

FunM²C: A Filter for Uncertainty Visualization of Multivariate Data on Multi-Core Devices

Gautam Hari*

Indiana University

Nrushad Joshi[†]

Indiana University

Zhe Wang[‡]

Oak Ridge National Laboratory

Qian Gong[§]

Oak Ridge National Laboratory

Dave Pugmire[¶]

Oak Ridge National Laboratory

Kenneth Moreland^{||}

Oak Ridge National Laboratory

Chris R. Johnson ^{**}

SCI Institute University of Utah

Scott Klasky^{††}

Oak Ridge National Laboratory

Norbert Podhorszki^{‡‡}

Oak Ridge National Laboratory

Tushar M. Athawale

Oak Ridge National Laboratory

ABSTRACT

Uncertainty visualization is an emerging research topic in data visualization because neglecting uncertainty in visualization can lead to inaccurate assessments. In this paper, we study the propagation of multivariate data uncertainty in visualization. Although there have been a few advancements in probabilistic uncertainty visualization of multivariate data, three critical challenges remain to be addressed. First, the state-of-the-art probabilistic uncertainty visualization framework is limited to bivariate data (two variables). Second, existing uncertainty visualization algorithms use computationally intensive techniques and lack support for cross-platform portability. Third, as a consequence of the computational expense, integration into production visualization tools is impractical. In this work, we address all three issues and make a threefold contribution. First, we take a step to generalize the state-of-the-art probabilistic framework for bivariate data to multivariate data with an arbitrary number of variables. Second, through utilization of VTK-m's shared-memory parallelism and cross-platform compatibility features, we demonstrate acceleration of multivariate uncertainty visualization on different many-core architectures, including OpenMP and AMD GPUs. Third, we demonstrate the integration of our algorithms with the ParaView software. We demonstrate the utility of our algorithms through experiments on multivariate simulation data with three and four variables.

Index Terms: Uncertainty, multivariate visualization.

1 INTRODUCTION

Visualizations play a vital role in large-scale data analysis and enable scientists to make informed decisions effectively and efficiently. All scientific data contains varying degrees of uncertainty due to instrument limitations, approximations in numerical and physical models, and data processing techniques [6]. Ignoring the data uncertainty in visualization can result in misleading analysis. Thus, visualizing data uncertainty has been an active area of research that has resulted in numerous breakthroughs [5, 11, 18].

*e-mail: gautamhari@outlook.com

[†]e-mail: nrushad2001@gmail.com

[‡]e-mail: godenwangzhe@gmail.com

[§]e-mail: gongq@ornl.gov

[¶]e-mail: pugmire@ornl.gov

^{||}e-mail: morelandkd@ornl.gov

^{**}e-mail: crj@sci.utah.edu

^{††}e-mail: klasky@ornl.gov

^{‡‡}e-mail: pnorbert@ornl.gov

e-mail: athawaletm@ornl.gov

The fiber surface [7] and feature level-set [10] techniques have proven to be powerful tools for the analysis of complex multivariate data from a large number of domains, including molecular dynamics, astrophysics, and fluid dynamics. State-of-the-art work by Zheng and Sadlo [24], Sane et al. [20], and Athawale et al. [3] have given scientists previously unattainable insight into the impact of uncertainty on the analysis of multivariate data. In this work, we identify and address three limitations in the state of the art.

First, the probabilistic frameworks proposed by Zheng and Sadlo [24] and Athawale et al. [3] are currently limited to bivariate data. Simulation or experimental data are often multivariate with more than two variables. Thus, in our first contribution, we take a step to extend the probabilistic framework for bivariate data to multivariate data with more than two variables. Although the previous work by Sane et al. [20] studied uncertainty in multivariate data with more than two variables, their work derived confidence intervals for depicting data uncertainty. Our work significantly differs from Sane et al.'s work in that we compute spatial probabilities of multivariate data features, similar to the prior work [3, 24].

Second, although the acceleration of uncertainty visualization has been previously demonstrated on GPUs for univariate data [22] and bivariate data [3], it has not been studied and evaluated yet for multivariate data with more than two variables. Thus, in our second contribution, we present platform-portable and GPU-accelerated implementation for uncertainty visualization of multivariate data with more than two variables using VTK-m [15]. We refer to our filter as FunM²C (filter for **u**ncertainty visualization of **m**ultivariate data on **multi**-core devices), given many similarities of our approach compared to the one by Wang et al. [22].

