

Comparative Visual Analysis of Vector Field Ensembles

Mihaela Jarema*

Ismail Demir†

Johannes Kehrert‡

Rüdiger Westermann§

Technische Universität München

ABSTRACT

We present a new visual analysis approach to support the comparative exploration of 2D vector-valued ensemble fields. Our approach enables the user to quickly identify the most similar groups of ensemble members, as well as the locations where the variation among the members is high. We further provide means to visualize the main features of the potentially multimodal directional distributions at user-selected locations. For this purpose, directional data is modelled using mixtures of probability density functions (pdfs), which allows us to characterize and classify complex distributions with relatively few parameters. The resulting mixture models are used to determine the degree of similarity between ensemble members, and to construct glyphs showing the direction, spread, and strength of the principal modes of the directional distributions. We also propose several similarity measures, based on which we compute pairwise member similarities in the spatial domain and form clusters of similar members. The hierarchical clustering is shown using dendrograms and similarity matrices, which can be used to select particular members and visualize their variations. A user interface providing multiple linked views enables the simultaneous visualization of aggregated global and detailed local variations, as well as the selection of members for a detailed comparison.

Keywords: Uncertainty Visualization, Vector Field Data, Coordinated and Multiple Views, Glyph-based Techniques.

Index Terms: I.3.m [Computer Graphics]: Miscellaneous—

1 INTRODUCTION

Scientists increasingly use computational resources to run a numerical simulation several times with different input parameterizations or specific model variations. This produces an ensemble of simulations, each representing a possible realization of the simulated phenomenon. Scientists then analyze the ensemble variability (or spread), for example, to find the most likely predictions or separate similar predictions into groups to give an overview of the major trends represented by the ensemble.

Especially in meteorology, analyzing the variability of ensembles of vector fields plays an important role. For example, a number of meteorological features, such as cyclones (low pressure systems), anticyclones (high pressure systems), and jet streams (strong upper level winds) are manifest in the wind field. Variations in wind fields are thus an important source of uncertainty. In particular, poor forecasts of high impact weather events, e.g., cyclones, can often be traced back to disturbances in wind direction. To analyze why and where weather forecasts break, it is helpful to determine where the wind field disturbances occur and which ensemble members, i.e., which physical model or input parameter setting, caused the disturbance. In this work, we propose a visual analysis framework

to support such an explorative investigation of ensembles of vector fields. In designing our framework, we have considered the following requirements specified by our collaborators from meteorology:

Similarity clustering: Ensemble members should be grouped by similarity, as this enables an effective identification of outliers.

Distribution-based visualization: Visualizations should convey the main directional trends at every domain location, to indicate where the ensemble members start to separate due to disturbances.

Member relationship: Visualizations should reveal relationships between pairs or groups of members regarding their membership to local trends, to reveal where certain members behave similarly.

Interactive exploration: Global variability information should be visualized first, together with means to interactively zoom into the data to guide the analysis towards regions with interesting behavior.

Our novel approach for the visual analysis of *directional* data in 2D vector field ensembles builds on the characterization of distributions based on their modality. It uses the resulting representation to derive new comparative measures for ensemble members, and visually encode directional distributions in an intuitive way. Its major components are:

- *Distribution-based classification:* We model directional distributions using mixtures of probability density functions (pdfs). Each component is represented by its mean direction, the variation around this mean, and its weight.
- *Clustering:* We employ the modes of the pdfs to develop similarity measures that describe how similar the vectors are locally, and use these measures to cluster the ensemble members according to the extent of their variation across the domain.
- *Glyph-based visualization:* We propose a glyph-based technique to visualize the directional pdfs. We map the modes to the shape and orientation of a glyph, comprising a number of primitives equal to the modality of the mixture model. Glyphs reveal relevant information, such as the variation in the number of modes or in spread over the domain of the ensemble.
- *Brushing and linking:* We propose an interactive user interface with multiple linked views to enable the simultaneous visualization of aggregated global and detailed local variations.

2 RELATED WORK

The visualization of uncertainty belongs to the top challenges in scientific visualization [4, 31]. Uncertainty in the data often arises from ensemble simulations, where each ensemble member is computed with varied parametric conditions. The sheer complexity of an ensemble hinders an efficient analysis, since the data is typically spatiotemporal, multivariate, and multivalued [20, 23]. Several techniques have been proposed to reduce this complexity, e.g., computing statistical summaries such as mean and standard deviation, and visualizing them via color maps, contours, surface deformation, opacity, boxplots, or glyphs [23, 26, 29]. Other methods [19, 30] represent 2D distributions as volumetric data and apply traditional 3D visualization methods. Full distributions are also used by Kao et al. [18] for exploring and comparing pdfs from lidar data. Thompson et al. [40] introduce hixels (per-voxel histograms) for feature detection. Potter et al. [28] use a contour display to reveal the normed difference between a user-specified pdf and the pdfs over a 2D domain. While these methods deal with scalar distributions, our approach focuses on directional distributions.

*e-mail: mihai@in.tum.de

†e-mail: ismail.demir@mytum.de

‡e-mail: johannes.kehrer@tum.de

§e-mail: westermann@tum.de

Accepted for publication in Proc. IEEE Conf. Visual Analytics Science and Technology (VAST 2015).

© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

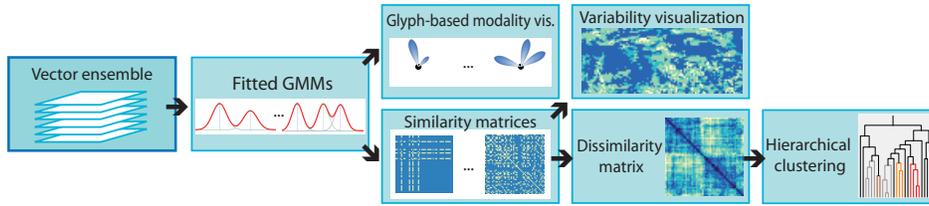


Figure 1: Directional distributions of vector-valued ensembles are modelled using Gaussian Mixture Models (GMMs). The resulting compact representations are used to design glyphs for visualizing local trends, and derive similarity matrices to classify members locally. The local similarity matrices are used to identify locations where ensemble members are dissimilar, and cluster members based on their similarity extent.

