Tugluk, Stefan Wild, Matthew Wolf, Justin M. Wozniak, Wei Xu, and Shinjae Yoo. 2017. Computing just what you need: Online data analysis and reduction at extreme scales. In *European conference on parallel processing*. Springer, 3–19.
- [15] Qian Gong, Xin Liang, Ben Whitney, Jong Youl Choi, Jieyang Chen, Lipeng Wan, Stéphane Ethier, Seung-Hoe Ku, R Michael Churchill, C-S Chang, et al. 2021. Maintaining trust in reduction: Preserving the accuracy of quantities of interest for lossy compression. In *Smoky Mountains Computational Sciences and Engineering Conference*. Springer, 22–39.
- [16] Qian Gong, Xin Liang, Ben Whitney, Jong Youl Choi, Jieyang Chen, Lipeng Wan, Stéphane Ethier, Seung-Hoe Ku, R Michael Churchill, C-S Chang, et al. 2021. Maintaining trust in reduction: Preserving the accuracy of quantities of interest for lossy compression. In *Smoky Mountains Computational Sciences and Engineering Conference*. Springer, 22–39.
- [17] Wassily Hoeffding. 1963. Probability Inequalities for Sums of Bounded Random Variables. *J. Amer. Statist. Assoc.* 58, 301 (1963), 13–30. <http://www.jstor.org/stable/2282952>
- [18] Lawrence Ibarria, Peter Lindstrom, Jarek Rossignac, and Andrzej Szymczak. 2003. Out-of-core compression and decompression of large n-dimensional scalar fields. In *Computer Graphics Forum*, Vol. 22. Wiley Online Library, 343–348.
- [19] Søren Keiser Jensen, Torben Bach Pedersen, and Christian Thomsen. 2018. Modelardb: Modular model-based time series management with spark and cassandra. *Proceedings of the VLDB Endowment* 11, 11 (2018), 1688–1701.
- [20] Pu Jiao, Sheng Di, Hanqi Guo, Kai Zhao, Jiannan Tian, Dingwen Tao, Xin Liang, and Franck Cappello. 2022. Toward Quantity-of-Interest Preserving Lossy Compression for Scientific Data. *Proceedings of the VLDB Endowment* 16, 4 (2022), 697–710.
- [21] Søren Keiser Jensen, Torben Bach Pedersen, and Christian Thomsen. 2019. Scalable Model-Based Management of Correlated Dimensional Time Series in Modelardb+. *arXiv e-prints* (2019), arXiv–1903.
- [22] Fabian Knorr, Peter Thoman, and Thomas Fahringer. 2021. ndzip: A high-throughput parallel lossless compressor for scientific data. In *2021 Data Compression Conference (DCC)*. IEEE, 103–112.
- [23] Fabian Knorr, Peter Thoman, and Thomas Fahringer. 2021. ndzip-gpu: efficient lossless compression of scientific floating-point data on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–14.
- [24] I. Lazaridis and S. Mehrotra. 2003. Capturing sensor-generated time series with quality guarantees. In *Proceedings 19th International Conference on Data Engineering (Cat. No.03CH37405)*. 429–440. <https://doi.org/10.1109/ICDE.2003.1260811>
- [25] Shaomeng Li, Peter Lindstrom, and John Clyne. 2023. Lossy scientific data compression with SPERR. In *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 1007–1017.
- [26] Panagiotis Liakos, Katia Papakonstantinou, and Yannis Kotidis. 2022. Chimp: efficient lossless floating point compression for time series databases. *Proceedings of the VLDB Endowment* 15, 11 (2022), 3058–3070.
- [27] Xin Liang, Sheng Di, Franck Cappello, Mukund Raj, Chunhui Liu, Kenji Ono, Zizhong Chen, Tom Peterka, and Hanqi Guo. 2022. Toward feature-preserving vector field compression. *IEEE Transactions on Visualization and Computer Graphics* 29, 12 (2022), 5434–5450.
- [28] Xin Liang, Sheng Di, Dingwen Tao, Zizhong Chen, and Franck Cappello. 2018. An efficient transformation scheme for lossy data compression with point-wise relative error bound. In *2018 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 179–189.
