VECNA: Visual Exploration, Comparison and Analysis of Reconstructed Spatiotemporal Scientific Simulation Data

A. Mishra[1], A. Biswas[2], and C. Bryan[3]

[1]Fujitsu Research, USA [2]Los Alamos National Laboratory, USA [3]Arizona State University, USA

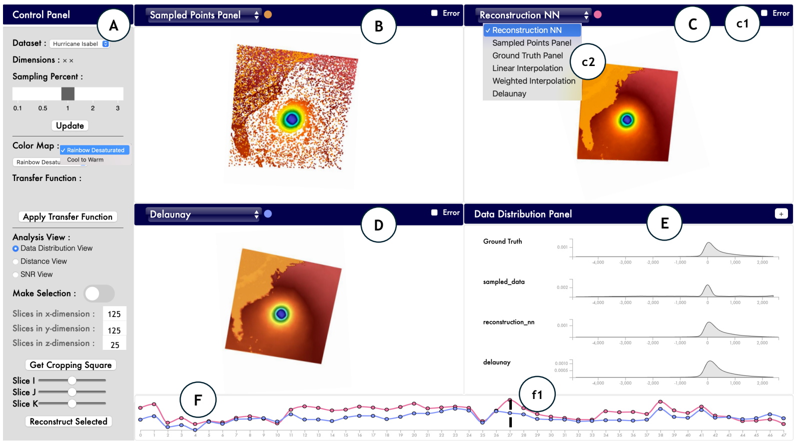

Figure 1: VECNA*'s interface consists of six linked panels which support (A) selecting datasets and rendering options, (B–D) exploring and comparing different dataset visualizations and reconstruction methods, (E) analyzing the data statistics and distributions of the datasets, and (F) reviewing contextual signal to noise ratios across timesteps.*

Abstract

Data-driven sampling and reconstruction techniques are increasingly being employed in scientific computing applications to aggressively reduce data volumes while retaining the crucial features of spatiotemporal datasets. Such data must be reconstructed for analysis, but it is difficult for domain experts to assess reconstruction quality, particularly given the pace at which new methods are being developed and a lack of support in existing tools. To help address this, we introduce VECNA*, a visual analytics system for exploring and comparing reconstructed scientific datasets. Developed through collaboration with high-performance computing researchers,* VECNA *enables intuitive qualitative and quantitative comparisons among diverse reconstruction methodologies. We validate* VECNA *via a usage scenario and empirical expert assessments to demonstrate its efficacy in empowering users to discern nuances in reconstruction quality and identify regions of interest within datasets, facilitating more informed subsequent analyses.*

CCS Concepts

• ***Human-centered computing*** → ***Visualization systems and tools;***

1. Introduction

Scientific simulations produce large volumes of data which create bottlenecks in network communication and disk I/O [BDPA18]. An increasingly prevalent strategy for addressing this congestion is to employ aggressive (sometimes called *extreme-scale*) data reduction techniques, which sample a very small subset (e.g., 1% or even 0.1%) of the data at a timestep for disk storage and post hoc analysis [BDL*21].

The necessary complement to sampling is *data reconstruction*, which reconstructs the full resolution of the dataset while aiming to preserve essential features [GBP*20]. Unfortunately, for extreme-scale sampling scenarios, this is a non-trivial task. The quality of a sampling/reconstruction can be influenced by several factors, such as the simulation type, sampling and reconstruction algorithm parameters (e.g., sampling rate), and the presence or complexity of specific subregions or features in the dataset. Moreover, there

DOI: 10.2312/evs.20251082

are now a variety of reconstruction methods available each with their own tradeoffs in terms of quality, time complexity, and accuracy [GBP*20]. Such confounds make it difficult for domain experts (e.g., scientists who perform data reduction and reconstruction in their workflows) to understand and compare reconstruction methods, and also to identify optimum reconstruction strategies for a given scenario. As such, several papers have highlighted a need for adaptable frameworks to accommodate not only diverse sampling and reconstruction pipelines [KH12, GBP*20, BDL*21], but also for tools to provide oversight, review, and comparison for reconstructed datasets and methods [GBP*20, BMM*24].

