

Compressive Modeling and Visualization of Multivariate Scientific Data using Implicit Neural Representation

Abhay Kumar Dwivedi, Shanu Saklani, and Soumya Dutta
abhaykd@cse.iitk.ac.in, shanu@cse.iitk.ac.in, soumyad@cse.iitk.ac.in
Indian Institute of Technology Kanpur (IIT Kanpur)
Kanpur, Uttar Pradesh, India

Abstract

The extensive adoption of Deep Neural Networks has led to their increased utilization in challenging scientific visualization tasks. Recent advancements in building compressed data models using implicit neural representations have shown promising results for tasks like spatiotemporal volume visualization and super-resolution. Inspired by these successes, we develop compressed neural representations for multivariate datasets containing tens to hundreds of variables. Our approach utilizes a single network to learn representations for all data variables simultaneously through parameter sharing. This allows us to achieve state-of-the-art data compression. Through comprehensive evaluations, we demonstrate superior performance in terms of reconstructed data quality, rendering and visualization quality, preservation of dependency information among variables, and storage efficiency.

CCS Concepts

• **Computing methodologies** → **Machine learning**; • **Human-centered computing** → **Visualization**.

Keywords

Multivariate Data Visualization, Data Compression, Deep Learning, Implicit Neural Representation, Volume Visualization.

ACM Reference Format:

Abhay Kumar Dwivedi, Shanu Saklani, and Soumya Dutta. 2025. Compressive Modeling and Visualization of Multivariate Scientific Data using Implicit Neural Representation. In *Proceedings of 16th Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP'25)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

Efficiently analyzing and visualizing multivariate scientific datasets is essential for gaining insights into various natural phenomena. Scientists from diverse domains utilize computational simulations to generate complex multivariate datasets, often comprising hundreds of variables. Instead of examining each variable independently, exploring the interaction among multiple variables concurrently is a powerful approach for studying intricate scientific phenomena [18].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICVGIP'25, Himachal Pradesh, India

© 2025 ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

For example, in weather simulation, while studying hurricanes, the interaction between pressure, wind velocity, and precipitation is often jointly examined to enable a comprehensive understanding of the complex interplay between these variables.

As the size, complexity, and number of variables in multivariate datasets increase, scientists across various domains face significant challenges in efficiently storing, managing, analyzing, and visualizing such data. This challenge becomes particularly pronounced when dealing with datasets containing hundreds of variables, such as those in climate and weather forecasting simulations. Consequently, data scientists are exploring various compression methods [2, 10, 18]. Recently, there has been considerable interest in using deep learning models to represent univariate data compactly [16, 27]. Although these methods have achieved impressive results, they have yet to address the direct compression of multivariate datasets. Moreover, it is crucial to comprehensively evaluate such models to ensure: (1) faithful on-demand reconstruction of all variables with high accuracy, (2) preservation of variable relationships (dependencies) in reconstructed data, and (3) meaningful results in various types of multivariate analysis tasks.

In this work, the primary focus is on thoroughly studying the viability of implicit neural representations (INRs) for modeling a large number of variables (tens to hundreds) jointly. We investigate whether using a single INR to model all variables jointly can efficiently learn such a diverse set of variables. We aim to evaluate the achievable reconstruction quality of this approach and compare it with alternative data reduction methods. To achieve this, we employ a residual SIREN architecture. By using a single network, we leverage parameter sharing across all data variables, enabling us to achieve a high compression ratio. We conduct a rigorous evaluation of the network's prediction quality from different perspectives, comparing it with other competitive approaches to demonstrate its superiority. In addition to comparing prediction quality using conventional data reconstruction and image-space metrics, we also perform a detailed comparative analysis using important multivariate application-specific tasks. We examine the accuracy of inter-variable relationships and assess their performance in conducting multivariate analyses. Finally, we investigate the performance of our model under varying conditions, including changes in the number of variables, variations in the amount of training data, and adjustments to the network depth. Our findings suggest that implicit neural representations are effective in capturing intricate patterns from multivariate scientific datasets, which contain numerous variables, and can efficiently represent complex multivariate data in a compact manner.

2 Related Works

Multivariate Data Visualization. Multivariate data analysis is a popular research area with many applications in various domains [6, 45]. Sauber et al. [31] propose a multifield-graph-based technique for determining the local correlation coefficients between various variables, facilitating the visualization and analysis of multivariate data. Gosink et al. [12] use local statistical distributions to improve it further for Query-driven analysis. To identify the interesting regions of multivariate datasets, local statistical complexities are utilized by Jänicke et al. [23]. For time-varying multivariate datasets, information theory is used to discover the causal relationship among variables [40]. Biswas et al. [4] use information theory to study variable interaction, and Dutta et al. [10] perform multivariate sampling.

Deep Learning for Scientific Visualization. Deep learning has found numerous applications in scientific visualization. Techniques for generating compact neural representations of scientific scalar data are proposed by Lu et al. [28] and Weiss et al. [43]. Hong et al. [22], He et al. [20], and Berger et al. [22] study the visualization of scalar field data using volume-rendered images. Weiss et al. [42] use isosurfaces for the same and further explored an adaptive sampling guided method for volume data visualization [44]. For compressing the volume data, new models for domain knowledge-aware latent space generation techniques for scalar data are also proposed [32]. Generation of high-resolution spatiotemporal volumes from low-resolution data [13, 14, 47] is another research area of focus. Also, for visualizing and exploring the parameter spaces for ensemble data, DNNs are utilized as surrogates [20, 33, 34]. GCNs are explored by Han and Wang [15] for learning surface representations. Han et al. [17] suggest a variable-to-variable translation technique for scientific multivariate data. There are many more applications of deep learning in scientific visualization, and for a more detailed review, please refer to the state-of-the-art survey [39].

Data Compression and Summarization. Data compression and summarization techniques have become increasingly pivotal as scientific datasets have grown significantly in recent years. Scientists have adopted several approaches to reduce the size of data with minimal loss of information. One such approach is to use compression techniques [2, 7, 26] to rapidly reduce the size of datasets and then decompress them during analysis and visualization. Woodring et al. [46] apply Wavelet-based compression to compress climate data. Statistical methods for data summarization and reduction have been extensively studied. More importantly, statistical Copula-model-based [18] and multivariate sampling [10] approaches have been investigated recently for summarizing multivariate datasets. For a detailed summary of existing data reduction techniques, please refer to [25].