- [29] Xin Liang, Sheng Di, Dingwen Tao, Sihuan Li, Shaomeng Li, Hanqi Guo, Zizhong Chen, and Franck Cappello. 2018. Error-Controlled Lossy Compression Optimized for High Compression Ratios of Scientific Datasets. In *2018 IEEE International Conference on Big Data*. IEEE.
- [30] Xin Liang, Hanqi Guo, Sheng Di, Franck Cappello, Mukund Raj, Chunhui Liu, Kenji Ono, Zizhong Chen, and Tom Peterka. 2020. Toward Feature-Preserving 2D and 3D Vector Field Compression.. In *PacificVis*. 81–90.
- [31] Xin Liang, Kai Zhao, Sheng Di, Sihuan Li, Robert Underwood, Ali M. Gok, Jiannan Tian, Junjing Deng, Jon C. Calhoun, Dingwen Tao, Zizhong Chen, and Franck Cappello. 2023. SZ3: A Modular Framework for Composing Prediction-Based Error-Bounded Lossy Compressors. *IEEE Transactions on Big Data* 9, 2 (2023), 485–498. <https://doi.org/10.1109/TBDDATA.2022.3201176>
- [32] Peter Lindstrom. 2014. Fixed-rate compressed floating-point arrays. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 2674–2683.
- [33] Peter G Lindstrom et al. 2017. *Fpzip*. Technical Report. Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States).
- [34] Chunwei Liu, Hao Jiang, John Paparrizos, and Aaron J Elmore. 2021. Decomposed bounded floats for fast compression and queries. *Proceedings of the VLDB Endowment* 14, 11 (2021), 2586–2598.
- [35] Jinyang Liu, Sheng Di, Kai Zhao, Xin Liang, Zizhong Chen, and Franck Cappello. 2022. Dynamic quality metric oriented error bounded lossy compression for scientific datasets. In *2022 SC22: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE Computer Society, 892–906.
- [36] Jinyang Liu, Sheng Di, Kai Zhao, Xin Liang, Sian Jin, Zizhe Jian, Jiajun Huang, Shixun Wu, Zizhong Chen, and Franck Cappello. 2024. High-performance effective scientific error-bounded lossy compression with auto-tuned multi-component interpolation. *Proceedings of the ACM on Management of Data* 2, 1 (2024), 1–27.
- [37] Kiyoshi Masui, Mandana Amiri, Liam Connor, Meiling Deng, Mateus Fandino, Carolin Höfer, Mark Halpern, David Hanna, Adam D Hincks, Gary Hinshaw, et al. 2015. A compression scheme for radio data in high performance computing. *Astronomy and Computing* 12 (2015), 181–190.
- [38] Tuomas Pelkonen et al. 2015. Gorilla: A Fast, Scalable, in-Memory Time Series Database. *Proc. VLDB Endow.* 8, 12 (Aug. 2015), 1816–1827.
- [39] X Carol Song, Preston Smith, Rajesh Kalyanam, Xiao Zhu, Eric Adams, Kevin Colby, Patrick Finnegan, Erik Gough, Elizabeth Hillery, Rick Irvine, et al. 2022. Anvil-system architecture and experiences from deployment and early user operations. In *Practice and experience in advanced research computing*. 1–9.
- [40] Zhaoyuan Su, Sheng Di, Ali Murat Gok, Yue Cheng, and Franck Cappello. 2022. Understanding impact of lossy compression on derivative-related metrics in scientific datasets. In *2022 IEEE/ACM 8th International Workshop on Data Analysis and Reduction for Big Scientific Data (DRBSD)*. IEEE, 44–53.
- [41] Vivienne Sze, Madhukar Budagavi, and Gary J Sullivan. 2014. High efficiency video coding (HEVC). In *Integrated circuit and systems, algorithms and architectures*. Vol. 39. Springer, 40.
- [42] Dingwen Tao, Sheng Di, Zizhong Chen, and Franck Cappello. 2017. Significantly improving lossy compression for scientific data sets based on multidimensional prediction and error-controlled quantization. In *2017 IEEE International Parallel and Distributed Processing Symposium*. IEEE, 1129–1139.