Third, partly due to the high computational costs, uncertainty visualization is often not available for use in production visualization tools such as VisIt [8] and ParaView [1]. Wang et al. [22] previously demonstrated how their VTK-m uncertainty visualization algorithms for univariate data integrate well with the ParaView software for wider accessibility. Motivated from the work by Wang et al., in our third contribution, we demonstrate the integration of our VTK-m multivariate uncertainty visualization algorithms in the ParaView software for accessibility to a wide audience.

2 BACKGROUND AND RELATED WORK

Here, we present an overview of fiber uncertainty visualization algorithm for bivariate data proposed by Zheng and Sadlo [24] since

This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a non-exclusive, paid up, irrevocable, world-wide license to publish or reproduce the published form of the manuscript, or allow others to do so, for U.S. Government purposes. The DOE will provide public access to these results in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

we extend their work to data with more than two variables. For a bivariate function $f: \mathbb{R}^n \rightarrow \mathbb{R}^2$, the spatial domain $\mathcal{D} \subset \mathbb{R}^n$ is mapped to an attribute space $\mathcal{A} \subset \mathbb{R}^2$. A *fiber* is then defined as the preimage of a point $a = (a_1, a_2) \in \mathcal{A}$, which represents an intersection of isosurfaces [13] for isovalues $\mathcal{A}_1 = a_1$ and $\mathcal{A}_2 = a_2$. A continuum of points (i.e., a curve) in \mathcal{A} corresponds to a continuum of fibers in \mathcal{D} , which forms a fiber surface. A polygon is generally used to specify the values of user interest for the extraction of fiber surfaces. Such a polygon is referred to as a trait \mathcal{T} .




Figure 1: Illustration of *fiber probability* computation at a single grid vertex for uncertain bivariate data. The *fiber probability* is computed as the intersection area of a trait \mathcal{T} and a blue rectangle representing uncertain data \mathcal{U} (with intersection shown by the dotted orange border) divided by the area of \mathcal{U} at a vertex. In Layman’s terms, the fiber probability corresponds to a fraction of uncertain data at a grid vertex residing inside the user-specified trait.

Fig. 1 summarizes the process for computing *fiber probability* or uncertainty [3, 24] at each grid vertex of domain. The fiber probability essentially denotes the probability of uncertain data at a grid vertex to lie inside a user-defined trait \mathcal{T} . Our work is currently limited to the uniformly-distributed data uncertainty with independence assumption between variables, but we discuss the ramifications of the uniform noise and independence assumptions later in Sec. 5. As represented in Fig. 1, the blue rectangle represents uniformly-distributed data uncertainty (denoted by \mathcal{U} with bounds $[a_1, b_1] \subset \mathcal{A}_1$ and $[a_2, b_2] \subset \mathcal{A}_2$) at a single grid vertex. In a closed-form method, the fiber probability is computed as the ratio of the area of intersection between trait \mathcal{T} (with bounds $[l_1, h_1] \subset \mathcal{A}_1$ and $[l_2, h_2] \subset \mathcal{A}_2$) and uncertain data support \mathcal{U} (i.e., the area of a rectangle enclosed by dotted orange lines in Fig. 1) and the total area of uncertain data support. This work expands the previous work on bivariate data [24] to general multivariate data $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ with m -dimensional attribute space. Further, we implement our multivariate uncertainty algorithms using the VTK-m library for faster performance using GPUs and seamless integration with ParaView, similar to the prior work for univariate [22] and bivariate [3] data.

3 METHODS

In this section, we describe our algorithm for visualizing uncertainty in multivariate data with more than two variables (Sec. 3.1) and its VTK-m implementation (Sec. 3.2).

3.1 Generalizing Bivariate Fiber Uncertainty to Data with More than Two Variables

We generalize the fiber uncertainty visualization framework originally designed for two variables [3] to m variables ($m \geq 2$) assuming uniform-distributed uncertainty per variable. Mathematically, each variable X_i representing uncertainty, where $i \in \{1, \dots, m\}$, is assumed to follow a uniform distribution $X_i \sim U[a_i, b_i]$. For each variable X_i , $[a_i, b_i] \subset \mathcal{A}_i$. Because of our data independence assumption, the joint distribution of all random variables X_i is given by the product of individual distributions. Thus, the joint distribution is also a uniform distribution over m -dimensional hyperrectangular support (\mathcal{U}) with bounds given by the ranges $\mathcal{U} = [[a_1, b_1], \dots, [a_m, b_m]]$. The m -dimensional trait \mathcal{T} is characterized by user-defined limits $[l_i, h_i]$ for each dimension $i \in \{1, \dots, m\}$, which corresponds to a m -dimensional hyperrectangle.

