Volume XXX (2018), Number XXX pp. 1–14 COMPUTER GRAPHICS *forum*

Making Sense of Scientific Simulation Ensembles with Semantic Interaction

M. Dahshan[1], N. F. Polys[1], R. S. Jayne[2], and R. M. Pollyea[2]

[1]Department of Computer Science, Virginia Tech, USA
[2]Department of Geosciences, Virginia Tech, USA

Abstract

In the study of complex physical systems, scientists use simulations to study the effects of different models and parameters. As they seek to understand the influence and relationships among multiple dimensions, they typically run many simulations and vary the initial conditions in what are known as 'ensembles'. Ensembles are then a number of runs that are each multidimensional and multivariate. In order to understand the connections between simulation parameters and patterns in the output data, we have been developing an approach to the visual analysis of scientific data that merges human expertise and intuition with machine learning and statistics. Our approach is manifested in a new visualization tool, GLEE (Graphically-Linked Ensemble Explorer), that allows scientists to explore, search, filter, and make sense of their ensembles. Our tool uses visualization and semantic interaction techniques to enable scientists to: find similarities and differences between runs, find correlation between different parameters, and explore relations and correlations between different runs and parameters. Our approach supports scientists in selecting interesting subsets of runs to investigate and summarizing factors and statistics to show variations and consistencies across different runs. In this paper, we evaluate our tool with experts to understand its strengths and weaknesses for optimization and inverse problems.

CCS Concepts

•***Scientific Visualization*** → *Ensembles , Sensemaking;*

1. Introduction

Recent advances in computing power and the availability of high-performance computing have led to the feasibility of running complex real-world simulations in an acceptable amount of time. Scientists usually need to run their simulations multiple times using different input conditions, simulation parameters, and simulation models. This supports the scientist in interpreting the variability in the system and gaining insights by alternating between models. Through these multiple runs, they can gain a more complete understanding of the simulated phenomenon and model, and refine their hypothesis and method for actual physical experiments. A set of simulation runs is known as an ensemble: it represents a parameter study or a set of studies using different computational models and paramters. Scientists from a variety of disciplines, such as aerodynamics, weather forecast climate, and computational fluid dynamics, use ensembles to simulate complex systems, explore unknowns in initial conditions, evaluate extreme cases, compare structural characteristics of their models, and investigate parameter sensitivity to assess the confidence in their findings. In other words, this guides the scientist in interpreting the distributions within the data, investigating the sensitivity of outputs to certain input parameters and understanding the similarities and dissimilarities between ensemble members.

The analysis of ensemble data is a challenging task due to its high multidimensionality, complexity, and size. Therefore, ensemble visualization is a crucial and essential component in the analysis process as it facilitates knowledge discoveries and helps the scientist see the characteristic features of the data through graphical representations. Such analysis of ensembles can help them find appropriate models and parameter ranges for hypothesized relationships and outcomes. Moreover, ensemble visualization helps in measuring the variability and sensitively of the model to its inputs and outputs and how output parameters react to input changes. Therefore, the focus of this paper is the visual exploration and comparison of the behaviors of simulations and their parameters.

Current research in the visual analysis of ensembles relies on multiple techniques for showing the variability of the ensemble members, major trends, and outliers. Some of these techniques focus on studying the parameter space and measuring the correlation between different parameters. Summary statistics [PKRJ10, BPFG11, PWB*09a, MWK14, WMK13, SEG*15], spaghetti plots [DNCP10, Det05], and probabilistic features such as multivariate Gaussian distributions, histograms and kernel density estimates (KDE) are examples of these techniques [PPH12, PW12]. Additionally, conventional visualization solutions such as glyphs [HLNW11,PKRJ10,PMW13,SZD*10] and visual variables

[BWE05, DKLP02], are used to model uncertainty in parameter space. However, some of these techniques have limited capabilities to show the intrinsic structures in the ensemble, and some of them are designed to work with 1D or 2D datasets.

Alternatively, other techniques study the shape and variability of the ensembles themselves. These techniques are divided into two broad groups, one that aggregates multiple ensembles creating an overview about the variability between ensembles while omitting potentially the details from the original ensembles [PWB*09b, CB12]. The second group supports the comparison between the ensemble members providing a better view but is limited to a certain number of ensembles [HHH15, FML16]. Most of these ensemble visualization techniques are tailored for time-varying ensembles for specific application areas (i.e., specifically weather and climate).

Our initial motivation was to develop a visual analysis tool that helps the scientists irrespective of their application discipline to make sense of their ensemble of runs. However, we found out the most of the techniques used focus on either parameter space or ensemble exploration with little intervention from humans. Although there is a disagreement about the optimal way to visually analyze ensemble data, we argue that integrating statistical, parameter, and ensemble displays in multiple coordinated views can help the scientists in exploring relations between ensemble members and parameters more efficiently. Displaying summary statistics only may suffer from misinterpretation due to size information. Alternatively, using ensemble display only may help in overweighting individual ensemble members leading to incorrect decisions. Therefore, integrating multiple displays using brushing and linking techniques for selecting or manipulating ensembles or regions of interest can help scientists gain a better understanding of their data as each visual display is designed to highlight certain aspects of the uncertainty within the data. As a result, we targeted a multi-view tool that helps scientists to make-sense of ensemble data during the interactive exploration process regarding both parameter space or ensemble runs while keeping human in the loop.

This paper proposes an interactive visual approach for exploring and analyzing the influence of parameters across the ensemble of simulations; we demonstrate a novel method for the exploration, comparison, and analysis of high-dimensional patterns and outliers. In other words, our tool not only enables scientists to identify the similarities and dissimilarities between ensemble members but also where and why these relationships exist. It also allows scientists to study sub-regions of interest in the simulation domain, helping them understand the correlations between parameters and result variables. Additionally, the statistical view would enable scientists to quantify and verify their hypotheses and validate their findings by presenting various summaries and descriptives of subsets. Our approach builds upon semantic interaction and statistical techniques with brushing and linking technique to visually analyze the high-dimensional parametric relationships within an ensemble. This enables our tool to work with any type of data, thus covering a broad range of application domains. In summary, the contributions of this paper are as follows:

- A new visual analysis approach that helps scientists make-sense of ensembles behaviors, patterns, and outliers for high dimensional data using both inputs as well as output parameters.
- Coupling interactive visual analysis and statistical summaries to detect and analyze characteristics of parameter spaces across multi-dimensional ensembles. In other words, determine the influence and sensitivity of different parameters of simulation results.

2. Related Work

2.1. Ensemble Visualization

Ensemble analysis and visualization belong to the area of uncertainty visualization, which is used for assessing uncertainty and variability of ensemble members. The common features of ensemble data are large size, high complexity, multi-dimensional, multivariate, and covers various scientific disciplines (i.e., such as ocean simulations [HMZ*14], weather and climate prediction and analysis [SZD*10, NFB07, PWB*09a, BLLS17], and high-energy physics [PPA*12]) that usually have temporal and spatial information, making its analysis and visualization a challenging task. This has motivated many researchers to develop a wide variety of techniques, approaches, and surveys [BHJ*14, MRH*05, BOL12, OJ14] to help scientists in analyzing the relationships between and within ensemble members.

Numerous approaches and tools have been devoted to fit existing visualization techniques to support the analysis of simulation ensembles. Some of them focused on the statistical distribution and properties of the ensemble members to provide quantitative information about the uncertainty in the data [HOGJ13, PWB*09b, SZD*10], while others visually explore ensemble runs, their parameter space, and the relationships and correlations between them. Summary based ensemble visualization techniques show the statistical distribution of the ensemble members (i.e., at least the mean, median, and some indication of the data spread) by offering different ways to visualize ensemble fields using color maps, contours, animation, opacity, boxplots, or glyphs [HLNW11, PMW13, ZWK10, CBDT11]. Box plots [MTL78] represents the distribution of the data by encoding the maximum and minimum values, mean, median, and other quartile information. Kao et al. [KLDP02] extended the traditional boxplot to support 2D ensemble data where the statistics are visualized on a 2D plane using color mapping and glyphs. They also used probability density function (PDF) as voxel values to handle statistics for 3D volume [KDP01]. Potter et al. [PKRJ10] proposed another extension of box plots that integrated multiple descriptive statistics (i.e., such as mean, standard deviation, kurtosis, skewness, and tailing information) into a a summary plot to convey additional information about the distribution of the data. Alternatively, Sanyal et al. [SZD*10] enhanced spaghetti plots by developing a tool, Noodles, that displays multiple isocontours using glyphs and confidence ribbons to highlight the spread of 2D contour ensembles.