Clustering is a standard method for analyzing large and complex data [15]. Bordoloi et al. [5] propose realization- and distribution-based hierarchical clustering of ensemble data to reduce the amount of information to be visualized. Bruckner and Möller [8] use density-based clustering to identify similar volumetric time sequences in physically-based ensemble simulations. Beham et al. [3] use hierarchical clustering to group similar geometric shapes. Hummel et al. [14] cluster using Minimum Spanning Trees to compare the material transport in flow ensembles. Reh et al. [33] cluster similar pores in industrial XCT data into mean objects (MObjects) and then visualize the per-voxel probability of belonging to an MObject using transfer functions. In contrast to these approaches, we cluster ensemble members based on the extent of their similarity, where the proposed similarity measures are based on the modes of fitted mixture models.

Fitting Gaussian Mixture Models (GMMs) to the data is a popular summarization method for pdfs, which provides more insight than mean and standard deviation alone. Correa et al. [9] use GMMs to model uncertainty in the visual analysis process, while Hollister and Pang [13] apply GMMs to perform pdf interpolation. Liu et al. [22] employ multiple Gaussian components for a volume rendering of stochastic fields. Pöthkow and Hege [27], however, use nonparametric models for probabilistic function extraction.

Coordinated multiple views are commonly used to study multivariate relations via linking and brushing [20], or the complex relations between the input and output of an ensemble simulation [38]. Matrices are useful for studying pair-wise relationships in data, an early example being scatterplot matrices [2], which plot data attributes against each other to study correlations. MatrixExplorer [12] uses node-link diagrams and matrices to explore social networks. Talbot et al. [39] study classifiers in machine learning using an interactive matrix visualization. Bruckner and Möller [7] use a matrix to show the similarities between isosurfaces computed with different isovalues.

Glyph-based techniques have attracted a lot of attention due to their ability to show many attributes at once [6, 34]. Van Pelt et al. [41] propose a glyph-based blood flow visualization to compare different treatment options at multiple zoom levels. Arrow glyphs are a standard approach to show directional information in vector fields. Wittenbrink et al. [43] use such glyphs for showing uncertainty in orientation and vector magnitude. Recently, Pfaffelmoser et al. [26] presented circular glyphs to convey the mean and spread of the orientation of uncertain 2D vectors. While these methods assume unimodal distributions, we address multimodal distributions. Superquadric glyphs are often used to visualize tensors [36]. Schultz and Kindlmann [35] propose higher-order tensor glyphs for diffusion imaging. Jiao et al. [17] propose glyphs that represent the fiber orientation distribution function and the associated uncertainty, and HiFiVE glyphs [37] show fiber directions by embedding their pdfs in a Hilbert space. Even though our glyphs have similarities to those used in other works, to the best of our knowledge, the mapping of modality information to glyphs has not been addressed before.

3 OVERVIEW

Our method starts with an ensemble of 2D vector fields given at the grid points of a Cartesian grid structure. At every point, we model the distributions of the directional data using a mixture of Gaussians (cf. Section 4). This enables us to reduce the distribution to its main features, the so-called modes, which we use in two different ways: Firstly, to perform a local similarity analysis of the ensemble members based on their membership to these modes (cf. Section 5.1). Secondly, to summarize the main features of these distributions using glyphs, and to use these glyphs to visually encode the main directions of the modes and their spread at every point (cf. Section 6.2). Next, we aggregate the local similarity measures over the whole domain to cluster ensemble members depending on the extent to which they behave similarly throughout the domain, and show the locations of their variability (cf. Section 5.2). Finally, we describe how the visual representation of the comparative analysis is realized (cf. Section 6.1). The different computational stages comprising our approach are illustrated in Fig. 1. Our interactive user interface comprising four linked views is shown in Fig. 4.

4 MODELLING DIRECTIONAL DATA

In this work, we model directional data using parametric mixture models. Popular parametric models for circular data are the Wrapped Normal (WN) and the von Mises (vM) distributions [10]. Initially we fitted mixtures of both vM and Normal distributions, and verified the extent to which the modalities (number of components) of these mixtures repeated at the grid points. This was necessary because, depending on the initial values of the parameters, the solution of a mixture model may converge to a local optimum and yield different structures. To alleviate this problem, we repeated the fitting process several times with different starting values and selected as the best solution the one with the highest average silhouette from those whose modalities were obtained most often during the repeated fittings. Because the Normal mixtures yielded more consistent results for the considered data sets, we decided, without loss of generality, to use the GMMs in the rest of the analysis. The procedures are detailed in Appendix A1.

4.1 Modelling Data using GMMs

Circular data consists of observations that can be regarded either as unit vectors in the plane or as points on the unit circle. Choosing an initial direction (the x -axis) and an orientation (counterclockwise) allows specifying the observations by the angle from the axis to the point on the unit circle. We model the directional data (given as angles) using GMMs [25]. The main steps of our algorithm are outlined below. For a sample $\theta_1, \dots, \theta_n$, the pdf is

$$f(\theta) = \sum_{i=1}^N \alpha_i \mathcal{N}(\mu_i, \sigma_i^2)(\theta), \quad \alpha_i > 0, \quad \sum_{i=1}^N \alpha_i = 1, \quad (1)$$

where N is the number of Gaussian components parameterized by the mean vectors μ_i and variances σ_i^2 , and α_i are the weights of the components. For a given N , the $3N$ parameters $\{\alpha_i, \mu_i, \sigma_i^2\}$ are estimated using the Expectation Maximization (EM) algorithm. The true number of mixture components is, however, unknown and must

be inferred. We determine the modality of a pdf and its structure automatically, without a priori knowledge on the number of modes. To this purpose, we first perform a statistical test of randomness, so that we do not attempt to find clusters when the data is distributed uniformly. If this is not the case, we follow a procedure similar to that of Hamerly and Elkan [11] for the k-means algorithm, which has been shown to work well to determine the correct number of modes. Instead of k-means, however, we use the less restrictive EM algorithm, and fit a GMM to the given data set. Specifically, we start with one component and test whether it can be modeled by a single Gaussian distribution. If this is not the case, we run the EM algorithm to fit two Gaussians to the data set and repeat this procedure recursively, until no component needs to be split anymore.

Note that GMMs cannot be applied directly to directional data, since, depending on the location where the circle is cut to unwrap it to an interval of length 2π on the real line, a mode on the circle may split into two modes on the line. Instead of cutting the circle, Mardia and Jupp [24] recommend repeating a complete cycle of the data, yielding an interval of length 4π on the real line. This allows using GMMs with the standard EM algorithm.

