Title
Interactive Visual Exploration and Analysis

Authors
Weber, Gunther H.
Hauser, Helwig

Publication Date
2014-10-21

Interactive Visual Exploration and Analysis

Gunther H. Weber[1,2] and Helwig Hauser[3]

Abstract Interactive exploration and analysis of multi-field data utilizes a tight feedback loop of computation/visualization and user interaction to facilitate knowledge discovery in complex datasets. It does so by providing both overview visualizations, as well as support for focusing on features utilizing iterative drill-down operations. When exploring multi-field data, interactive exploration and analysis relies on a combination of the following concepts: (i) *physical views* that show information in the context of the spatiotemporal domain (domain perspective), (ii) *range views* show relationships between multiple fields (range perspective), and (iii) selecting/marking data subsets in one view (e.g., regions in a physical view) leading to a consistent highlighting of this subset in all other views (brushing and linking). Based on these principles, interactive exploration and analysis supports building complex feature definitions, e.g., using Boolean operations to combine multiple selections. Utilizing derived fields, statistical methods, etc., adds a further layer of flexibility to this approach. Using these concepts, it is also possible to integrate feature detection methods from the other chapters of this part, as well as application-specific feature extraction methods into an joint framework. This methodology of interactive visual data exploration and analysis has proven its potential in a larger number of successful applications. It has been implemented in a larger number of systems and is already available for a wide spectrum of different application domains.

[1]Computational Research Division, Lawrence Berkeley National Laboratory, Email: GHWeber@lbl.gov
[2]Department of Computer Science, University of California, Davis, Email: ghweber@ucdavis.edu
[3]Department of Informatics, University of Bergen, Email: Helwig.Hauser@UiB.no

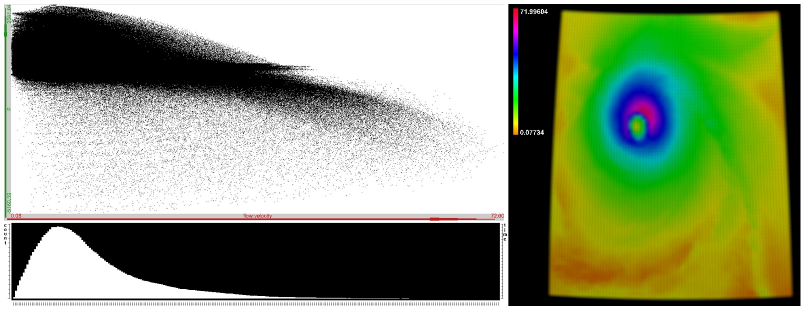

Fig. 1 Interactive Visual Analysis (IVA) uses two types of views: *Spatiotemporal views* (like the image on the right), e.g., false color plots, show the distribution of a quantity within the domain. Range views (like the images on the left), including scatter plots (top left) and histograms (bottom left), show the correlation between multiple fields or additional information about a single field, respectively.

1 Basic Concepts

At its basis, interactive visual analysis (IVA) builds on the combination of different views on data with the ability to emphasize data subsets interactively (most commonly features of interest). In the context of multi-field data exploration and analysis, two aspects of data are of primary interest: (i) the spatiotemporal distribution of one or more fields, and (ii) the relationship of one or multiple fields with respect to each other. For example, examining multiple fields in a simulation of a hurricane, one may be interested in the spatial location of regions of high velocity, but also in learning how velocity correlates with pressure. To provide this information, IVA utilizes two types of views displaying complementary information. (i) *Spatiotemporal views*, such as false color plots or volume rendered images provide a domain-centric perspective on the data. For example, in the hurricane example we can map velocity to color and display a false color plot that shows the spatial distribution of velocity in the simulation domain (Fig. 1 (right)). (ii) *Range views*, such as scatter plots [3] or parallel coordinate plots [4, 13], show the correlation between two or more fields and show the data from a range perspective. For example, for the hurricane example, a scatter plot of pressure and velocity shows their correlation (Fig. 1 (top left)). Individually, the use of these types of views has a long history in science and statistics. Considering only one aspect at a time limits data analysis capabilities. The fundamental idea underlying IVA is to combine different views on the same data in such a way that a user can correlate the different views. One way to achieve this correlation is to enable the interactive selection of data subsets, and highlight such a data subset in other views in a consistent manner, i.e., ensuring the same data items are visually emphasized over their context in all views. Selection is often performed directly on a view by interactive visual means, similar to those in a drawing program, and therefore are usually called *brushing*. Highlighting in this