To compute the probability of trait satisfaction at each grid point, we calculate the intersection between the m -dimensional hyperrectangles corresponding to the uncertain data distribution support \mathcal{U} and trait \mathcal{T} . Let \mathcal{I} denote the hyperrectangle corresponding to the intersection of \mathcal{U} and \mathcal{T} . The lowest and highest limits of \mathcal{I} then correspond to m -dimensional vectors $[\max(a_1, l_1), \dots, \max(a_m, l_m)]$ and $[\min(b_1, h_1), \dots, \min(b_m, h_m)]$, respectively. The logic underneath the computation of the lowest and highest limits of \mathcal{I} is directly derived from a 2-dimensional case shown in Fig. 1. The volume of the hyperrectangle \mathcal{I} is then computed using a product $P = \prod_{i=1}^m [\min(b_i, h_i) - \max(a_i, l_i)]$. Finally, the probability of trait satisfaction at a grid vertex is computed as the ratio of the volume of hyperrectangle \mathcal{I} (i.e., P) and the volume of hyperrectangle \mathcal{U} [i.e., $\prod_{i=1}^m (b_i - a_i)$].




Figure 2: Illustration of uncertainty analysis for trivariate data at a single grid point. The red box represents bounds indicating uncertain data, while the green box highlights a user-specified trait. The orange shading illustrates the intersection volume, showing how much uncertain data coincides with the specified traits.

Figure 2 demonstrates our approach for a three variable case. The green box in the figure represents the user-selected trait (\mathcal{T}) corresponding to the ranges $[l_i, h_i]$ with $i \in \{1, 2, 3\}$. Each variable X_i follows a uniform distribution $X_i \sim U[a_i, b_i]$ with $i \in \{1, 2, 3\}$, creating a joint uniform distribution over the support (\mathcal{U}) depicted in red. The intersection volume (\mathcal{I}) illustrated in orange has the lowest limits as $\max(a_i, l_i)$ and the highest limits as $\min(h_i, b_i)$ in the i ’th dimension. Thus, the probability of trait satisfaction at a grid point can be computed as $\frac{\text{Volume}(\mathcal{I})}{\text{Volume}(\mathcal{U})}$.

3.2 Integration Into VTK-m and Paraview Software

In a prior work by Wang et al. [22], an isosurface algorithm for uncertain data (FunMC²) was implemented in VTK-m [15] and integrated into ParaView. Our approach extends the capabilities of univariate uncertainty visualization proposed by Wang et al. to multivariate data with more than two variables through the development of a specialized worklet within VTK-m called *MultivariateUncertaintyClosedForm*. The *MultivariateUncertaintyClosedForm* worklet leverages a deterministic, closed-form solution proposed in Sec. 3.1, which calculates probabilities based on analytical expressions derived from the distribution parameters. The worklet

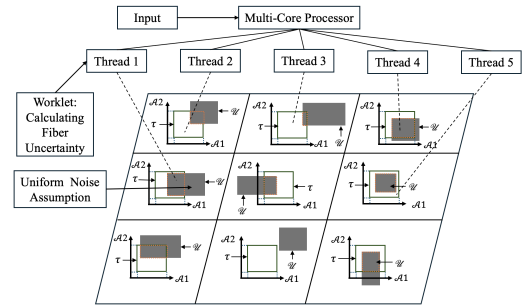


Figure 3: Illustration of VTK-m’s multi-threading process that reduces the algorithm’s execution period. Each parallelogram above corresponds to a grid vertex and indicates data uncertainty (\mathcal{U}) with respect to a trait (\mathcal{T}). The process generalizes to multivariate data with an arbitrary number of dimensions.