Once a GMM has been fitted to the augmented data set, the next step is to restrict the extended interval to the initial one and summarize the final components on the circle via sample trigonometric moments. For a component with m angles, we consider the corresponding unit vectors and compute the mean direction μ as the angle of the resultant vector after vector addition. The mean resultant length ρ associated to the mean direction μ is the length of the resultant vector, normalized by m . ρ takes values in the range $[0, 1]$, higher values showing increased concentrations around the mean direction [24]. The two parameters characterize the Wrapped Normal distribution $WN(\mu, \rho)$, the result of wrapping a Normal distribution $\mathcal{N}(\mu, \sigma^2)$ (given on the line) around the circle:

$$\rho = \exp\left(-\frac{1}{2}\sigma^2\right), \text{ where } \sigma^2 = -2\log(\rho). \quad (2)$$

5 COMPARATIVE ANALYSIS OF ENSEMBLE MEMBERS BASED ON FITTED GMMs

In this section, we show how the mixture models of the directional distributions can be used to develop local comparative measures for the ensemble members and reveal the locations of their variations in the spatial domain. We then present how the local measures can be accumulated and used to cluster members depending on the extent of their similarity. Finally, we give insight into how the behavior of the members at one grid point can be compared across the domain.

5.1 Local Comparative Measures

Various measures can be used to assess the directional similarity of two ensemble members at a certain location, e.g., the angle between the vectors or their cosine similarity. Two members can be classified as angularly similar if, for example, the measure is below a prescribed threshold. Instead of imposing an arbitrary threshold, we use the previously derived mixture models to assess pairwise similarities relative to the variability of all the ensemble members (explicit formulae are given in Appendix A2).

GMMs model the n angular observations at every grid point of the domain with a mixture of Gaussian components. Observations are assigned to these modes based on the estimated posterior probabilities that an observation belongs to a certain mode. The clustering is called *hard* when observations are assigned to the mode which they are most likely to belong to a posteriori, and *soft* when observations have various probabilities of belonging to each mode.

For our first local similarity measure – the *modality* measure – we use soft clustering and consider two ensemble members as similar at a certain location if their angles have a posterior probability over a given threshold of belonging to the same mode. By using soft clustering, we do not force members that lie at the borders of two

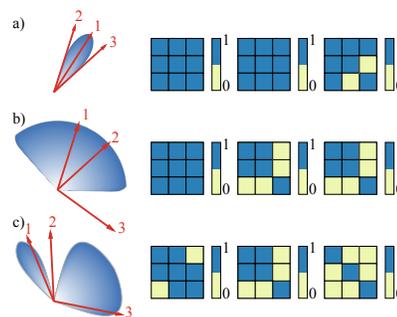


Figure 2: Similarity measures for selected members of three synthetic data sets. First column: selected members (red) and glyph representations of a unimodal pdf showing (a) low variation and (b) high variation, and (c) a bimodal pdf. The next columns depict the similarity matrices corresponding to the three types of measures.

modes to be classified as similar to the members of only one mode, and the members that lie well within their clusters are not affected. At every location, the binary similarity values (1 for directions with the same mode membership and 0 otherwise) can be summarized for n members into an $n \times n$ symmetric similarity matrix.

While this measure differentiates between members that clearly belong to different modes, it does not give any insight into how similar the members are within the modes. For instance, members modelled by unimodal distributions are all classified as similar, even if some observations are located far from the main data mass. We thus propose a second similarity measure – the *scaled angular* measure: For every pair of ensemble members that are similar using the modality measure, we select the mode with the highest sum of posterior probabilities and consider the members as similar if the smallest angle between them is less than or equal to the maximum sample circular standard deviation (cf. Eq. 2) in the domain. In this way, members with large deviations are no longer considered similar, even if they belong to the same mode.

If users are also interested in the relatively larger deviations occurring in regions with low variation, the angular deviation can be computed relative to the local sample circular standard deviation. Thus, for our third measure – the *locally scaled angular* measure – we consider the angular deviation of two members relative to local variation, locations exhibiting higher variations allowing larger deviations than those with lower variations. For uniform distributions, where all directions are equally likely, we set a threshold, e.g., $\pi/2$, from which two realizations are considered dissimilar.

Fig. 2 illustrates the proposed similarity measures using three synthetic examples. The first column contains glyph representations (cf. Section 6 for details on the glyph construction) of the corresponding pdfs, together with three selected members; the following columns show the similarity matrices (corresponding to the three similarity measures) for the selected members. Assessing members based on their cluster membership (second column), classifies the members of both unimodal pdfs as similar. For the bimodal pdf (Fig. 2(c)), members 1 and 3 belong to different modes and are classified as dissimilar. Member 2, which is located at the border of the two clusters, is similar to both 1 and 3. Considering the angular deviations relative to the highest variability (third column), new dissimilarities emerge: in Fig. 2(b), member 3 (an outlier) is now dissimilar to both 1 and 2 (both inliers); furthermore, in Fig. 2(c), members 2 and 3, located towards the opposite ends of the wider lobe, are also classified as dissimilar. Relative to the local variations (fourth column), members 2 and 3 in Fig. 2(a), and 1 and 2 in (c) are additionally described as dissimilar. Even though members 1 and 2 in Fig. 2(b) have approximately the same angle as 2 and 3 in (a), they are not dissimilar relative to their local variation.

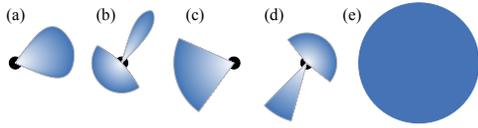


Figure 3: Glyphs for various directional distributions: lobular glyphs showing (a) unimodal pdf, medium variation and (b) bimodal pdf, main mode (low variation), second mode (high variation), and (c)-(d) corresponding pie glyphs; (e) uniform pdf.

5.2 Global Comparative Measures

Summarizing similarity measures via similarity matrices at every grid point reveals the locations in the spatial domain where two or more members disagree. For instance, to find the regions of different behavior for a pair of members (i, j) , we extract from every local similarity matrix the binary value at row i and column j to construct a binary field of similar and dissimilar locations. For more than two members, the values at the every grid point range between 0 and 1, depending on how many pairs of members are similar, out of the total number of pairs $(1 + k(k - 1))/2$ for k members).

The binary fields of any two members can be used to define the global similarity of these members, as the percentage of similar locations out of the total number of grid points. To summarize all global similarities into one representation, the pairwise global measures can be assembled in a *global similarity matrix*. Then, the binary fields and the global similarity matrix can be used to group together ensemble members showing similar behavior.

We cluster ensemble members hierarchically [16], starting with pairs of similar members and joining clusters at one level into larger clusters at a next level until a hierarchical cluster tree is formed. For the linkage criterion, the similarity between sets of observations is computed by accumulating the similarity values (between 0 and 1) at the grid points and normalizing by the total number of points.

