DOI: 10.1111/cgf.14029

COMPUTER GRAPHICS *forum*
Volume 39 (2020), number 6 pp. 325–343

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
npolys@vt.edu
[2]Department of Geosciences, Virginia Tech, USA
rjayne@vt.edu, rpollyea@vt.edu

Abstract

In the study of complex physical systems, scientists use simulations to study the effects of different models and parameters. Seeking 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 multi-dimensional and multi-variate. 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. GLEE uses visualization and semantic interaction (SI) techniques to enable scientists to find similarities and differences between runs, find correlation(s) between different parameters and explore relations and correlations across and between different runs and parameters. Our approach supports scientists in selecting interesting subsets of runs in order to investigate and summarize the factors and statistics that 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.

Keywords: ensemble, scientific visualization, visual analytics, human computer interaction (HCI), sensemaking

ACM CCS: • Scientific Visualization → Ensembles, Sensemaking

1. Introduction

Recent advancements in data acquisition, storage and computing power have led to the feasibility of running complex simulations in an acceptable amount of time. Scientists from multiple disciplines, such as meteorology, computational fluid dynamics [HMZ*14, PWB*09a, HOGJ13, HLNW11] usually run their simulations multiple times using different initial conditions, input parameters or simulation models to understand the uncertainty in the data. A set of simulation runs is known as an ensemble. An ensemble supports scientists in gaining deeper insights about the simulated phenomena, exploring unknowns in initial conditions and understanding their influence on the simulation output, evaluating extreme cases and investigating parameter sensitivity. As a result, scientists are capable of assessing the confidence in their findings and refining their hypothesis before physical experiments.

Simulation ensemble analysis is a challenging task due to its high-dimensionality, complexity and size. Ensemble visualization plays a major role in guiding scientists towards a better and more intuitive understanding of the data. Current research in interactive ensemble visualization has produced new modes for scientists to visualize their data by exploring either parameter space or ensemble space [MWK14, WMK13, SEG*15, PMW13, FML16]. This makes the visual analysis of ensembles a long-standing problem that needs careful study. Therefore, the focus of this paper is the visual analysis, exploration and comparison of simulation ensembles.

Simulation ensembles provide scientists with ample data to explore processes of interest where they can ask questions including but not limited to: which attribute(s) influence(s) the ensemble the most/least? How do input parameters influence the ensemble results? What are the optimal parameter settings for a given scenario? Which simulation runs within the ensemble are similar and/or how are they similar? Which ensemble runs follow a particular pattern or distribution? Answering these questions is critical to any modelling study; however, answering some of these questions is difficult without the proper tools to analyse high-dimensional datasets.

The initial motivation of the paper was developing a visual analysis tool to assist scientists, regardless of their domain expertise, to explore their simulation ensemble. However, we found that current tools focus on either exploring the parameter space or the ensemble space. In this paper, we attempt to close the gap between parameter space and ensemble space: we developed a visualization tool, GLEE (Graphically-Linked Ensemble Explorer), that helps scientists explore the ensemble without in-depth knowledge of the underlying machine learning algorithms. GLEE aims to answer questions that address the following ensemble visualization problems: (1) exploring the parameter space by combining input parameters and simulation output(s) into the same space; (2) exploring the ensemble space by interpreting and understanding the similarities between ensemble members; (3) finding patterns, distributions and optimal settings between various attributes (i.e. inputs and outputs).

GLEE represents our interactive visual approach for exploring high-dimensional ensembles when prior knowledge about features and associations is unknown and/or unavailable. It is based on combining parameter space, ensemble space and summary statistics into coordinated views. Relying on a single space for analysis can offer an incomplete picture of the ensemble precluding the efficiency of the exploration process. GLEE not only empowers scientists to identify the similarities between ensemble members but also where and why similarities exist. Moreover, GLEE allows scientists to study sub-regions of interest, helping them understand global and local correlations between parameters and result variables. GLEE's statistical view assists scientists in quantifying and verifying their hypotheses and validating their findings through various statistical summaries and descriptions.

In this paper, we describe GLEE, a new visualization tool that uses semantic interaction (SI), dimensionality reduction, statistical visualizations and brushing and linking techniques to visually explore simulation ensembles. GLEE demonstrates the power of the SI approach by enabling the visual analysis of the high-dimensional parametric relationships both within an ensemble as well as the similarity and dissimilarity among individual members in a unified visualization layout. This makes our tool more generalizable than previous tools and significantly advances SI for simulation ensembles: the task of analysing the nuanced relations between simulation parameters and output ensemble runs simultaneously. The contributions of this paper are as follows:

- A visual analysis approach that combines SI and multi-dimensional projection techniques to support sensemaking: a platform where machine and user learn from each other in order to understand and analyse a high-dimensional ensemble's patterns, associations, similarities and uncertainties, considering both the input parameters as well as the simulation outputs.
- Coupling SI with parallel coordinates and statistical summaries to support analysis into the characteristics of high-dimensional parameter spaces, finding optimal parameter settings and for determining the sensitivity and impact of different parameters in simulation ensemble results.
- We demonstrate how our analysis approach enables scientists to explore simulation ensembles, find interesting and relevant patterns with little to no prior knowledge about the underlying simulation model; we use real-world simulation ensembles and gather domain expert feedback.

2. Related Work

In this section, we provide a brief survey of ensemble visualization, multi-dimensional data visualization, and visual analysis techniques to highlight the lack of connection between parameter space and ensemble space analysis.

2.1. Ensemble visualization and analysis

Various techniques have been developed to visualize the relationships between and within ensemble members mostly focusing on parameter space or ensemble space [BHJ*14, BOL12, OJ14, WHLS18]. Parameter space visualizations have studied the correlation between parameters using different methodologies including but not limited to, summary statistics [PKRJ10, BPFG11, PWB*09a, MWK14, WMK13, SEG*15], spaghetti plots [DNCP10, Det05], glyphs [HLNW11, PMW13, SZD*10] and probabilistic features [PPH12, PW12]. However, some techniques fall short in showing the intrinsic structures in the ensemble while others are not designed to handle high-dimensional data sets. On the other hand, ensemble space has been visualized either as an aggregation of multiple members omitting potential details of the original ensemble [PWB*09b, CB12] or as a comparative representation of a limited number of ensemble members [HHH15, FML16].

Summary-based ensemble visualizations use a variety of techniques to show the statistical distribution and properties of the ensemble members, ranging from simplistic ways (e.g. colour maps, contours, animation or glyphs [PMW13, PWB*09b, CBDT11]) to more complex and informative ways (i.e. spaghetti plots, contour box plots or curve box plot [WMK13, MWK14, SZD*10]). These techniques usually statistically aggregate the ensemble members before visualization, visually composite them after visualization or combine both techniques [WHLS18]. Although these techniques show different representations of statistical distributions, they are not designed to handle high-dimensional ensembles with large number runs and different types of parameters. So, they cannot work well with dense 2D/3D data and cannot efficiently differentiate some distributions.

To overcome these limitations, traditional boxplots [KDP01, PKRJ10] were extended to show the depth and the centrality of ensemble members using contours [WMK13] and curves [MWK14]. Another set of ensemble visualization techniques was developed to show variation in the ensemble. Bensema *et al.* [BGOJ16] used modality ensemble members to define the high-variance locations while Chen *et al.* [CZC*15] differentiated the distributions of similar mean ensemble members using uncertainty-aware projection scheme. On the other hand, Demir *et al.* [DJW16] analysed the central tendency of 2D and 3D ensembles using mixture models. Similarly, Ferstl *et al.* [FBW16] statistically modelled the distribution of streamlines by deriving clusters. The currently proposed visualization tools have helped scientists to understand the statistical distributions of simulation output data. However, they have fallen short in considering the input parameters during the analysis process. Additionally, they have supported basic interaction techniques (i.e. zooming, selection,..etc.) that have not conveyed much information about underlying similarities and associations between ensemble members, thus hindering a comprehensive analysis of the simulation features in the ensemble.

Various approaches have been developed to visualize high-dimensional data ranging from scatter plot matrix [CLNL87], star coordinates [Kan00] to more complex and informative ways such as dimensionality reduction techniques [Pea01, KW78] or interactive widgets (i.e, parallel coordinates) [DCK12, ID87]. Existing high-dimensional ensemble visualization includes EnsembleLens [XXM*18], Ensemble-Vis [PWB*09a], Drag and Track [OKB*18] and Noodles [SZD*10], which mainly focuses on parameter sensitivity analysis. EnsembleLens utilizes ensemble visualizations to facilitate anomaly detection for multi-dimensional data. Drag and Track projected high-dimensional input parameters and simulation outputs in 2D space, and allowed direct manipulate data points within a continuous parameter space and observe the change in the output space and vice versa. Noodles used glyph-based techniques to quantify uncertainty in a high-dimensional ensemble. Our approach is focused on exploring high-dimensional parameter and ensemble spaces simultaneously.

Ensemble visual analysis tools focus on characterizing the uncertainty encoded in data to infer new information. Several parameter space analysis tools have been developed, such as density-based clustering of animation sequences [BM10] and Cupid system for geometry generators [BHGK14]. In spite of incorporating input parameters and output in the analysis, these tools are designed to analyse a single ensemble member at a time and do not fit all scientific data. On the other hand, ensemble space visual analytic tools have been developed with different levels of complexity. The least complex tool shows a side-by-side comparison of 3D surfaces [AWH*12]. Increasing in complexity, we see the use of pairwise sequential animation and screen door tinting to show the differences between ensemble members using value changes to field points [PPA*12], the comparison of 2D function ensembles on three levels of details (i.e. surface plot, domain-oriented and member-oriented) [PPBT12], the comparison of 3D scalar field ensembles using mean isosurface [DKW16] and grouping ensemble members using hierarchical clustering [HHBY16]. Although these tools offer different ways to visualize and compare different data types, they are limited in the number of ensemble members that they visualize.