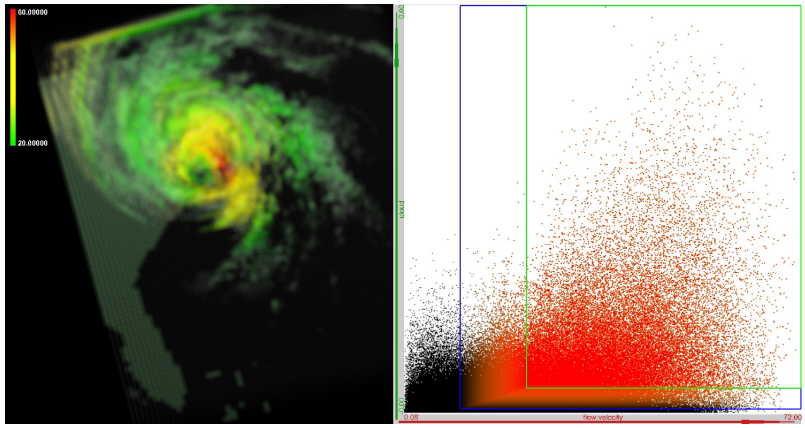

Fig. 2 *Brushing-and-linking* correlates multiple views by making it possible to select a data subset in one view, e.g., selected ranges in attribute/field space, and consistently highlighting this subset in all other views. Changing color (e.g., showing the selected points in a different color (left) or using saturation (right)) is a common way to achieve this effect of highlighting the emphasized data subset over its context (commonly referred to as focus plus context (F+C).

instance serves as means of *linking* views together, and this technique is referred to as *brushing-and-linking*[1] [1, 19].

For example, in the hurricane case, one might be interested in spatial regions corresponding to fast moving clouds. Using brushing-and-linking it is possible to select such a feature in the scatter plot (Fig. 2 (right)) and highlight the corresponding regions in other views, such as in a physical view of the hurricane (Fig. 2 (left)). Using brushing-and-linking, it is possible to formulate simple queries interactively, such as “where are regions of high temperature and low velocity” and visualize the results in a physical view. While there are many instances, where such features of interest are known a-priory and analysis is driven by known queries [29], the full power of IVA lies in the fact that a user can discover features of interest during interaction and pose or refine queries during interactive analysis.

As a consequence, IVA often defines “features” as data subsets of interest to the user, be it due to prior knowledge or because a data subset has caught the user’s attention. Common user interactions include brushing for outliers (e.g., to determine why a subset is behaving differently from the rest), regions of strong correlation (e.g., to verify if this correlation holds in the entire data set or only in particular subsets), or a spatial region of interest, like an inlet or outlet of a flow simulation (e.g., to determine if a correlation exists in that region). In general, we distinguish three patterns of explorative/analytical procedures:

[1] We note that the visualization community uses the order “linking-and-brushing” more commonly, while the database community uses the order “brushing-and-linking”. We use the term “brushing-and-linking” here as brushing is usually performed before linking.

1. From the domain to the range perspective, we select a subset of data items in a physical view and examine the selection in range views. This type of analysis serves to *localize the investigation* to a region of interest such as an inlet or outlet.
2. From the range to the domain perspective, we select a subset of data items in a range view, such as in a scatter plot, and examine the result in a physical view. This type of analysis enables *localizing features*. In this case brushing defines a feature, usually as a set of thresholds, and highlighting in a spatiotemporal view shows whether the selection corresponds to a localized feature.
3. Within the range perspective, we select subset of data items in a range view and observe the selection in another range view. This type of analysis provides a means of performing an interactive *multivariate analysis*, e.g., by brushing in one scatter plot and examining the selection in another scatter plot of different variables. This pattern was originally introduced in the field of information visualization [1, 31].

Using one or more of these patterns is the simplest form of IVA, more recently referred to as "Show & Brush." It utilizes multiple views, usually at least one range view for visually correlating multiple fields and one domain view to show properties in a physical domain context. Though being the simplest form of IVA, this method already covers a large percentage of use cases in multi-field analysis and serves as powerful basis for more advanced types of exploration and analysis. This type of IVA has proven valuable in many application areas, including aeronautical design [12], climate research [15, 20], biomedical visualization [8, 25], the analysis of gene expression data [32], the analysis of combustion engines [7, 22], and the analysis of simulations of particle accelerators [27].

2 Additional Concepts

Based on the simple "Show & Brush" paradigm, a few extensions can greatly enhance the expressiveness of IVA. First, in many cases it is useful to define brushes not as binary classifiers into two categories "of interest" and "not of interest" but as a means to map each data item to a *degree of interest* [6]. It is possible to define this degree by specifying two selections (e.g., regions in a scatter plot). All items inside an *inner range* have a degree of interest of 100% (i.e., are definitely of interest), and all items outside an *outer range* have a degree of interest of 0% (i.e., not of interest). Between those regions, a transfer function maps the distance of a sample from inner and outer range to a degree of interest between 100% and 0%. A linear ramp is a common choice for this transfer function. More generally, we can utilize fuzzy logic operators to combine multiple smooth brushes.