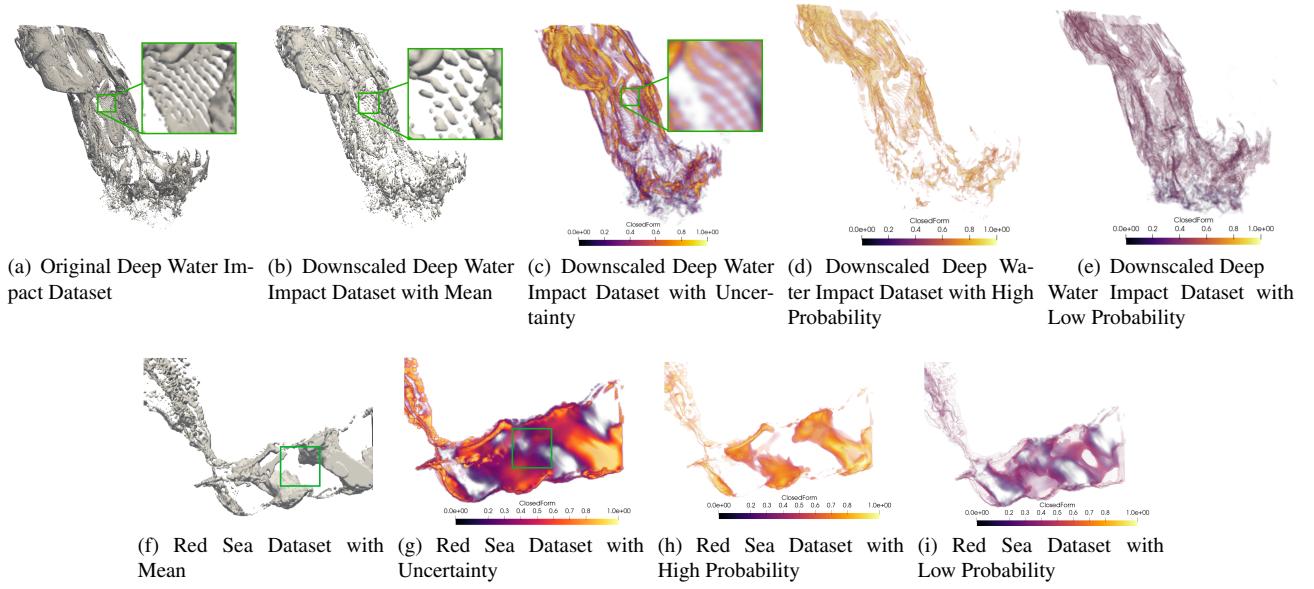


Figure 4: The top row visualizes the Deep Water Impact dataset with three variables for the trait (\mathcal{T}) v_2 :(0.1, 0.5), tev :(0.2, 0.25), and prs :(30000, 400000). Fig.(a) visualizes the original resolution ($460 \times 280 \times 120$) dataset, which is treated as the reference. Fig.(b) shows the mean-field surface, and Fig.(c) shows our probabilistic visualization. The probabilistic visualization recovers important topological structures (e.g., rib-like structure in the inset views) that appears to be broken in the mean-field visualization. The bottom row depicts results for the Red Sea dataset with four variables for the trait $curlZ$:(-15.0, -0.1), Vorticity:(0.0, 15.0), Velocity:(0.0, 0.4), and Temperature:(10.0, 29.0). The results for the Red Sea are similar to the Deep Water Impact dataset, however, the reference is not known in the Red Sea dataset. The Green rectangles point out feature inconsistencies between the mean-field in Fig.(f) and our probabilistic visualization in Fig.(g). Both Fig.(d) and Fig.(h) illustrate the regions with probability greater than 0.5, and Fig. (e) and Fig. (i) show the regions with probability smaller than 0.5.

inherits from VTK-m’s foundational class, `WorkletMapField` [14], which facilitates the systematic application of computations across multi-dimensional datasets. At each grid point, these worklets operate in conjunction with uncertain multivariate data representations. Figure 3 visually illustrates the execution process within VTK-m, showcasing its capability to manage computations concurrently across multiple grid points. The figure shows bivariate data being parallelized but the same process still applies to multivariate data with more than two variables (described in Sec. 3.1).

Moreover, the integration of `MultivariateUncertaintyClosed-Form` worklet into VTK-m extends beyond mere algorithmic implementation. Utilizing VTK-m’s plugin architecture, this worklet seamlessly integrates into ParaView [1], a widely-used visualization tool. Within ParaView, users gain access to new filters that execute our multivariate uncertainty visualization algorithms. The graphical user interface (GUI) of ParaView enables users to have direct control over algorithmic parameters, facilitating interactive exploration, and refinement of visualization results.