Ensemble and parameter space analysis are often combined into a coordinated multi-view visualization. Various multi-view visualization tools that have been developed include, but are not limited to families of data surfaces representing pairs of independent data dimensions [MGKH09, MGJ*10], interactive interaction plots [SEG*15], multi-chart visualization of a 3D volumetric ensemble [DDW14], series of parallel coordinates plots (PCPs) [LS16] and nested PCP [WLSL17]. Similar to our proposed approach, Höllt *et al.* [HMZ*14, HdMRHH16], Aboulhassan *et al.* [AWH*12], Cibulski *et al.* [CKS*17] and Luciani *et al.* [LBS*18] proposed multi-linked views that integrate ensemble visualization with statistical plots to facilitate the understanding of ensemble characteristics. However, all these techniques did not consider the effect of both inputs parameters and simulation outputs on the simulation ensemble. In our work, we tried to address the limitation of the previous research and incorporate human in the loop through SIs to power the visual analysis process.

2.2. Sensemaking and SI

This paper is inspired by interactive visual analytic tools that support sensemaking in enabling users to explore possible connections, investigate hypotheses and eventually create new knowledge about their data. Sensemaking is an exploration process that helps users find meaning in their data, and this involves establishing implicit connections between the user's intuition and information based on domain expertise. The sensemaking process can be broken down into two parts: foraging and synthesis [PC05]. Foraging refers to gathering and filtering relevant or interesting information, while synthesis uses forged information to construct and test hypotheses using human intuition. This makes the forging process lend itself more to computational support, while the synthesis process uses human intuition to establish relationships between information. To gain better insights into the data, sensemaking is usually integrated with SI.

SI is an approach to user interaction designed for visual exploration and the analysis of data [EHM*11, EFN12, EBN13]. It occupies a new design space for interaction that couples computational models and human reasoning. Using SI, users' direct manipulation of visual objects within the workspace incrementally builds the computational model, which describes the weighting of different (high-dimensional) attributes. This mapping of interaction to model implicitly captures scientists' insights and expertise, steering the underlying semi-supervised machine learning algorithms to provide scientists with feedback on what the computational models have learned. SI uses human interactions as a data source: incrementally formalizing their sensemaking intentions, and domain expertise. From these data, visual analytics systems use SI to capture and understand scientists' actions and make inferences guiding the computational models to react and even take initiative. SI merges the foraging abilities of statistical models with the synthesizing process to keep the sensemaking loop tight. This coupling helps to strengthen the cognitive connection formed between the user and the visualization layout [HBM*13, EFN12, EBN13].

Prior research has shown some tools for spatializations that externalize knowledge to steer the sensemaking process [AEN10]. However, this externalization requires the use of control panels outside the spatial metaphor [TG07]. Moreover, these tools do not scale well with high-dimensional data. To resolve this problem, Dust & Magnet [SYMSJ05] uses parametric interaction to adjust the model's parameters within the same spatial metaphor. Although this provides an intuitive way to control parameter(s) within the spatialization, they used attributes of the data (i.e. 'magnets' representing keywords), not data itself. On the other hand, tools such as ForceSPIRE [EFN12], Dis-function [BLBC12], Andromeda [SH] and StarSPIRE [BNH14] focus on using human cognition to steer the underlying computations by directly manipulating the spatializations, giving users the chance to interact with data points and translate this feedback through a dimension-reduction algorithm to a new view reflecting the user's interaction. This helps provide 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 analysis to make the data more manageable [SH, CD18, SKBE17, FGS19]. The interaction techniques used in this paper are built on those used by Andromeda [SH].

3. Approach

Ensemble visualizations are intended to capture the variation in the simulation model with respect to different input parameters, initial settings and outcomes. However, visual analysis of ensembles is still needed in order to analyse and explore the similarities and relationships among the ensemble members and their various parameters. Visual analysis tools should be designed to support users in gaining insights into their data by understanding the cognitive processes of scientists as they reason about data [Nor06]. Such understanding implies considering the ‘human-in-the-loop’ not only during the design of the visualization but also during the analysis process. This requires capturing human expertise and intuition through the interactive analysis processes. Most of the prior research uses brushing and linking techniques as a way to interact with multiple coordinated view visualizations. We believe this is not sufficient: scientists’ expertise and intuition should be part of the computational and analysis processes. Therefore, our novel contribution is using SI techniques for the exploratory analysis of ensembles.

3.1. SI in GLEE

SI represents an opportunity to engage scientists with machine learning, metric learning and computational models that can learn about their data without needing to translate their cognitive artefacts into computational actions. For example, after projecting high-dimensional ensemble runs into a two-dimensional space, scientists’ expertise may suggest a strong similarity between two or more runs that were not initially clustered together. So, scientists visually group those runs in the workspace. In response to their direct manipulation, the system builds a model of similarity that describes this relationship. The similarity model can then be applied to the workspace layout, leading to a new visualization that reflects what the machine learned through that interaction. Therefore, SI can support scientists in incrementally developing a computational model of the domain that embodies their domain expertise.

Scientists can use SI to test and explore hypotheses without having to think about manipulating underlying statistical and computational models, which allows them to focus on the analytic process. Typically, scientists interact with visual analysis tools using external controls such as menu(s) or text field(s) to control underlying model parameters. However, in high-dimensional datasets, the number of controls can quickly exceed human ability. In addition, in exploratory visualization, the user may not have *a priori* syntactic formulation of the problem; for example, during the early stages of analysis, the scientist may not have gained a comprehensive understanding of the ensemble space and so their insights are still informal. SI can help scientists who do not yet have the basis for expressing their inputs as specific model parameters and parameter combinations.

3.2. Method overview

Our method starts with an ensemble $\mathrm{E} = \{s_1, \ldots, s_N\}$ of N members. Each member is visualized using an image representing the output(s) of the simulation run. Initially, we spatialize the high-dimensional ensemble members in 2D space using the input parameters, simulation output(s) and weights associated with them via a low-dimensional projection technique, such as PCA [Pea01], MDS [KW78], etc. In the SI workspace spatialization, close proximity reflects relative similarity; specifically, two ensemble members close to each other in the spatialization have more similar attributes than far apart ones. Thus, ensemble members closely positioned in the SI workspace (low-dimensional layout) indicate similar attributes between these ensemble members in the high-dimensional space. Initially, all weights are set to be equal, but with user interactions, these weights change to reflect the importance of the corresponding attributes, so attributes with large weights are considered more heavily in the spatialization than those with low weights. Thus, a user can deepen his or her interpretation of a complex system considering the weighted attributes of the projections that preserve the relative distances among members. SI and GLEE thus support scientists in understanding the influence of attributes on the ensemble, the parameter view displays the weights corresponding to attributes on an attribute slider. This gives scientists the freedom to explore and directly modify the importance of different attributes.

Our tool supports two interactions for interacting with ensemble members and their attributes: observation level interaction (OLI) and parametric level interaction (PLI). OLI is a user interaction technique based on the principles of SI. It is an interaction occurring within the spatialization enabling scientists to directly manipulate ensemble members. Using OLI, scientists can move and cluster ensemble members within the spatialization to express knowledge about the data that could disagree with the generated visualization or test hypotheses. The new position of the clustered members only denotes their similarity/dissimilarity, so the specific location within the spatialization is less important. OLI transforms user interactions within a spatialized layout (i.e. visual feedback) into knowledge that is fed as input into underlying statistical and machine learning algorithms, leading to updates in the model’s attributes.

Scientists interact with the input parameters and simulation outputs using PLI by manipulating the slider. PLI allows the scientists to directly adjust or manipulate the dimension weight(s) of the underlying mathematical model based on the weight values on the slider. Adjusting the importance of either inputs or outputs gives scientists the chance to provide parametric feedback to a model regarding which dimension they believe is important, which in turn updates the weight vector and results in new projections of the data. Additionally, PLI enables them to study the correlation(s) between input parameters and simulation outputs and their influence on the ensemble members. In both OLI and PLI interactions, the weights and the spatial coordinates of the ensembles are manipulated to reflect a user’s interaction. Moreover, GLEE presents several statistical displays to help scientists understand the correlation between different attributes in multiple dimensions (1D, 2D and high-D). It also helps to identify areas of interest, asking quantitative questions about the ensemble behaviour, and exploring the distribution associated with the data between the different linked views.

4. System Design

We worked closely with scientists from various domains during GLEE's design. Geoscience and physics graduate students helped us in defining and achieving GLEE's design goals. Involving scientists from a diverse group preclude building *ad hoc* tailored solutions that work with specific application domains. Initially, we started by understanding scientists' analysis workflow through observations and unstructured interviews. We noticed that most of the scientists follow a manual analysis process that uses the visualizations of the simulation results and some statistical displays for input parameters or/and simulation outputs. This process relies heavily on trial and error, which is time-consuming and can easily lead to mistakes. Our preliminary visual design relied only on using SI as a way to help scientists analyse and understand the hidden patterns and similarities between ensemble members.

Iterating over the design during meetings, semi-formal interviews, focus groups and discussions with scientists, we found that SI alone would not be sufficient. Scientists need both the ensemble members and statistical measures in the same visualization to decrease the cognitive load taken by scientists when using different tools or scripts to analyse their data. Moreover, we found out that scientists are interested in analysing both parameter and ensemble domains simultaneously. Therefore, we altered our design to include a statistical view with multiple displays and a parameter view displaying the different simulation parameters to explore the effect of simulation parameters on the ensemble.

We observed that ensemble outputs are either 2D or 3D objects. So, instead of displaying the ensemble members as points within the spatial space, we used the output image of the simulation run. The images can help scientists interpret the similarities between ensemble members visually, and which may lead to better data analysis in terms of time and efficiency. Moreover, we added a viewpoint control slider for the Cinema thumbnails [OAJ*16] to allow scientists to view the 3D objects from different camera perspectives for a better understanding of their data and for emulating real-time rendering of and interaction with 3D objects.

- **Goal 1: Understanding the similarities and differences among and between different ensemble members**: Finding the distribution and similarities between similar groups of ensemble members, finding differences between dissimilar groups of ensemble members and exploring regions of variability within and across ensemble runs. Determining these similarities and differences will help scientists have a better understanding of the ensemble space.
- **Goal 2: 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 affect different ensemble members? Answering these questions will help scientists in predicting how changes in one or more of the parameter values influence the simulation output.
- **Goal 3: Understanding the correlations between parameters in single and multiple runs**: What is the correlation between two or more specific input parameters? How does the correlation between these input parameters change across the ensemble runs? Are there any global or local correlations between these input parameters across ensemble runs? Given a set of input parameter(s) and ensemble outputs, how do certain parameters correlate to ensemble output(s)? Answering these questions will help scientists gain a better understanding of the parameter space.