This smooth drop-off of a degree of interest makes it possible to transition seamlessly between data items of interest and those of not interest and use generalized focus plus context (F+C) methods [5, 26, 9] to reduce cluttering in resulting visualizations and draw a user's attention to the most important details. Traditionally, focus and context methods use space distortion such as a fish-eye lens to assign more space to data of interest while presenting the remainder as context for orientation. However, in a generalized setting, various visual attributes can serve to emphasize or deemphasize data items, including, for example:

- Color (hue, saturation, brightness, or an alternative representation) and opacity: a typical example would be to present the data subset in focus in color and its context in gray scale by mapping the degree of interest to saturation [8]. Alternatively, the degree of interest can be mapped to opacity, rendering the focus opaque and the context semi-transparent [21].
- Style: different visualization modalities (isosurfaces, volume rendering, etc.) can be used to discriminate focus and context. Alternatively, rendering styles, in particular non-photorealistic/illustrative styles (halos, outlines, cross-hatched/dotted lines/polygonal primitives) can serve this purpose (for example in a *two-level volume rendering* approach [11]).
- Frequency: Only use the full spectrum of spatial frequencies for the data subsets in focus and render the context band-limited. This approach is called *Semantic Depth of Field* [18] and results in a blurred style for the context, directing the user's attention to the sharply rendered data subsets in focus.
- Space: This approach refers to the traditional notion of F+C visualization, i.e., that the visualization space is distorted in order to give more space (or time) to the visualization of data subsets in focus.

3 Levels of IVA

So far, with "Show & Brush", we have seen the base level of IVA. Based on the complexity of feature definitions, we distinguish additional, more complex (and thereby also more powerful) levels of IVA. It is our experience, however, that in many cases—if not in most cases—the simple Show & Brush technique already provides sufficient functionality to enable an effective data analysis; the more complex levels of IVA, as introduced below, are only advanced solutions for more complicated cases which cannot be served with the base-level IVA.

1. **Show & Brush (level 1)**: This level captures the analysis as described so far. It utilizes at least two linked views, usually one physical and one range view. The interactive selection of features of interest is accomplished by brushing in one view, leading to a focus plus context visualization in the linked view(s).
2. **Relational analysis (level 2):** This level supports the combination of brushes using logical opearations and a simple feature definitions language.
3. **Complex analysis (level 3):** This level integrates computational analysis, e.g., derived fields, statistical methods, machine learning [28], etc., into the interactive visual approach, thus adding a new dimension of possible procedures. A typical scenario would be that, prompted by insights gained during visual exploration and analysis, the user decides to initiate a certain computational analysis procedure, such as clustering. This procedure results in at least one additional (syn-

thetic) data attribute, such as membership in a cluster, which can subsequently be used together with all other data to improve the analysis.

4. **Proprietary analysis (level 4):** This class is a container for everything beyond complex analysis and includes, e.g., the integration of application-specific feature definitions (such as flow feature detectors [2]) or could entail the integration of higher-level feature definition languages. Identifying common concepts and refining IVA beyond this level is a subject of future research.

We note that this terminology is potentially controversial and that "relational analysis" and "complex analysis" have other possible meanings. Consequently, we present this classification as a starting point that can evolve as research in IVA progresses. In the following, we describe the higher levels of IVA in greater detail.

4 Relational Analysis

Relational analysis takes the selections in form of brushes and provides means to combine these brushes (or selections) into more complex feature definitions. A simple feature definition language uses Boolean expressions, for example, to combine brushes into more complex feature definitions. Fig. 3 shows an example from the analysis of three-dimensional gene expression data. Here, positions correspond to the locations of cells in an organism, and the multiple fields represent expression values of genes, i.e., they specify whether a certain gene is expressed in a given cell. Individual brushes select expression patterns based on single genes. Combining these brushes using Boolean operations, it is possible to define complex selections. The example in the figure uses this capability, to combine patterns based on a-priori knowledge about how genes interact, and verify whether the genes involved completely explain the arising pattern.

It is possible to generalize logical operations to smooth brushes [6, 5] and enable F+C visualization in relational analysis. One associated challenge is to extend the visual means, which discriminate data subsets in focus from their context, in such a way that takes this more complex form of feature definition into account. Within each view, an appropriate F+C visualization is necessary to reflect the brush(es) applied to this view. Another level of F+C visualization must reflect the overall feature specification, possibly also involving multiple features. One possible solution to this problem is a four-level F+C visualization approach proposed by Muigg et al. [23], which, as one particular aspect, is based on an intelligent color combination scheme.