3 Learning Multivariate Data using Implicit Neural Representation

3.1 Implicit Neural Representation

INRs with periodic activation functions have emerged as a promising approach for learning coordinate-based datasets such as images, scientific data, shape-based data, etc. Sitzmann et al. [35] showed




Figure 1: Schematic architecture of MVNet. It uses a residual SIREN architecture to enhance its learning capability.

that a feed-forward neural network with sinusoidal activation function, when initialized with a carefully chosen weighting scheme, can be used to model various types of coordinate-based data efficiently. Such networks are known as SIREN (sinusoidal representation network) [35]. Several variations of SIREN have been employed in the visualization research community to learn scalar field data [27], and model representations of various visualization types [16]. In a recent research, Tang and Wang used implicit neural networks to generate spatiotemporal super-resolution of scientific data from low-resolution representations [38].

3.2 Motivation

In this work, we employ a modified SIREN for modeling multivariate datasets with a large number of variables. In multivariate data, different variables contain different features and patterns, adding to the complexity of the data. Subsets of variables also often show the existence of complex local and global dependencies [4]. Hence, the primary motivation of our work is to investigate whether a SIREN-based neural network can accurately learn complex variable dynamics with a very large number of variables in the dataset. A potential reason for reduced accuracy on this task could be that the data variability across a large number of variables negatively impacts training, causing performance deterioration. Such a network may also require many hidden layers to successfully learn the multivariate data patterns, increasing the model complexity. Besides these possible pitfalls, we also aim to investigate the amount of compression that can be achieved by using such a neural representation compared to existing state-of-the-art data reduction techniques, such as tensor compression [2]. Finally, we also seek to evaluate (a) how accurately such a network preserves the dependency among the variables and (b) the effectiveness of the reconstructed data in performing important multivariate analysis tasks.

3.3 Model Architecture

We build our MultiVariate data representation Network (MVNet), based on a SIREN architecture [35]. Our goal is to learn the representation of multivariate data as a function, which is represented by the parameters of a neural network. We want the size of the neural network to be significantly smaller than the size of the raw multivariate data so that a compressed multivariate data model is produced for efficient analysis and visualization. Our model represents a function that takes a d dimensional location vector, i.e., for 3D volumetric data, $d = 3$, and we pass the coordinate of the grid point as input and predict a v dimensional vector output where v is the number of variables in the multivariate dataset. Essentially,

our neural network learns a function $\mathcal{F}(\theta) : \mathbb{R}^d \mapsto \mathbb{R}^v$, where θ represents the parameters of the neural network.

We build a multilayer perceptron consisting of d neurons as inputs, followed by l hidden layers, and an output layer containing v neurons. For boosting the model's learning capability and training a deep network stably, we enhance this architecture by incorporating residual blocks and skip connections [19]. By formulating the learning of v variables using a single neural network, we pose this as a simple multi-task learning problem [5]. As the variables in the scientific datasets are often correlated [4, 18], we believe that such inter-dependencies can help our network to learn the diverse patterns of all the variables efficiently, where learning representations of one variable is considered as a single task, and the network learns all the variable representations jointly. The applicability of multi-task learning where the dataset has several related variables has been shown to be effective by researchers in the past [5, 21, 29]. Furthermore, using a single network to learn all the variables jointly for a multivariate dataset also helps us to obtain a compact representation of the entire data since the network parameters are shared, and hence the network size is minimal.

In Fig. 1, a schematic of our model architecture is shown. We use a sinusoidal activation function for its superior performance over alternative activation functions such as ReLU and Sigmoid [35]. It can be seen that the network takes a d dimensional location vector as input and predicts a v dimensional vector output, where each element of the output vector indicates the value of an individual variable. To train this network, we use the Adam optimizer [24] and the conventional mean squared error (MSE) loss ($\mathcal{L}_{mse}$) function. We conduct a thorough hyperparameter search to find the best parameters and learning rate for the Adam optimizer. While computing the MSE loss, we use equal weight on loss values for all the variables. To train MVNet in a stable manner, we normalize the input coordinates between 0 and 1 and rescale the values of each variable between -1 and 1.

4 Comparative Study

To show the efficacy of MVNet over the existing data compression and summarization approaches, we present a thorough comparative study. We use four multivariate datasets to demonstrate the results. The dimensionality, number of variables, and spatial resolution of the datasets are reported in Table 1. We use a GPU server with NVIDIA GeForce GTX 1080Ti GPUs with 12GB GPU memory for all the experimentation. All the models are implemented in PyTorch [30]. Hurricane Isabel data was produced by the Weather Research and Forecast model, courtesy of NCAR and the U.S. National Science Foundation. The Turbulent Combustion dataset is made available by Dr. Jacqueline Chen at Sandia Laboratories through the U.S. Department of Energy's SciDAC Institute for Ultrascale Visualization. The Climate50 and Climate100 datasets were obtained from the Earth System Grid Federation's online data repository. The dataset was generated by Energy Exascale Earth System Model (E3SM). First we compare MVNet with three different data compression approaches. The first one is the baseline, where the multivariate dataset is first sub-sampled into a lower resolution grid using linear interpolation, and then again linear interpolation is used to up-sample the data back to the original grid

Table 1: Description of datasets used in the experimentation.

Dataset	Dimensionality	#Variables	Resolution
Combustion	3D	5	240×360×60
Isabel	3D	14	250×250×50
Climate50	2D	53	2880×1440
Climate100	2D	100	2880×1440

resolution. Typically such sub-sampling is often used as a means of data reduction. We call this method as *LERP*. The second method is the state-of-the-art tensor compression, *TTHRESH* [2]. The third method is another widely used floating point compression method *Zfp* [26]. We compare the reconstruction quality of all the variables using Peak Signal to Noise Ratio (PSNR). To conduct a comparison in feature space, we use isocontour similarity to estimate the accuracy of isocontours generated by each method. Next, to evaluate the visualization quality in image space, we render images of data variables and compare their accuracy by comparing them against the ground truth. Finally, we compare MVNet with the statistical Copula-based summarization method proposed in [18].