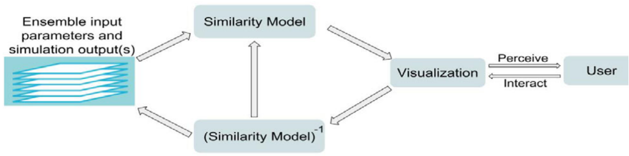

Figure 1: *GLEE's visualization pipeline has three main components: data source, similarity models and multi-view visualization. Data source used simulation's input parameters and outputs as well as corresponding weights as an input source to the similarity model. Similarity models pair forward and inverse computation through the translation of semantic and parametric interactions into manipulations of model parameters that are transformed into new visualization. The multi-view visualization allows scientists to interact with ensemble members and their attributes for exploring the ensemble.*

5. System

Our work is focused on visual exploration and analysis of high-dimensional simulation ensemble(s). Our system addresses the needs of domain experts, and the visualization components used are targeted towards GLEE's design goals. In this section, we discuss in detail the main components of our visualization pipeline (Figure 1).

5.1. Ensemble simulation attributes

GLEE's pipeline uses simulations' input parameters and outputs or derived output(s) as the data basis for representation in GLEE. We argue that considering both inputs and outputs in representing the ensemble visualization gives a more comprehensive representation of the domain and thus provides scientists with a more accurate picture of the data during the exploration process. Initially, scientists upload ensemble attributes and the corresponding images for each ensemble member. During the pipeline initialization, to avoid any distortions in the projection caused by raw attribute values, attribute values are z-score normalized before visualization. Additionally, a weight vector corresponding to all attributes is created. Each attribute is assigned an initial weight of $1/k$, where k is the number of attributes. All the weights inside the weight vector (l) are constrained to two conditions: sum to one and be within 0 and 1. The weight vector (l), along with normalized values of attribute(s), are then passed down the pipeline to the similarity model for processing.

5.2. Similarity model

Below, we discuss in detail the two components of the similarity model: forward model and backward model. The forward algorithm defines how data are processed for projection in the visualization layout using ensemble attributes and their weights. Conversely, the backward algorithm responds to user interactions by updating the

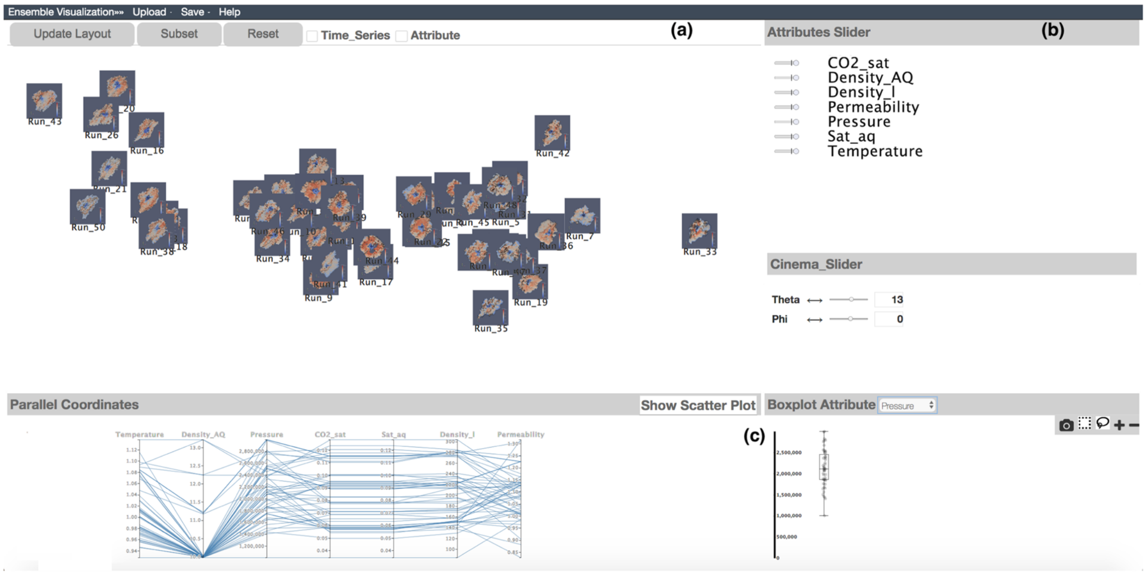

Figure 2: *The main interface of our visual exploration tool for ensemble simulation analysis: (a) ensemble view, each image represents an ensemble member laid out spatially using WMDS; (b) parameter view, shows weight slider for ensemble inputs and outputs; (c) statistical view, displays statistical properties and distributions of ensemble using different graphs (i.e. boxplot, scatter plot and parallel coordinates).*

attributes' weights and manipulating the ensembles spatially. 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 onto the visualization layout.

5.2.1. *Forward similarity model*

The forward similarity model projects high-dimensional data using the normalized values of ensemble input parameters and simulation output(s) and the corresponding weights. The forward model begins by calculating the weighted pairwise distance between all pairs of ensemble members. Many distance functions could be used for projection, such as Euclidean distance, Manhattan distance. However, the choice of the distance functions is determined based on the supported task and data type while taking the compatible of the projection technique into consideration. These pairwise distances are then fed to the projection technique, which determines the location of each ensemble member in the low-dimensional space by optimizing the following equation:

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

where $|\cdot|$ indicates the absolute value, $dist_L$ is the weighted Euclidean distance between low-dimensional points, and $dist_H$ is weighted high-dimensional distance function (in GLEE, we use weighted Euclidean distance) to calculate similarities between ensemble members normalized values. r represents the low-dimensional position of each ensemble member, r_i represents the i-th low dimensional value of each ensemble member and d_i represents the i-th high-dimensional value that represents the i-th ensemble member.

5.2.2. *Backward similarity model*

The backward similarity model is triggered when the scientist performs an OLI or a PLI or interacts with the statistical view. When performing an OLI, the new low-dimension positions of the moved ensemble members only are fed into an optimization algorithm that tries to find a new set of weights that describes the scientist interaction. The optimization algorithm starts an initial set weight for all attributes and iterates until it finds the set of weights that reflect the new positioning of the moved ensemble members. The following equation represents the optimization algorithm where w represents the newly calculated set of weights corresponding to the attributes, while $r_i{}^*$ represents the i-th low-dimensional value for the moved ensemble member i-th:

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

Alternatively, when the weight of a certain attribute is changed via a PLI, other weights are updated so that they sum to one. The updated weight vector from OLI or PLI is then fed to the forward similarity model for projecting the ensemble members. Similarly, when the scientist selects interesting regions or patterns in statistical views, corresponding runs are extracted and re-projected.

5.3. Visual encoding methods and interactions

In the previous sections, we described the main components of the pipeline but not the communication between the multiple coordinated views. GLEE consists of three linked views: ensemble view, parameter view and statistical view (Figure 2). In the following, we detail the three main data views and present interactions and linking capabilities.

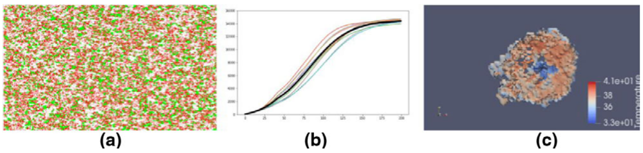

Figure 3: *Examples of simulation output images displayed in ensemble view: (a) the population of prey and predator on 2D lattice using the Lotka Volterra model, (b) mortality rates of different populations and (c) isosurface of saturated aqueous fluid for injecting CO_2 in a rock coloured by temperature.*

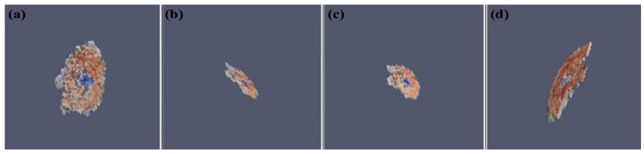

Figure 4: *Different camera positions for an ensemble member using different camera parameters: (a) phi = 0 and theta = 1, (b) phi= 1 and Theta = 0, (c) phi = 1 and theta = 2, (d) phi = 5 and theta = 5; this 3D rotation is driven by the Cinema slider in the ensemble view.*

5.3.1. *Ensemble view*

The ensemble view displays the low-dimensional projection of ensemble members. We used weighted multi-dimensional scaling (WMDS) [SSRY81] for the projection as it is easy to interpret weighted dimensions, making it accessible for scientists from various disciplines. Additionally, SI can loosely overcome WMDS linearity assumption. Each ensemble member has its own 2D image (Figure 2(a)) that can be 2D or 3D depending on the type of data (Figure 3). The image could be directly produced from the simulation run (Figure 3(a)) or processed using visualization tool (i.e. Paraview or VisIT) (Figures 3(b) and (c)). Scientists interact with ensemble members within the ensemble view by directly manipulating ensemble members through OLI, changing camera position to view the ensemble from different camera angles, zooming, lasso selection and multi-selection.

OLI has two modes of interaction: exploratory and expressive. During exploratory interaction, scientists can drag an ensemble member towards a cluster of ensemble members and update the layout. If the ensemble member is similar to the cluster members, it will be attracted to the cluster in the re-projected layout; otherwise, it will be repelled. This helps scientists in gaining insight into the structure of the ensemble by learning about a single run and how it relates to clusters of runs in the ensemble. On the other hand, expressive interaction allows scientists to express their knowledge or test an assumption or hypothesis. For example, scientists may disagree with the layout of runs within the spatialization based on their domain knowledge or can visually observe some interesting patterns from runs' images. Hence, they wish to determine what is similar between these runs. So, they drag these runs together forming a cluster and update the layout. The underlying computational models try to find the set of attribute weights that correspond to this similarity and update the spatial layout accordingly. Therefore, OLI supports the scientist in gaining more insight about the correlations and structure in the ensemble, testing a hypothesis or generally exploring the ensemble space achieving the first design goal.