Combining brushes usually defines a relation between multiple fields. Early work on query-driven visualization (QDV) [29] used similar concepts in that it defined features as a Boolean combination of relational expressions. However, in this QDV work, the features and expressions were known a priori and not refined during analysis. An important aspect of QDV visualization is the use of indices, such as FastBit [33], to accelerate data selection based on queries. However, there is also work on combining QDV concepts with IVA, e.g., using parallel coordinates [27].

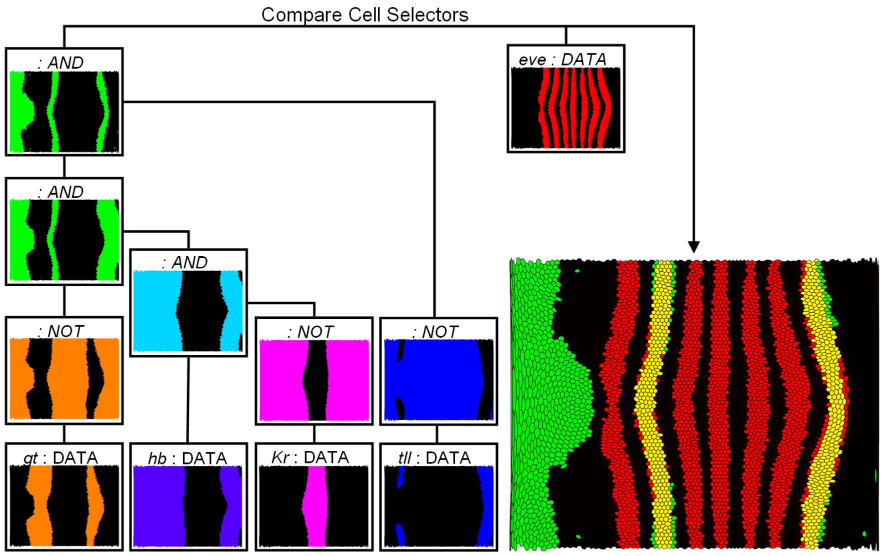

Fig. 3 Building complex feature definitions from individual brushes using Boolean operations in an example from three-dimensional gene expression. Genes are expressed in spatial patterns that control specialization of cells into different tissue types. More complex patterns, such as the seven stripes of the gene *even skipped (eve)* (red pattern in the image), arise from simpler expression patterns when expression of one gene controls (enhances or suppresses) the expression of other genes. The image shows the use of brushes to verify known relations that create *eve* stripes two and seven. The expression patterns of the genes *giant (gt)*, *hunchback (hb)*, *Krüppel (Kr)*, and *tailless (tll)* are first classified by defining an independent brushes in scatter plots. Subsequently, the brushes defining the *gt*, *Kr*, and *tll* patterns are inverted using a NOT operation (to model suppression of gene expression). Afterwards these brushes as well the brush defining the *hb* pattern are combined using a sequence of AND operations. In this way the overlap of the hb expression pattern, and the inverted *gt*, *Kr*, and *tll* expression patterns can be determined. The result (green) is compared to the *eve* expression pattern (red) identified by another brush.

5 Complex Analysis

The levels of IVA described so far are an extremely versatile and powerful framework for enabling effective and efficient visual data analysis. Certain aspects of complex datasets, however, cannot be captured with these mechanisms. In such situations, the integration of computational data analysis tools, like those known from statistics, data mining, or machine learning, can help, leading to a solution which is tightly aligned with the currently modern *visual analytics* methodology [30, 17]). Alternatively, the implementation of extended interaction mechanisms, such as brushes that are capable of grasping aspects of the data that are not explicitly represented in a visualization, can also help in these situations. In the following, we exemplify both approaches to achieve *complex analysis* in the context of IVA.

A very powerful extension of the IVA methodology as described up to here is adding the capability to interactively derive new, user-defined data attributes, based on computational data analysis procedures. While in principle there are no limits to the set of potentially useful data derivation mechanisms, it is the author's opinion that it is worthwhile to emphasize a few more general examples:

Interactive Spatiotemporal Data Derivation: Interactive estimation of gradients with respect to the usually spatiotemporal domain is a generally useful data derivation mechanisms. Spatial and temporal derivatives, including higher order derivatives obtained by repeated application of an interactive derivation operator, are often useful for defining features definition since features are often based on some notion of change. Using temporal derivatives, for example, supports a more advanced analysis of time-dependent aspects of such datasets, where the consideration of first- and second-order derivatives (wrt. time) leads to a massively parallel data analysis similar to how curve sketching is performed for individual time series.