5.3 Comparing Mixture Models

In addition to comparing ensemble members locally and clustering them based on their global similarities, we also examine how the behavior of the members at one point compares to that at other locations. Several questions arise in this context. For instance, which other locations exhibit the same modality? This can be answered by a simple query-based exploration of the modality scalar field. Queries can also be run to determine the locations where the pdfs have similar shapes. While there is no best measure to assess the similarity of two directional pdfs, we looked for a measure that has a closed form for GMMs and is thus easily computable on the fly. We thus built upon the concordance coefficient [32] (cf. Appendix A3). Finally, we want to assess the extent to which ensemble members vary in the same manner at other locations. For instance, given a bimodal distribution, are the sub-sets of members contributing to the two modes the same at other bimodal pdfs? To answer this question, we compare the similarity matrices at the other points with the matrix at the selected location and determine the percentage of common similarities from the total number of similarities.

6 VISUAL ANALYSIS OF VECTOR ENSEMBLES

Our interactive user interface (cf. Fig. 4) comprises four linked views, two over the spatial domain and another two abstract representations of the members similarities.

6.1 Similarity Visualization

The hierarchical clustering of the ensemble members for any selected similarity measure is summarized graphically as a dendrogram (cf. Fig. 4(C)), with the members along the bottom horizontal axis and the levels of clustering along the vertical axis. The U-shaped lines join subclusters (consisting of one or more members)

to each other until all members have been linked into a single cluster. Similar members are joined at lower-level clusters, the variability among the members increasing as subclusters are merged. The ordering of the members on the horizontal axis is such that the sum of the similarities between adjacent leaves is maximized (but without dividing the clusters). To separate the groups of members more easily, main neighboring branches are depicted in different colors.

Dendrograms provide the users with more than a single set of clusters, because the nested partitions show how subclusters relate to one another. Nevertheless, because the merging step is greedy and constrained by previous choices, rather similar members may end up in different main branches. Thus, the global similarity of any two members is shown in the second abstract representation (cf. Fig. 4(D)) – the global dissimilarity matrix (the complement of the global similarity matrix). The dissimilarity values are shown using a sequential ColorBrewer scheme (also used in the two spatial views). The ordering of the members in the matrix is that of the dendrogram, the subclusters appearing as block-diagonal forms. Users can also have the dissimilarity matrix sorted by member ID.

The main view over the domain color codes various derived scalar fields, such as the modality (default) or the dissimilarity field of the selected members, showing at each location the number of dissimilar members out of the total number of pairs. (Groups of) members may be selected by clicking on the respective branch in the dendrogram or on a cell in the dissimilarity matrix. Brushing in the dissimilarity matrix allows the simultaneous selection of several members. Specific members can also be specified as text input.

6.2 Glyph-based Modality Visualization

Since displaying pdfs over the entire domain would lead to massive clutter and occlusion, users can interactively zoom into a region (of fixed dimensions) in a detail view (cf. Fig. 4(B)). To visualize circular distributions, we propose a lobular glyph (cf. Fig. 3(a)-(b)) that reveals the main characteristics of a multi-modal directional pdf: the number of components, as well as their directions, weights, and widths. A glyph consists of a number of lobes equal to the number of modes, where the direction, weight, and width of each mode is mapped visually to the orientation, length, and opening angle of the corresponding lobe. The width is given by twice the sample circular standard deviation. For the sake of clarity, lobes have a minimum opening angle and length. Moreover, for multimodal pdfs, weights are normalized by the maximum weight before being mapped to the length of the lobes. Uniform distributions are shown by a full disk (cf. Fig. 3(e)). To ease the matching process, glyphs have the same color as the corresponding locations in the main spatial view.

The lobular glyphs convey the main directions in which the majority of the unit vectors point, and the variation around these directions. Meteorologists, however, frequently show the directions from which the winds blow as well. To accommodate this type of visualization, showing the directions from rather than into which the members point, we need to mirror the glyphs, the direction being now read from the head of the glyph to the tail. We also modify the lobular form to a pie (cf. Fig. 3(c)-(d)). Comparing lobular to pie glyphs, their pointy heads emphasize the main directions and are thus effective at conveying the direction, especially for narrow lobes (where showing a clear direction is meaningful). It is, nevertheless, less efficient at showing the variability, because narrow modes will look narrower displayed via lobes than via pies. Wider modes will not suffer from this drawback, though, since lobular glyphs converge to pie glyphs as the width increases. The main direction is more difficult to read in pie glyphs. However, when mirrored to show the directions from which the vectors point, the direction is read from the head to the tail of the pie, which is pointy and thus clearly indicates the direction. To be able to distinguish the directions in (from) which most of the members point (if existent), we draw small anchor circles at the tail (head) of each lobe (pie).

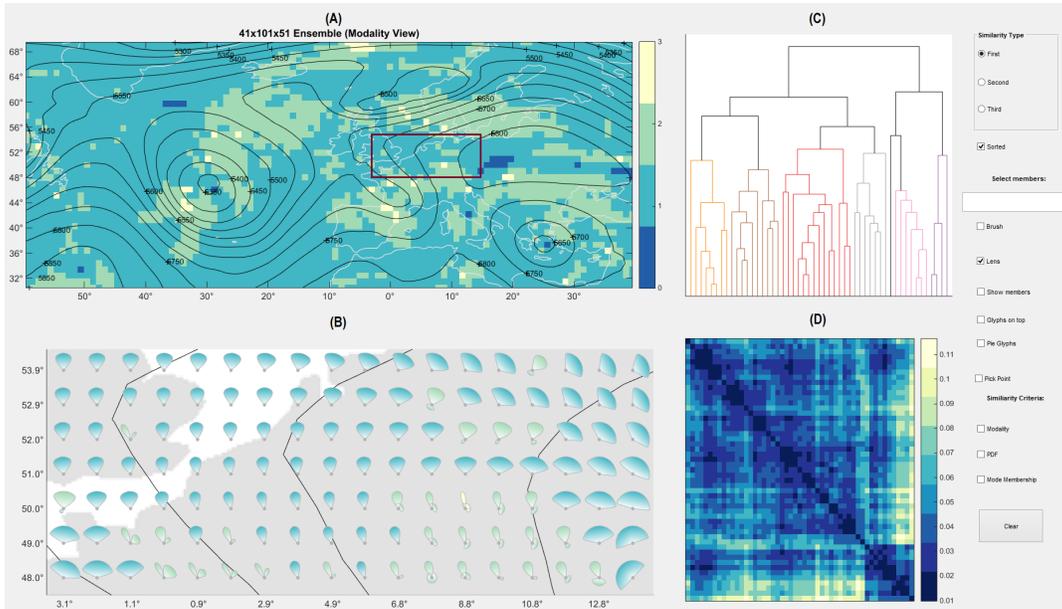


Figure 4: Multiple linked views of a wind ensemble. View (A) color codes geolocated scalar measures over the entire domain (here the modality of the distributions), while (B) shows a detail view over the region selected by the user (marked in A by a rectangle), where directional distributions are displayed using glyphs. The isocontours show the geopotential height. The hierarchical clustering of the ensemble members is shown using a dendrogram in (C), and the pairwise dissimilarities between ensemble members are shown in the global dissimilarity matrix in (D).