- [43] Dingwen Tao, Sheng Di, Hanqi Guo, Zizhong Chen, and Franck Cappello. 2019. Z-checker: A framework for assessing lossy compression of scientific data. *The International Journal of High Performance Computing Applications* 33, 2 (2019), 285–303. <https://doi.org/10.1177/1094342017737147>
- [44] David S Taubman and Michael W Marcellin. 2002. JPEG2000: Standard for interactive imaging. *Proc. IEEE* 90, 8 (2002), 1336–1357.

- [45] Jiannan Tian, Sheng Di, Xiaodong Yu, Cody Rivera, Kai Zhao, Sian Jin, Yunhe Feng, Xin Liang, Dingwen Tao, and Franck Cappello. 2021. cuSZ (x): Optimizing Error-Bounded Lossy Compression for Scientific Data on GPUs. *CoRR* (2021).
- [46] Jiannan Tian, Sheng Di, Kai Zhao, Cody Rivera, Megan Hickman Fulp, Robert Underwood, Sian Jin, Xin Liang, Jon Calhoun, Dingwen Tao, and Franck Cappello. 2020. cuSZ: An Efficient GPU-Based Error-Bounded Lossy Compression Framework for Scientific Data. In *Proceedings of the ACM International Conference on Parallel Architectures and Compilation Techniques* (Virtual Event, GA, USA) (PACT '20). Association for Computing Machinery, New York, NY, USA, 3–15. <https://doi.org/10.1145/3410463.3414624>
- [47] Robert Underwood, Jon C Calhoun, Sheng Di, Amy Apon, and Franck Cappello. 2022. OptZConfig: Efficient Parallel Optimization of Lossy Compression Configuration. *IEEE Transactions on Parallel and Distributed Systems* (2022).
- [48] Robert Underwood, Sheng Di, Jon C Calhoun, and Franck Cappello. 2020. Fraz: A generic high-fidelity fixed-ratio lossy compression framework for scientific floating-point data. In *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 567–577.
- [49] Roman Vershynin. 2018. *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge university press.
- [50] Martin J Wainwright. 2019. *High-dimensional statistics: A non-asymptotic view-point*. Vol. 48. Cambridge university press.
- [51] Gregory K Wallace. 1991. The JPEG still picture compression standard. *Commun. ACM* 34, 4 (1991), 30–44.
- [52] Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. 2003. Overview of the H. 264/AVC video coding standard. *IEEE Transactions on circuits and systems for video technology* 13, 7 (2003), 560–576.
- [53] Xuan Wu, Qian Gong, Jieyang Chen, Qing Liu, Norbert Podhorszki, Xin Liang, and Scott Klasky. 2024. Error-controlled Progressive Retrieval of Scientific Data under Derivable Quantities of Interest. In *2024 SC24: International Conference for High Performance Computing, Networking, Storage and Analysis SC*. IEEE Computer Society, 1368–1383.
- [54] Mingze Xia, Sheng Di, Franck Cappello, Pu Jiao, Kai Zhao, Jinyang Liu, Xuan Wu, Xin Liang, and Hanqi Guo. 2024. Preserving Topological Feature with Sign-of-Determinant Predicates in Lossy Compression: A Case Study of Vector Field Critical Points. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 4979–4992.
- [55] Lin Yan, Xin Liang, Hanqi Guo, and Bei Wang. 2023. TopoSZ: Preserving topology in error-bounded lossy compression. *IEEE Transactions on Visualization and Computer Graphics* (2023).
- [56] Feng Zhang, Zaifeng Pan, Yanliang Zhou, Jidong Zhai, Xipeng Shen, Onur Mutlu, and Xiaoyong Du. 2021. G-TADOC: Enabling efficient GPU-based text analytics without decompression. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 1679–1690.
- [57] Feng Zhang, Jidong Zhai, Xipeng Shen, Onur Mutlu, and Wenguang Chen. 2018. Efficient document analytics on compressed data: Method, challenges, algorithms, insights. *Proceedings of the VLDB Endowment* 11, 11 (2018), 1522–1535.
- [58] Kai Zhao, Sheng Di, Maxim Dmitriev, Thierry-Laurent D. Tonellot, Zizhong Chen, and Franck Cappello. 2021. Optimizing Error-Bounded Lossy Compression for Scientific Data by Dynamic Spline Interpolation. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. 1643–1654. <https://doi.org/10.1109/ICDE51399.2021.00145>