4 RESULTS

Evaluation results are presented in three parts. Sec. 4.1 describes the platforms and datasets used for evaluation. Sec. 4.2 presents uncertainty visualization results for multivariate data obtained with the proposed `FunM2C` filter. Sec. 4.3 validates the correctness and demonstrates the enhanced performance of our closed-form approach (Sec. 3.1) using OpenMP and AMD GPUs.

4.1 Platform and Datasets

Implementing multivariate uncertainty visualization algorithms using VTK-m provides flexibility to run the algorithms on different backends based on various multi-core devices. The Serial, OpenMP, and HIP backends are evaluated on the “Frontier” supercomputing system [2]. The details of Datasets are listed as follows:

Deep Water Impact is created by a cosmology simulation [16] that simulates the Impact of an asteroid. The original data resolution is $460 \times 280 \times 240$. The data is partitioned into blocks of

size $2 \times 2 \times 2$, and uncertainty in each block is modeled with the uniform distribution using the `hixel` idea [21]. We assess the visualizations for three fields, including “`v02`” representing the volume of water fraction, “`tev`” representing the temperature in electron-volt (eV), and the field “`prs`” representing the pressure in microbars (μbar).

Red Sea dataset [23] is publicly accessible through IEEE SciVis contest 2020. The size of each ensemble member is $500 \times 500 \times 50$ in a uniform grid, and there are 20 ensemble members in total. We assess the four attributes representing vorticity, curl, velocity, and temperature [20] for visualization of flow simulations.

Supernova is created by a cosmology simulation [19] that simulates the core-collapse of a supernova. The size of each ensemble member is resampled to $100 \times 100 \times 100$ in a uniform grid, and there are 64 ensemble members in total. We assess the fields “Nickel” and “Iron” in the dataset.

4.2 Visualization of Multivariate Probability

Fig. 4 visualizes the results of the application of our multivariate uncertainty filter (`FunM2C`) to the Deep Water Impact (top row) and Red Sea (bottom row) datasets. The uncertainty visualizations shown are of volume renderings [9] of probability volumes derived using techniques proposed in Sec. 3 with the `inferno` color map. In particular, transparent deep purple regions indicate the positions of lower probability, whereas a less transparent bright yellow regions indicate the positions of higher probabilities.

In the case of the Deep Water Impact dataset, the high-resolution dataset ($460 \times 280 \times 240$) is visualized in Fig. 4(a), which is treated as the reference. Fig. 4(b) visualizes the results for the mean field, which is computed per data block of size $2 \times 2 \times 2$. Fig. 4(c) visualizes probabilistic field computed with our algorithms by modeling uncertainty with uniform distribution over block for each variable. In visualizing of the probabilistic field, low-probability regions are dark purple, and high-probability regions are yellow.

For the Deep Water Impact dataset, the mean-field visualization (Fig. 4(b)) produces a significant number of cracks and loses im-

portant features compared to the reference (Fig. 4(a)). In contrast, the probabilistic field computed with our methods (Fig. 4(c)) recovers important topological features (e.g., rib-like features shown in the inset views) that are lost in the mean field when compared to the reference. Fig. 4(d) and Fig. 4(e) visualize the high probability (probability ≥ 0.5) and low probability (probability < 0.5) features, respectively. The results for the Red Sea dataset in Fig. 4 are similar to those for the asteroid dataset. In the Red Sea dataset, we do not know the reference. However, we could see prominent feature differences between the mean-field (Fig. 4(f)) and our probabilistic visualizations (Fig. 4(g)), as indicated by the green boxes.

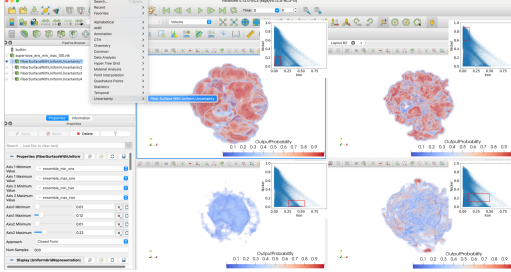


Figure 5: Uncertainty visualization of bivariate Supernova [19] using the proposed FunM²C filter as a ParaView plugin. Scatter plots associated with each render view are added through post image processing, which is not the capability of the plugin.