4.1 Comparison of Reconstructed Data Quality

The quality of the reconstructed data is quantitatively compared by computing the PSNR for each data variable for each method. Our goal is to compare MVNet against TTHRESH, Zfp, and LERP by comparing the reconstructed data quality given comparable storage footprint. We find that MVNet with 10 residual blocks with 120 neurons in each layer produces the best quality vs. storage trade-off. In Table 2, the results for all the four methods are provided. We use Zfp with both fixed bitrate (FBR) and fixed absolute error (FAE) as compression guideline. Since the datasets contain a large number of variables and due to space constraints, we present the average PSNR of all the variables in Table 2. We observe that TTHRESH, Zfp, and LERP methods produce reconstructed data with lower PSNR than MVNet, even with slightly higher storage. Hence, MVNet gives the best trade-off between storage and PSNR. We further observe that, between Zfp and TTHRESH, TTHRESH gives better trade-off. Hence, we use only TTHRESH for the subsequent analyses and visualization. Next, Fig 2 shows a joint boxplot comparing PSNR for MVNet, LERP, and TTHRESH across various datasets. In each boxplot, we show the mean, median and further map the maximum and minimum variable PSNR values to the whiskers. A boxplot is drawn for each dataset for each method. The boxplots illustrate that MVNet achieves higher PSNR, compared to LERP and TTHRESH, showing its superior and consistent reconstruction performance. Finally, in Table 3, we report the average absolute maximum error, 95th percentile absolute error and the fraction above the tolerance level (0.05) for all the variables for MVNet. Note that, we compute these error statistics on variable values scaled between -1 to 1 to obtain consistent and comparable error statistics.

4.2 Comparison of Reconstructed Features

Isocontours are considered important feature representations in scientific datasets. Hence, we evaluate the accuracy of the extracted isocontours for the three methods. Given a dataset, we first uniformly randomly select 20 isovalues for each variable and then extract those isosurfaces. Then, we compute the Hausdorff distance




Figure 2: Comparison of reconstruction quality for MVNet, LERP, and TTHRESH. We provide a boxplot for each method for each dataset. We observe that MVNet consistently delivers superior PSNR compared to LERP and TTHRESH.

Table 2: Quantitative comparison of average reconstruction quality for all the variables vs. storage overhead for MVNet and other methods. All raw datasets are stored using VTK data format. It is observed that MVNet produces higher reconstruction quality with the minimum storage.

Dataset	Raw Stor. (MB)	MVNet		TTHRESH		LERP		ZFP (FBR)		ZFP (FAE)	
		Stor. (KB)	PSNR (dB)	Stor. (KB)	PSNR (dB)	Stor. (KB)	PSNR (dB)	Stor. (KB)	PSNR (dB)	Stor. (KB)	PSNR (dB)
Comb.	108	1160	52.51	1608	49.89	1816	32.47	1400	30.45	4284	46.88
Isabel	158	1164	47.32	2072	44.94	2164	37.74	1960	34.87	1856	39.07
CLM50	842	1184	51.27	2792	47.10	2052	33.01	21836	18.79	8808	17.97
CLM100	1568	1208	49.68	4568	48.62	3076	33.25	47200	23.99	9160	22.2

Table 3: Maximum absolute error, 95th percentile absolute error, and fraction above tolerance (0.05) for MVNet for all the datasets.

DataSet	Max. Error	95th Percentile error	Frac. above tolerance (0.05)
Combustion	0.056	0.012	0.000012
Isabel	0.371	0.026	0.004
climate50	0.103	0.021	0.013
climate100	0.144	0.029	0.06




Figure 3: Visualization of isosurface of Velocity Magnitude at isovalue=41 of Isabel dataset with ground truth (GT) shown in Fig. 19a. Isosurface produced by LERP is least accurate while MVNet and TTHRESH produce visually comparable results.




Figure 4: Reconstructed data visualization for Pressure (P) variable of Isabel dataset. Images generated by ground truth (GT), MVNet, TTHRESH, and LERP are shown. We observe that MVNet produces the most accurate result.

and Chamfer distance [3] for each isocontour against the ground truth isocontours extracted from the raw data. This study is conducted for the 3D datasets. The average isocontour distance over 20 isocontours is computed for each variable, and finally, the average of such variable-wise average values is computed for each dataset. Table 4 shows the results of this feature-based comparison. We observe that, on average, the accuracy of the isocontours is the highest for MVNet compared to TTHRESH and LERP. In Fig. 19, we show the visualization of a representative isosurface generated from Isabel dataset using the Velocity field with isovalue=41. We observe that the MVNet and TTHRESH produce visually comparable results, while the isosurface generated by the LERP fails to capture detailed surface features.

Table 4: Quantitative feature quality comparison using average Chamfer and Hausdorff distances for all data variables for two 3D datasets. The findings indicate that for both the datasets, MVNet outperforms TTHRESH and LERP regarding isocontour reconstruction quality.

Dataset	MVNet		TTHRESH		LERP	
	Chamfer ↓	Hausdorff ↓	Chamfer ↓	Hausdorff ↓	Chamfer ↓	Hausdorff ↓
Combustion	0.250	11.937	0.344	13.813	5.715	32.187
Isabel	0.955	23.016	2.814	36.131	12.959	60.941

Table 5: Visualization quality comparison in image space. We use SSIM, LPIPS, and DISTS measures to quantify perceptual image similarity between ground truth and the images generated by MVNet, TTHRESH, and LERP methods.

Dataset	Var Name	SSIM ↑			LPIPS ↓			DISTS ↓		
		MVNet	TTHRESH	LERP	MVNet	TTHRESH	LERP	MVNet	TTHRESH	LERP
Isabel	P	0.991	0.936	0.675	0.049	0.248	0.356	0.053	0.224	0.243
	Vel	0.956	0.95	0.858	0.085	0.075	0.174	0.081	0.115	0.299
Climate50	U10	0.975	0.914	0.744	0.003	0.023	0.155	0.028	0.221	0.15
	CLDTOT	0.987	0.967	0.88	0.005	0.014	0.137	0.036	0.191	0.198

4.3 Comparison of Visualization Quality

Next comparative evaluation is conducted in image space to study the accuracy of the visualizations generated. We use selected variables from Isabel and Climate50 data to conduct this study. While generating images, we keep all the rendering settings fixed for each dataset. To quantify the image similarity between the ground truth and images generated by each method, we use (1) Structural Similarity (SSIM) [41], (2) Learned Perceptual Image Patch Similarity (LPIPS) [48], and (3) Deep Image Structure And Texture Similarity




Figure 5: Reconstructed data visualization for the U-Velocity (U10) variable of the Climate100 dataset. Images generated by ground truth data, MVNet, TTHRESH, and LERP methods are shown from left to right with a region zoomed below to show the differences. We observe that MVNet produces the most accurate result, while the result from TTHRESH shows artifacts as highlighted by the red dotted box in Fig. 5c. The LERP method-generated image is the least accurate.