To help address this problem of reviewing and analyzing reconstructed datasets during post hoc analysis, we introduce a visual analytics tool called VECNA. VECNA facilitates the exploration of various reconstruction methods for 3D/4D voxel datasets, enables quantitative and qualitative comparison of reconstructed images, and provides tools for analyzing the reconstructed data. The system is algorithm agnostic, and can support various state-of-the-art sampling and reconstruction methods, including recent machine learning driven reconstruction approaches [BMM*24]. VECNA's design is based on a domain analysis and discussions with domain experts who work in extreme-scale sampling and reconstruction, and integrates several coordinated visualizations for examining single and multiple time steps. Specifically, the tool is designed for analysts who need to (i) compare multiple reconstruction pipelines or parameter settings side-by-side, (ii) blend quantitative quality metrics with qualitative inspection, and (iii) localize reconstruction failures across space and time. To validate VECNA, we present a use case and a set of expert review sessions, demonstrating its efficacy as a focused tool for assessing reconstruction methods.

2. Related Work

Sampling and Reconstruction Approaches. Sampling methods are widely used in the scientific and high-performance computing community to reduce the size of large-scale data sets; two recent surveys include discussions of sampling and reconstruction methods [WH23, SZD*23]. Efforts in extreme-scale sampling have focused on optimizing the in situ sampling process for post hoc analysis. For example, Woodring et al. [WAF*11] proposed a stratified random sampling based algorithm to enable downstream interactive visualization tasks. More recent works by Biswas et al. [BDL*21, BDPA18] have proposed techniques that prioritize important data features and gradient properties, ensuring the extraction of important data features given a storage constraint.

Similarly, the reconstruction of volumetric data is also a focus of study, particularly by the scientific visualization community. While there are many "straightforward" algorithms available, such as interpolation schemes based on Delaunaey triangulation, nearest neighbors, and radial basis functions, other approaches have proposed various machine learning methods to upscale or reconstruct volumetric data [WCTW21, GYH*20, HWG*20, BMM*24] (also see Wang et al.'s survey on deep learning in scientific visualization for a discussion on this [WH23]). Unfortunately, such methods provide little in the way of interactive tools to allow users to assess and compare reconstruction quality.

Visual Analytics for Scientific Datasets. For analyzing large 3D/4D scientific datasets, tools like VisIt [CBW*12], ParaView [par], Z-checker [TDG*19] and Foresight [GBP*20] represent mature applications that are widely used. In particular, ParaView represents a *de facto* standard for the community and is widely used across simulation and computing domains. In contrast, Z-checker and Foresight primarily focus on analyzing compression methods on scientific data. One drawback for these tools is they are not optimized for reconstruction assessment. For example, ParaView lacks streamlined side-by-side error overlays, distribution-based comparison views, and timeline-linked metrics that were considered essential for rapid, task-focused analysis in our requirements analysis. This lack of a purpose-built tool was a primary motivation for developing VECNA.

3. Design and Implementation

Requirements Analysis. VECNA was motivated by discussions with scientific computing researchers and a meta-analysis of limitations in current scientific visualization and analysis tools (including ParaView, Z-checker, and Foresight), as well as recent papers that discuss open issues for reconstructed spatiotemporal datasets (e.g. [GBP*20, BMM*24]). Based on this process, we distilled a set of four high-level design goals **DG1–DG4**:

DG1: Provide avenues for users to easily explore various reconstruction methods. Because the number of reconstruction methods is only increasing, users should be able to *easily and interactively import, explore, and compare across methods and datasets*, facilitating comprehensive analysis and informed decision-making.

DG2: Enable both quantitative and qualitative analysis. Two reconstruction methods may yield visually similar outputs at first glance, but it is imperative for users to discern differences not only among the methods themselves, but also in comparison to the ground truth. Therefore, tools *should facilitate holistic assessments*, i.e., both quantitative analysis and qualitative inspection.

DG3: Support identifying regions of interest (ROIs). Comparing different reconstructed images can help users identify high-level differences between various methods, but *fine-grained analysis can help users identify regions where reconstruction fails (or performs poorly)*. Such granularity can empower users to make iterative adjustments for subsequent downstream analyses.