Scientists usually observe interesting features in the images of some ensemble members and want to explore them more. GLEE offers multiple selection mechanisms in the ensemble view: lasso and multiple selections. Selecting multiple runs using multiple selection then clicking 'subset' button or lasso will automatically re-project the selected runs inside the ensemble view. The advantage of having selection mechanisms is not only inspecting and analysing regions of interests in a subset of the ensemble, but also determining local or global uncertainties in the data. This helps in eliminating misinterpretations and false assumptions about the ensemble correlations and relationships to parameters/attributes. Additionally, zooming is supported inside the ensemble view to assist scientists in getting finer details about each ensemble member in its thumbnail. Moreover, scientists could be interested in exploring different time steps. So, GLEE supports a time-series slider in order to help scientists navigate through the different time steps. The time-series slider is activated by checking the 'Time Series' checkbox. Each time step in the slider is treated separately from other steps. Manipulating the time slider will reset the pipeline with a new cinema database for this step, a new set of weights, simulation input parameters and the simulation output for this step.

5.3.2. *Cinema user interface*

To emulate real-time rendering and interaction with 3D ensemble members, we use image-based rendering from Paraview's Cinema exporter [OAJ*16, WAP*17]. Cinema is a framework for generating images from a structured sampling of camera positions of the visualization objects. Cinema facilitates capturing and exploration of important features in the simulation by taking many pictures from various camera positions around the dataset. GLEE's *Cinema Slider* enables scientists to view thumbnail images representing the ensemble runs from different camera positions. The camera positions of the runs form a database of pre-rendered images following the Cinema concept. This offers scientists different options for viewing angles on the simulation output (Figure 2 centre right).

To explore the image database, GLEE has a *Cinema Slider*, that has two sliders, each representing a camera parameter (i.e. phi and theta) to view different angles of the simulation output. Manipulating the slider(s) will change the camera angle resulting in new images for all the ensemble members in the ensemble view. This allows scientists to browse 3D data in real-time, facilitating the exploration of features and properties in the simulation ensemble. Moreover, it enables scientists to compare and contrast different ensemble members and to discover the view of interest within the ensemble that they want to explore and analyse (Figure 4).

5.3.3. *Parameter view*

Exploring the influence of parameters' space on the simulation ensemble has equal importance to the visual comparison and analysis of ensemble members. Therefore, GLEE has a parameter view

(Figure 2(b)) that represents the ensemble's high-dimensional input parameters and simulation outputs. The parameter view visualizes ensemble attributes using a horizontal slider where the value on the slider is not the raw data for this dimension but rather its weight. The weight represents the importance of this dimension. We believe that using weights instead of the actual numerical value is more representative because how the weights relate to one another is more meaningful. For example, if a scientist increases the weight of one dimension over the others, this implies that she orhe believes that this dimension is more important than others. Since we constrain all weights to sum to 1, this entails that increasing the weight of one more dimension requires the decrease of all other weights and vice versa.

Scientists interact with ensemble attributes using PLI. PLI allows scientists to directly manipulate the dimension weight(s) of the underlying mathematical model based on the weight values on the slider. This gives scientists the chance to provide parametric feedback to model about which dimension s/he believes is important, which in turn update the weight vector (l) resulting in new projections. This suggests that the similarities and differences of the data points on the manipulated dimension are amplified when that dimension weight is increased or lessened when that dimension weight is decreased.

The parameter view is linked with the ensemble view and vice versa. The result of PLI interaction is an updated weight vector that leads to a change of weights on the attribute slider and a new projection in the ensemble view. Scientists can use PLI to explore the sensitivity of attributes on ensemble, find the correlation between different inputs and outputs and find dimension(s) affecting the ensemble members clusters in the ensemble view. This helps in achieving the second and third design goal. Similarly, OLI interaction results in a new projection in the ensemble view based on the similarities/differences between ensemble members. It also results in a new weight vector that updates the slider in parameter view. OLI helps in identifying the most common and influential attributes between different clusters of ensembles.

5.3.4. *Statistical view*

One of the most challenging tasks when analysing high-dimensional data is identifying the regions of variability across all ensemble members and determining the associations between interrelated variables. Parameter and ensemble views enable scientists to find the most influential parameters, associations between parameters and similarities between ensemble members. However, a statistical view is still needed to allow 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.

Statistical displays can provide scientists with ways to catch mistakes, validate and refine the assumptions concluded from the other views, find new patterns and correlations not discovered by other views and generate new hypotheses that can be tested by other views. The statistical view can be used to get an overview about the distributions of the data as well as the correlation between parameters before manipulating the parameter and ensemble views. This helps in improving the accuracy and understanding of findings discovered by other views, which translates into a better understanding of the whole simulation model achieving the second and third design goals. Determining which statistical displays to use is a difficult question. In our choice for the statistical displays, we tried to balance between what is typically used by scientists in their analysis and the usability, readability and understandability of statistical displays for univariate, bivariate and multi-variate attributes. Therefore, we decided to use boxplots, scatter plots and parallel coordinates.

- Boxplot is a statistical display that characterizes univariate attributes using five-number summary. It is used to determine the central tendency, spread, variability, quartiles, outliers and the skewness of the data. GLEE's boxplot enables scientists to view the distribution of raw data for one single attribute across all runs at a time. Scientists can also use the dropdown list to navigate through different attributes.
- Scatterplot represents the correlation and relationship between bivariate data assisting scientists in finding patterns, trends, shape and distribution between them. This makes outlier detection an easy task, as regions with a higher density of points will be grouped perceptually, which helps in eliminating any bias produced. GLEE's scatterplot is designed to enable scientists to explore the relationships and correlations between any two attributes across all runs.
- PCP is a geometric projection method that visualizes multivariate data in a two-dimensional space. PCP displays multivariate attributes as parallel axes, where each attribute is mapped onto one axis and each ensemble run is mapped onto a polyline with vertices on these parallel axes. PCP supports plotting much information about complex multi-variate relationships simultaneously, which in turn, helps in getting an overview of all the data or, at least, a subgroup of attributes. It also assists in understanding and finding features within the data such as clusters, correlations, distributions across attributes or/and outliers. This is really beneficial in the analysis process when scientists have little to no prior knowledge about the ensemble. Moreover, scientists can interact with the PCP through axes changes and data filtering by imposing constraints on the value range of some attribute(s). This helps in comparing two or more ensemble members, comparing the variations of values between different attributes, investigating the influence of a single attribute or a group of attributes on the rest of attributes or/and estimating the degree of similarity between ensemble members. PCP could be an unfriendly graphical interface to novices; however, with little training, scientists can use it more easily [DCK12, ID87].

All our statistical displays support interactive brushing of the data points to select or highlight regions exhibiting specific statistical properties or interesting patterns, which provides immediate visual feedback to other views and statistical displays. In all three displays, scientists can select interesting patterns within the display and the corresponded runs are subsetted in the ensemble view and other statistical displays. Scientists can also select and highlight an interesting point within any statistical display that, in turn, highlights the selected run in all other views and displays and vice versa. Although statistical displays can suffer from visual clutter in cases of large datasets making identification of data structures and relationships

Figure 5: *(a) Population levels in an agent-based simulation ensemble that study epidemic dynamics and social behaviour. (b) Stochastic spatially extended Lotka–Volterra models that study noise-induced pattern formation and phase transitions in non-equilibrium systems for a predator–prey population in a Monte Carlo simulation ensemble.*

much harder, we believe this is not the case with GLEE as all of the domain experts we have been collaborating with have been working with ensembles that have less than 100 ensemble members. In summary, PCP is suitable for providing a general display of a large number of attributes at the same time, while the scatterplot offers a detailed comparison for pairs of attributes. On the other hand, the boxplot provides additional details about the distribution of univariate attributes that are somethings hard to detect in PCP, especially with a large number of runs and parameters. Therefore, cooperative employment of the three displays might not only enhance the exploration and analysis of the data but also supply detailed information about the data.

6. Implementation, Performance and Evaluation

Our web-based visual analytics tool uses JavaScript and the main visualization modules were built using the Data Driven Documents (D3) Library [BOH11] while the backend algorithm was implemented using Python. Visualization preparation and Cinema image database generation were done using Paraview. Ensemble data fed to the visualization pipeline is stored in a comma-separated values format, which is fetched directly to plot the visualization when the web page is loaded. Once the data are uploaded, scientists can freely explore and interact with GLEE.

6.1. Use cases

To validate the importance of GLEE in ensemble analysis across different fields of science, we demonstrated two potential applications of GLEE: population health and ecology. The first demonstration application is an agent-based simulation used to evaluate the 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. As the simulation is stochastic, ensembles are used to bound the uncertainly of results. In our initial example (Figure 5(a)), scientists used GLEE to explore the influence of variables for influenza immunization scenarios, which include the agents' compliance to public safety notices, the thresholds of triggering risk and the duration of the intervention. Dependent variables are simulation results like infection rates, mortality and productivity loss [VLC*18].

The second application is a Monte Carlo simulation ensemble that examines non-equilibrium relaxation of features in a stochastic Lotka-Volterra predator–prey model based on a 2D lattice (Figure 5(b)). Physicists studied the biodiversity in ecology and population dynamics through pattern formation and phase transitions in order to protect endangered species in threatened ecosystems. They used GLEE's visualization and interactions to help them find the reason for the system's relaxation. They found 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 neighbourhood of the critical predation rate [CT16].

6.2. Experiments

We conducted three experiments to assess the benefits and the drawbacks of GLEE in analysing and making sense of simulation ensemble. We mainly focused on measuring performance, evaluating functionality and usability and identifying features that require further research. GLEE's performance was assessed in terms of speed and accuracy. Speed was measured in terms of the time taken to react to users' interaction, synchronization of different views after each interaction and projection of high-dimensional data to 2D space. Accuracy was based on scientific correctness, insights and scientists' ability to derive meaningful relations and conclusions. On the other hand, GLEE's usability and usefulness were measured based on the ability to facilitate the reasoning process and achieve the designated design goals. Additionally, we investigated GLEE's power to fits into the research work, how it would help scientists improve their research, what scientists liked and disliked, and what features they wished to have in GLEE.