Although these approaches showed different representations of statistical distributions, they can not work well with dense 2D/3D data and can not efficiently differentiate some distributions within one or two values. To overcome these limitations, Bensema et al. [BGOJ16] focus on the modality of data distributions of ensemble members to define high-variance locations in the ensemble while Chen et al. [CZC*15] differentiate the distributions of the ensemble

members having similar mean using uncertainty-aware projection scheme. Moreover, Whitaker et al. [WMK13] proposed the idea of contour boxplot, as an extension to traditional boxplot, to visualize the spread, outliers, and variability in ensembles of spatial contours. Inspired by contour box plots, curve box plot [MWK14] was proposed to handle spatial curves in the ensemble focusing on revealing the variability in ensemble pathlines. Both tools are based on the concept of statistical data depth that helps in showing the centrality of members of an ensemble and overcoming the visual cluttering limitations of spaghetti plots.

Inspired by contour variability plots, Demir et al. [DJW16] proposed a tool to analyze the central tendency of ensembles of 2D and 3D shapes using mixture models. Similarly, Ferstl et al. [FBW16] proposed an approach to statistically model the distribution of streamlines through deriving clusters that show the significant trends in the ensemble. Although these tools offer different visualization and visual analysis tools to understand the statistical distributions of output data, they did not consider the exploration methods for the input parameters, i.e., they do not allow the users to use input parameters in the analysis process. Moreover, most of these approaches only display statistical information about ensemble members (i.e, usually users are not able to interact with the statistical representations, they are just used for display) without connecting it to ensemble members which hinder a comprehensive analysis of the simulation features in the ensemble.

The visual analysis of ensemble parameter space to explore the relationships between the simulation parameters is a very challenging problem that attracted significant attention in the recent years. Bruckner et al. [BM10] proposed a density-based clustering of animation sequences to visually analyze parameter space to help visual effect designers in finding the appropriate parameters for the desired results. Similarly, Beham et al. [BHGK14] proposed the Cupid system for geometry generators that combines abstract parameter space of geometry generators with the output space of the resulting shapes to help users in identifying similar behavior in different ensemble members. Although these approaches incorporated analysis of input and output parameters, they are not designed to work with scientific data. Additionally, they are designed to analyze one simulation run at a time. On the other hand, Alabi et al. [AWH*12] compare ensemble members using a side-by-side comparison of surfaces in 3D to illustrate ensemble geometries and show the differences and similarities between surfaces. Phadke et al. [PPA*12] presented another technique for ensemble exploration and comparison using pairwise sequential animation and screen door tinting. Pairwise sequential animations use shape, color, and size for comparing data between pairs of ensemble members while screen door tinting shows the differences between ensemble members using value changes to field points. Similarly, Piringer et al. [PPBT12] compared 2D function ensembles on three levels of details (i.e., surface plot, domain-oriented, and member-oriented) to enables scientists to have domain overview as well as a detailed view of ensemble members. Alternatively, Hao et al. [HHBY16] presented a static ensemble visualization system that helps scientists to find interesting subsets of ensemble members using hierarchical clustering. Moreover, Demir et al. [DKW16] compared and visualized 3D scalar field ensembles using mean isosurface surrounded by a spaghetti plot of silhouettes of individual ensemble members while Hazarika et al. [HDS16] visualized ensemble isosurfaces using color-mapping to represent the distance from the median surface. Although these approaches offer different ways to visualize and compare different data types(i.e., 2D and 3D), they are limited in the number of ensemble members that they visualize.

Other visualization techniques for ensemble and parameter space representations embed brushing and linking technique into coordinated multiple views to show the different facet of ensemble members in different displays. MatkoviÄĞ et al. [MGKH09] proposed multi-linked views to visualize ensemble data as families of data surfaces with respect to pairs of independent data dimensions. Later, they extended their work to model three different levels of detail (using scatterplot, curve view, or a histogram) [MGJ*10] and to support interactive interaction plots for better exploration and analysis of high-dimensional parameters [SEG*15]. Similarly, Demir et al. [DDW14] proposed multi-chart visualization with brushing and linking to explore summaries of a 3D volumetric ensemble in different regions. Alternatively, Liu et al. [LS16] used a series of parallel coordinates plots (PCPs) in multiple coordinated views to explore relationships between scalar values in a multivariate context. Extending the idea of parallel coordinate plots to investigate parameter correlations, Wang et al. [WLSL17] proposed a nested PCP that combines superimposition and parallel design to explore intra-set and inter-set correlations between different parameters.

Similar to our proposed approach, Hollt et al. [HMZ*14, HdMRHH16], aboulhassan et al. [ASB*17], and Cibulski et al. [CKS*17] proposed multi-linked views that integrate ensemble visualization with statistical plots to facilitate the understanding of ensemble characteristics. However, all these techniques focus on a specific application domain that works well with spatial data. They also did not consider the effect of both inputs and outputs parameters on the simulation ensemble. Some of them focus on either ensemble members analysis or parameter space analysis. In contrast to those works, we focus on the effect of both ensemble parameters (i.e., inputs and outputs) and ensemble members during the simulation ensemble analysis. We also incorporate human in the loop through semantic interactions to power the visual analysis process.

2.2. Spatial Sensemaking and Semantic Interaction

Visual analytics tools help analysts to get insights about the data through a sensemaking loop [PC05]. Spatial sensemaking environment shows how users interact with a spatialization for exploratory data analysis. To gain better insights about the data, spatial sensemaking is usually integrated with semantic interactions. Semantic interaction merges the foraging abilities of statistical models with synthesizing process to keep the sensemaking loop tight [HBM*13,EFN12a,EBN13]. This helps users to test hypotheses without the need to understand the underlying models. Prior research has shown some tools for spatializations which externalize knowledge to steer the sensemaking process [AEN10].

However, this externalization requires the use of control panels outside the spatial metaphor [TG07]. Additionally, these tools do not scale well,so to resolve this problem, the user interaction should be within the spatial metaphor. Dust & Magnet [SYMSJ05]

uses parametric interaction to adjust the model's parameters within the same spatial metaphor. It uses dimension weights as parameters for their model. These weights are represented with magnets in the visualization layout. The user can interact with these weights (parameters) to update their importance. Although this provides an intuitive way to control parameter within the spatialization, these magnets are attributes of the data, not data itself.

On the other hand, tools such as ForceSPIRE [EFN12a, EFN12b], Dis-function [BLBC12], Andromeda [SH] and StarSPIRE [BNH14] allow the user to interact with data points and translate this feedback through dimensionality reduction algorithm to a new view that reflects user's interaction. This helps in providing an intuitive space for strengthening insight creation and data understanding. High dimensional data, in particular, is hard for users to understand because humans are limited in the number of dimensions that they can think of simultaneously. Therefore, a number of dimension reduction techniques have been developed and incorporated into interactive visual analytics to make the data more manageable [SH]. The work done in this paper builds on Andromeda [SH]. Andromeda is a visual analytics interface for exploring high-dimensional data using observation-level interaction (OLI)(i.e., type of semantic interaction), parametric interaction and multi-dimensional scaling (MDS) (i.e., dimension reduction technique) enabling users to gain more complex insights and accomplish new types of tasks.

3. Approach

The fundamental objective of any ensemble visualization is to analyze and make sense of the relationships among the ensemble members and their various parameters. In this paper, we demonstrate a multi-view visualization tool that uses semantic interactions, statistical visualizations, and linking and brushing technique to help scientists visually explore and analyze simulation ensembles. Our tool, GLEE (Graphically-Linked Ensemble Explorer), demonstrates the power of our approach. The main advantage of our approach is enabling scientists to understand and make sense of both runs and parameters in the same visualization application. Existing research seems to focus on supporting insight into one or the other. Moreover, most of the prior research uses only linking and brushing for human interaction with the visualization. Our novel contribution is using Semantic Interaction so that the ensemble view is not just a display view, users can directly manipulate runs of an ensemble spatially, in order to make sense of their relationships and explore how various parameters are related.

3.1. Visualization Domains

3.1.1. Parameter Domain

A parameter setting is a combination of values used for different parameters in one run of the simulation. Multiple runs of the simulation model use different parameter settings. A collection of all parameter settings used in an ensemble is a parameter domain. Considering each parameter as a dimension, multiple parameters with different settings form a high-dimensional space.

3.1.2. Ensemble Domain

The output from one execution of the simulation model constitutes one ensemble member. Multiple simulation runs produce multiple ensemble members. All members of one ensemble usually contain the same number of parameters but with different values.

3.2. Task Analysis

Working closely with scientists from different domains, we observed that most of them have mutual interests in analyzing parameter and ensemble domains. They almost share the same analysis procedure and have the same requirements. Based on these observations, we define their requirements into three basic tasks:

- Understanding the correlations between parameters in single and multiple runs: what is the correlation between two or more specific parameters? How does the correlation change over the ensemble runs? Are there any global or local correlations between parameters and ensemble runs? Answering these questions will help scientists gain a better understanding of the parameter space.