DG4: Support temporal analysis. Scientific simulations are often complex spatially and temporally. Enabling users to navigate through temporal snapshots enhances their understanding of reconstruction methodologies and amplifies the contextual understanding of their reconstruction methodologies. However, existing visualization frameworks and tools often do not facilitate the easy time-varying analysis of datasets. *Hence, tools should incorporate seamless transitioning between timesteps, enabling users to delve into exploration and analysis with ease.*

The VECNA System. VECNA is designed to support **DG1–DG4** by allowing users to review, analyze, and compare reconstructions of 3D/4D voxel datasets. Figure 1 shows the interface, which is composed of six linked panels (A)–(F). The application interface is primarily written in D3.js and VTK.js, build atop a web server for data processing and handling. The codebase is available at: Anonymized for submission.

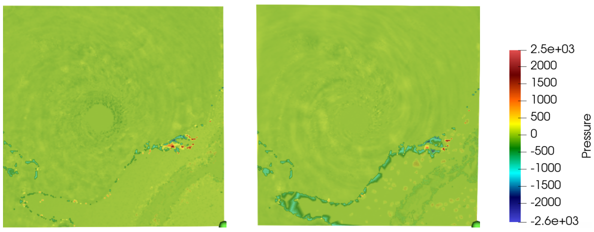

(a) FCNN (b) Delaunay

Figure 2: *Using* VECNA *to examine how reconstruction error for the Hurricane Isabel dataset (from a 1% sampling percentage) differs between FCNN and Delaunay triangulation methods.*

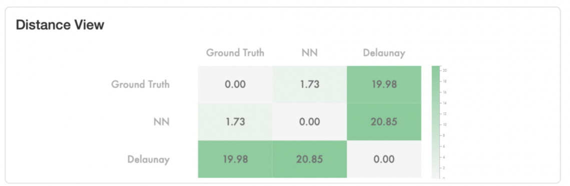

(a) Pairwise distances between Neural network and Delaunay reconstruction with Ground Truth

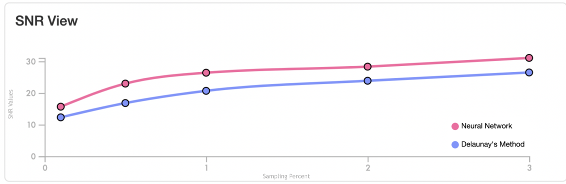

(b) SNR values over sampling percentages for Neural network (pink) and Delaunay (blue) reconstructions

Figure 3: *Examples of analysis metrics provided in* VECNA.

(A) The **Control Panel** supports loading datasets, adjusting sampling percentages for analysis, and editing colormap and transfer functions. In addition, users can toggle the charts shown in the Data Distribution Panel (E, described below), and apply slice operations (which are uniformly applied across the visualization panels).

(B, C, D) To support **DG1**, visualizations of the dataset and its reconstructions can be loaded into a trio of **Image Panels**. Users can select the dataset for each panel (see (c2); options include the dataset's ground truth if available (i.e., the full resolution data), showing the sample points only, and of the available dataset reconstructions based on the methods that have been applied thus far. For example, Figure 1 shows a visualization of the sampled points (at a 1% sampling percentage) and two reconstructions: one using a neural network and one based on interpolation using Delaunay triangulation. Each panel independently supports zooming, panning, and rotating, and an error checkbox for each panel (e.g., (c1)) overlays the reconstruction error on the dataset (if applicable, see Figure 2 for an example), supporting **DG3**.

(E) The **Data Distribution Panel** helps support **DG2** by letting users tab between three analysis charts. A distribution view (shown in Figure 1(E)) aligns a small multiples plot of kernel density functions mapping the scalar value distributions across the selected panels (plus the ground truth, if available). Alternatively, users can switch to an adjacency matrix or line chart (see examples in Figure 3). The matrix summarizes the pairwise distances between different datasets (we currently show Wasserstein distance, though VECNA can supports other methods such as KL divergence), and the line chart plots the signal-to-noise ratio (SNR) across available sampling percentages for the datasets.