The experiments were performed on a laptop equipped with a 2.4 GHz Intel Core i5, with 8 GB of memory, and an Intel Iris Pro 1536 MB and NVIDIA GeForce card. GLEE was evaluated over two stages. During the first stage, we conducted a preliminary evaluation with 10 scientists (i.e. two assistant professors and eight graduate students) from different domains (i.e. Biology, Physics, Geoscience) to assess the viability of the concepts of OLI and PLI and to identify opportunities for refining the tool. In the second stage, GLEE was evaluated by two physics research scientists, three geoscience domain experts (i.e. an assistant professor and two graduate students). One of the geoscience graduate students and physicists provided the ensemble data used in the experiments. Before starting the evaluation, scientists were given a training session introducing them to GLEE's main functionalities and interactions. Later, they were asked to complete a set of analysis tasks. Finally, they were

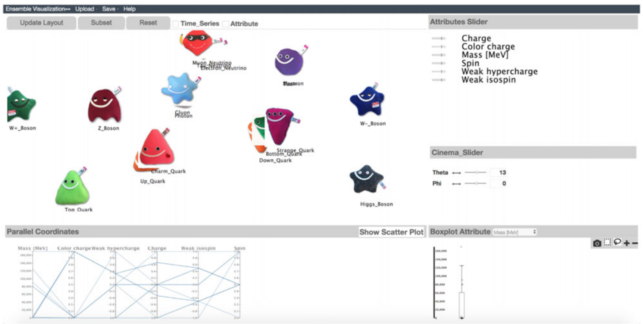

Figure 6: *The initial projection of elementary particle zoo. The subparticles are projected in ensemble view forming several clusters: electrically charged leptons, non-electrically charged leptons and two groups of quarks. All attributes have the same importance on the attribute slider in the parameter view and raw data are displayed in the statistical view.*

asked to report the weaknesses, usefulness, usability, the degree of difficulty and guidance provided by GLEE.

6.2.1. *Elementary particle zoo dataset*

The purpose of the first experiment was to measure the effectiveness of OLI and PLI and to discover how these interactions affect analysis tasks and the types of insights gained. For this experiment, we used an elementary particle dataset. Elementary particles are the smallest fundamental building blocks of matter that constitute the universe. They have two types depending on their spin: fermions and bosons. Fermions are particles that make up all matter, have particle spin equal to a half-integer value and come in two types: leptons and quarks. On the contrary, bosons are particles that carry the force and have spin equal to an integer. In this experiment, we treated each subparticle as an ensemble member. Each ensemble member has a six-dimensional vector describing its characteristics including: Mass [MeV], Charge, Spin, Colour charge, Weak Isospin and Weak Hypercharge.

From the initial projection of the elementary particle zoo (Figure 6), the scientist noticed several clusters of subparticles: electrically charged leptons, non-electrically charged leptons, and two groups of quarks. S/he started confirming his or her understanding of the data by clicking and dragging fermions and bosons to specific locations in the spatialization, forming two clusters. The location of the clusters could be anywhere in the spatialization, since the location just denotes the desired similarity/dissimilarity of the moved subparticles. After forming the clusters, the scientist clicked 'Update Layout', performing an OLI (Figure 7(a)). The resulting projection from the SI triggered by OLI coincided with the scientist's understanding that the spin attribute is what differentiates fermions from bosons (Figure 7(b)). This result gave the scientist trust in the correctness of the performed interaction in GLEE and freedom to interact with subparticles at an object-level without the need to manipulate external control menus that require an understanding of the underlying model.

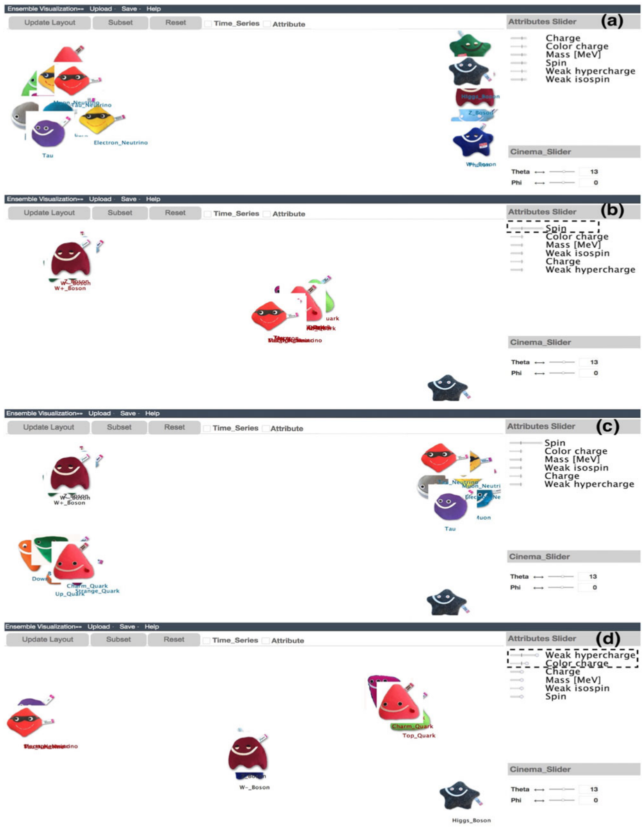

Figure 7: *(a) Scientists group fermions subparticles together in one corner to express their desired similarities and bosons subparticles in the other corner expressing their difference with fermions. (b) After that clicked 'Update Layout', the data are re-projected with new attribute weights showing that 'spin' attribute is what differentiates between fermions and bosons. (c) Scientists moved quarks subparticles together and leptons subparticles together to express their difference. (d) After that clicked 'Update Layout', the data are re-projected forming three clusters for leptopns, quarks and bosons differentiated by colour charge and weakisospin.*

The scientist was interested in gaining deeper insights about the data based on a couple of questions, including but not limited to: What attributes separate the quarks from leptons? What attribute(s) has the most influence on the data? Is there any relationship or correlation(s) between attributes that differentiate different ensemble members? To answer these questions, the scientist started by grouping quarks and leptons into two clusters performing an OLI (Figure 7(c)). Based on the reprojection, the scientist observed three new clusters for quarks, leptons and bosons. This led to an insight that colour charge and weak hypercharge are the attributes that describe the differences between the leptons, quarks and bosons (Figure 7(d)). Additionally, s/he noticed that Higgs boson is an outlier. This motivated him or her to determine what other attributes have an influence on the data, so s/he performed a PLI by increasing the importance of the mass attribute. According to the resulting projection, s/he observed two clusters (i.e. fermions and bosons) and an outlier (i.e. top quark) (Figure 8(a)). So, the scientist used

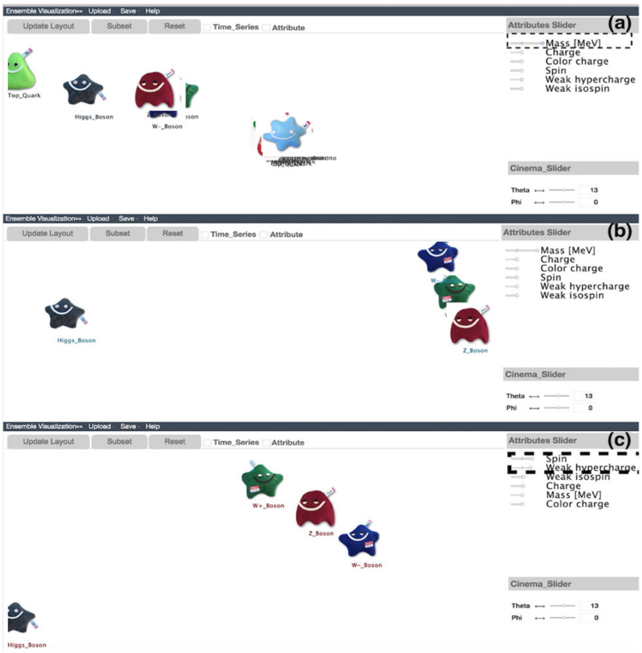

Figure 8: *(a) Increasing the weight of the mass attribute leads to an insight that fermions and bosons can be differentiated by mass attribute with top quark as an outlier. (b) Use OLI to determine what attribute(s) differentiates between high mass ensemble members and (c) the result projection showed that weak hypercharge and spin are what differentiate Higgs bosons from other high-mass subparticles.*

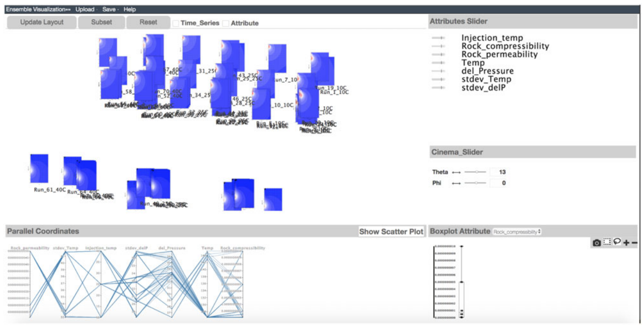

Figure 9: *The initial projection of the oilfield wastewater disposal dataset. Runs are clustered where deeper and lower magnitude pressure build-up runs are sorted on together and shallow and high magnitude pressure build-up runs are sorted together.*

the PCP– (Figure 6—statistical view)—to interpret and confirm this correlation. Based on the observation and PCP interpretation, the scientist gained a new insight that almost all fermions have low mass, most of bosons have a high mass and the top quark is an outlier that has the highest mass among all subparticles.

The scientist was interested in high mass subparticles and wanted to explore more correlations or similarities between them. So s/he subsetted high mass subparticles (W+ boson,W- boson, Z boson, Higgs boson) using multi-selection/lasso. By observing the mass attribute on the PCP, the scientists noticed that the Higgs boson had the highest weight. So, s/he separated the Higgs boson from other subparticles and performed an OLI. The new reprojection led to the insight that weak hypercharge and spin are the attributes that differentiate the Higgs bosons from other high mass subparticles (Figure 8(b)). The scientist was interested in knowing the exact correlation between these attributes across the different subparticles. So, s/he used the PCP to confirm and investigate more about this correlation. S/he concluded that Higgs bosons have an inverse relationship with other high mass subparticles in terms mass, spin and weak hypercharge—where it has the highest mass and weak hypercharge but lowest spin(Figure 8(c)).