- Understanding and determining the sensitivity and influence of parameters on different ensemble members. Does changing the parameter values influence the ensemble members? Does one or more parameter affects different ensemble members? Answering these questions will help the scientist in predicting how changes in one or more of the parameter values influence the simulation output.

- Understanding the correlations between different ensemble members: Finding the distribution and similarities between similar ensemble members. Finding differences between dissimilar ensemble members. Understanding these similarities and differences will help scientists have a better understanding of the ensemble space.

3.3. Design

The final design of this tool went through several stages. Initially, we started to observe and examine the steps used by scientists to understand and explore relations in their ensemble. Involving scientists from a diverse group (i.e., Physics and Geoscience) in the design process and gathering their feedback helps in preventing from building ad-hoc tailored solutions that work with specific application domains. We noticed that most of them usually apply a manual analysis process that relies heavily on trial and error, which consumes a lot of time and can easily lead to mistakes. So, we began a preliminary visual design for the system that used only Semantic Interaction to help scientists analyze and understand the hidden patterns and relationships between ensemble members.

However, during our meetings, semi-formal interviews, focus groups and discussions with scientists, we find that semantic interactions alone will not be sufficient. We find that scientists are

interested not only in viewing the similarities and differences between the ensemble members but were also interested in considering the statistics underlying these ensemble members. Therefore, we decided to alter our design to have both the ensemble members and statistical measures in the same visualization to decrease the cognitive load taken by scientists in using different tools to analyze their data.

Additionally, we observed that most of the ensemble outputs are 2D/3D objects, so instead of displaying the ensemble members as points in the spatial space, we decided to use the output image of the simulation run to represent the ensemble members in ensemble view. These images would help the scientists to interpret the relationships between ensemble members visually and this may lead to a better analysis of the data in terms of time and efficiency. Moreover, we added a camera slider to allow users to view 3D objects from different camera perspective for a better understanding of their data. In order to make our tool applicable to a wide range of scientific disciplines, we included a variety of statistical measures. We found that a combination of boxplots, scatterplots, and parallel coordinates provided sufficient statistical measures to cover different levels of detail that aided the scientists in understanding the patterns and outliers within their datasets.

3.4. Method Overview

Our method starts with an ensemble $(s_1, ..., s_N)$ of N 2D images, visualized outputs from each run. Initially, we spatialize the ensemble using its input and output parameters, where each parameter represents a dimension forming high dimensional data, in two dimensional space using weighted multidimensional scaling (WMDS). In the spatialization, distance reflects relative similarity; e.g., two ensemble members (i.e., runs) close to each other in the spatialization have more similar parameters than ensemble members far away from each other. To set the spatial coordinates of the runs, WMDS relies on the ensemble parameters. We refer to the parameters as variable weights because parameters with large weights are considered more heavily in the spatialization than those with low weights.

Thus, one can deepen their interpretation of a visualization tool by considering both the distance and the weights. For example, two ensemble members sitting near each other in a spatialization are more similar to each other in terms of parameters values and their corresponding weights than two points far away from each other. In other words, ensemble runs closely positioned in the low dimensional layout (i.e., 2D) indicate similar parameters between these ensemble runs in the high dimensional space.

Our tool also uses several statistical visualization tool such as boxplot, scatter plot, and parallel coordinates to understand the correlation between different parameters. Users can use these statistical measures to identify areas of interest, ask quantitative questions about the ensemble behavior, and explore the distribution associated with the data between the different linked views.

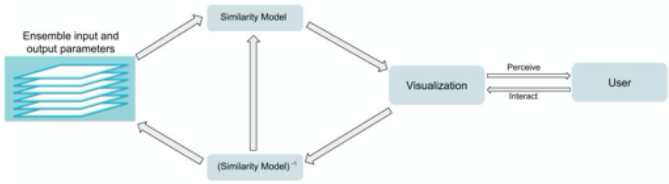

Figure 1: *Pipeline of our comparative visualization*

4. System

4.1. Ensemble simulation parameters

Scientific data in most emendable simulations is considered, to a certain degree, as a table with rows and columns. Each row represents a simulation run, and the columns represent the simulation parameter. In our model, we are using both inputs and output parameters as an access point for providing data to our visualization tool. We argue that considering both inputs and outputs for representing the ensemble visualization will give more accurate results in analyzing the relationships between the different ensembles and their parameters. Initially, users uploads all their data, which represents the input and output of all simulation runs, that is passed to the pipeline. During the pipeline initialization, a weight vector corresponding to all parameters from the ensemble is created. Each parameter inside the weight vector is assigned an initial default value of $1/k$, where k is the number of parameters. All the weights inside the weight vector sum to one. The weight vector is then passed down the pipeline to the similarity model for processing.

4.2. Similarity Model

The Similarity model has a forward and inverse (i.e., backward) implementation. The forward algorithm defines how data is processed for projection in the visualization layout. Therefore, it relies on the ensemble parameters and their weights. In contrast, the inverse algorithm responds to user interactions by updating the ensemble parameters or manipulating the ensembles spatially. This implies that that inverse algorithms receive parameter changes from the visualization layout. In this case, an iteration in the pipeline typically begins by running the inverse algorithms from the inverse similarity model through to the ensemble data. After that, the forward model is executed and the results are projected on the visualization layout.

4.2.1. Forward Similarity Model

The forward similarity model is concerned with creating a 2D projection of the runs for visualization. The 2D dimensional projection is created using Weighted Multidimensional Scaling (WMDS), which aims to preserve pairwise dissimilarities among ensemble members. For instance, if two ensembles are far apart in the high dimensional space, then they should be far apart in the low dimensional projection. Using the normalized values for ensemble input and output parameters, it begins by calculating the pairwise similarities between all pairs of ensembles using Euclidean distance. The Euclidean distance between ensemble members u and v, incorporating attribute weights w, where the weights from the visualization

Figure 2: *The main interface of our visual exploration tool for ensemble simulation analysis (a) Ensemble view: each image represents an ensemble member laid out spatially using WMDS (b) Parameter view: shows weight slider for ensemble inputs and outputs (c) Statistical view: displays statistical properties and distributions of ensemble using different graphs.*

weight vector are used to mark the importance of the parameters to the current projection. These pairwise similarities are then fed into WMDS which determines the location of each ensemble in the low dimensional (i.e., 2D) space by optimizing the following equation:

$$r = \min_{r_1,\ldots,r_n} \sum_{i=1}^{n} \sum_{j>i}^{n} \left(|dist_L(r_i, r_j) - dist_H(w, d_i, d_j|), \right) \quad (1)$$

where $|\cdot|$ indicates the absolute value, $dist_L$ is the Euclidean distance between low dimensional points, and $dist_H$ is given by Equation

4.2.2. Backward Similarity Model

The backward similarity is triggered when the scientist either moves an ensemble member within the visualization layout to assert some knowledge the scientist has about the similarity or dissimilarity between ensembles using Observation-Level Interaction (OLI) [EHM*11] or when he increases/decreases the weight of a certain parameter using parameter sliders. When using OLI, the new low dimension positions of the moved ensemble are fed into an optimization algorithm that looks for the corresponding set of weights for the parameters that best match the new positions of the moved ensemble. The optimization algorithm is represented by the following equation:

$$w = \min_{w_1,\ldots,w_k} \sum_{i=1}^{n} \sum_{j>i}^{n} \left(|dist_L(r_i{}^*, r_j{}^*) - dist_H(w, d_i, d_j)| \right) \quad (2)$$

Alternatively, when the the weight of a certain parameter is changed (i.e., the scientist moves the slider to increase or decrease the importance of a certain parameter using parametric interaction(PI)), other weights are updated so that they sum to one. The updated weight vector is then fed to the forward similarity model for projecting the ensemble members based on the new weights.

4.3. Visual Encoding Method

In the previous sections, the main components of the pipeline are described, but there is still a need to show the communication between the coordinated multiple views. The system consists of several linked views. In the following we detail the three main data views and present interaction and linking capabilities.

4.3.1. Ensemble View

The ensemble view displays the low-dimensional (2D) projection of the runs where each ensemble member is its own 2D graph. Each image represents the final output of the simulation run which could be an output of a static dataset or a final step in a times series. Depending on the type of data, the image can be 2D or 3D (i.e., figure 3). The image could be directly produced from the simulation run (figure 3a) or processed using visualization tool (i.e., Paraview [Aya15]) or libraries (i.e., Matplotlib or R graphics packages) as in (figure 3b and 3c). Additionally our tool has a Cinema slider that shows images (i.e., 3D) in the ensemble view, based on several camera positions, allowing scientists to interactively navigate, analyze, and make sense of ensemble runs (figure 4).

The projected distances among ensemble members encode the similarities of the runs in high-dimensional space including input and output parameters; thus, similar runs are placed near each other and dissimilar ones are set farther apart. The scientists can then interact with the runs to gain a better understanding of hidden patterns and outliers. These interactions include moving an ensemble member spatially within the visualization, zooming, lasso selection, multi-selection, and changing camera position to view the ensemble from different camera angles.