(DISTS) [9]. The results are presented in Table 5. While MVNet produces the most accurate visualizations, the LERP consistently yields the least accurate visualizations. For a qualitative comparison, in Fig 16 and Fig 5, we provide visualizations of selected variables from Isabel and Climate100 dataset respectively. From Fig 16, we observe that MVNet produces the most accurate result while TTHRESH and LERP-generated images contain visual artifacts. Fig 5 shows a similar finding as Fig 16. By observing the zoomed inset, we see that MVNet produces the most accurate result for the U10 variable and TTHRESH-generated image contains visual artifacts. The LERP method fails to preserve the features and produces the least accurate result.

4.4 Comparison with Copula-based Method

Representing multivariate datasets with a large number of variables compactly using statistical Copula-based summaries was proposed by Hazarika et al. [18]. Since our method is addressing the similar problem, we use the Climate50 and Climate100 datasets to compare MVNet with the Copula-based summarization method [18]. To keep the storage footprint minimal for the Copula-based model, we model each variable in a block using Gaussian distribution, which only requires two floating points per variable per block. We also store the pairwise correlation coefficient for all pairs of variables for each data block. To make the storage size comparable to MVNet, we use 80×80 block size for the Climate50 dataset and 120×120 block size for the Climate100 dataset for Copula modeling. For both MVNet and Copula-based methods, we reconstruct all the variables at their original resolution. In Table 6, the quality (average PSNR across all variables) vs. storage trade-off is presented. We observe that MVNet produces significantly higher-quality reconstruction with lower storage (as shown in boxplot of Fig. 6). Due to space limitation, we show visualization of a representative variable, Surface Temperature (TS), in the supplementary material.

5 Multivariate Application Study

5.1 Multivariate Dependency Analysis

Variables in multivariate scientific datasets are often correlated with other variables in the dataset. The nature of this correlation or

Table 6: Quantitative comparison of the data reconstruction quality between MVNet and Copula-based method using Climate dataset. It is observed that MVNet produces higher reconstruction quality with lower storage overhead.

dataset	#Variables	MVNet		Copula-based Method		
		Storage ↓ (KB)	PSNR ↑ (dB)	Block Size	Storage ↓ (KB)	PSNR ↑ (dB)
Climate50	53	1184	51.274	80X80	5116	33.618
Climate100	100	1208	49.679	120X120	7108	32.192




Figure 6: Boxplot-based comparison of PSNR for MVNet and Copula method for Climate50 and Climate100 datasets.

interdependence can be both linear and non-linear [4]. By analyzing and visualizing such correlated (dependent) variables, complex multivariate features are often studied [4, 18, 31, 37]. Hence, it is essential to preserve such correlations for any multivariate data modeling technique. In this section, we compare the accuracy with which correlation (linear and non-linear) is preserved in datasets reconstructed by the MVNet, TTHRESH, and LERP methods. We use Pearson's Correlation Coefficient to quantify linear dependence and Mutual Information to measure non-linear dependence between all variable pairs for each method. Next, to compute the error in these estimated correlation and mutual information values, we compute the absolute difference between each estimated value and the ground truth value for each variable pair. In Table 7, we show the average absolute error value computed over the three methods' variable pairs for each dataset.

Table 7: Comparison of average absolute error in linear correlation and mutual information for different methods.

Dataset	Error in Linear Correlation ↓			Error in Mutual Information ↓		
	MVNet	TTHRESH	LERP	MVNet	TTHRESH	LERP
Combustion	0.002139	0.000919	0.170004	0.115754	0.226789	0.258035
Isabel	0.082532	0.188319	0.284613	0.195037	0.316863	0.227967
Climate50	0.032771	0.057524	0.268297	0.601893	1.535684	0.787852
Climate100	0.109252	0.137031	0.737098	1.554588	2.497101	1.963269

It is observed that, in terms of linear correlation error (deviation), for Combustion data, TTHRESH incurs the least average error, and for all other datasets, MVNet produces the minimum average error. Next, it is found that MVNet produces the least average absolute error in mutual information for all the datasets. LERP produces the highest error. Therefore, from Table 7, we conclude that MVNet recovers the dependencies among variables most accurately. In Fig. 7 and Fig. 8, we show all pairwise correlation and mutual information error matrices for Isabel and Climate50 datasets, respectively. Since we show the absolute error values in these matrices, if a matrix has fewer darker cells, then it can be interpreted that the corresponding method produces fewer errors. We use a consistent color map while generating these plots across the three methods so that the same color intensity in two different plots indicates the same absolute error value. We observe that between Figs. 7(a)-(c), Fig. 7(a), and between Figs. 7(d)-(f), Fig. 7(d) produces a minimum error indicating MVNet's superiority over TTHRESH and LERP in estimating linear correlation. Similarly, from Fig. 8, by comparing Figs. 8(a)-(c) and Figs. 8(d)-(f), we conclude that MVNet produces the most accurate mutual information estimation as compared to TTHRESH and LERP.

Table 8: Comparison of the accuracy of multivariate QDV for different methods. Query accuracy between a method and the ground truth result is quantified using the Dice Similarity Coefficient (DSC). It is observed that MVNet produces the most accurate results for both datasets.