Moreover, the preferred direction is emphasized using a color gradient, the luminance value decreasing as the distance from the tail increases. According to the Gestalt laws of pattern perception [42], this enhances the direction information, since the luminance contrast makes the head more distinct than the tail. Furthermore, as our collaborators from meteorology wanted to be able to see the outliers in addition to the main trends, we allow displaying the members as gray sticks, either on top of or under the glyphs. Selected members are shown as black sticks and are always drawn on top.

Finally, users can run queries to see how members at selected locations behave compared to the other locations (cf. Section 5.3). The selection of a glyph can be done by clicking into either one of the spatial views. Once a glyph has been selected, nonsimilar locations (and corresponding glyphs) fade out and only similar locations and glyphs maintain their colors. The fading out depends on the similarity criteria chosen by the user. Thus, when looking for glyphs where the pdfs have the same modality and/or a similar pdf shape (up to a predefined threshold), all dissimilar locations in the main spatial view and the glyphs in the detail view fade out to the same extent. When searching for locations where ensemble members vary in the same manner, the degree of fading at each location is proportional to the extent to which members disagree in their behavior. Since the color scheme goes over yellow, in order to discern between similar yellow glyphs and non-similar gray glyphs, the contours of the dissimilar glyphs become light gray, while the contours of the similar glyphs maintain their dark gray color. To enhance those similar glyphs at locations where members exhibit comparable behavior, the contours of these glyphs become black. The contour of the selected glyph is drawn in a contrasting color.

7 RESULTS

We demonstrate the potential of our techniques on three ECMWF (European Centre for Medium-Range Weather Forecasts) ensembles. The first two data sets are wind ensembles, given on a 101×41 Cartesian grid, and comprising 50 members and a control run. Our third data set is an ensemble of temperature gradient fields derived from an ECMWF forecast, given on a 221×101 Cartesian grid and comprising 50 members. Details on the system are available in [21].

7.1 Similarity Analysis

The first ensemble is the 120 hour wind direction forecast at a pressure level of 500 hPa, valid on October 19, 2012. The geolocated scalar field containing the modalities is shown in the upper spatial view in Fig. 4, along with the isocontours of the geopotential height field of the control run. A hard clustering based on an empirical orthogonal function analysis of the 500 hPa geopotential field was available from ECMWF, identifying three main groups of cardinality 21, 20, and 10. While a detailed analysis of this clustering solution was not possible (the intra- and inter-cluster similarities were not available), since at this pressure level winds blow approximately parallel to the contours of the geopotential heights, we were interested to verify whether the ECMWF clustering was consistent with our results. Examining clustering solutions obtained using several methods and measures is highly important in meteorology, because there is no single “best” clustering method and the different techniques influence the clustering solutions, so that only consistency across techniques helps gain confidence in the validity of the results [1].

We used silhouette coefficients (internal cluster assessment criteria) on the ECMWF clusters to check the extent to which the members in these clusters are both closely related and well-separated from the members in the other clusters. The analysis revealed that the members in the third ECMWF cluster (nine of them placed in the last main group in both the dendrogram and the dissimilarity matrix in Fig. 4) form a coherent cluster (with coefficients around 0.4 – 0.6) and are well-separated from the majority of the other members. This is also noticeable in the two abstract views in Fig. 4, where the merge with the last main branch occurs at the highest level in the dendrogram and the bulk of the dissimilarity values to the other members are at the upper end of the color scheme. The members in the other two ECMWF clusters were not well-separated based on the silhouette analysis. However, several nested subclusters consisted primarily of subgroups of these ECMWF clusters.

The two abstract views showing the relationships between the ensemble members aid the users to quickly identify similar members, as these are joined at low levels in the dendrogram and their pairwise low dissimilarities are shown in the global dissimilarity

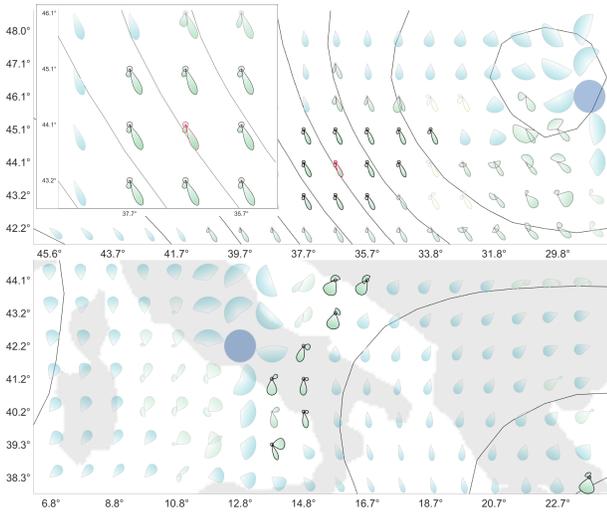


Figure 5: Querying for locations similar to the one selected (corresponding glyph outlined in red). The locally scaled similarity measure finds only neighboring locations as similar (top figure), whereas the scaled measure is less restrictive and also indicates remote locations (bottom figure) as similar.

matrix in a dark blue color (primarily on the diagonal of the matrix). Outliers are also rapidly identified, since these members are typically merged at a higher level in the dendrogram and the dissimilarity values vary primarily in the middle and upper range of the color scheme. An example is member 11, which is the last member to be merged in the first (orange-colored) subcluster in the dendrogram.

Various measures of the variability of the members over the domain are encoded in the main spatial view, Fig. 4(A) showing the modality of the pdfs. Users can then select interesting regions to examine the pdfs at these locations in a detail view (cf. Fig. 4(B)). We noticed that regions of high variability (due to high spreads or multimodality) occur primarily at and around the pressure centers (the smallest circles in the centers of the concentric closed contours of the geopotential height) and at the boundaries between centers.