Moving ensemble members: scientists can move ensemble members within the spatialization layout (i.e., perform an Observation-Level Interaction, or OLI) to express knowledge or test hypotheses using the âĂIJUpdate LayoutâĂİ button. OLI is an automated procedure that transforms user interactions with data visualizations (visual feedback) to the parametric feedback that in turn adjusts an entire visual space. After moving the points to their desired locations, an update message is sent down the pipeline with the coordinates of the moved ensemble members. This information is used by the backward similarity model to calculate the new set of weights that closely reflect the relative pairwise distances between the trans-

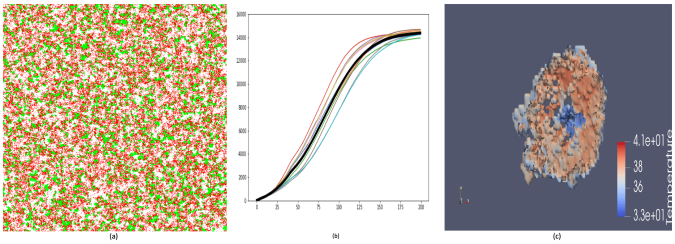

Figure 3: *Examples of simulation output images displayed in ensemble view a) shows the population of prey and predator at the final time step using Lotka Volterra model on 2D lattice b) shows mortality rates of different populations c) shows isosurface of saturated aqueous fluid colored by temperature for the final time step of a simulation injecting CO_2 in a rock*

ferred members. The updated weights along the with the ensembles parameters are fed to the forward similarity model for projection.

Selection: Scientists usually observe regions of interesting patterns and want to explore them more. Our tool offers different selection mechanisms such as lasso and multiple selections to give them the freedom to select these sub-regions. The advantage of having a selection mechanism is not only to further inspection and analysis for regions of interests, but also in determining local or global uncertainties in the data that helps in eliminating misinterpretations and false assumptions about the ensembles correlations and relationships to parameters.

4.3.2. Parameter View

Exploring parameters' influences on the simulation ensemble has equal importance to the visual comparison and analysis of the ensemble members to identify their features and behavior patterns. So, we added a parameter view. Our parameter view displays all the ensemble parameters using sliders. Each slider represents the weight of this parameter. This weight represents importance of this parameter. Scientists use parametric interaction to allow the scientists to move the slider so that they can increase or decrease the importance of specific parameters. The parameter view is linked with ensemble view, so when the user moves the slider (either increasing or decreasing), an update message with the new weight is sent to the backward similarity model. The weight vector is updated by assigning the new weight to the designated parameter, and the weights of other parameters are updated so they sum to one. These weights along with parameters are fed into forward similarity model for projection on the ensemble view. By moving the slider, the ensembles in the ensemble view are affected, and their spatial position is changed based on the 2D projection of the updated weights. In this case, scientists can use parametric interactions to see what is the most influential parameters and how these parameters are affecting the ensemble groups in the ensemble view. Similarly, when scientists do an OLI operation by spatially moving ensemble members to find the hidden relations between ensembles, using the link with parameter view, the weights on sliders are changed and this helps in identifying the most common and most influential parameters between different clusters of ensembles.

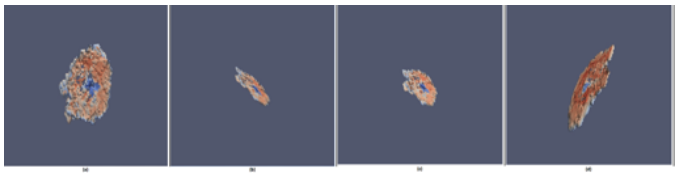

Figure 4: *Different camera positions for an ensemble member driven with a from Cinema viewpoint slider a) Phi= 0 and Theta= 1 b) Phi= 1 and Theta= 0 c) Phi= 1 and Theta= 2 d) Phi= 5 and Theta= 5*

4.3.3. Statistical View

One of the most challenging tasks when analyzing and making sense of high-dimensional data is identifying the regions of variability in the domain across all ensemble members, in addition to, determining the associations between interrelated variables. Although parameter and ensemble views enable the scientists to find the most influencing parameters and associations between parameters and ensemble members, a statistical view is still needed as it allows the scientists to determine the regions of variability in their data. Scientists usually have some understanding of the relationships between parameters, but unexpected discoveries are hard to find using parameter and ensemble views only.

Moreover, scientists can use the statistical view to have an initial idea about the distributions of the data as well as the correlation between parameters before manipulating the parameter and ensemble views. In other words, the statistical view can be used to improve the accuracy and understanding of findings discovered by parameter and ensemble views and can detect hidden relations not recognized by other views which translate into better understanding of the whole simulation model. Based on our observations of the analysis process taken by scientists, our statistical view is composed of three main graphs: box plot, scatter plot, and parallel coordinates.

- Box plot: is used to display the main features of ensemble attributes such as median, quartiles, and outliers which helps in showing the distribution range in data. Although box plot can suffer from visual clutter in case of large datasets, we believe this is not the case with our tool as our tool is designed to work with hundreds of ensemble runs.
- ScatterPlot: is used to explore the correlations and trends between different parameters. The scatter plot can help in giving more information about the relationship between attribute and eliminate any bias produced by summary displays [CH17].
- Parallel Coordinates: represents all parameters (i.e., inputs and outputs) in a single display as polylines on the parallel axes. This helps in displaying an overview of the whole data, exploring high dimension data and showing relations between multiple parameters at the same time.

All our statistical displays support interactive brushing of the data points to select regions exhibiting specific statistical properties (i.e., interesting patterns) and this provides immediate visual feedback to other views and statistical displays.

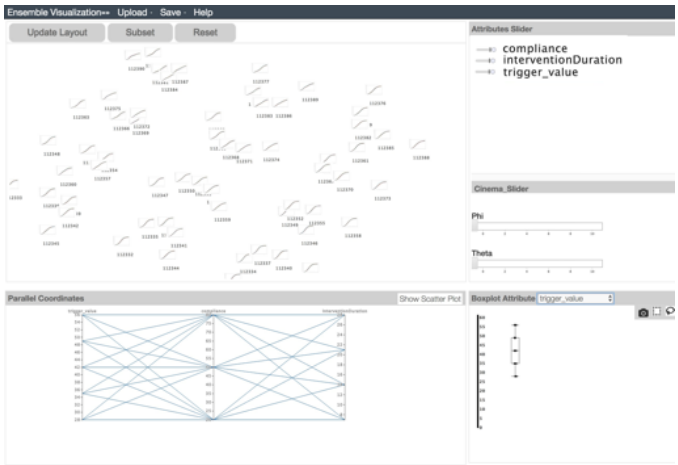

Figure 5: *Example 1: population levels in an agent-based simulation ensemble*

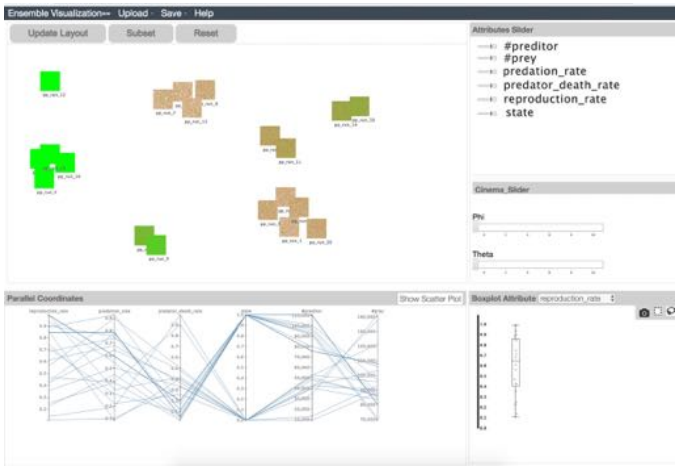

Figure 6: *Example 2: Prey and predator population in Monte Carlo simulation ensemble*

5. Implementation and Use Cases

The tool is implemented as a Web application that can be used remotely (client-server mode) or locally. The user interface is implemented by using DataDriven Documents (D3) [BOH11] while the backend algorithm (i.e., pipeline) is implemented using python. Visualization preparation and Cinema image database generation [OAJ*16] is done using Paraview [AGL*05].

To demonstrate the importance and usefulness of our tool to analyze and make sense of scientific data, independently of the underlying field of science, we will demonstrate three applications: geoscience, population health, and ecology. In addition, we evaluated our tool in detail with domain experts in geoscience simulation of fluid transport through porous media. These examples and the evaluation results illustrate how scientists can use our tool to interpret and analyze their data. The results emphasize the effectiveness of our tool in showing trends in the data and helping scientists to find correlations between parameters and ensembles irrespective of the field.