(F) Finally, the **Timeline Selector** panel shows how the SNR (for the current sampling percentage) changes across simulation timesteps for the shown reconstruction methods, and allows users to click across the timeline to update the data that is shown in the other panels (see (f1)). Line colors in this plot correspond to the Data Distribution Panel's line charts, and also to the colored circles beside each dataset label in the header bar of the Image Panels.

4. Results

To help demonstrate VECNA's utility, we first present a brief usage scenario, and then report an empirical evaluation conducted with domain experts. For testing data, we use datasets generated from two well-known simulations: Hurricane Isabel [Har02] and ExaAm [TBB*22]. (Note that VECNA supports larger datasets: we choose these due to their popularity as testing datasets in scientific computing, and because they are representative of the types of simulations that require aggressive sampling and reconstruction while accounting for the presence and interactions of complex spatiotemporal features.)

For data sampling, we employ a state-of-the-art value-based sampling method introduced by Biswas et al. [BDL*21]; for reconstruction, we include a number of reconstruction methods that were recently studied in a reconstruction experiment paper [BMM*24], including simple interpolation using Delaunay triangulation, weighted interpolation using pairwise distances, and training a fully-connected neural network (FCNN). Like above, we note that VECNA is flexible to supporting other types of sampling and reconstruction methods.

Gary's Usage Scenario. Gary, a scientific computing researcher, wants to assess and compare various reconstruction methodologies using a Hurricane Isabel simulation dataset, which simulates the evolution of a hurricane off the coast of Florida, comprising thirteen attributes across 48 timesteps. He focuses on the pressure attribute, which is a key indicator of hurricane intensity [Har02]. For example, the hurricane's eye (a low pressure zone) is depicted as a prominent blue circle in Figure 1.

Gary's work is shown in Figures 1–3. He begins by selecting the 1% sampling level and populating the three Image Panels. Clicking the error checkboxes on the two reconstruction images (showing Delaunay triangulation and FCNN results) overlays the error values on the reconstruction datasets (Figure 2). He can immediately see that Delaunay has more significant errors in the regions that trace the Florida coast, compared to the FCNN. Looking at the Data Distribution panels, he can see that the neural network is significantly more accurate compared ot the ground truth (Figure 3(a)), and that this trend holds across SNR values at different sampling percentages (Figure 3(b)). Further analysis across timesteps (not shown due to space constraints) confirms that Delaunay-based reconstruction consistently under-performs compared to the FCNN model.

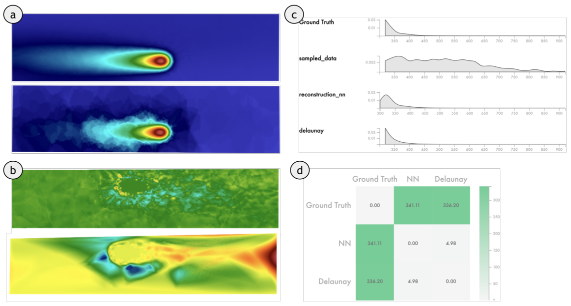

Figure 4: *(a) Example reconstructions of the ExaAM dataset by FCNN (top) and Delaunay (bottom). (b) Error renderings of the FCNN (top) and Delaunay (bottom) reconstructions. (c) Data distributions for ground truth, sampled data, FCNN, and Delaunay reconstruction datasets; note how the sampling results in a different distribution compared to the ground truth, but the reconstructions successfully capture (most of) the ground truth distribution. (d) Pairwise distances between distributions of the ground truth and reconstructed datasets.*

Expert User Study. To further evaluate VECNA, we conducted pair analytic sessions [AHKGF11] and semi-structured interviews with three domain experts, denoted as e1–e3. The goal of this evaluation was to better understand the benefits and limitations of VECNA for assessment, analysis, and comparison of reconstructed datasets and methods, and to learn how the system could be extended to better support their workflows in the future.

Expert e1 is a senior PhD student at Anonymized for submission; e2 and e3 are staff research scientists at Anonymized for submission with at least 5 years post-PhD research experience. All the experts had 3+ years of experience working on the reconstruction of large scientific datasets. One participant was also a co-author and one of the primary developers of Foresight [GBP*20].