6.2.2. *Oilfield wastewater disposal*

The second experiment used a 2-D simulation ensemble of 72 models of oilfield wastewater disposal that reproduce the operational and geologic characteristics of the Anadarko Shelf in northern Oklahoma and southern Kansas. Each simulation in the ensemble comprises the same model and injection scenario, but they have unique combinations of depth-decaying basement permeability, homogeneous basement compressibility and injection fluid temperature. The simulation ensemble is used to understand how measurable parameters affect the depth to which fluid pressure transients can migrate after 10 years of wastewater disposal operations. Because fluid pressure transients cause earthquakes in this region [PMTC18], the scientist is interested in the relationship between fluid pressure build-up and wastewater fluid temperature, rock permeability and rock compressibility [PMTC18, PCJW19]. He wants to gain insights based on the following questions: (1) Are there parametric similarities between runs in which fluid pressure migrates deeper into the formation? (2) What role, if any, does injection fluid temperature have on fluid pressure build-up? (3) What parameter(s) govern the runs with the highest fluid pressure build-up?

From the initial projection, the scientist observed that runs are sorted in a way that deeper and lower magnitude pressure build-up are sorted on one side, while shallow and high magnitude pressure build-up are sorted on the other side (Figure 9). Moreover, from the PCP, he observed an inverse relationship between rock permeability and outputted temperature. This leads him to an insight that high permeability allows the thermal fluid mass to mix with naturally occurring fluids. In contrast, low permeability inhibits mixing between the wastewater and natural fluids, so the thermal mass maintains it's injection temperature. Based on these observations, s/he started grouping runs' images visually based on how deep fluid pressure migrates in the image, performing as SI of OLI (Figure 10(a)). Following this re-projection, the scientist gains an insight that rock permeability is the most important parameter that controls the depth of pressure migration (Figure 10(b)), which answers the first question.

To explore the second question, the scientist increases the importance of injection fluid temperature on the parameter slider, performing a PLI. The resulting projection grouped simulation runs such that fluid pressure depth increases from left-to-right and fluid pressure magnitude increases from top-to-bottom (Figure 10(c)). Because this grouping sorted the simulation runs by pressure depth (left-to-right) with more subtle effects due to pressure magnitude

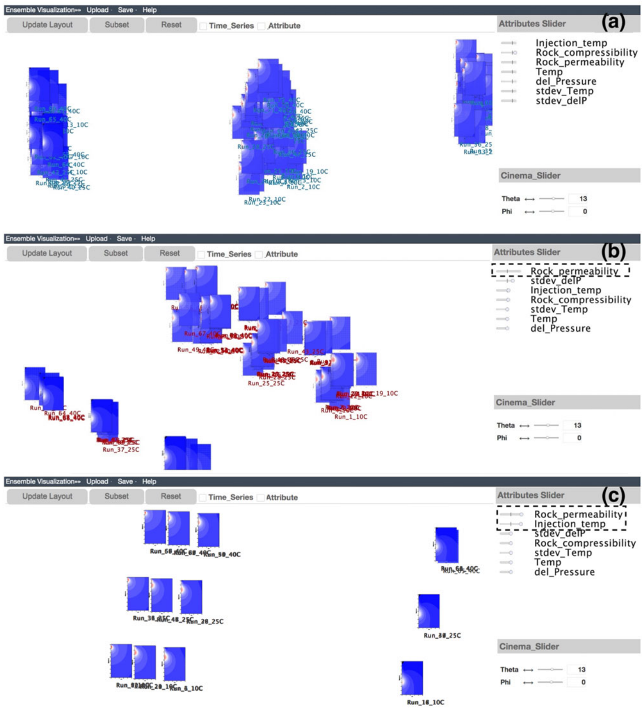

Figure 10: *(a) Scientists investigate how fluid pressure migrates deeper into the formation by dragging and clustering runs into three groups based on fluid pressure migration in the runs' images. (b) After that, he clicked 'Update Layout' gaining an insight that rock permeability is the most important parameter that controls the depth of pressure migration. (c) Increasing the importance of injection fluid temperature attribute through PLI results in grouping runs in a way that fluid pressure depth increases from left-to-right and fluid pressure magnitude increases from top-to-bottom. This leading to an insight that the temperature of fluid injection plays a role in the magnitude of fluid pressure build-up.*

Figure 11: *(a) Scientists investigate what differentiates the highest fluid pressure build-up runs by dragging and clustering runs based on deepness level, then performing OLI. (b) After clicking 'Update Layout', they gain an insight that pressure magnitude of these deep runs is governed by the temperature of injection fluid.*

(top-to-bottom), the scientist learns that the temperature of fluid injection plays a role in the magnitude of fluid pressure build-up, which answers the second question. The scientist then performed an OLI interaction on a subset of the simulations with the deepest fluid pressure perturbations (Figure 11(a)) and found that the pressure magnitude of these deep runs is governed by the temperature of injection fluid (Figure 11(b)). This both confirmed the answer second question and answered the third question because rock compressibility did not exert any quantifiable effects on the pressure magnitude.

6.2.3. *Geologic CO_2 sequestration*

The third experiment used a geoscience dataset of 50 numerical simulations ensemble that model the effects of geologic CO_2 sequestration at a site located in Richland, WA, USA. For this project, the permeability of the target formation is known only at the injection well, so there is substantial uncertainty with respect to CO_2 migration away from the well. To quantify the effects of this spatial uncertainty, the permeability distribution is randomly generated for each ensemble member to reproduce the known permeability values within the injection borehole, as well as the known cumulative distribution and spatial correlation of permeability in the region [JP18]. The models are comprised of identical geometry, CO_2 injection pressure, initial and boundary conditions and bulk permeability statistics; however, the spatial configuration of permeability in each ensemble member is both unique and equally probable. The CO_2 injection simulation is then completed for each ensemble member. Results are utilized to understand how the permeability distribution affects CO_2 migration pathways, fluid pressure propagation, temperature, CO_2 saturation levels and density of CO_2 water mixture [JWP19].

In this experiment, scientists have little prior knowledge about the simulated phenomena. So, they were interested in: (1) comparing and contrasting different groups of runs by finding out which inputs or outputs best describe these groups; (2) determining the sensitivity of input parameters on simulation outputs and exploring correlations between inputs and outputs; and (3) finding the best parameter settings for significant parameters, using distributions of data to investigate interesting patterns and subsets and confirming conclusions derived from both parameter and ensemble views.

The initial projection of the ensemble is shown in (Figure 12(a)), where each image shows the CO_2 plume at 1% saturation coloured by temperature. From this projection, the scientist recognized

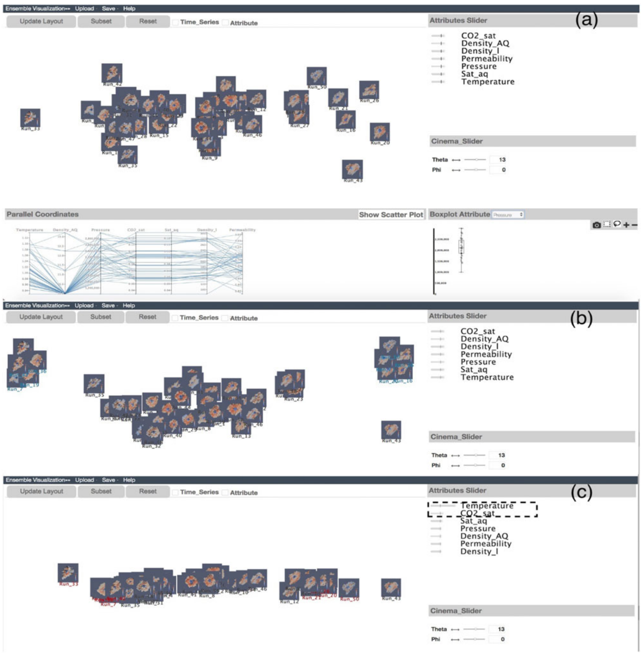

Figure 12: *(a) The initial projection of high-dimensional geologic ensemble inputs and outputs, (b) scientist observed interesting temperature patterns in ensemble images. So, they dragged high temperature patterns together in one corner and low temperature in the opposite corner performing an OLI. (c) After clicking 'Update Layout', they gained an insight that larger CO_2 plumes seem to exhibit more variability in temperature which in turn affects the ability of CO_2 to expand.*

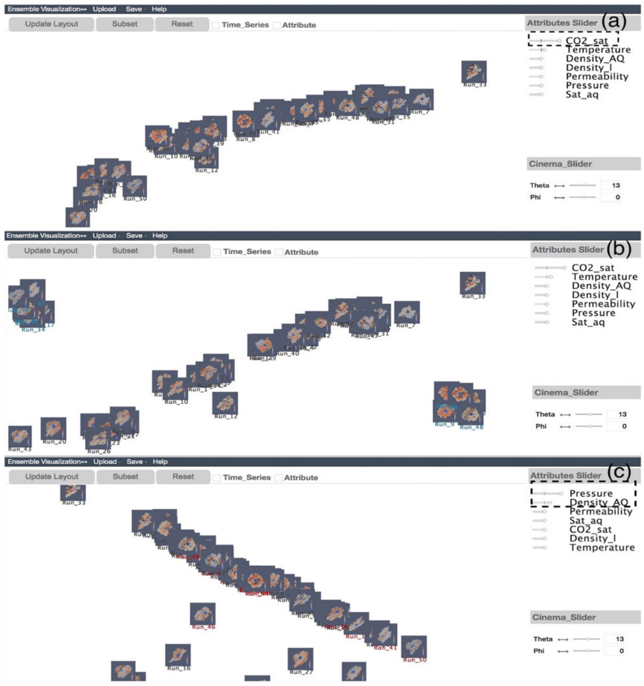

Figure 13: *(a) Scientists explore the influence of CO_2 saturation on the ensemble by increasing its importance using PLI. Based on the reproduction, they noticed that smaller CO_2 plumes appeared to be more circular than the more oblique large plumes. (b) Exploring the observation from resulted PLI, scientists grouped smaller CO_2 plumes on one side and larger CO_2 plumes on the other side performing an OLI. (c) After clicking 'Update Layout', scientists gained an insight 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.*

several interesting temperature patterns. To determine the distinction between these patterns, the scientist moved runs with high temperature patterns together in one corner within the spatialization and runs with low temperature in the opposite corner performing an OLI by hitting the 'Update Layout' button (Figure 12(b)). This interaction results in increasing the weights of CO_2 saturation and temperature (Figure 12(c)) leading to an insight that larger CO_2 plumes seem to exhibit more variability in temperature. The scientist then postulated this phenomenon is likely the result of interconnected high-permeability pathways that lessen fluid pressure, which in turn affects the ability of CO_2 to expand.