The first demonstration application is an agent-based simulation used to evaluate allocation of resources in emergency situations. The agents are the demographically-generated citizens of a real city, who spend their days pursuing activities on a transportation grid of nodes and edges. In this case, the independent variables being explored in the ensemble can quickly multiply into the dozens. As the simulation is stochastic, ensembles are used to bound the uncertainly of results; sometimes numbering into 10,000s of runs. In our initial example (Figure 3), the variables of interest are about an influenza immunization scenario and are things like: the agent compliance to public safety notices, thresholds of triggering risk, and the duration of the intervention. Dependent variables are simulation results like infection rates, mortalilty, and productivity loss for example [VLC*18], [CKL*14], [VCG*17].

The second application was a Monte Carlo simulation that examines non-equilibrium relaxation features in a stochastic LotkaâĂŞ-Volterra predator-prey model based on a two-dimensional lattice. Physicists/Researchers were studying the biodiversity in ecology and population dynamics through pattern formation and phase transitions and how can this protect the endangered species in threatened ecosystems. Depending on different simulation parameters (number of preys, number of predators, predation rate, etc.), they used our tool to analyze and make sense of their data. Using the different multi-linked views and interactions, our tool helped them to find the reason for the systemâĂŹs relaxation. Additionally, they were able to observe that there was a critical slowing-down in predator density at the extinction critical point in the case of non-equilibrium relaxation of the predator density in the neighborhood of the critical predation rate. Our tool also helped them to have a better understand the relationship between the inputs and outputs in the ensemble and to affirm their physical properties. Additionally, it helped them to find correlations between ensemble runs and parameters [CT16, Vol26, Lot20].

6. Evaluation

To verify the effectiveness and usefulness of our visualization tool for providing scientists with multiple ways to interact and make sense of their data, we performed a structured user study to assess the benefits and the drawbacks of our multi-linked visual analysis tool. We believe that tools with a multiplicity of views and interactions might allow for more complete analyses, therefore, in our study, we are mainly focusing on performance and usability of our tool. In other words, we need to determine how the different linked views can help the scientists in analyzing their high dimensional data.

In this study, we collaborated with experts in the geoscience field to evaluate our tool. One of them is the scientist who provided the ensemble data used in the study and was deeply involved in the design of the tool. To familiarize them with the tool, they were given a training session using a different dataset where we explain the purpose of each view and how to interact with each view and how the multiple views are linked. After that, they were asked to

Figure 7: *(a) Most dominant attributes (b) Runs grouped by levels of CO2*

complete a set of tasks that asked them to analyze the data and answer specific questions.

The geoscience dataset consists of an ensemble of results from 50 constant pressure CO2 injections, which simulate a commercial-scale CO2 injection into the Columbia River Basalt Group located in Richland, WA, USA. Due to the large scale of these reservoir simulations, there is large uncertainties related to how rock properties vary in space. Specifically, permeability distributions are studied within these simulations as permeability has a first-order control on fluid flow. In order to account for a wide range of permeabilities, sequential indicator simulation (sisim) is employed to create equally-probable, stochastically generated permeability distributions. In this approached the sisim creates a permeability value for each grid cell within the injection zone by utilizing (1) the cumulative distribution function of regional hydrologic data, (2) known data points within the borehole, (3) previously simulated grid cells, and (4) a chosen spatial correlation model, which for this study is an anisotropic semivariogram. The injection zone is composed of 40,000 grid cells with the dimensions 50m x 50m x 25m. The results from these simulations are utilized to understand how the permeability distribution affects the variability of total CO2 injected, pressure, temperature, CO2 saturation, density of CO2, and density of the aqueous species within each the synthetic reservoirs [JJP16].

The study included questions that are directly related to our measures (i.e., performance and usability of our tool). We divided the questions into categories to serve as benchmark tasks for each measure. Regarding the performance, we were focusing on tool's performance in terms of time (i.e., the time taken by the tool to react to user's interaction, the time taken for synchronizing the different views after each interaction, and time taken for projecting the data). We also consider users' performance in terms of accuracy in their answers (i.e., does their answers make sense scientifically? Were they able to analyze their data correctly and derive meaningful relations and conclusions?). On the other hand, to measure the usability and understandability of our tool, three to four low-level questions about each view was given to the scientists, and we expected the scientists would use these views to answer the question without asking them which view to use. For example, questions targeting parameter view asked if a particular parameter (i.e., temperature) has any influence on the ensemble runs, if there is any correction between any two or more parameters, and what attributes have the least and the most effect on ensemble runs.

Questions for users in the ensemble view were focused on finding and comparing characteristics between different clusters of runs (e.g., users were asked to find common attributes between clusters, form clusters of runs that they believe could be similar and find commonalities between them, and characterize and compare differences between similar and dissimilar groups of runs). Statistical view questions were focusing on finding the best parameter space for significant attributes, using distributions of data to investigate interesting patterns and subsets, and confirming conclusions derived from both parameter and ensemble views. Additionally, users were given some general questions about the usability and usefulness of the tool to their analysis process like if the tool fits into their research work and how would it help them to improve their research, what they like and disliked in the tool, and what features that liked or wished to have in each view.

In summary, the study questions were focused on detecting the simulation features, identifying behavioral patterns, comparing a number of simulation runs to draw conclusions about initial parameter influence and estimating corresponding dependencies, and searching for possible outliers. For example, they were asked to find the most dominant attributes between ensembles (Figure 7 a), to find runs that have the highest levels of CO_2 saturation (Figure 7 b), to characterize and compare correlation patterns between runs(Figure 8 a), and to find the distribution and properties of their attributes (Figure 8 b)

6.1. Discoveries and Results

Our tool helped the domain experts in making several specific discoveries. Initially, the scientist uploads a file containing the input and output parameters for each run and gets initial projection. Each image in this projection shows an isosurface of saturated aqueous fluid colored by temperature. Users can easily change the colored contour to reflect other output parameters (Figure 9a). Grouping high temperature runs together and low temperature runs together (i.e., an OLI interaction) in the ensemble (Figure 9b) view revealed that larger CO_2 plumes seem to exhibit more variability (Figure 9c) in temperature as this has to do with permeability affecting

Figure 8: *(a) Finding correlations between similar runs using OLI through ensemble view (b) Finding distributions and trends in data using statistical view*

fluid pressure, which in turn affects the ability of CO_2 to expand because CO_2 cools as it expands. The connection between the ensemble view and parameter view revealed that grouping of high and low-temperature runs leads to increasing the importance (i.e., weight) of CO_2 saturation.

This further led the scientists to explore the effect of CO_2 saturation on the whole ensemble by examining the distribution of the CO2 using the box plot, checking the correlation CO_2 with other parameters, giving more weight to CO2 saturation on the attribute slider (PI interaction) (Figure 9d). They noticed that smaller CO_2 plumes appeared to be more circular than the more oblique large plumes. This suggests that the anisotropic permeability correlation structure of this geologic formation exerts more control on plume geometry when the CO_2 reaches longer radial distances from the injections well.

Users grouped smaller CO_2 plumes together and larger ones together (OLI) (Figure 9e) leading them to find out that the density of aqueous fluid increases linearly with the pressure and the slopes at different concentrations are almost the same at a certain temperature (Figure 9f). This is an interesting discovery for them that would require more analysis. Additionally, they were interested in finding any relationship between temperature and CO_2, so they reset the pipeline and clustered different runs by temperature, CO_2 plume size, and CO_2 plume shape (Figure 10a). They noticed that there is a dominant relationship of permeability (Figure 10b), which was a discovery that they did not expect.

To further investigate it, they used the Cinema slider to see their runs from different views (Figure 10c and d). They then decided to check if there is any correlation between permeability and CO_2 saturation or not using PI interaction on parameter view (Figure10e). As expected they found that lower permeability leads to lower CO_2 where lower CO_2 runs were grouped together and high CO_2 runs were grouped together. To check if this correlation holds globally and locally between the parameters and the runs, they selected interesting patterns in the parallel coordinates, which leads to an iteration in the pipeline returning the runs corresponding to selected areas (Figure 10f). From this new projection, they found out that this correlation holds locally and globally.

From our observations and records, we have noticed that the performance of our tool is reasonable. It takes an average of 2 to 3 seconds for projecting the data in the ensemble view, less than a second for synchronizing between views and 1 to 3 seconds to react to user interactions (i.e., 1 second or less for slider interactions in parameter view or manipulating the graphs in statistical view and 2 to 3 seconds for updating the view in ensemble view after OLI operation). Additionally, users takes from 1 to 5 minutes to answer the questions and these times take into consideration only the time spent interacting with our tool and do not consider the time that users spent reading the questions and writing responses.