Dataset	Multivariate Query	MVNet DSC ↑	TTHRESH DSC ↑	LERP DSC ↑
Combustion	Mixfrac >0.3 & <0.7 & Y_OH >0.006 & <0.1	0.9914	0.988	0.8966
Isabel	CLOUD >0.0001 & <0.002 & PRECIP >0.0001 & <0.0093 & QVAPOR >0.01 & <0.0235	0.7815	0.7083	0.3569

5.2 Multivariate Query-Driven Visual Analysis

Multivariate variable interaction analysis using query-driven visualization (QDV) is an effective approach for large scientific data exploration [10, 11, 18, 36]. Scientists often perform queries on a group of variables jointly to filter a subset of related data points for detailed feature analysis. Such query-driven analyses help reduce the explorable data size and scientists' cognitive load while accelerating the scientific discovery process. Therefore, we study the efficacy of MVNet in performing multivariate QDV. To compare the accuracy of the QDV results, we compute the Dice Similarity Coefficient (DSC) [8] between the query result obtained from ground truth data and the other methods. The value of DSC varies between

0 and 1, where 1 indicates a perfect match. The results and the exact queries are summarized in Table 8.

We execute a query on the Hurricane Isabel dataset involving variables CLOUD, PRECIPITATION, and QVAPOR to study the cloud structure and regions where high rainfall is possible. We jointly query moderate to high value ranges for each variable. The exact query is shown in Table 8. In another experiment, we perform QDV for Combustion data. Turbulent Combustion data is studied to understand the complex combustion process and the interplay among fuel, oxidizer, and other chemical components [1]. The interaction between mixture fraction (Mixfrac) and mass fraction of Hydroxyl (Y_OH) variables can be studied to understand the flame structure. Hence, we conduct a QDV using these two variables to isolate data points that indicate the potential flame structure [1]. From Table 8, we observe that MVNet produces the most accurate query results among all the methods for both datasets. In Fig. 17 and Fig. 18, the results of these QDV are shown. We observe that compared to the ground truth, MVNet produces visually similar results, whereas the LERP method produces the least accurate results.

6 Evaluation & Parameter Study

6.1 Varying Number of Variables

To evaluate the performance of MVNet in this regard, we conduct experiments where we gradually increase the number of variables while keeping the architecture of MVNet fixed. Only the number of neurons in the output layer is adjusted to accommodate the varying number of variables. All other hyperparameters are set to the best configuration determined from experimentation (details are in supplementary document). The results are depicted in Fig. 11 where different colors representing different datasets. As the number of variables increases, the learning task becomes more challenging, leading to a gradual decrease in the average PSNR value across all variables. However, it's worth noting that even when trained with the maximum number of variables, MVNet maintains a high average PSNR for each dataset. Furthermore, Table 9 highlights that employing a deeper MVNet model can further enhance PSNR, providing the option for users to obtain higher-quality reconstructions if desired. These findings offer valuable insights for users regarding the trade-off between expected quality and storage overhead across different datasets and variable configurations.

6.2 Compression Ratio vs. Prediction Quality

The architecture of MVNet relies on residual blocks, which allows for increasing the network's depth to potentially enhance prediction quality and generalizability. However, this also increases model complexity, leading to higher storage requirements. To understand the trade-off between storage and prediction quality, we conducted an experiment varying the number of residual blocks from 4 to 14, while keeping the number of neurons per hidden layer constant at 120. This analysis also sheds light on MVNet's compression capabilities. The results, presented in Table 9, reveal that with only 4 residual blocks, MVNet achieves maximum compression ratio while maintaining good average PSNR across datasets. This trade-off is also provided as a form of grouped bar chart in Fig. 13. Notably, as the number of variables increases in multivariate datasets, MVNet offers higher compression ratios due to its consistent network size.




Figure 7: Linear Correlation Error Matrices for Isabel and Climate50 dataset for MVNet, TTHRESH, and LERP method.




Figure 8: Mutual Information Error Matrices for Isabel and Climate50 dataset for MVNet, TTHRESH, and LERP method.

Table 9: Average PSNR, storage footprint, and compression ratio (CR) under varying numbers of residual blocks to study the trade-off between compression ratio and PSNR. We observe that MVNet maintains a high PSNR even when fewer residual blocks are used.

Dataset	Num. of Res. Block = 4			Num. of Res. Block = 6			Num. of Res. Block = 8			Num. of Res. Block = 10			Num. of Res. Block = 12			Num. of Res. Block = 14		
	Storage (KB)	CR	PSNR (dB)	Storage (KB)	CR	PSNR (dB)	Storage (KB)	CR	PSNR (dB)	Storage (KB)	CR	PSNR (dB)	Storage (KB)	CR	PSNR (dB)	Storage (KB)	CR	PSNR (dB)
Combustion	468	236.2:1	45.567	700	157.9:1	48.386	928	119.14:1	50.701	1160	95.3:1	52.521	1388	79.65:1	53.984	1620	68.25:1	55.411
Isabel	472	343.13:1	44.207	704	230.05:1	45.464	932	173.77:1	46.381	1164	139.13:1	47.634	1396	116.01:1	47.947	1624	99.72:1	49.168
Climate50	492	1753.38:1	44.038	720	1198.14:1	46.697	952	906.15:1	48.740	1184	728.6:1	51.367	1412	610.95:1	52.805	1644	524.73:1	54.152
Climate100	512	3172.01:1	42.784	744	2182.89:1	45.303	976	1664:1	47.380	1204	1348.89:1	49.991	1436	1130.96:1	51.515	1664	976:1	52.624




Figure 9: Query-driven visualization for the Isabel dataset. A multivariate query: (CLOUD >0.0001 AND <0.002) AND (PRECIPITATION >0.0001 AND <0.0093) AND (QVAPOR >0.01 AND <0.0235) is shown.

Table 10: Hyperparameters, training and inference time. All experiments are done on a GPU server with NVIDIA GeForce GTX 1080Ti GPUs with 12GB GPU memory

Dataset	Learning Rate	Batch Size	Decay Rate	Decay Frequency	Training Time (Hrs.)	Inference Time (Secs.)
Combustion	0.00005	2048	0.8	Every 15th Epoch	3.89	2.04
Isabel	0.0001	2048	0.8	Every 15th Epoch	2.32	1.46
Climate50	0.0001	2048	0.8	Every 15th Epoch	2.72	3.96
Climate100	0.0001	2048	0.8	Every 15th Epoch	2.93	10.38




Figure 10: Query-driven visualization for the Combustion dataset. A multivariate query: (Mixfrac >0.3 AND <0.7) AND (Y_OH >0.006 AND <0.1) is shown.