Studies were conducted over Zoom. In accordance with pair analytics protocols, a study administrator "drove" the system according to the expert's directions. Sessions lasted as long as each expert desired, who could switch between the Hurricane Isabel and ExaAM datasets. After the pair analytics session concluded, each participant were asked to provide freeform commentary and feedback on the system, such as what features they liked, disliked, and found useful. Verbal feedback was qualitatively coded by the authors to identify major themes and takeaways; below, we briefly discuss four main ones.

An effective and focused tool for reconstruction analysis. All three experts stated that the interface was a useful analytic tool that would proactively help them to both qualitatively and quantitatively assess reconstructed images. All participants remarked that while they use commercially available tools (such as ParaView and VisIt), they preferred VECNA's focus in supporting the viewing and analyzing the nuances of reconstructed datasets. For example, e2 stated, *"This interface would be such a great tool to use, I don't think I have seen any interface for our use cases before."*. Likewise, e3 noted a lack of similar functionalities in ParaView, *"[t]his interface is really useful for people like us. It's very difficult and kind of unintuitive to do any kind of analysis on ParaView. This however gets straight to the point."* In particular, viewing the reconstruction error (e.g., Figures 2 and 4(b)) was considered highly useful. *"Honestly, it's harder to see any visible changes just from the reconstructed images a lot of times and not just for these datasets. The error rendering is pretty useful to check these differences out as we don't really care about every region being perfectly reconstructed but just the important ones"* (e2).

Analysis views provide deeper insights for reconstructed data diagnosis. All three experts found the charts in the Data Distribution Panel a useful complement to the Image Panel's 3D visualizations. *"The first thing I generally do is plot a distribution plot of my reconstructions to see if there are any bugs, so I definitely liked this view and the matrix view. It basically reinforces what I just saw."* (e1). In particular, e2 and e3 found this panel particularly useful when analyzing the ExaAM dataset (see Figure 4 for examples of charts they created): *"The distribution in the ExaAM dataset is very interesting. The Delaunay reconstruction distribution looks pretty similar to the ground truth data distribution, but it's error rendering gives it away"* (e2).

Timeline panel provided contextual analysis. Surprisingly, the timeline panel was one of the most positively spoken of features in the interface, particularly as a way to retain contextual information. *"The bottom timeline panel is the best thing about this interface. It gives me an overview of SNR values right away"* (e3). *"Not a lot of work has been done on sampling over time and generally just doing it on a time instance loses contextual information and how each time step is correlated to one another. So this SNR definitely is good"* (e2). All three experts suggested temporal analysis was an area that could be expanded in the future (e.g., showing more than just SNR values, or analyzing how errors evolve over time).

5. Discussion and Conclusion

This paper presents an initial design study on developing visual analytics interfaces for interactive assessment and comparison of reconstructed spatiotemporal datasets. Expert evaluations highlight its advantages over tools like ParaView. The system is dataset and method-agnostic, and we have future plans to incorporate additional reconstruction and sampling methods, while also expanding the interface and its features, based on study feedback.

One potential limitation in the current paper's evaluation is that we test on two "modest" sized datasets (Hurricane Isabel and ExaAM). While these are sufficient to validate the design goals and UI/UX aspects of VECNA, they certainly do not match the scale of simulations being run on today's exascale architectures. Supporting such datasets will require significant backend engineering, though we intend to develop such capabilities as we mature and expand VECNA's codebase and enhance the user interface (for example, by allowing users to manipulate the number of reconstruction panels they wish to show). Ultimately, as we continue to move through the exascale era, flexible and scalable visual analytics strategies will be critical for effectively analyzing reconstruction pipelines. VECNA provides a foundation for tackling these challenges, offering insights into how interactive, user-centric design can enhance large-scale scientific data analysis.