We export our FunM²C filter as a plugin to showcase its usage inside ParaView and its potential utility to the broader scientific community. Fig. 5 illustrates the ParaView GUI interface utilizing FunM²C filter as a backend through the plugin. Users can conveniently explore different regions of the data by specifying bounds denoting traits through one interface and study respective probabilistic multivariate visualization.

4.3 Validation and Performance of FunM²C Filter

In this experiment, we initially validate the correctness of the FunM²C implementation by comparing results from the proposed closed-form framework (Sec. 3.1) with the Monte Carlo technique. The Monte Carlo technique involves generating uniformly distributed random samples within an m -dimensional hyperrectangle representing uncertain data ranges ($\mathcal{U} = [[a_1, b_1], \dots, [a_m, b_m]]$) and computing the ratio of samples that fall within a m -dimensional hyperrectangle corresponding to a user-defined trait ($\mathcal{T} = [[l_1, h_1], \dots, [l_m, h_m]]$), providing a statistical estimate of likelihood of trait satisfaction at a grid point. Due to the simplicity of the Monte Carlo approach, it serves as a benchmark to verify the correctness of closed-form implementation.

To validate our closed-form implementation, we ensure that the Monte Carlo solution converges to the results from the closed-form filter as the number of samples increases. In particular, increasing the number of Monte Carlo samples from 100 to 1000 reduces the mean squared difference between our closed-form and Monte Carlo solutions from 2.26×10^{-4} to 1.60×10^{-4} for the Red Sea dataset, and from 0.03 to 0.017 for the Deep Sea Impact dataset with four variables. The maximum error between any given grid point also decreases to negligible numbers, which can be attributed to the randomness of the sampling. These findings confirm that as the sample size grows, the Monte Carlo solution converges toward our closed-form solutions, validating our implementation’s accuracy.

We assessed the performance of our VTK-m algorithms for both the Deep Water Impact and the Red Sea datasets with three and four variables, respectively, utilizing a single node of the Frontier Supercomputer. For the OpenMP benchmarks, we tested the performance with 32 threads, while for the HIP benchmarks, the performance was evaluated on a single AMD GPU. The performance results are depicted in Fig. 6.




Figure 6: The speedup for the proposed closed-form solutions Sec. 3.1 for the evaluated datasets using different acceleration backends for probabilistic computations.

Our algorithms scale well with increase in the number of variables. On the Deep Water Impact dataset with three variables, execution times improved significantly: serial execution took approximately 13.64ms, reducing to 4.99ms with OpenMP (a 3 \times improvement), and further to 0.37ms with HIP (a 37 \times improvement). Performance gains were even more pronounced with four variables, with serial times of 29.17ms decreasing to 4.91ms with OpenMP (a 6 \times improvement), and to 0.50ms with HIP (a 58 \times improvement). These speedup results are depicted in Fig. 6(a). Additionally, our algorithms show better scalability with an increase in the data size. Particularly, performance improvements are greater for the Red Sea dataset (Fig. 6(b)) compared to the Deep Water Impact dataset (Fig. 6(a)) because of its larger data size.

Utilizing accelerators enables FunM²C to produce results in near-real-time, minimizing disruption during interactive result analysis. The speedup figures vary across datasets and backend architectures, highlighting the importance of systematically evaluating different datasets and workloads. Comparing acceleration architectures requires such systematic evaluation [12], which lies beyond the scope of this paper.

5 CONCLUSION & FUTURE WORK

This paper presented the design, implementation, and evaluation for efficiently computing multivariate data uncertainty using multi-core devices. We propose generalization of the previous work on bivariate uncertainty visualization [24] to multivariate data with arbitrary number of variables. Specifically, we showcase the results of our generalized framework for the Deep Water Impact dataset with three variables and the Red Sea dataset with four variables. Our results demonstrate a significant performance increase on multi-core CPUs and AMD GPUs compared to the serial implementation. Our implementation and integration with the VTK-m software makes it platform-portable and can be used on additional types (both current and future) architectures. Further, we demonstrate the seamless integration of our work with the ParaView software using the plugin mechanism. The integration into a production visualization tool provides access to uncertainty visualization techniques to a wider scientific community for data analysis.