6.3 Varying Amounts of Training Data

As computational capabilities advance, storing entire high-resolution multivariate datasets becomes increasingly costly. In such cases, storing sub-sampled data becomes a viable option to reduce storage requirements for offline analysis and visualization. Consequently, training MVNet on sub-sampled data becomes necessary. To investigate the effectiveness of MVNet under these circumstances, we conducted a comprehensive experiment by training MVNet using 25%, 50%, 75%, and 100% data samples, randomly selected uniformly. The Fig. 12 illustrates the average quality of the reconstructed data variables for various datasets. We observe a steady increase in reconstruction quality with more training samples. However, it's noteworthy that the rate of improvement in PSNR values isn't




Figure 11: The variation of average PSNR when the number of variables increases. We keep the MVNet architecture fixed and only change the number of neurons in the final layer.




Figure 12: Average PSNR values under different sampling percentages for different datasets.




Figure 13: Change of average PSNR when number of res. blocks are varied. The corresponding compression ratio is also shown at the top of each bar.

significantly high with increasing sample points. Even at a low sampling rate of 25%, MVNet can still achieve very good average PSNR, suggesting its efficacy when trained on sub-sampled data.

6.4 Hyperparameters, Training, and Inference Timing of MVNet

The hyperparameters, timings for both training and inference of MVNet are presented in Table 10. Consistency is maintained by

training MVNet for 300 epochs across all datasets. It's demonstrated that MVNet achieves rapid reconstruction (inference) of all variables, with inference time correlating to the number of variables in the multivariate dataset. Table 10 illustrates training time using the optimal hyperparameter combinations for each dataset. Through comprehensive evaluation, it's noted that increasing the batch size decreases training time, and training time scales with both the number of data points and variables.

7 Discussion

This work demonstrates the effectiveness of utilizing residual implicit neural representations with sinusoidal activation functions to model complex multivariate scientific datasets when initialized with carefully selected weights and hyperparameters. Our primary focus is on datasets featuring a large number of variables with dynamic characteristics, enabling us to assess the model's performance comprehensively. We compare MVNet against established methods such as tensor compressor TTHRESH, Zfp, and Copula modeling, and employ various metrics to assess reconstructed data, feature, and image quality. We further delve into MVNet's ability to recover both linear and nonlinear dependencies among variable combinations, which is important for multivariate data analysis. We then comparatively study the efficacy of MVNet in carrying out important multivariate tasks such as query-driven visual analysis. We conduct experiments to report the trade-off between compression ratio and prediction quality, allowing users to use this as a guideline for selecting the appropriate model architecture for their dataset. We also depict how increasing variables in the dataset impact MVNet's overall prediction quality. Depending on the need for reconstruction quality and the number of variables to be modeled, the users can select an appropriate depth for MVNet to get a desirable accuracy. Through our thorough evaluation, we aim to encourage domain scientists to integrate INR-based models, such as MVNet, into their analysis workflows. These models provide a compact yet versatile representation of multivariate data, facilitating streamlined analytics and visualization, and accelerating the scientific discovery process.

8 Conclusions and Future Work

In summary, we introduce MVNet to efficiently capture complex multivariate scientific data containing a large number of variables. We undertake a comprehensive investigation to showcase the effectiveness of MVNet compared to state-of-the-art compression and statistical methods, focusing on the balance between storage and the reconstruction quality, as well as its performance in several scientific applications. Looking ahead, we aim to expand this work by applying MVNet to model time-varying scientific datasets including spatio-temporal vector and tensor data.

References

- [1] Hiroshi Akiba, Kwan-liu Ma, Jacqueline H. Chen, and Evatt R. Hawkes. 2007. Visualizing Multivariate Volume Data from Turbulent Combustion Simulations. *Computing in Science & Engineering* 9, 2 (2007), 76–83. doi:10.1109/MCSE.2007.42
- [2] Rafael Ballester-Ripoll, Peter Lindstrom, and Renato Pajarola. 2020. TTHRESH: Tensor Compression for Multidimensional Visual Data. *IEEE Transactions on Visualization and Computer Graphics* 26, 9 (2020), 2891–2903. doi:10.1109/TVCG.2019.2904063