Local understanding can also be enhanced by running queries to determine, for a selected glyph, glyphs with the same modality, a similarity pdf under a threshold τ , and a similar behavior of the ensemble members. Fig. 5 shows an example, where the bimodal glyph of the selected location is drawn in a red outline in the top figure. The top figure reveals locations where members behave similarly in the neighborhood of the selected glyph. The glyphs that are most faded out and drawn with a light gray contour have different modalities and are thus dissimilar. Regarding the pdf shape, the amount of allowed angular deviation in the means depends on the existent variation: A mode with a higher variation and a larger deviation is assessed as more similar than a mode with a lower variation and a smaller deviation, if its deviation is larger relative to the existent variation. As the bimodal glyphs have the main modes pointing in a direction similar to that of the selected glyph, the similarity in shape (for a threshold $\tau = 1$) does not exclude any of these glyphs. The final color of the glyphs that represent pdfs of similar modality and shape depends on whether the members at these locations vary in the same manner. Because we use the locally scaled angular similarity measure, the percentage of members classified as similar at other locations is around 70% only in the immediate neighborhood of the selected point. Glyphs at these locations are visually enhanced in the figure. Other locations do not exceed 30%. If, however, we use the scaled angular measure instead (relative to the maximum global variation), several remote locations reach values of even 80%, while the neighboring locations are over 90%. This

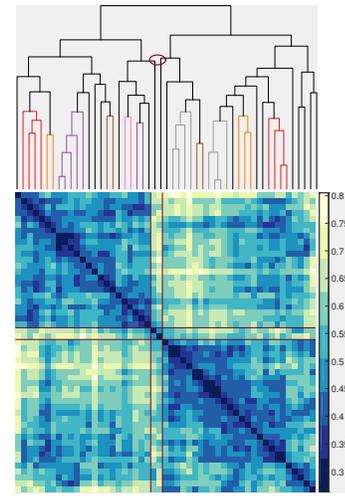


Figure 6: Clustering representations of the 51 ensemble members at 850 hPa (shown here) reveal better separated groups than at 500 hPa: dendrogram view (top) and dissimilarity matrix (bottom). The color coding of the dissimilarity varies from blue (low) to yellow (high).

happens because this measure considers members as similar if they belong to the same mode and their angular deviation is within the maximum limit allowed. An example of a remote region exhibiting similar locations is shown in the bottom figure of Fig. 5. Here, it can also be noticed that several glyphs showing bimodal pdfs have been classified as dissimilar due to their different pdf shapes.

The second ensemble is the 120 hour wind direction forecast at a pressure level of 850 hPa. Closer to the ground, the wind direction is influenced by the surface (due to friction). Regions of high variability occupy a much larger area than before, emerging at pressure centers, boundaries between centers, and over irregular terrain.

The groups of ensemble members are now better separated than at the 500 hPa level, (cf. Fig. 6 for the locally scaled angular similarity measure). Only a few ensemble members exhibit similarities to members in other clusters. The dendrogram points out several outliers merging at high levels, members 16 and 39 (their dissimilarity values to all members are outlined in brown in the dissimilarity matrix) being the last members to be merged. Their outlier status is confirmed by the dissimilarity matrix, which nonetheless also reveals that the two members are rather similar to one another.

Fig. 7 (top) illustrates the dissimilarity field for a group of three similar members (according to the locally scaled angular similarity measure). The scalar dissimilarity field is displayed in the same spatial view as the modality field. There are four possible values, depending on the number of dissimilar pairs out of the three possible (besides the 0 value occurring when all three members are similar). Adding a fourth member, dissimilar to the previous three according to the global dissimilarity matrix, reveals that the latter disagrees with the previous three at the majority of the locations that were previously shown as similar (colored dark-blue in Fig. 7 (bottom)). This is indicated by the dissimilarity value at these locations, which shows that three out of the six pairs possible are dissimilar.

The third data set is the temperature gradient ensemble, at a pressure level of 200 hPa, valid on January 12, 2011. Fig. 8 (top) shows the geolocated scalar field containing the modalities, together with the isotherms of the mean temperature field. It can be noticed that uniform distributions (modality zero), indicative of no preferred direction in the greatest rate of change in temperature, occur primarily around the centers of the air masses of the mean field, while multimodal distributions tend to occur across the boundaries of warm and cold air masses. A detail view of such a region located at the boundaries is shown in Fig. 8 (bottom). Its location within the whole domain is marked in brown in the top figure.

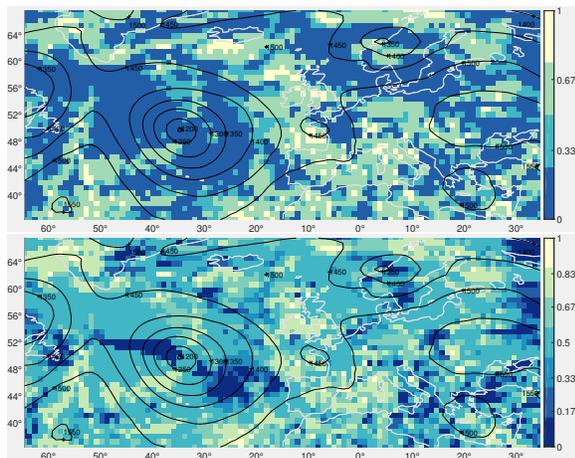


Figure 7: Spatial view highlighting dissimilar locations for a group of three similar members (top). Updated dissimilar field after a fourth member, dissimilar to the first three, is added (bottom).

7.2 Implementation, Performance and Scalability

Our tests were run on a standard desktop PC, equipped with an Intel i7-4790 quad-core processor running at 3.6GHz and with 12GB RAM. All computational operations were performed in Matlab in a preprocessing step, using available Matlab functionality to fit GMMs on the real line and for the normality tests. For each of the two wind ensembles, it took roughly 12 minutes to perform all computations (including the ten repetitions), while the larger temperature gradient ensemble required about 80 minutes. The most time consuming operation was the fitting (most of the processing time), followed by the similarity computations and the global assembly. Other operations, such as computing the clustering or the glyph representations, last no longer than a few seconds for all data sets. Due to the compact representation of the fitted pdfs, the memory requirement of all precomputed quantities is moderate. We are able to handle the analysis session with multiple linked views at interactive frame rates (see the accompanying video).