There are few important limitations of our work, which we plan to address in the future. Our current implementation is limited to the uniform noise models and assumes independence among variables. Most real datasets often exhibit distributions that may significantly deviate from the uniform distribution. Thus, we plan to expand our generalized framework for Gaussian and histogram models in the future for more accurate quantification of probabilities. Also, variables often exhibit correlation in real datasets. Thus, our independent uniform assumption can lead to overestimation of probabilities compared to the correlated models with different probability distributions (as shown in the previous studies [4, 17]). Thus, we would like to account for correlations in our generalized framework. Lastly, our work currently assumes hyperrectangular traits for analysis of multivariate data, but we plan to investigate enhanced representation of traits based on scatterplots and extend our work to improve uncertainty analysis using flexible trait shapes.

ACKNOWLEDGMENTS

This work was supported in part by the U.S. Department of Energy (DOE) RAPIDS-2 SciDAC project under contract number DE-AC0500OR22725, the Intel OneAPI CoE, and the DOE Ab-initio Visualization for Innovative Science (AIVIS) grant 2428225. This research used resources of the Oak Ridge Leadership Computing Facility (OLCF), which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

REFERENCES

- [1] J. Ahrens, B. Geveci, and C. Law. *ParaView: An End-User Tool for Large Data Visualization*, chap. 36, pp. 717–731. Elsevier, 2005. 1, 3
- [2] S. Atchley et al. Frontier: Exploring exascale. In *SC '23: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, number 52, Nov. 2023. doi: 10.1145/3581784.3607089 3
- [3] T. M. Athawale, C. R. Johnson, S. Sane, and D. Pugmire. Fiber uncertainty visualization for bivariate data with parametric and nonparametric noise models. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):613–623, 2022. doi: 10.1109/TVCG.2022.3209424 1, 2
- [4] T. M. Athawale, S. Sane, and C. R. Johnson. Uncertainty Visualization of the Marching Squares and Marching Cubes Topology Cases. In *IEEE VIS 2021 (to appear)*. IEEE Xplore, 2021. 4
- [5] G. Bonneau, H. Hege, C. R. Johnson, M. Oliveira, K. Potter, P. Rheingans, and T. Schultz. Overview and state-of-the-art of uncertainty visualization. In M. Chen, H. Hagen, C. Hansen, C. R. Johnson, and A. Kauffman, eds., *Scientific Visualization: Uncertainty, Multifield, Biomedical, and Scalable Visualization*, pp. 3–27. Springer London, 2014. doi: 10.1007/978-1-4471-6497-5_1 1
- [6] K. Brodlie, R. A. Osorio, and A. Lopes. A review of uncertainty in data visualization. In J. Dill, R. Earnshaw, D. Kasik, J. Vince, and P. C. Wong, eds., *Expanding the Frontiers of Visual Analytics and Visualization*, pp. 81–109. Springer Verlag London, 2012. doi: 10.1007/978-1-4471-2804-5_6 1
- [7] H. Carr, Z. Geng, J. Tierny, A. Chattopadhyay, and A. Knoll. Fiber surfaces: Generalizing isosurfaces to bivariate data. In *Computer Graphics Forum*, vol. 34, pp. 241–250. Wiley Online Library, 2015. doi: 10.1111/cgf.12636 1
- [8] H. Childs, E. Brugger, B. Whitlock, J. Meredith, S. Ahern, K. Bonnell, M. Miller, G. Weber, C. Harrison, D. Pugmire, T. Fogal, C. Garth, A. Sanderson, E. W. Bethel, M. Durant, D. Camp, J. Favre, O. Rübel, P. Navratil, and F. Vivodtzev. VisIt: An end-user tool for visualizing and analyzing very large data. *Proceed SciDAC*, pp. 1–16, 01 2011. 1
- [9] K. Engel, M. Hadwiger, J. M. Kniss, A. E. Lefohn, C. R. Salama, and D. Weiskopf. Real-time volume graphics. In *ACM SIGGRAPH 2004 Course Notes*, SIGGRAPH '04, p. 29–es. Association for Computing Machinery, New York, NY, USA, 2004. doi: 10.1145/1103900.1103929 3