References

[AHKGF11] Arias-Hernandez R., Kaastra L., Green T., Fisher B.: Pair analytics: Capturing reasoning processes in collaborative visual analytics. In *2011 44th Hawaii International Conference on System Sciences* (2011), IEEE, pp. 1–10. 4

[BDL*21] Biswas A., Dutta S., Lawrence E., Patchett J., Calhoun J., Ahrens J.: Probabilistic data-driven sampling via multi-criteria importance analysis. *IEEE Transactions on Visualization and Computer Graphics 27*, 12 (2021), 4439–4454. 1, 2, 3

[BDPA18] Biswas A., Dutta S., Pulido J., Ahrens J.: In situ data-driven adaptive sampling for large-scale simulation data summarization. In *Proceedings of the Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization* (2018), ACM, p. 13–18. 1, 2

[BMM*24] Biswas A., Mishra A., Majumder M., Hazarika S., Most A., Castorena J., Bryan C., McCormick P., Ahrens J., Lawrence E., Hagberg A.: Filling the void: Data-driven machine learning-based reconstruction of sampled spatiotemporal scientific simulation data. In *SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis* (2024), pp. 290–299. 2, 3

[CBW*12] Childs H., Brugger E., Whitlock B., Meredith J., Ahern S., Pugmire D., Biagas K., Miller M., Harrison C., Weber G., Krishnan H., Fogal T., Sanderson A., Garth C., Bethel E., Camp D., Rübel O., Durant M., Favre J., Navrátil P.: VisIt: An end-user tool for visualizing and analyzing very large data. In *High Performance Visualization–Enabling Extreme-Scale Scientific Insight*. 2012, pp. 357–372. 2

[GBP*20] Grosset P., Biwer C., Pulido J., Mohan A., Biswas A., Patchett J., Turton T., Rogers D., Livescu D., Ahrens J.: Foresight: Analysis that matters for data reduction. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis* (2020), IEEE, pp. 1–15. 1, 2, 4

[GYH*20] Guo L., Ye S., Han J., Zheng H., Gao H., Chen D., Wang J., Wang C.: SSR-VFD: Spatial super-resolution for vector field data analysis and visualization. In *IEEE Pacific Visualization Symposium (PacificVis)* (2020), pp. 71–80. 2

[Har02] Harper B.: Tropical cyclone parameter estimation in the Australian region. *Systems Engineering Australia Pty Ltd for Woodside Energy Ltd, Perth 83* (2002). 3

[HWG*20] He W., Wang J., Guo H., Wang K.-C., Shen H.-W., Raj M., Nashed Y., Peterka T.: InSituNet: Deep image synthesis for parameter space exploration of ensemble simulations. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2020), 23–33. 2

[KH12] Kehrer J., Hauser H.: Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Transactions on Visualization and Computer Graphics 19*, 3 (2012), 495–513. 2

[par] ParaView - Open-source, multi-platform data analysis and visualization application. https://www.paraview.org. 2

[SZD*23] Sarton J., Zellmann S., Demirci S., Güdükbay U., Alexandre-Barff W., Lucas L., Dischler J.-M., Wesner S., Wald I.: State-of-the-art in large-scale volume visualization beyond structured data. *Computer Graphics Forum 42*, 3 (2023), 491–515. 2

[TBB*22] Turner J., Belak J., Barton N., Bement M., Carlson N., Carson R., DeWitt S., Fattebert J.-L., Hodge N., Jibben Z., et al.: ExaAM: Metal additive manufacturing simulation at the fidelity of the microstructure. *The International Journal of High Performance Computing Applications 36*, 1 (2022), 13–39. 3

[TDG*19] Tao D., Di S., Guo H., Chen Z., Cappello F.: Z-checker: A framework for assessing lossy compression of scientific data. *The International Journal of High Performance Computing Applications 33*, 2 (2019), 285–303. 2

[WAF*11] Woodring J., Ahrens J., Figg J., Wendelberger J., Habib S., Heitmann K.: In-situ sampling of a large-scale particle simulation for interactive visualization and analysis. *Computer Graphics Forum 30*, 3 (2011), 1151–1160. 2

[WCTW21] Weiss S., Chu M., Thuerey N., Westermann R.: Volumetric isosurface rendering with deep learning-based super-resolution. *IEEE Transactions on Visualization and Computer Graphics 27*, 6 (2021), 3064–3078. 2

[WH23] Wang C., Han J.: DL4SciVis: A state-of-the-art survey on deep learning for scientific visualization. *IEEE Transactions on Visualization and Computer Graphics 29*, 8 (2023), 3714–3733. 2