Interactive and Targeted Data Normalization: Data analysis commonly adopts two types of perspective: an *absolute* perspective that considers absolute data values (or derived attribute values), and a *relative perspective* that examines relative values. One mechanism that enables a relative perspective in IVA is to support interactive data normalization. A powerful aspect of performing this normalization as part of IVA is that it not only allows for global normalization procedures, which usually do not add too much in terms of opportunities to understand data aspects that otherwise would not be accessible, but to also enables more localized normalization operations. Examples are normalization per time step, normalization per height-level, etc. Useful normalization operators include the scaling to the unit interval, z-standardization, or the normalization against other data statistics like the median and the MAD.

Interactive Derivation of Data Statistics: Statistics are powerful means to summarize and characterize data. Having data statistics, in particular localized data statistics, available for subsequent computations and interactive feature specifications, enriches the spectrum of possibilities in IVA substantially. A very good starting point are the standard descriptive statistics *mean*, *standard deviation*, *skewness*, and *kurtosis*. Interesting complements include more robust estimates such as the *median*, *MAD*, etc., as well as ranking-based statistics (e.g., based on quartiles or octiles). Interesting applications for IVA have been demonstrated, for example, in the context of multi-run data analysis for climatology [14].

Considering correlation information, data clustering, etc.: Data analysis techniques from statistics, data mining, machine learning, etc., are very rich in terms of history and available related work, and the potential set of useful mechanisms that are promising candidates for integration into IVA is almost unlimited. Particularly interesting candidates for extending the power of IVA are: the interactive derivation of *correlation* information between data attributes (e.g., based on the standard Pearson correlation, or Spearman's correlation measure), techniques for *attribute selection* or *dimension reduction* (such as PCA or LDA, for example), the consideration of data *clustering* (e.g., based on supervised or unsupervised clustering techniques), the integration of measures of *outlyingness* (e.g., based on the Mahalanobis distance of the data points, or derived from normality tests such as Shapiro's p-test), etc. One example for combining IVA with clusterings is the analysis of three-dimensional gene expression data with integrated clustering [28].

Advanced brushing mechanisms can be integrated in IVA as an alternative or in addition to these data derivation approaches. Brushes developed for special purposes include *angular brushing* [10] of parallel coordinates to access the slopes of the lines, or *similarity brushing* [24, 23], which utilizes a more advanced similarity measure between data and brush to determine the data items that are selected by a certain brushing interaction.

In principle, it is possible to design advanced brushes for any of the data aspects that otherwise could be made accessible (to standard brushing) via the further above described data derivation mechanism. The more indirections, however, in terms of implicitly considered data derivations, are built into an advanced brush, the more challenging the additional cognitive load becomes when using such a brush. It therefore stands to reason that highly complicated relations in the data, which only can be accessed through a number of concepts as described above (some statistics, some dimension reduction, some outlyingness measure, etc.), are better made available to interactive feature specification in a step-by-step procedure (a certain sequence of data derivation steps, for example) than packing too much into a single advanced brushing tool.

Fig. 4 shows an example of a Complex Analysis—in this case an outlier analysis in a multi-run climate simulation dataset. As part of a coupled atmosphere–ocean–biosphere simulation model, temperature values in the world's big oceans, represented by three 2D cross-sections (longitude vs. depth), are analyzed, which are given over a 500 year period at about 6000 BC. The goal of this analysis was to identify spatiotemporal locations where the simulated temperature values exhibit large differences (as compared to the main trend) in some simulation runs. Using the interactive data derivation mechanism, first the overall number of outliers per space-time location was computed (this step uses a mild univariate outlyingness measure, i.e., all values which lie more than $3 \cdot \mathrm{IQR}/2$ above q_3 (the 3$^{\mathrm{rd}}$ quartile) or below q_1 (with IQR being the interquartile range $q_3 - q_1$). The scatter plot in Fig. 4 (a) identifies all locations according to how many such outliers exist (x-axis) and to which degree they are large- or small-value outliers (y-axis). A smooth brush was then used to highlight all locations with a substantial number of outliers, and the glyph-based visualization in Fig. 4 (b) shows these locations emphasized (larger, less transparent glyphs). In a next step, the analysis was confined to lower-value outliers. This restriction was achieved by first using the data derivation mechanism, again, to "normalize" the y-axis wrt. its vertical extent per x-location. This step enables a selection—with a standard rectangular brush—of those outliers, which are mainly lower-value outliers. The scatter plot after loading this new attribute and the according brush are illustrated in Fig. 4 (c).