- [10] J. Jankowai and I. Hotz. Feature level-sets: Generalizing iso-surfaces to multi-variate data. *IEEE Transactions on Visualization and Computer Graphics*, 26(2):1308–1319, Feb. 2020. doi: 10.1109/TVCG.2018.2867488 1
- [11] A. Kamal, P. Dhakal, A. Y. Javaid, V. K. Devabhaktuni, D. Kaur, J. Zientz, and R. Marinier. Recent advances and challenges in uncertainty visualization: a survey. *Journal of Visualization*, 24(5):861–890, May 2021. doi: 10.1007/s12650-021-00755-1 1
- [12] N. Kondratyuk, V. Nikolskiy, D. Pavlov, and V. Stegailov. GPU-accelerated molecular dynamics: State-of-art software performance and porting from Nvidia CUDA to AMD HIP. *The International Journal of High Performance Computing Applications*, 35(4):312–324, 2021. doi: 10.1177/10943420211008288 4
- [13] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *SIGGRAPH Computer Graphics*, 21(4):163–169, Aug. 1987. doi: 10.1145/37402.37422 2
- [14] K. Moreland. The vtk-m user's guide version 2.1. Technical Report ORNL/TM-2023/3182, Oak Ridge National Laboratory, 2023. 3
- [15] K. Moreland, C. Sewell, W. Usher, L.-t. Lo, J. Meredith, D. Pugmire, J. Kress, H. Schroots, K.-L. Ma, H. Childs, M. Larsen, C.-M. Chen, R. Maynard, and B. Geveci. VTK-m: Accelerating the visualization toolkit for massively threaded architectures. *IEEE Computer Graphics and Applications*, 36(3):48–58, 2016. doi: 10.1109/MCG.2016.48 1, 2
- [16] J. Patchett, F. Samsel, K. C. Tsai, G. R. Gisler, D. H. Rogers, G. D. Abram, T. L. Turton, and L. Alamos. Visualization and analysis of threats from asteroid ocean impacts. *Los Alamos National Laboratory*, 5, 2016. 3
- [17] K. Pöthkow, B. Weber, and H.-C. Hege. Probabilistic marching cubes. *Computer Graphics Forum*, 30(3):931–940, June 2011. doi: 10.1111/j.1467-8659.2011.01942.x 4
- [18] K. Potter, P. Rosen, and C. R. Johnson. From quantification to visualization: A taxonomy of uncertainty visualization approaches. In A. M. Dienstfrey and R. F. Boisvert, eds., *Uncertainty Quantification in Scientific Computing*, pp. 226–249. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. doi: 10.1007/978-3-642-32677-6_15 1
- [19] M. A. Sandoval, W. R. Hix, O. B. Messer, E. J. Lentz, and J. A. Harris. Three-dimensional core-collapse supernova simulations with 160 isotopic species evolved to shock breakout. *The Astrophysical Journal*, 921(2):113, 2021. doi: 10.1051/0004-6361/201425025 3, 4
- [20] S. Sane, T. M. Athawale, and C. R. Johnson. Visualization of Uncertain Multivariate Data via Feature Confidence Level-Sets. In M. Agus, C. Garth, and A. Kerren, eds., *Proc. EuroVis - Short Papers*. The Eurographics Association, 2021. doi: 10.2312/evs.20211053 1, 3
- [21] D. Thompson, J. A. Levine, J. C. Bennett, P.-T. Bremer, A. Gyulassy, V. Pascucci, and P. P. Pébay. Analysis of large-scale scalar data using hixels. In *2011 IEEE Symposium on Large Data Analysis and Visualization*, pp. 23–30, Oct. 2011. doi: 10.1109/LDAV.2011.6092313 3
- [22] J. Wang, T. Athawale, K. Moreland, J. Chen, C. Johnson, and D. Pugmire. FunMC²: A filter for uncertainty visualization of marching cubes on multi-core devices. Technical report, Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States), 2023. doi: 10.2312/pgv.20231081 1, 2
- [23] P. Zhan, A. C. S. F. Yao, and I. Hoteit. Eddies in the Red Sea: A statistical and dynamical study. *JGR Oceans*, (6):3909–3925, June 2014. doi: 10.1002/2013JC009563 3
- [24] B. Zheng and F. Sadlo. Uncertainty in continuous scatterplots, continuous parallel coordinates, and fibers. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1819–1828, Feb. 2021. doi: 10.1109/TVCG.2020.3030466 1, 2, 4