- [3] H. G. Barrow, J. M. Tenenbaum, R. C. Bolles, and H. C. Wolf. 1977. Parametric correspondence and chamfer matching: two new techniques for image matching. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence - Volume 2* (Cambridge, USA) (IJCAI'77). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 659–663.
- [4] Ayan Biswas, Soumya Dutta, Han-Wei Shen, and Jonathan Woodring. 2013. An Information-Aware Framework for Exploring Multivariate Data Sets. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2683–2692. doi:10.1109/TVCG.2013.133
- [5] Michael Crawshaw. 2020. Multi-Task Learning with Deep Neural Networks: A Survey. *arXiv preprint arXiv:2009.09796* (2020).
- [6] MC Ferreira De Oliveira and Haim Levkowitz. 2003. From visual data exploration to visual data mining: A survey. *IEEE transactions on visualization and computer graphics* 9, 3 (2003), 378–394.
- [7] Sheng Di and Franck Cappello. 2016. Fast Error-Bounded Lossy HPC Data Compression with SZ. In *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 730–739. doi:10.1109/IPDPS.2016.11
- [8] Lee R Dice. 1945. Measures of the amount of ecologic association between species. *Ecology* 26, 3 (1945), 297–302.
- [9] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P. Simoncelli. 2022. Image Quality Assessment: Unifying Structure and Texture Similarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 5 (2022), 2567–2581. doi:10.1109/TPAMI.2020.3045810
- [10] Soumya Dutta, Ayan Biswas, and James Ahrens. 2019. Multivariate Pointwise Information-Driven Data Sampling and Visualization. *Entropy* 21, 7 (2019). doi:10.3390/e21070699
- [11] Luke Gosink, John Anderson, Wes Bethel, and Kenneth Joy. 2007. Variable Interactions in Query-Driven Visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1400–1407. doi:10.1109/TVCG.2007.70519
- [12] Luke J Gosink, Christoph Garth, John C Anderson, E Wes Bethel, and Kenneth I Joy. 2010. An application of multivariate statistical analysis for query-driven visualization. *IEEE Transactions on Visualization and Computer Graphics* 17, 3 (2010), 264–275.
- [13] Jun Han and Chaoli Wang. 2020. TSR-TVD: Temporal Super-Resolution for Time-Varying Data Analysis and Visualization. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 205–215.
- [14] Jun Han and Chaoli Wang. 2022. SSR-TVD: Spatial Super-Resolution for Time-Varying Data Analysis and Visualization. *IEEE Transactions on Visualization and Computer Graphics* 28, 6 (2022), 2445–2456.
- [15] Jun Han and Chaoli Wang. 2022. SurfNet: Learning surface representations via graph convolutional network. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 109–120.
- [16] Jun Han and Chaoli Wang. 2023. CoordNet: Data Generation and Visualization Generation for Time-Varying Volumes via a Coordinate-Based Neural Network. *IEEE Transactions on Visualization and Computer Graphics* 29, 12 (2023), 4951–4963. doi:10.1109/TVCG.2022.3197203
- [17] Jun Han, Hao Zheng, Yunhao Xing, Danny Z Chen, and Chaoli Wang. 2020. V2V: A deep learning approach to variable-to-variable selection and translation for multivariate time-varying data. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1290–1300.
- [18] Subhashis Hazarika, Soumya Dutta, Han-Wei Shen, and Jen-Ping Chen. 2019. CoDDA: A Flexible Copula-based Distribution Driven Analysis Framework for Large-Scale Multivariate Data. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 1214–1224. doi:10.1109/TVCG.2018.2864801
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [20] Wenbin He, Junpeng Wang, Hanqi Guo, Ko-Chih Wang, Han-Wei Shen, Mukund Raj, Youssef SG Nashed, and Tom Peterka. 2019. InSituNet: Deep image synthesis for parameter space exploration of ensemble simulations. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 23–33.
- [21] Xiao He, Francesco Alesiani, and Ammar Shaker. 2019. Efficient and scalable multi-task regression on massive number of tasks. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence* (Honolulu, Hawaii, USA) (AAAI'19/IAAI'19/EAAI'19). AAAI Press, Article 462, 9 pages. doi:10.1609/aaai.v33i01.33013763
- [22] Fan Hong, Can Liu, and Xiaoru Yuan. 2019. DNN-VolVis: Interactive volume visualization supported by deep neural network. In *2019 IEEE Pacific Visualization Symposium (PacificVis)*. IEEE, 282–291.
- [23] Heike Janicke, Alexander Wiebel, Gerik Scheuermann, and Wolfgang Kollmann. 2007. Multifield visualization using local statistical complexity. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1384–1391.
- [24] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
- [25] S. Li, N. Marsaglia, C. Garth, J. Woodring, J. Clyne, and H. Childs. 2018. Data Reduction Techniques for Simulation, Visualization and Data Analysis. *Computer Graphics Forum* 37, 6 (2018), 422–447. doi:10.1111/cgf.13336
- [26] Peter Lindstrom. 2014. Fixed-Rate Compressed Floating-Point Arrays. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2674–2683. doi:10.1109/TVCG.2014.2346458
- [27] Y. Lu, K. Jiang, J. A. Levine, and M. Berger. 2021. Compressive Neural Representations of Volumetric Scalar Fields. *Computer Graphics Forum* 40, 3 (2021), 135–146. doi:10.1111/cgf.14295
- [28] Yuzhe Lu, Kairong Jiang, Joshua A Levine, and Matthew Berger. 2021. Compressive neural representations of volumetric scalar fields. In *Computer Graphics Forum*, Vol. 40. Wiley Online Library, 135–146.
- [29] Umberto Michelucci and Francesca Venturini. 2019. Multi-Task Learning for Multi-Dimensional Regression: Application to Luminescence Sensing. *Applied Sciences* 9, 22 (2019). doi:10.3390/app9224748
- [30] A. Paszke et al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Proc. of NIPS*. 8024–8035.
- [31] Natascha Sauber, Holger Theisel, and Hans-peter Seidel. 2006. Multifield-Graphs: An Approach to Visualizing Correlations in Multifield Scalar Data. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 917–924. doi:10.1109/TVCG.2006.165
- [32] Jingyi Shen, Haoyu Li, Jiayi Xu, Ayan Biswas, and Han-Wei Shen. 2022. IDLat: An Importance-Driven Latent Generation Method for Scientific Data. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 679–689.
- [33] Neng Shi, Jiayi Xu, Haoyu Li, Hanqi Guo, Jonathan Woodring, and Han-Wei Shen. 2022. VDL-Surrogate: A View-Dependent Latent-based Model for Parameter Space Exploration of Ensemble Simulations. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 820–830.
- [34] Neng Shi, Jiayi Xu, Skylar W Wurster, Hanqi Guo, Jonathan Woodring, Luke P Van Roekel, and Han-Wei Shen. 2022. Gnn-surrogate: A hierarchical and adaptive graph neural network for parameter space exploration of unstructured-mesh ocean simulations. *IEEE Transactions on Visualization and Computer Graphics* 28, 6 (2022), 2301–2313.
- [35] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. 2020. Implicit Neural Representations with Periodic Activation Functions. In *Proc. NeurIPS*.
- [36] K. Stockinger, J. Shalf, K. Wu, and E.W. Bethel. 2005. Query-driven visualization of large data sets. In *VIS 05. IEEE Visualization, 2005*. 167–174. doi:10.1109/VISUAL.2005.1532792
- [37] Jeffrey Sukharev, Chaoli Wang, Kwan-Liu Ma, and Andrew T. Wittenberg. 2009. Correlation study of time-varying multivariate climate data sets. In *2009 IEEE Pacific Visualization Symposium*. 161–168. doi:10.1109/PACIFICVIS.2009.4906852
- [38] Kaiyuan Tang and Chaoli Wang. 2024. STSR-INR: Spatiotemporal super-resolution for multivariate time-varying volumetric data via implicit neural representation. *Computers & Graphics* 119 (2024), 103874. doi:10.1016/j.cag.2024.01.001
- [39] Chaoli Wang and Jun Han. 2022. DL4SciVis: A state-of-the-art survey on deep learning for scientific visualization. *IEEE transactions on visualization and computer graphics* (2022).
- [40] Chaoli Wang, Hongfeng Yu, Ray W Grout, Kwan-Liu Ma, and Jacqueline H Chen. 2011. Analyzing information transfer in time-varying multivariate data. In *2011 IEEE Pacific Visualization Symposium*. IEEE, 99–106.
- [41] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. doi:10.1109/TIP.2003.819861
- [42] Sebastian Weiss, Mengyu Chu, Nils Thuerey, and Rüdiger Westermann. 2019. Volumetric isosurface rendering with deep learning-based super-resolution. *IEEE Transactions on Visualization and Computer Graphics* 27, 6 (2019), 3064–3078.
- [43] Sebastian Weiss, Philipp Hermüller, and Rüdiger Westermann. 2022. Fast neural representations for direct volume rendering. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 196–211.
- [44] Sebastian Weiss, Mustafa Işık, Justus Thies, and Rüdiger Westermann. 2020. Learning adaptive sampling and reconstruction for volume visualization. *IEEE Transactions on Visualization and Computer Graphics* 28, 7 (2020), 2654–2667.
- [45] Pak Chung Wong and R Daniel Bergeron. 1994. 30 years of multidimensional multivariate visualization. *Scientific Visualization* 2 (1994), 3–33.
- [46] Jonathan Woodring, Susan Mniszewski, Christopher Brislawn, David DeMarle, and James Ahrens. 2011. Revisiting wavelet compression for large-scale climate data using JPEG 2000 and ensuring data precision. In *2011 IEEE Symposium on Large Data Analysis and Visualization*. 31–38. doi:10.1109/LDAV.2011.6092314
- [47] Skylar W Wurster, Hanqi Guo, Han-Wei Shen, Tom Peterka, and Jiayi Xu. 2022. Deep hierarchical super resolution for scientific data. *IEEE Transactions on Visualization and Computer Graphics* (2022).
- [48] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.