To investigate this insight further, the scientist used the box plot to examine the distribution of the CO_2 and then increased the weight of CO_2 saturation on the attribute slider performing a PLI. By observing the re-projected runs' images, the scientist visually determined that smaller CO_2 plumes appeared to be more circular than the more oblique large plumes. This observation guided the scientist to an insight that the anisotropic permeability correlation structure of the geologic formation exerts more control on plume geometry when the CO_2 reaches longer radial distances from the injection well (Figure 13(a)). Digging deeper into this insight, s/he grouped smaller CO_2 plumes on one side and larger CO_2 plumes on the other side performing an OLI (Figure 13(b)). The resulted projection leads to an insight 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 13(c)). This was an interesting new discovery that would require more experiments from the scientist.

The scientist was interested in finding if there is a correlation between temperature and CO_2, so he reset GLEE's pipeline and cluster ensemble runs based on temperature, CO_2 plume size and CO_2 plume shape then performed an SI of OLI (Figure 14(a)). The reprojected ensemble shows a dominant relationship of permeability (Figure 14(b)), which was a discovery that he did not expect. Trying to understand this relationship, the scientist examined the permeability of the ensemble runs from different angles using the cinema slider (Figures 14(c) and (d)). Based on the observations from different angles, he suspected a correlation between permeability and CO_2 saturation. So, he increased the weight of CO_2 through PLI (Figure 15(a)). The resulted projection leads to an insight that lower permeability leads to lower CO_2 saturation, and thus lower CO_2 storage capacity. Trying to confirm that this relation holds globally and locally across all ensemble runs, the scientist selected some runs based on interesting patterns in the runs' image, PCP and scatterplot. From the new projection, the scientist found out from runs' images that this correlation holds both locally and globally (Figure 15(b)).

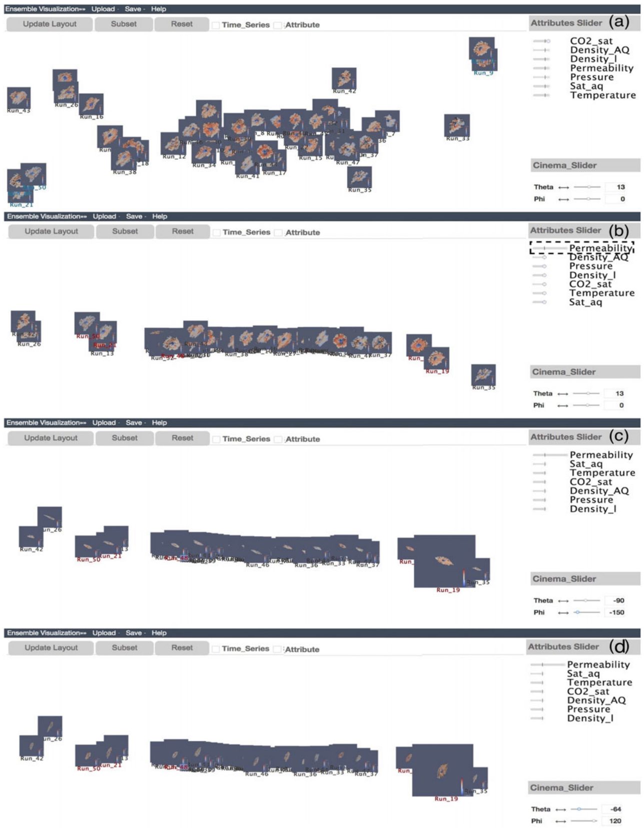

Figure 14: *Scientists explore the ensemble using OLI and Cinema slider in order to discover new findings by (a) grouping ensemble members by temperature, CO_2 plume size and CO_2 plume shape performing an OLI . (b) The re-projection showed a dominant relationship of permeability. They take advantage of the Cinema slider to interpret the results of projection and view data from different angles (c), (d).*

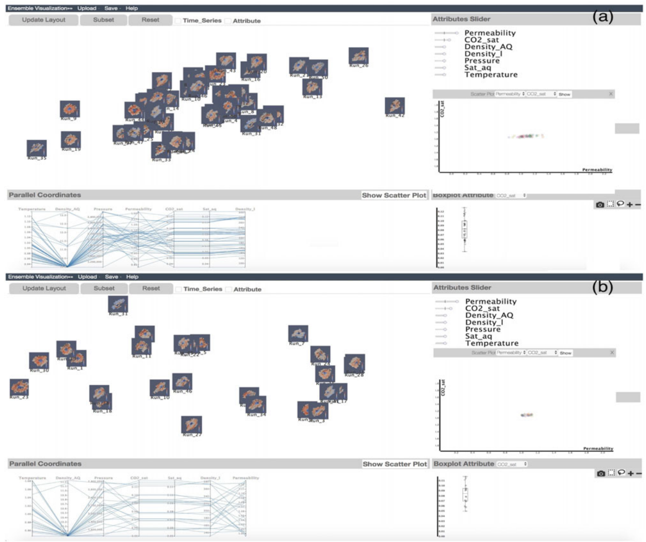

Figure 15: *(a) Scientists explore the correlation between permeability and CO_2 saturation by increasing the importance of CO_2 saturation using PLI. Based on the reprojection, scientists gained an insight that lower permeability leads to lower CO_2 saturation, and thus lower CO_2 storage capacity. (b) Confirming that the relationship between permeability and CO_2 saturation holds globally between all runs and locally between subset of runs, Scientists selected a subset of runs based on patterns in runs' images, PCP, and scatterplot.*

6.3. Domain expert evaluation

Domain experts feedback was satisfying. They mentioned that using OLI and PLI offers a different prescriptive for analysing their ensemble leading to several domain-specific discoveries. They believed that GLEE would help them in analysing datasets that they have little to no prior knowledge about. Additionally, they confirmed that the statistical view was beneficial in exploring the uncertainty and showing the distributions of the attributes, relationships between multiple attributes, trends and outliers. Similarly, the Cinema slider provides a means to view different angles of the data set, emulating real-time 3D rendering. This feature helped in viewing the data from different perspectives that open other angles for exploring parameter settings and ensemble members. Moreover, they confirmed that each of the three linked views conveys information about the ensemble helping in having a complete picture of the ensemble. They also mentioned that the GLEE's performance was reasonable and confirmed the ease of use of GLEE and its applicability to different datasets.

6.4. Discussion

This paper presents an exploratory visualization approach that helps scientists maximize their knowledge about the ensemble when they only have a big picture of their model. This implies that they do not know all the correlations or similarities held in the data. In such workflows, scientists are interested in gaining deeper insight into their ensemble. Our visualization approach applies to domain expert workflows that start with a basic understanding of the underlying model and then seek a broader understanding of the relationships and similarities in the model. In summary, GLEE is designed to support scientists in analysing their data when they are unsure which questions or hypotheses are worth exploring.

GLEE starts with an 'Overview first' mantra [Shn03], where a high-dimensional ensemble is projected into 2D space, scientists can then filter down to relevant ensemble runs during the analysis based on interesting patterns or features. This supports the situation where users' main activities are oriented towards exploratory analysis. However, if the scientists have a well-defined model where they know exactly what they need to analyse, they can use GLEE's 'selection' feature to select a certain run or subset(s) of runs of interest. GLEE's selection functionality offer scientists two modes of details based on either features observed in the images [CFV*16] or interesting patterns observed in statistical displays. Additionally, using the 'reset' button, scientists can reset the pipeline, recalling the whole ensemble and then select a new subset of runs that have a different feature. This gives scientists the opportunity to explore different features in the data. Using the 'selection' or 'reset' functionalities, scientists can alternate between 'Overview first' and 'Details first'. This makes GLEE capable of handling the two analysis modes. Additionally, GLEE enables scientists to benefit from the integration of the parameter space, ensemble space and statistics onto one screen without the need to change screens, programs or scripts that visualize the data—accelerating the analysis process.

During the experiments, we examined how the three linked views convey different information about the ensemble helping scientists to get a complete picture of the simulated model. The qualitative analysis of the experiments' results showed that using PLI, OLI and/or statistical displays enabled scientists to find optimal ranges, correlations between parameters, the sensitivity of parameters on ensemble members and similarities and differences between ensemble members. Additionally, GLEE's performance was reasonable as it took less than 3 s to react to user interactions and less than 1 s for synchronizing between views. GLEE's accuracy in answering the exploration questions varies, where some questions were answered with one interaction, while others took more interactions and time to be answered. Scientists used the appropriate view in GLEE to answer most of the questions. However, sometimes, when we expected them to use the ensemble view directly, they preferred using the statistical view first to have an overview of the distributions and the patterns in the data.

Current research in the ensemble visualization focuses on visualizing high-dimensional ensembles through parameter space or ensemble space. While some recent research addressed this problem, most of the solutions are oriented towards domain-specific problems without giving much attention to the importance of exploratory analysis. Moreover, human expertise and intuition are not integrated as part of the computational model during the interactive analysis process. Our visualization approach tries to fill this gap by designing a visualization tool that incorporates parameter space and ensemble space into one screen taking both the input parameter(s) and the simulation output(s) into consideration.

Different ensemble visualization approaches have used dimension reduction methods to visualize high-dimensional ensembles. What differentiates our work is the use of 2D images instead of glyphs [PMW13, HLNW11] to represent runs. Runs' images and Cinema slider support scientists in visually interpreting the similarities between runs (unlike other techniques that used points or circles to represent the runs [CZC*15]). Additionally, unlike interaction techniques supported by other visualization tools, GLEE's interactions provide scientists with the ability to get their expertise and knowledge as part of the visual analysis computational pipeline by manipulating the data on the object level, reinforcing the importance of exploratory analysis. This provides a bridge between user intention and the underlying mathematical models, which helps scientists not only understand their data but also find new information that could lead to discoveries.