Up to this point, the entire analysis was solely focused on delimiting locations that have outliers of a particular characteristic. In the next step, the focus was di-

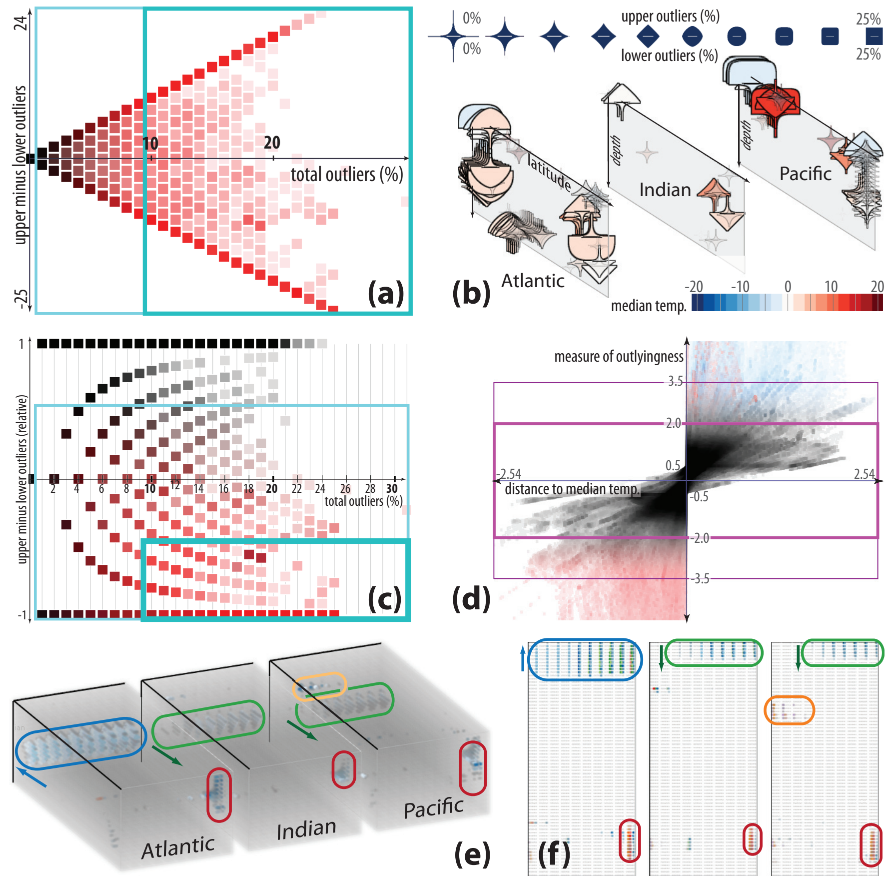

Fig. 4 Selected steps within a Complex Analysis example (outlier analysis in a multi-run climate simulation dataset—more details in the main text). After having used the interactive data derivation mechanism to compute an IQR-based outlyingness measure, spatiotemporal locations are identified with outliers. This identification is achieved by brushing scatter plot (a) and observing the selected locations in the linked visualization (b). In a next step, the analysis was confined to lower-value outliers (using the data derivation mechanism, again) by brushing scatterplot (c). Subsequently, to see the actual outliers themselves, a new scatter plot was used, with detrended and accordingly normalized temperature values on y, to focus on the actual outliers, then observed in views (e) and (f). More details about this study are available in the main text and in a paper by Kehrer et al. [16].

rected to the outliers themselves. To select them, another data derivation steps was performed, computing detrended and normalized temperature values per location (the performed operation was to first subtract the median temperature wrt. all simulation runs, per location, and then divide by IQR). A new scatter plot, shown in Fig. 4 (d), was used to show all data points wrt. their distance to the median (x-axis) and this detrended and normalized temperature measure (y-axis). Consistent with Fig. 4 (a) and (b), all points with y-values beyond ± 2 are also considered as outliers (and brushed accordingly). This brushing leads to their identification in the views Fig. 4 (e) and (f), where each ocean section is repeated 100 times (once for every computed simulation run). This analysis resulted in an interesting deep-water pattern of some ”outliers” in the north of the simulation, translating from the Atlantic slice into the Arctic basin (which actually look much more like a distinct pattern than just outliers) as well as some surface-water outliers (warm water, half-way north in the Pacific, marked orange) and some other outliers near Antarctica (circled red). More details about this study have been presented by Kehrer et al. [16].

6 Conclusions and Future Directions

IVA has already proven valuable in a wide range of application areas, including engineering, climate research, biomedical research and economy. The ability to define features interactively and refine feature definitions based on insights gained during visual and exploration and analysis provides an extremely powerful and versatile tool for knowledge discovery. Future challenges lie in the integration of alternate feature detection methods and their utilization in intelligent brushes. Furthermore, integrating IVA and simulations, thus supporting computational steering, offers a wide range of new possibility.

Acknowledgements

This work was supported by the Director, Office of Advanced Scientific Computing Research, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. We thank the members of the LBNL visualization group and the Berkeley *Drosophila* Transcription Network. Special thanks are extended to Helmut Doleisch and colleagues from SimVis GmbH in Vienna, Austria, for the cooperation over many years. We thank Johannes Kehrer and Peter Filzmoser as well as Thomas Nocke, Michael Flechsig, and colleagues from the Potsdam Institute for Climate Impact Research in Germany for the collaboration on the climate data analysis. We also thank Krešimir Matković and colleagues from the VRVis Research Center in Vienna, Austria, for many years of fruitful colaboration on IVA research, as well as many others from Vienna, Bratislava, Magdeburg, Zürich, and Bergen.