Regarding scalability, an increase in the grid dimensions reflects only on the spatial view, where the size of the pixel for each grid point decreases (cf. Fig. 8 (top)). The detail view showing the directional distributions is not affected, since the number of shown glyphs was adjusted to fit the displayed view and is constant for all grid dimensions (cf. Fig. 8 (bottom)). The two abstract views are also not influenced by the grid size. A significant increase in the number of ensemble members, however, would affect the readability of the abstract views (but not that of the domain views). Since the data sets used by our collaborators generally comprise approximately 50 members, legibility is not problematic in our case. Ensembles with considerably more members would nevertheless require focus and context techniques, or zooming and panning.

8 CONCLUSION AND EVALUATION

In this paper, we proposed a novel technique that allows an interactive comparative exploration of 2D ensembles of vector fields. To this purpose, we modelled the directional distributions at every grid point using mixtures of Gaussians. Fitting GMMs employs an EM algorithm and is thus prone to converging to local optima. Nevertheless, the used repetition process helped to significantly reduce the percentage of non-repetitive structures. Based on the GMM modelling, we introduced similarity measures to reveal the locations where two or more ensemble members disagree locally (with respect to the behavior of the whole ensemble). These measures were then used to identify the most similar groups of members. We also proposed a novel glyph-based technique that permits an intu-

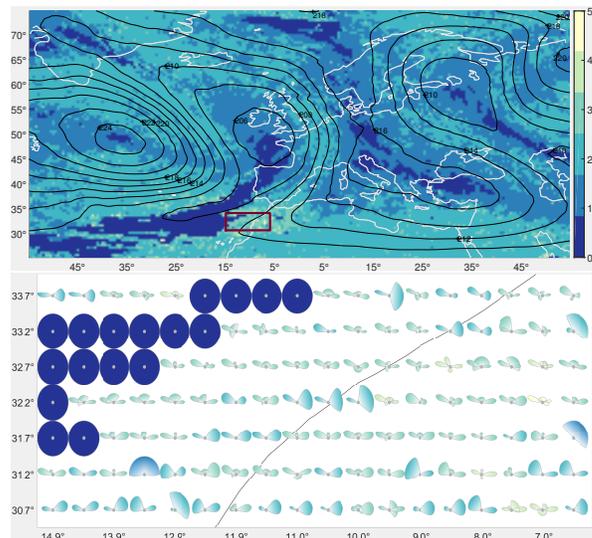


Figure 8: ECMWF temperature gradient ensemble: spatial view with 221×101 modality field (top) and 18×7 detail view (bottom).

itive visualization of distributions of 2D directional data.

In developing our techniques, we collaborated with domain experts from meteorology, from whom we have already received informal feedback. The experts highly appreciated the interaction facilities that allowed them to select a spatial region and investigate the corresponding directional distributions in the detail view. They also liked the color-coding of the modalities, which helped them identify the locations with multimodal distributions. However, they were initially hesitant about the matrix, since they were not familiar with such a representation and thus required some effort to understand it. On the other hand, they found it interesting to quickly grasp which members disagree with most of the other members (as the corresponding columns/rows were mainly yellow). They also found the possibility to select ensemble members in the matrix and dendrogram, and inspect the locations where these members agree/disagree with other members very useful. Moreover, they liked the option to blend in the individual ensemble members, which allowed them to see outliers in addition to the modes. They found that the lobular glyphs nicely convey the modality of the distributions, but preferred the pie glyphs, which resembled other graphical representations used in meteorology.

While this first feedback was positive, we will evaluate the full potential of our method in a more detailed ensemble analysis in the future. To gain an initial insight into how the lobular glyphs compare to the pie glyphs, we conducted a user study with 65 inexperienced participants (54 males and 11 females, aged 18 to 30 years). The participants were given six main tasks (three for the lobular glyphs and another three for the pie glyphs) in a limited amount of time, namely to estimate, for each primitive in every figure, either its direction, its spread, or the ratio between its length and the maximum length in the respective figure. They were also asked to assess the difficulty of the task on a five-level Likert scale. The questions were identical for both lobes and pies, but in a different order. Furthermore, the subtasks of each main task were split on two different pages and delivered in two separate batches, so that we could also record the answers of the participants to the two types of visualizations after they had got to interact with both. Upon a preliminary examination of the results, we noticed that participants were equally capable of solving the tasks with either method, the difference in accuracy (measured via mean and standard deviation) being insignificant. In the perceived difficulty, however, we noticed that the majority of the participants (81%) found the lobes more intuitive when estimating directions, while they preferred the pies for

the other two types of tasks (87% for the spread and 81% for the ratios). In the future, we would like to gain a deeper understanding of the ability of the users to assess the different characteristics of directional data using the two proposed methods.

Moreover, we plan to extend our analysis to 3D directional data. The mathematical analysis in 3D is already available: we model the directional data using a mixture model of von Mises-Fisher components – the equivalent of the Gaussian distribution on the sphere. However, even though the proposed methods for 2D extend straightforwardly to 3D, the 3D representation requires further improvements to minimize clutter and occlusion. Another aspect will be to include time-varying vector ensembles and reveal relevant information on the ensemble dynamics. Possible solutions could include animation, illustrative techniques, or alpha-blending of glyphs from different time steps.

ACKNOWLEDGEMENTS

Access to ECMWF prediction data has been kindly provided in the context of the ECMWF special project “Support Tool for HALO Missions”. We are grateful to the special project members Marc Rautenhaus and Andreas Dörnbrack for providing the ECMWF ENS dataset used in this study. The work was partly funded by the European Union under the ERC Advanced Grant 291372: Safer-Vis - Uncertainty Visualization for Reliable Data Discovery.