7. Limitations and Future Work

There are several limitations to the current design of GLEE. First, our weighting scheme and inverse model currently works only with derived datasets or summary statistics of spatial data. Scientists still need to analyse spatial and temporal ensembles. So, we consider these important extensions for future work. However, the summary statistics of some spatial data show promising results, as illustrated in experiment three. Second, our tool does not scale well if there are thousands of ensemble members—as it could result in visual clutter; large displays can increase this number. Given that the size of most scientific data that 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 iteration. Future work should include research into scalability concerns, including data and display sizes. Moreover, our time-series slider treats each time step independently from other steps. We believe this limits the analysis of the ensemble as scientists could be interested in exploring the influence of parameters across different time or/and comparing runs from time steps. Future work includes more features for exploring with time-series ensembles.

In the current implementation, moved ensemble members are highlighted in a different colour after performing OLI or PLI, so scientists can track them for better understanding. Future work should include a backend technique that tracks the moved members to determine if they are moved to different clusters or not. Additionally, the current statistical view displays can be modified to incorporate a wider variety of charts (i.e. histogram, line charts, … etc.) 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 ensemble members. Although GLEE supports Lasso selection to facilitate ensemble members manipulation during OLI, more advanced mechanisms for selection will are needed as part of our future work.

8. Conclusion

In this paper, we presented GLEE, an interactive approach for visual exploration and analysis of high-dimensional simulation ensembles. Based on SI, dimensionality reduction, and brushing and linking techniques, GLEE enables interactive exploration of parameter space, ensemble space and statistical characteristics simultaneously while fitting human expertise and intuition into the computational model and analysis process. GLEE's multi-linked views offer different levels of details to support scientists in the analysis workflow for a complete picture of the ensemble. This helps in filling the gap between the analysis techniques used by domain scientists and the approaches available from visualization research, which could lead to new discoveries. We applied GLEE to case studies in physics, ecology and geoscience and got encouraging feedback from domain scientists who believe that GLEE could help them in improving the ensemble analysis process. Finally, we illustrated the utility of GLEE through several experiments with domain experts. The next steps in the development of GLEE will include supporting spatial and time-series ensembles, integrating machine learning to learn from user interactions to suggest new interactions that may help scientists in their analytic process, and scaling up to support thousands of ensemble members.

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.

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

[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.

[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.

[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. Hansen, C. D., Chen, M., Johnson, C. R., Kaufman, A. E. & Hagen, H. (Eds.), In *Scientific Visualization*. London: Springer, 2014, pp. 3–27.

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

[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.

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

[BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D^3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 2301–2309.

[BOL12] BRODLIE K., OSORIO R. A., LOPES A.: A review of uncertainty in data visualization. John, D., Rae, E., David, K., John, V. & Pak, C. W. (Eds.), In *Expanding the Frontiers of Visual Analytics and Visualization*. London: Springer, 2012, pp. 81–109.

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

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

[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.

[CD18] CAVALLO M., DEMIRALP Ç.: A visual interaction framework for dimensionality reduction based data exploration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018), ACM, p. 635.

[CFV*16] CHEN M., FEIXAS M., VIOLA I., BARDERA A., SHEN H.-W., SBERT M.: *Information Theory Tools for Visualization*. AK Peters/CRC Press, Boca Raton, FL, 2016.

[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.

[CLNL87] CARR D. B., LITTLEFIELD R. J., NICHOLSON W., LITTLEFIELD J.: Scatterplot matrix techniques for large n. *Journal of the American Statistical Association 82*, 398 (1987), 424–436.

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

[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.

[DCK12] DASGUPTA A., CHEN M., KOSARA R.: Conceptualizing visual uncertainty in parallel coordinates. *Computer Graphics Forum 31* (2012), 1015–1024.

[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.

[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), 1–15.

[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.

[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.

[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.

[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.

[EFN12] 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.

[EHM*11] Endert A., Han C., Maiti D., House L., North C.: Observation-level interaction with statistical models for visual analytics. *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2011), IEEE, pp. 121–130.

[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.

[FGS19] Faust R., Glickenstein D., Scheidegger C.: Dimreader: Axis lines that explain non-linear projections. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2019), 481–490.

[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.

[HdMRHH16] Höllt T., de Matos Ravanelli F., Hadwiger M., Hoteit I.: Visual analysis of reservoir simulation ensembles. Karsten, R. & Dirk, Z. (Eds.), In *Workshop on Visualization in Environmental Sciences* (2016) Groningen Netherlands, pp. 1–4.

[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.

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

[HLNW11] Hlawatsch M., Leube P., Nowak W., Weiskopf D.: Flow radar glyphs—static visualization of unsteady flow with uncertainty. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 1949–1958.

[HMZ*14] Höllt T., Magdy A., Zhan P., Chen G., Gopalakrishnan G., Hoteit I., Hansen C. D., Hadwiger M.: Ovis: A framework for visual analysis of ocean forecast ensembles. *IEEE Transactions on Visualization and Computer Graphics 20*, 8 (2014), 1114–1126.

[HOGJ13] Hummel M., Obermaier H., Garth C., Joy K. I.: Comparative visual analysis of Lagrangian transport in CDF ensembles. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2743–2752.

[ID87] Inselberg A., Dimsdale B.: Parallel coordinates for visualizing multi-dimensional geometry. Tosiyasu, L.K., In *Computer Graphics 1987*. Tokyo: Springer, 1987, pp. 25–44.

[JP18] Jayne R. S., Pollyea R. M.: Permeability correlation structure of the Columbia River Plateau and implications for fluid system architecture in continental large igneous provinces. *Geology 46*, 8 (2018), 715–718.

[JWP19] Jayne R. S., Wu H., Pollyea R. M.: Geologic CO_2 sequestration and permeability uncertainty in a highly heterogeneous reservoir. *International Journal of Greenhouse Gas Control 83* (2019), 128–139.

[Kan00] Kandogan E.: Star coordinates: A multi-dimensional visualization technique with uniform treatment of dimensions. In *Proceedings of the IEEE Information Visualization Symposium* (2000), vol. 650, Citeseer, p. 22.

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

[KW78] Kruskal Joseph B. Multidimensional scaling. No. 11. Sage, 1978. https://doi.org/10.4135/9781412985130

[LBS*18] Luciani T., Burks A., Sugiyama C., Komperda J., Marai G. E.: Details-first, show context, overview last: Supporting exploration of viscous fingers in large-scale ensemble simulations. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2018), 1–11.

[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.

[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.

[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), 1351–1358.

[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.

[Nor06] North C.: Toward measuring visualization insight. *IEEE Computer Graphics and Applications 26*, 3 (2006), 6–9.

[OAJ*16] O'Leary 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.

[OKB*18] Orban D., Keefe D. F., Biswas A., Ahrens J., Rogers D.: Drag and track: A direct manipulation interface for contextualizing data instances within a continuous parameter space. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2018), 256–266.

[PC05] Pirolli P., Card S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. Mark, L. & Mark, M. (Eds.), In *Proceedings of International Conference on Intelligence Analysis* (2005), vol. 5, McLean, VA, USA, pp. 2–4.

[PCJW19] Pollyea R. M., Chapman M. C., Jayne R. S., Wu H.: High density oilfield wastewater disposal causes deeper, stronger, and more persistent earthquakes. *Nature Communications 10*, 1 (2019), 1–10.

[Pea01] Pearson K.: Principal components analysis. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 6*, 2 (1901), 559.

[PKRJ10] Potter K., Kniss J., Riesenfeld R., Johnson C. R.: Visualizing summary statistics and uncertainty. *Computer Graphics Forum 29* (2010), 823–832.

[PMTC18] Pollyea R. M., Mohammadi N., Taylor J. E., Chapman M. C.: Geospatial analysis of Oklahoma (USA) earthquakes (2011–2016): Quantifying the limits of regional-scale earthquake mitigation measures. *Geology 46*, 3 (2018), 215–218.

[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.

[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.

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

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

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

[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. *Journal of Physics: Conference Series 180* (2009), 012089.

[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 *IEEE International Conference on Data Mining Workshops, 2009. ICDMW'09* (2009), IEEE, pp. 233–240.

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

[SH] Self J. Z., House L.: *Andromeda: Observation-Level and Parametric Interaction for Exploratory Data Analysis*. Technical report, Department of Computer Science, Virginia Tech, Blacksburg, Virginia, 2015.

[Shn03] Shneiderman B.: The eyes have it: A task by data type taxonomy for information visualizations. In *The Craft of Information Visualization*. Proceedings 1996 IEEE Symposium on Visual Languages. Boulder, CO, USA: Elsevier, 2003, pp. 364–371.

[SKBE17] Saket B., Kim H., Brown E. T., Endert A.: Visualization by demonstration: An interaction paradigm for visual data exploration. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2017), 331–340.

[SSRY81] Schiffman S. S., Schiffman S. B., Reynolds M. L., Young F. W.: *Introduction to Multidimensional Scaling: Theory, Methods and Applications*. Academic Press Incorporated, New York, NY, 1981.

[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.

[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.

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

[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.

[WAP*17] Woodring J., Ahrens J. P., Patchett J., Tauxe C., Rogers D. H.: High-dimensional scientific data exploration via cinema. In *2017 IEEE Workshop on Data Systems for Interactive Analysis (DSIA)* (2017), IEEE, pp. 1–5.

[WHLS18] Wang J., Hazarika S., Li C., Shen H.-W.: Visualization and visual analysis of ensemble data: A survey. *IEEE Transactions on Visualization and Computer Graphics 25*, (2018), 2853–2872.

[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.

[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.

[XXM*18] XU K., XIA M., MU X., WANG Y., CAO N.: Ensemblelens: Ensemble-based visual exploration of anomaly detection algorithms with multidimensional data. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2018), 109–119.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Data video S1

Data video S2

Data video S3

Data video S4

Data video S5

Data video S6