Regarding the accuracy of userss answers, some answers were more complete than others; however, this did not have a significant affect on the analysis process and users were still able to explore the data and gain multiple insights. In most of the questions, they were able to answer all low-level questions using the appropriate view correctly. However, sometimes when we expect them to use ensemble view directly, they preferred using statistical view first to have an overview of the distributions and the patterns in the data.

6.2. Discussion

During our study, we examined how the three linked views helped the scientists in analyzing their high dimension ensembles. The qualitative analysis of the study results showed that the scientists were able to explore and make sense of the ensemble using our tool. Although the interactions supported by our tool were new to the scientists and they did not have a complete understanding of the underlying algorithm, this did not affect their exploratory analyses. Our study proves that our multi-linked views and interactions (i.e., OLI, PI, and brushing and linking) provide enough context to the scientists to understand the relations and correlations between the ensemble and its parameters.

OLI gives scientists the ability to pose what-if questions to test hypotheses about relation and correlation between and within ensemble runs. From their background and existing knowledge about the nature of the simulation, they know and understand preexisting relationships between runs and parameters. OLI can help in refining the claims that support this understanding, in addition to

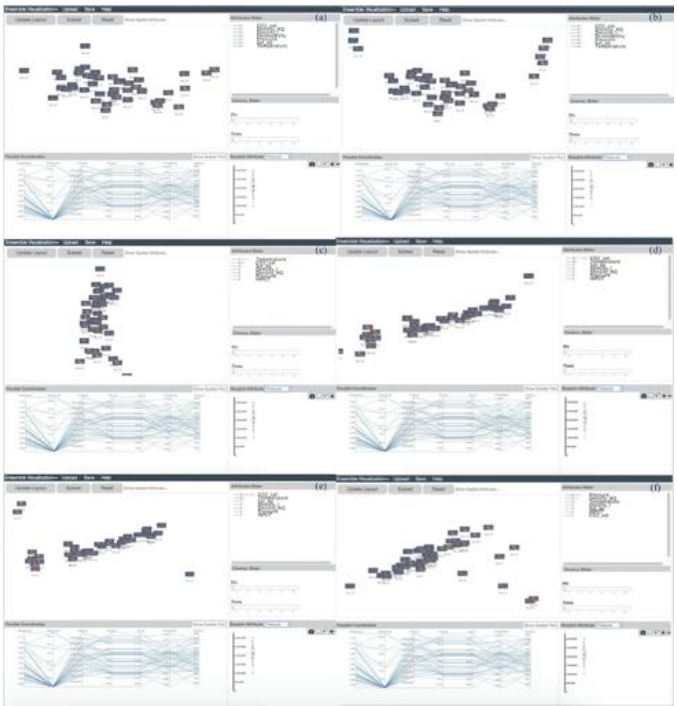

Figure 9: *shows the results of analyzing the geoscience dataset using our tool. a) initially WMDS projection of the dataset inputs and outputs; b,c,e,f) scientists uses OLI interaction in ensemble view by dragging similar runs (i.e., highlighted) to understand effect of temperature and CO_2 respectively on the ensemble ;d) scientists uses uses PI interaction in parameter view to check the influence of CO_2 on the whole ensemble*

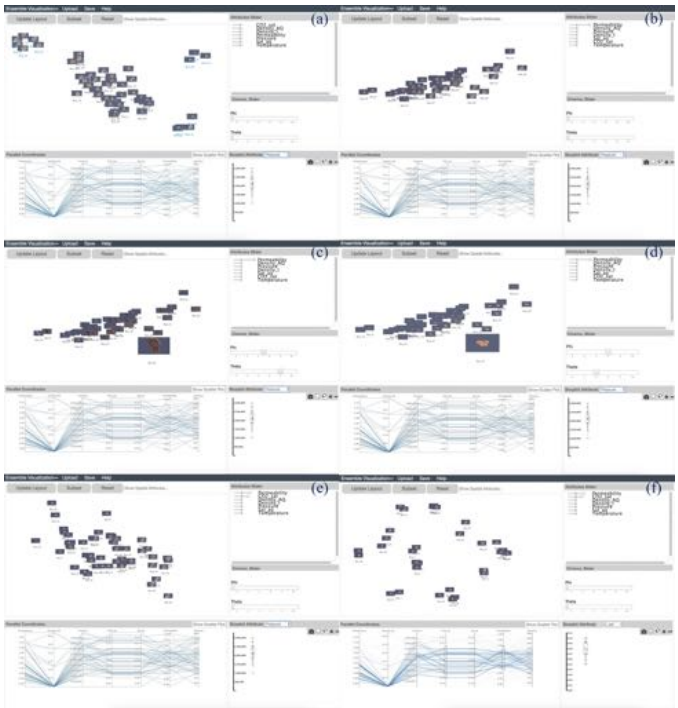

Figure 10: *shows further analysis of the dataset using Cinema slider and brushing and linking after finding new discoveries a,b)cluster runs by temperature, CO_2 plume size, and CO_2 plume shape and use OLI interaction c,d) take advantages of the Cinema slider to interpret the results of projection and view data from different angles e,f) make use of the brushing and linking to select interesting patterns in the parallel coordinates*

finding discoveries that are hard to find by manual analysis methods or would require many trial-and-error iterations using parametric interaction, especially with high dimensional data. OLI also gives users the opportunity to manipulate the data on the object level, providing a bridge between user intention and the underlying mathematical models. Although OLI interaction offers a meaningful space for interaction, depending solely on OLI for analyzing simulation ensemble will limit a comprehensive analysis of the simulation features in the ensemble. Each of the three linked views conveys different information about the ensemble will helps scientists to have a broad understanding of their ensemble.

The feedback we received from domain experts was satisfying. They confirmed that our multi-linked views visualization tool enabled them to explore the uncertainty in the ensemble and helped them to build connections and correlations between parameter settings and ensemble members. They also mentioned that using semantic interaction by spatially clustering interesting groups (i.e., OLI) and using attribute slider (PI) gives them a different prescriptive for analyzing their ensemble. Additionally, they confirmed that statistical view was beneficial in showing the distributions of the attributes, relationships between multiple attributes, trends, and outliers. Our tool provided them the methodology to compare all the different attributes, runs, and their statistical properties all at the same time without changing screens, programs, or changing scripts to visualize the data. They also did not have any existing visualization tool that allows them to explore the uncertainty of the input and output parameters at the same time.

Moreover, they confirmed that each of the three linked views conveys information about the ensemble helping them to have a complete picture about the ensembles for a better analysis process. The Cinema slider provides a means to view different angles of the data set. It helped them to view their data from different perspectives that open other angles for exploring parameter settings and ensemble members. They also mentioned that the performance of tool with respect to time was reasonable and confirmed the ease of use of the tool and its applicability to different datasets. In summary, they affirmed that our tool met their goals for analyzing and making sense of their data and all the views used in the tool were needed during their analysis process.

7. Limitations and Future work

There are several limitations in the current design of our tool. First, our weighting scheme and inverse model currently only works with final outcomes (the last timestep), thus the Sematic Interaction is working on a static snapshot of the system. While our current tool

implementation shows promise to help scientists to gain insights about the correlations between parameters and ensemble members, scientists also need to work and analyze across temporal features in their ensembles. We consider these important extensions as future work. Second, our tool does not scale well up to thousands of ensemble members as it could result in visual clutter; large displays can increase this number, and currently we work well with hundreds of ensemble members (simulation runs). Given that the size of most of the scientific data we were dealing with during the design of our tool was not extremely large, we put less emphasis on the scalability issue in the current work. Future work should include research into scalability concerns including data and display sizes.

Additionally, the domain experts provided us with valuable suggestions for improving our tool. They suggested modifying the current statistical view to incorporate a wider variety of charts (i.e., histogram, bar charts, and line charts) and give the scientists more control over them by selecting specific runs and parameters. This allows them to see the statistics from an individual or multiple selected runs without having to subset these runs. They also suggested filtering/subsetting the space by raw data values and gives them the chance to save this raw data. Moreover, they wanted us to extend our tool to incorporate the in situ workflow. Finally, we plan to take these suggestions and limitations as our future work.

8. Conclusion

In this paper, we have demonstrated multi-linked visualization approach that merges human expertise and intuition with machine learning and statistics for sense-making, analyzing, and exploring any multi-dimensional scientific ensembles irrespective of the application domain. At the core of our visual analysis tool (GLEE: Graphically-Linked Ensemble Explorer) is a novel way to represent simulation ensembles by integrating ensemble runs, their input and output parameters, and statistical characteristics into one screen, where each view/display represents meaningful interpretation for certain aspect of the ensemble. We believe that these three linked views fill the gap between the different analysis techniques used by domain scientists and the approaches available from visualization research.