REFERENCES

- [1] A. Alhamed, S. Lakshmirarahan, and D. J. Stensrud. Cluster analysis of multimodel ensemble data from samex. *Monthly weather review*, 130(2):226–256, 2002.
- [2] R. Becker and W. Cleveland. Brushing scatterplots. *Technometrics*, 29(2):127–142, 1987.
- [3] M. Beham, W. Herzner, M. E. Gröller, and J. Kehrler. Cupid: Cluster-based exploration of geometry generators with parallel coordinates and radial trees. *IEEE TVCG*, 20(12):1693–1702, 2014.
- [4] G.-P. Bonneau, H.-C. Hege, C. Johnson, M. Oliveira, K. Potter, P. Rheingans, and T. Schultz. Overview and state-of-the-art of uncertainty visualization. In *Scientific Visualization*, pages 3–27. Springer, 2014.
- [5] U. Bordoloi, D. Kao, and H.-W. Shen. Visualization techniques for spatial probability density function data. *Data Sci. J.*, 3:153–162, 2004.
- [6] R. Borgo, J. Kehrler, D. H. S. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen. Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics State of the Art Reports*, pages 39–63, 2013.
- [7] S. Bruckner and T. Möller. Isosurface similarity maps. *Comput. Graph. Forum*, 29(3):773–782, 2010.
- [8] S. Bruckner and T. Möller. Result-driven exploration of simulation parameter spaces for visual effects design. *IEEE TVCG*, 16(6):1468–1476, 2010.
- [9] C. Correa, Y.-H. Chan, and K.-L. Ma. A framework for uncertainty-aware visual analytics. In *Proc. IEEE VAST*, pages 51–58, 2009.
- [10] N. Fisher. *Statistical analysis of circular data*. Cambridge Univ. Press, 1995.
- [11] G. Hamerly and C. Elkan. Learning the k in k-means. In *Neural Information Processing Systems*. MIT Press, 2003.
- [12] N. Henry and J.-D. Fekete. MatrixExplorer: A dual-representation system to explore social networks. *IEEE TVCG*, 12(5):677–684, 2006.
- [13] B. E. Hollister and A. Pang. Interpolation of non-gaussian probability distributions for ensemble visualization. Technical report, 2013.
- [14] M. Hummel, H. Obermaier, C. Garth, and K. Joy. Comparative visual analysis of Lagrangian transport in CFD ensembles. *IEEE TVCG*, 19(12):2743–2752, 2013.
- [15] A. K. Jain. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651 – 666, 2010.
- [16] A. K. Jain and R. C. Dubes. *Algorithms for Clustering Data*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1988.
- [17] F. Jiao, J. Phillips, Y. Gur, and C. Johnson. Uncertainty visualization in HARDI based on ensembles of ODFs. In *Proc. PacificVis*, pages 193–200, 2012.
- [18] D. Kao, M. Kramer, A. Love, J. Dungan, and A. Pang. Visualizing distributions from multi-return lidar data to understand forest structure. *Cartographic J.*, 42(1):35–47, 2005.
- [19] D. Kao, A. Luo, J. L. Dungan, and A. Pang. Visualizing spatially varying distribution data. In *Proc. Int’l. Conf. Information Visualization (IV)*, pages 219–225, 2002.
- [20] J. Kehrler and H. Hauser. Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE TVCG*, 19(3):495–513, 2013.
- [21] M. Leutbecher and T. Palmer. Ensemble forecasting. *Journal of Computational Physics*, 227(7):3515–3539, 2008.
- [22] S. Liu, J. Levine, P. Bremer, and V. Pascucci. Gaussian mixture model based volume visualization. In *Proc. Lдав*, pages 73–77, Oct 2012.
- [23] A. Love, A. Pang, and D. Kao. Visualizing spatial multivalued data. *IEEE Comput. Graph. Appl.*, 25(3):69–79, 2005.
- [24] K. Mardia and P. Jupp. *Directional statistics*. Wiley & Sons, 2000.
- [25] G. McLachlan and D. Peel. *Finite mixture models*. Wiley & Sons, 2004.
- [26] T. Pfaffelmoser, M. Mihai, and R. Westermann. Visualizing the variability of gradients in uncertain 2D scalar fields. *IEEE TVCG*, 19(11):1948–1961, Nov 2013.
- [27] K. Pöthkow and H.-C. Hege. Nonparametric models for uncertainty visualization. *Comput. Graph. Forum*, 32(3):131–140, 2013.
- [28] K. Potter, M. Kirby, D. Xiu, and C. Johnson. Interactive visualization of probability and cumulative density functions. *Int’l. J. Uncertainty Quantification*, 2(4):397–412, 2012.
- [29] K. Potter, J. Kniss, R. Riesenfeld, and C. Johnson. Visualizing summary statistics and uncertainty. *Comput. Graph. Forum*, 29(3):823–832, 2010.
- [30] K. Potter, J. Krüger, and C. Johnson. Towards the visualization of multi-dimensional stochastic distribution data. In *Proc. Int’l. Conf. Comput. Graph. and Vis. (IADIS)*, 2008.
- [31] K. Potter, P. Rosen, and C. Johnson. From quantification to visualization: A taxonomy of uncertainty visualization approaches. In *Uncertainty Quantification in Scientific Computing*, pages 226–249. 2012.
- [32] S. Ray. *Distance-based model-selection with application to the analysis of gene expression data*. PhD thesis, PSU, 2003.
- [33] A. Reh, C. Gusenbauer, J. Kastner, M. E. Gröller, and C. Heinzl. Mobjects: A novel method for visualization and interactive exploration of defects in industrial XCT data. *IEEE TVCG*, 19(12):2906–2915, 2013.
- [34] T. Ropinski, S. Oeltze, and B. Preim. Survey of glyph-based visualization techniques for spatial multivariate medical data. *Comput. & Graph.*, 35(2):392–401, 2011.
- [35] T. Schultz and G. Kindlmann. A maximum enhancing higher-order tensor glyph. *Comput. Graph. Forum*, 29(3):1143–1152, 2010.
- [36] T. Schultz and G. Kindlmann. Superquadric glyphs for symmetric second-order tensors. *IEEE TVCG*, 16(6):1595–1604, 2010.
- [37] T. Schultz, L. Schlaffke, B. Schölkopf, and T. Schmidt-Wilcke. Hi-FIVE: A hilbert space embedding of fiber variability estimates for uncertainty modeling and visualization. *Comput. Graph. Forum*, 32(3):121–130, 2013.
- [38] M. Sedlmair, C. Heinzl, S. Bruckner, H. Piringer, and T. Möller. Visual parameter space analysis: A conceptual framework. *IEEE TVCG*, 20(12):2161–2170, 2014.
- [39] J. Talbot, B. Lee, A. Kapoor, and D. S. Tan. EnsembleMatrix: Interactive visualization to support machine learning with multiple classifiers. In *Human Factors in Computing Systems (CHI)*, pages 1283–1292, 2009.
- [40] D. Thompson, J. Levine, J. Bennett, P.-T. Bremer, A. Gyulassy, V. Pascucci, and P. Pebay. Analysis of large-scale scalar data using hixels. In *Proc. Lдав*, pages 23–30, 2011.
- [41] R. van Pelt, R. Gasteiger, K. Lawonn, M. Meuschke, and B. Preim. Comparative blood flow visualization for cerebral aneurysm treatment assessment. *Comput. Graph. Forum*, 33(3):131–140, 2014.
- [42] C. Ware. *Information visualization: Perception for design*. Elsevier, 3rd edition, 2013.
- [43] C. Wittenbrink, A. Pang, and S. Lodha. Glyphs for visualizing uncertainty in vector fields. *IEEE TVCG*, 2(3):266–279, 1996.