A Supplementary Material

A.1 Comparison of Gradient Quality


The gradient of scientific data variables has been shown to be an important quantity to carry out several data analysis tasks, such as importance-driven sampling and multi-dimensional transfer function design. Therefore, we conduct a comparison based on the quality of the gradient computed per grid point from the reconstructed volume data among the three methods. After reconstructing all the variables, we compute spatial gradient magnitudes for each grid location and then compute the PSNR value between the gradient magnitudes estimated from the ground truth data and the gradients from the reconstructed data. Table 11 shows the average PSNR values for all the variables for the three methods. It is observed that the proposed MVNet results in the most accurate gradients using the reconstructed data variable values for all the datasets. In Fig. 14, we show the visualization of the reconstructed gradient magnitude for a representative variable, Pressure (P), of the Isabel dataset. We observe that the TTHRESH (Fig. 14c) and LERP (Fig. 14d) generated gradient fields show visual artifacts, while MVNet (Fig. 14b) produces the most accurate gradient reconstruction when compared against ground truth (Fig. 14a).

A.2 Comparison with Copula-based Method: Visualization Result

In order to compare the quality of the reconstructed data variables visually, we show a rendering of a representative variable, Surface Temperature (TS), using MVNet and Copula-based summaries in Fig. 15. It is seen that MVNet produces a visualization (Fig. 15b) identical to the ground truth (Fig. 15a). In contrast, the visualization generated by Copula-based summaries (Fig. 15c) fails to preserve the detailed features of the TS variable as indicated by black dotted circles in the image.


Table 11: Quantitative comparison of average gradient magnitude values for all the variables between MVNet and other methods. It is observed that the PSNR (dB) value computed between raw data gradient and reconstructed data gradient is highest for MVNet compared to TTHRESH and LERP.

Dataset	MVNet PSNR ↑	TTHRESH PSNR ↑	LERP PSNR ↑
Combustion	52.930	49.566	35.179
Isabel	46.732	45.204	39.460
Climate50	38.453	31.909	38.015
Climate100	38.911	33.896	38.363




(a) Ground truth grad. mag. of P variable. (b) MVNet generated grad. mag. of P variable. (c) TTHRESH generated grad. mag. of P variable. (d) LERP generated grad. mag. of P variable.

Figure 14: Reconstructed gradient magnitude visualization for Pressure (P) variable of Isabel dataset. From left to right, ground truth, MVNet-generated result, TTHRESH-generated result, and LERP-generated result are provided. We observe that MVNet produces the most accurate gradient reconstruction, while the TTHRESH and LERP-generated images have artifacts, resulting in reduced accuracy.



(a) Ground truth visualization of TS variable. (b) MVNet generated visualization of TS variable. (c) Copula-based summary generated visualization of TS variable.

Figure 15: Visualization of Surface Temperature (TS) as a representative variable from the Climate100 dataset. Fig. 15a shows the ground truth, Fig. 15b shows the result reconstructed by MVNet, and Fig. 15c depicts the visualization reconstructed using the Copula-based summary data. It is observed that MVNet accurately reconstructs the TS variable. In contrast, the copula-based method fails to preserve the detailed features in the data, as highlighted by black dotted circles in the image.



(a) GT (b) MVNet (c) TTHRESH (d) LERP

Figure 16: Reconstructed data visualization for Pressure (P) variable of Isabel dataset. Images generated by ground truth (GT) data, MVNet, TTHRESH, and LERP are shown from left to right. We observe that the generated image is the most similar to the ground truth, while the TTHRESH and LERP-generated images have artifacts (as shown by the red dotted box).




Figure 17: Query-driven visualization for Hurricane Isabel dataset. A multivariate query: (CLOUD >0.0001 AND <0.002) AND (PRECIPITATION >0.0001 AND <0.0093) AND (QVAPOR >0.01 AND <0.0235) is shown. We observe that MVNet produces the most accurate result among the three methods.




Figure 18: Query-driven visualization for the Combustion dataset. A multivariate query: (Mixfrac >0.3 AND <0.7) AND (Y_OH >0.006 AND <0.1) is shown. We observe that while MVNet produces the most accurate result quantitatively, visually, MVNet and TTHRESH produce similar results, whereas LERP is the least accurate.




Figure 19: visualization of isosurface of Velocity Magnitude at isovalue=41 of Isabel dataset with ground truth (GT) shown in Fig. 19a. Isosurface produced by LERP is the least accurate, while MVNet and TTHRESH produce visually comparable results.