Our approach gives scientists the chance to explore interesting patterns to confirm their conclusions about both local and global uncertainties in the ensemble. The machine learns from these semantic interactions, converging on a shared model. Moreover, the images of ensemble members displayed in ensemble view and the Cinema slider give scientists a new way of visually analyzing their data by having the ability to view ensemble members from a different perspective. Additionally, we demonstrated the effectiveness of our tool through a user study with domain experts and showed the potential of our design with two different use cases. Our future steps include: modifying the current statistical view to incorporate a wider variety of charts, extending our tool to incorporate in-situ workflows, supporting time-series ensembles, scaling to support larger ensembles(i.e, although our tool is tailored to work with hundreds of members), and integrating machine learning to learn from user interactions and suggest new interactions that may help scientists in their analytic process.

9. ACKNOWLEDGMENTS

The authors wish to thank Sheng Chen from Physics department at Virginia Tech for his contributions and support during the development of this project. We also thank Virginia Tech Advanced Research Computing for their support in this research.

References

[AEN10] ANDREWS C., ENDERT A., NORTH C.: Space to think: large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI conference on human factors in computing systems* (2010), ACM, pp. 55–64. 3

[AGL*05] AHRENS J., GEVECI B., LAW C., HANSEN C., JOHNSON C.: 36-paraview: An end-user tool for large-data visualization. *The visualization handbook 717* (2005). 8

[ASB*17] ABOULHASSAN A., SICAT R., BAUM D., WODO O., HADWIGER M.: Comparative visual analysis of structure-performance relations in complex bulk-heterojunction morphologies. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 329–339.

[AWH*12] ALABI O. S., WU X., HARTER J. M., PHADKE M., PINTO L., PETERSEN H., BASS S., KEIFER M., ZHONG S., HEALEY C., ET AL.: Comparative visualization of ensembles using ensemble surface slicing. In *Proceedings of SPIE* (2012), vol. 8294, NIH Public Access. 3

[Aya15] AYACHIT U.: The paraview guide: a parallel visualization application. 6

[BGOJ16] BENSEMA K., GOSINK L., OBERMAIER H., JOY K. I.: Modality-driven classification and visualization of ensemble variance. *IEEE transactions on visualization and computer graphics 22*, 10 (2016), 2289–2299. 2

[BHGK14] BEHAM M., HERZNER W., GRÖLLER M. E., KEHRER J.: Cupid: Cluster-based exploration of geometry generators with parallel coordinates and radial trees. *IEEE transactions on visualization and computer graphics 20*, 12 (2014), 1693–1702. 3

[BHJ*14] BONNEAU G.-P., HEGE H.-C., JOHNSON C. R., OLIVEIRA M. M., POTTER K., RHEINGANS P., SCHULTZ T.: Overview and state-of-the-art of uncertainty visualization. In *Scientific Visualization*. Springer, 2014, pp. 3–27. 2

[BLBC12] BROWN E. T., LIU J., BRODLEY C. E., CHANG R.: Disfunction: Learning distance functions interactively. In *Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on* (2012), IEEE, pp. 83–92. 4

[BLLS17] BISWAS A., LIN G., LIU X., SHEN H.-W.: Visualization of time-varying weather ensembles across multiple resolutions. *IEEE transactions on visualization and computer graphics 23*, 1 (2017), 841–850. 2

[BM10] BRUCKNER S., MOLLER T.: Result-driven exploration of simulation parameter spaces for visual effects design. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 1468–1476. 3

[BNH14] BRADEL L., NORTH C., HOUSE L.: Multi-model semantic interaction for text analytics. In *Visual Analytics Science and Technology (VAST), 2014 IEEE Conference on* (2014), IEEE, pp. 163–172. 4

[BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D^3 data-driven documents. *IEEE transactions on visualization and computer graphics 17*, 12 (2011), 2301–2309. 8

[BOL12] BRODLIE K., OSORIO R. A., LOPES A.: A review of uncertainty in data visualization. In *Expanding the frontiers of visual analytics and visualization*. Springer, 2012, pp. 81–109. 2

[BPFG11] BERGER W., PIRINGER H., FILZMOSER P., GRÖLLER E.: Uncertainty-aware exploration of continuous parameter spaces using multivariate prediction. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 911–920. 1

[BWE05] BOTCHEN R. P., WEISKOPF D., ERTL T.: Texture-based visualization of uncertainty in flow fields. In *Visualization, 2005. VIS 05. IEEE* (2005), IEEE, pp. 647–654. 2

[CB12] CORCHADO E., BARUQUE B.: Wevos-visom: An ensemble summarization algorithm for enhanced data visualization. *Neurocomputing 75*, 1 (2012), 171–184. 2

[CBDT11] CONINX A., BONNEAU G.-P., DROULEZ J., THIBAULT G.: Visualization of uncertain scalar data fields using color scales and perceptually adapted noise. In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization* (2011), ACM, pp. 59–66. 2

[CH17] CORRELL M., HEER J.: Regression by eye: Estimating trends in bivariate visualizations. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), ACM, pp. 1387–1396. 7

[CKL*14] CHAKRABORTY P., KHADIVI P., LEWIS B., MAHENDIRAN A., CHEN J., BUTLER P., NSOESIE E. O., MEKARU S. R., BROWNSTEIN J. S., MARATHE M. V., ET AL.: Forecasting a moving target: Ensemble models for ili case count predictions. In *Proceedings of the 2014 SIAM international conference on data mining* (2014), SIAM, pp. 262–270. 8

[CKS*17] CIBULSKI L., KLARIN B., SOPOUCH M., PREIM B., THEISEL H., MATKOVIĆ K.: Super-ensembler: interactive visual analysis of data surface sets. In *Proceedings of the 33rd Spring Conference on Computer Graphics* (2017), ACM, p. 19. 3

[CT16] CHEN S., TÄUBER U. C.: Non-equilibrium relaxation in a stochastic lattice lotka–volterra model. *Physical biology 13*, 2 (2016), 025005. 8

[CZC*15] CHEN H., ZHANG S., CHEN W., MEI H., ZHANG J., MERCER A., LIANG R., QU H.: Uncertainty-aware multidimensional ensemble data visualization and exploration. *IEEE transactions on visualization and computer graphics 21*, 9 (2015), 1072–1086. 2

[DDW14] DEMIR I., DICK C., WESTERMANN R.: Multi-charts for comparative 3d ensemble visualization. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 2694–2703. 3

[Det05] DETTINGER M. D.: From climate-change spaghetti to climate-change distributions for 21st-century california. *San Francisco Estuary and Watershed Science 3*, 1 (2005). 1

[DJW16] DEMIR I., JAREMA M., WESTERMANN R.: Visualizing the central tendency of ensembles of shapes. In *SIGGRAPH ASIA 2016 Symposium on Visualization* (2016), ACM, p. 3. 3

[DKLP02] DJURCILOV S., KIM K., LERMUSIAUX P., PANG A.: Visualizing scalar volumetric data with uncertainty. *Computers & Graphics 26*, 2 (2002), 239–248. 2

[DKW16] DEMIR I., KEHRER J., WESTERMANN R.: Screen-space silhouettes for visualizing ensembles of 3d isosurfaces. In *2016 IEEE Pacific Visualization Symposium (PacificVis)* (2016), IEEE, pp. 204–208. 3

[DNCP10] DEMERITT D., NOBERT S., CLOKE H., PAPPENBERGER F.: Challenges in communicating and using ensembles in operational flood forecasting. *Meteorological applications 17*, 2 (2010), 209–222. 1

[EBN13] ENDERT A., BRADEL L., NORTH C.: Beyond control panels: Direct manipulation for visual analytics. *IEEE computer graphics and applications 33*, 4 (2013), 6–13. 3

[EFN12a] ENDERT A., FIAUX P., NORTH C.: Semantic interaction for sensemaking: inferring analytical reasoning for model steering. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2879–2888. 3, 4

[EFN12b] ENDERT A., FIAUX P., NORTH C.: Semantic interaction for sensemaking: inferring analytical reasoning for model steering. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2879–2888. 4

[EHM*11] ENDERT A., HAN C., MAITI D., HOUSE L., NORTH C.: Observation-level interaction with statistical models for visual analytics. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on* (2011), IEEE, pp. 121–130. 6

[FBW16] FERSTL F., BÜRGER K., WESTERMANN R.: Streamline variability plots for characterizing the uncertainty in vector field ensembles. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2016), 767–776. 3

[FML16] FOFONOV A., MOLCHANOV V., LINSEN L.: Visual analysis of multi-run spatio-temporal simulations using isocontour similarity for projected views. *IEEE transactions on visualization and computer graphics 22*, 8 (2016), 2037–2050.