References

1. Becker, R.A., Cleveland, W.S.: Brushing scatterplots. Technometrics **29**(2) (1987)
2. Buerger, R., Muigg, P., Ilčík, M., Doleisch, H., Hauser, H.: Integrating local feature detectors in the interactive visual analysis of flow simulation data. pp. 171–178. Eurographics Association, Norrköping, Sweden (2007). DOI 10.2312/VisSym/EuroVis07/171-178. URL http://www.eg.org/EG/DL/WS/VisSym/EuroVis07/171-178.pdf
3. Chambers, J.M., Cleveland, S., W., Kleiner, B., Tukey, P.A.: Graphical Methods for Data Analysis. Wadsworth/Brooks Cole, Monterey, California (1983)
4. d'Ocagne, M.: Coordonnées Parallèles et Axiales: Méthode de transformation géométrique et procédé nouveau de calcul graphique déduits de la considération des coordonnées parallèlles. Gauthier-Villars, Paris, France (1885)
5. Doleisch, H., Gasser, M., Hauser, H.: Interactive feature specification for focus+context visualization of complex simulation data. In: G.P. Bonneau, S. Hahmann, C.D. Hansen (eds.) Data Visualization 2003 (Proceedings of the Eurographics/IEEETCVG Symposim on Visualization) (2003)
6. Doleisch, H., Hauser, H.: Smooth brushing for focus+context visualization of simulation data in 3d. In: Proceedings WSCG'02, pp. 147–154 (2002)
7. Doleisch, H., Mayer, M., Gasser, M., Wanker, R., Hauser, H.: Case study: Visual analysis of complex, time-dependent simulation results of a diesel exhaust system. pp. 91–96. Eurographics Association, Konstanz, Germany (2004). URL http://www.eg.org/EG/DL/WS/VisSym/VisSym04/091-096.pdf
8. Gresh, D.L., Rogowitz, B.E., Winslow, R.L., Scollan, D.F., Yung, C.K.: WEAVE: A system for visually linking 3-d and statistical visualizations, applied to cardiac simulation and measurement data. In: T. Ertl, B. Hamann, A. Varshney (eds.) Proceedings IEEE Visualization 2000, pp. 489–492. IEEE Computer Society Press, Los Alamitos, CA, USA (2000)
9. Hauser, H.: Generalizing Focus+Context Visualization, pp. 305–327. Springer (2006). URL http://www.springerlink.com/index/v2521525q6148447.pdf
10. Hauser, H., Ledermann, F., Doleisch, H.: Angular brushing for extended parallel coordinates. In: Proceedings of the IEEE Symposium on Information Visualization 2002 (InfoVis 2002), pp. 127–130
11. Hauser, H., Mroz, L., Bischi, G.I., Gröller, M.E.: Two-level volume rendering. IEEE Transactions on Visualization and Computer Graphics **7**, 242–252 (2001). DOI http://doi.ieeecomputersociety.org/10.1109/2945.942692
12. Henze, C.: Feature detection in linked derived spaces. In: D. Ebert, H. Rushmeier, H. Hagen (eds.) Proceedings IEEE Visualization '98, pp. 87–94. IEEE Computer Society Press, Los Alamitos, CA, USA (1998)
13. Inselberg, A.: The plane with parallel coordinates. Visual Computer **1**(4), 69–91 (1985)