[HBM*13] HU X., BRADEL L., MAITI D., HOUSE L., NORTH C.: Semantics of directly manipulating spatializations. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2052–2059. 3

[HdMRHH16] HÖLLT T., DE MATOS RAVANELLI F., HADWIGER M., HOTEIT I.: Visual analysis of reservoir simulation ensembles. In *Workshop on Visualization in Environmental Sciences* (2016), pp. 1–4. 3

[HDS16] HAZARIKA S., DUTTA S., SHEN H.-W.: Visualizing the variations of ensemble of isosurfaces. In *2016 IEEE Pacific Visualization Symposium (PacificVis)* (2016), IEEE, pp. 209–213. 3

[HHBY16] HAO L., HEALEY C. G., BASS S. A., YU H.-Y.: Visualizing static ensembles for effective shape and data comparison. *Electronic Imaging 2016*, 1 (2016), 1–10. 3

[HHH15] HAO L., HEALEY C. G., HUTCHINSON S. E.: Ensemble visualization for cyber situation awareness of network security data. In *Visualization for Cyber Security (VizSec), 2015 IEEE Symposium on* (2015), IEEE, pp. 1–8. 2

[HLNW11] HLAWATSCH M., LEUBE P., NOWAK W., WEISKOPF D.: Flow radar glyphsâĂŤstatic visualization of unsteady flow with uncertainty. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 1949–1958. 1, 2

[HMZ*14] HÖLLT T., MAGDY A., ZHAN P., CHEN G., GOPALAKRISHNAN G., HOTEIT I., HANSEN C. D., HADWIGER M.: Ovis: A framework for visual analysisof ocean forecast ensembles. *IEEE Transactions on Visualization and Computer Graphics 20*, 8 (2014), 1114–1126. 2, 3

[HOGJ13] HUMMEL M., OBERMAIER H., GARTH C., JOY K. I.: Comparative visual analysis of lagrangian transport in cfd ensembles. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2743–2752. 2

[JJP16] JAYNE JR R., POLLYEA R.: Constraining the effects of permeability uncertainty for geologic co2 sequestration in a basalt reservoir. In *AGU Fall Meeting Abstracts* (2016). 9

[KDP01] KAO D., DUNGAN J. L., PANG A.: Visualizing 2d probability distributions from eos satellite image-derived data sets: A case study. In *Visualization, 2001. VIS'01. Proceedings* (2001), IEEE, pp. 457–589. 2

[KLDP02] KAO D., LUO A., DUNGAN J. L., PANG A.: Visualizing spatially varying distribution data. In *Information Visualisation, 2002. Proceedings. Sixth International Conference on* (2002), IEEE, pp. 219–225. 2

[Lot20] LOTKA A. J.: Undamped oscillations derived from the law of mass action. *Journal of the american chemical society 42*, 8 (1920), 1595–1599. 8

[LS16] LIU X., SHEN H.-W.: Association analysis for visual exploration of multivariate scientific data sets. *IEEE transactions on visualization and computer graphics 22*, 1 (2016), 955–964. 3

[MGJ*10] MATKOVIC K., GRACANIN D., JELOVIC M., AMMER A., LEZ A., HAUSER H.: Interactive visual analysis of multiple simulation runs using the simulation model view: Understanding and tuning of an electronic unit injector. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 1449–1457. 3

[MGKH09] MATKOVIC K., GRACANIN D., KLARIN B., HAUSER H.: Interactive visual analysis of complex scientific data as families of data surfaces. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (2009). 3

[MRH*05] MacEachren A. M., Robinson A., Hopper S., Gardner S., Murray R., Gahegan M., Hetzler E.: Visualizing geospatial information uncertainty: What we know and what we need to know. *Cartography and Geographic Information Science 32*, 3 (2005), 139–160. 2

[MTL78] McGill R., Tukey J. W., Larsen W. A.: Variations of box plots. *The American Statistician 32*, 1 (1978), 12–16. 2

[MWK14] Mirzargar M., Whitaker R. T., Kirby R. M.: Curve boxplot: Generalization of boxplot for ensembles of curves. *IEEE transactions on visualization and computer graphics 20*, 12 (2014), 2654–2663. 1, 3

[NFB07] Nocke T., Flechsig M., Böhm U.: Visual exploration and evaluation of climate-related simulation data. In *Proceedings of the 39th conference on Winter simulation: 40 years! The best is yet to come* (2007), IEEE Press, pp. 703–711. 2

[OAJ*16] OÂĂŹLeary P., Ahrens J., Jourdain S., Wittenburg S., Rogers D. H., Petersen M.: Cinema image-based in situ analysis and visualization of mpas-ocean simulations. *Parallel Computing 55* (2016), 43–48.

[OJ14] Obermaier H., Joy K. I.: Future challenges for ensemble visualization. *IEEE Computer Graphics and Applications 34*, 3 (2014), 8–11. 2

[PC05] Pirolli P., Card S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis* (2005), vol. 5, pp. 2–4. 3

[PKRJ10] Potter K., Kniss J., Riesenfeld R., Johnson C. R.: Visualizing summary statistics and uncertainty. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 823–832. 1, 2

[PMW13] Pfaffelmoser T., Mihai M., Westermann R.: Visualizing the variability of gradients in uncertain 2d scalar fields. *IEEE transactions on visualization and computer graphics 19*, 11 (2013), 1948–1961. 1, 2

[PPA*12] Phadke M. N., Pinto L., Alabi O., Harter J., Taylor R. M., Wu X., Petersen H., Bass S. A., Healey C. G.: Exploring ensemble visualization. In *Visualization and Data Analysis 2012* (2012), vol. 8294, International Society for Optics and Photonics, p. 82940B. 2, 3

[PPBT12] Piringer H., Pajer S., Berger W., Teichmann H.: Comparative visual analysis of 2d function ensembles. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1195–1204. 3

[PPH12] Petz C., Pöthkow K., Hege H.-C.: Probabilistic local features in uncertain vector fields with spatial correlation. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1045–1054. 1

[PW12] Pfaffelmoser T., Westermann R.: Visualization of global correlation structures in uncertain 2d scalar fields. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1025–1034. 1

[PWB*09a] Potter K., Wilson A., Bremer P.-T., Williams D., Doutriaux C., Pascucci V., Johhson C.: Visualization of uncertainty and ensemble data: Exploration of climate modeling and weather forecast data with integrated visus-cdat systems. In *Journal of Physics: Conference Series* (2009), vol. 180, IOP Publishing, p. 012089. 1, 2

[PWB*09b] Potter K., Wilson A., Bremer P.-T., Williams D., Doutriaux C., Pascucci V., Johnson C. R.: Ensemble-vis: A framework for the statistical visualization of ensemble data. In *Data Mining Workshops, 2009. ICDMW'09. IEEE International Conference on* (2009), IEEE, pp. 233–240. 2

[SEG*15] Splechtna R., Elshehaly M., Gračanin D., ÆL'uras M., Bühler K., Matković K.: Interactive interaction plot. *The Visual Computer 31*, 6-8 (2015), 1055–1065. 1, 3

[SH] Self J. Z., House L.: *Andromeda: Observation-level and parametric interaction for exploratory data analysis*. Tech. rep. 4

[SYMSJ05] Soo Yi J., Melton R., Stasko J., Jacko J. A.: Dust & magnet: multivariate information visualization using a magnet metaphor. *Information visualization 4*, 4 (2005), 239–256. 3

[SZD*10] Sanyal J., Zhang S., Dyer J., Mercer A., Amburn P., Moorhead R.: Noodles: A tool for visualization of numerical weather model ensemble uncertainty. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 1421–1430. 1, 2

[TG07] Tesone D. R., Goodall J. R.: Balancing interactive data management of massive data with situational awareness through smart aggregation. In *Visual Analytics Science and Technology, 2007. VAST 2007. IEEE Symposium on* (2007), IEEE, pp. 67–74. 3

[VCG*17] Venkatramanan S., Chen J., Gupta S., Lewis B., Marathe M., Mortveit H., Vullikanti A.: Spatio-temporal optimization of seasonal vaccination using a metapopulation model of influenza. In *Healthcare Informatics (ICHI), 2017 IEEE International Conference on* (2017), IEEE, pp. 134–143. 8

[VLC*18] Venkatramanan S., Lewis B., Chen J., Higdon D., Vullikanti A., Marathe M.: Using data-driven agent-based models for forecasting emerging infectious diseases. *Epidemics 22* (2018), 43–49. 8

[Vol26] Volterra V.: Fluctuations in the abundance of a species considered mathematically, 1926. 8

[WLSL17] Wang J., Liu X., Shen H.-W., Lin G.: Multi-resolution climate ensemble parameter analysis with nested parallel coordinates plots. *IEEE transactions on visualization and computer graphics 23*, 1 (2017), 81–90. 3

[WMK13] Whitaker R. T., Mirzargar M., Kirby R. M.: Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2713–2722. 1, 3

[ZWK10] Zehner B., Watanabe N., Kolditz O.: Visualization of gridded scalar data with uncertainty in geosciences. *Computers & Geosciences 36*, 10 (2010), 1268–1275. 2