14. Kehrer, J., Filzmoser, P., Hauser, H.: Brushing moments in interactive visual analysis. Computer Graphics Forum (Proc. of EuroVis) **29**(3), 813–822 (2010)
15. Kehrer, J., Ladstädter, F., Muigg, P., Doleisch, H., Steiner, A., Hauser, H.: Hypothesis generation in climate research with interactive visual data exploration. IEEE Transactions on Visualization and Computer Graphics **14**, 1579–1586 (2008). DOI http://doi.ieeecomputersociety.org/10.1109/TVCG.2008.139
16. Kehrer, J., Muigg, P., Doleisch, H., Hauser, H.: Interactive visual analysis of heterogeneous scientific data across an interface. IEEE Transactions on Visualization and Computer Graphics **17**(7), 934–946 (2011). URL http://dx.doi.org/10.1109/TVCG.2010.111
17. Keim, D., Kohlhammer, J., Ellis, G., Mansmann, F. (eds.): Mastering The Information Age - Solving Problems with Visual Analytics. Eurographics (2010). URL http://diglib.EG.org/
18. Kosara, R., Miksch, S., Hauser, H.: Semantic depth of field. IEEE Symposium on Information Visualization pp. 97–104 (2001). DOI http://doi.ieeecomputersociety.org/10.1109/INFVIS.2001.963286
19. Kosara, R., Sahling, G.N., Hauser, H.: Linking scientific and information visualization with interactive 3D scatterplots. In: Short Communication Papers Proceedings of the 12th International Conference in Central Europe on Computer Graphics, Visualization, and Computer Vision (WSCG), pp. 133–140 (2004)
20. Ladstädter, F., Steiner, A.K., Lackner, B.C., Pirscher, B., Kirchengast, G., Kehrer, J., Hauser, H., Muigg, P., Doleisch, H.: Exploration of climate data using interactive visualization. Journal of Atmospheric and Oceanic Technology **27**(4), 667–679 (2010). URL http://dx.doi.org/10.1175/2009JTECHA1374.1
21. Lokuge, I., Ishizaki, S.: Geospace: An interactive visualization system for exploring complex information spaces. In: Proceedings of the ACM CHI '95 Conference on Human Factors in Computing Systems, pp. 409–414 (1995)
22. Matković, K., Gračanin, D., Jelović, M., Ammer, A., Lež, 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), 1449–1457 (2010). URL http://dx.doi.org/10.1109/TVCG.2010.171
23. Muigg, P., Kehrer, J., Oeltze, S., Piringer, H., Doleisch, H., Preim, B., Hauser, H.: A four-level focus+context approach to interactive visual analysis of temporal features in large scientific data. Computer Graphics Forum (Proc. EuroVis 2008) **27**, 775–782 (2008)
24. Novotný, M., Hauser, H.: Similarity brushing for exploring multidimensional relations. Journal of WSCG **14**(1–3), 105–112 (2006)
25. Oeltze, S., Doleisch, H., Hauser, H., Muigg, P., Preim, B.: Interactive visual analysis of perfusion data. IEEE Transactions on Visualization and Computer Graphics **13**, 1392–1399 (2007). DOI http://doi.ieeecomputersociety.org/10.1109/TVCG.2007.70569
26. Piringer, H., Kosara, R., Hauser, H.: Interactive focus+context visualization with linked 2D/3D scatterplots. In: Proceedings of the Second International Conference on Coordinated & Multiple Views in Exploratory Visualization (CMV'04), pp. 49–60. IEEE Computer Society, Washington, DC, USA (2004). DOI http://dx.doi.org/10.1109/CMV.2004.11
27. Rübel, O., Prabhat, Wu, K., Childs, H., Meredith, J., Geddes, C.G.R., Cormier-Michel, E., Ahern, S., weber, G.H., Messmer, P., Hagen, H., Hamann, B., Bethel, E.W.: High performance multivariate visual data exploration for extemely large data. In: SuperComputing 2008 (SC08) (2008)
28. Rübel, O., Weber, G.H., Huang, M.Y., Bethel, E.W., Biggin, M.D., Fowlkes, C.C., Hendriks, C.L., Keränen, S.V.E., Eisen, M.B., Knowles, D.W., Malik, J., Hagen, H., Hamann, B.: Integrating data clustering and visualization for the analysis of 3d gene expression data. IEEE Transactions on Computational Biology and Bioinformatics **7**(1), 64–79 (2010)
29. Stockinger, K., Shalf, J., Wu, K., Bethel, E.W.: Query-driven visualization of large data sets. In: Proc. IEEE Visualization 2005, pp. 167–174. IEEE Computer Society Press (2005)
30. Thomas, J.J., Cook, K.A. (eds.): Illuminating the Path: The Research and Development Agenda for Visual Analytics (2005). Available at http://nvac.pnl.gov/agenda.stm

31. Wang Baldonado, M.Q., Woodruff, A., Kuchinsky, A.: Guidelines for using multiple views in information visualization. In: AVI '00: Proceedings of the working conference on Advanced visual interfaces, pp. 110–119. ACM Press, New York, NY, USA (2000). DOI http://doi.acm.org/10.1145/345513.345271
32. Weber, G.H., Rübel, O., Huang, M.Y., DePace, A.H., Fowlkes, C.C., Keränen, S.V., Luengo Hendriks, C.L., Hagen, H., Knowles, D.W., Malik, J., Biggin, M.D., Hamann, B.: Visual exploration of three-dimensional gene expression using physical views and linked abstract views. IEEE Transactions on Computational Biology and Bioinformatics **6**(2), 296–309 (2009). Doi: 10.1109/TCBB.2007.70249. LBNL-63776
33. Wu, K., Otoo, E., Shoshani, A.: Optimizing bitmap indices with efficient compression. ACM Transactions on Database Systems **31**, 1–38 (2006)