HHS Public Access

Author manuscript

2018 Int Symp Big Data Vis Immers Analyt (BDVA) (2018). Author manuscript; available in PMC 2019 November 15.

Published in final edited form as:

2018 Int Symp Big Data Vis Immers Analyt (BDVA) (2018). 2018 October ; 2018: . doi:10.1109/BDVA.2018.8534028.

Visual Pattern-Driven Exploration of Big Data

Michael Behrisch,

Harvard University, Cambridge, USA

Tobias Schreck,

Graz University of Technology, Graz, Austria

Robert Krüger,

Harvard University, Cambridge, USA

Nils Gehlenborg,

Harvard Medical School, Cambridge, USA

Fritz Lekschas,

Harvard University, Cambridge, USA

Hanspeter Pfister

Harvard University, Cambridge, USA

Abstract

Pattern extraction algorithms are enabling insights into the ever-growing amount of today's datasets by translating reoccurring data properties into compact representations. Yet, a practical problem arises: With increasing data volumes and complexity also the number of patterns increases, leaving the analyst with a vast result space. Current algorithmic and especially visualization approaches often fail to answer central overview questions essential for a comprehensive understanding of pattern distributions and support, their quality, and relevance to the analysis task. To address these challenges, we contribute a visual analytics pipeline targeted on the pattern-driven exploration of result spaces in a semi-automatic fashion. Specifically, we combine image feature analysis and unsupervised learning to partition the pattern space into interpretable, coherent chunks, which should be given priority in a subsequent in-depth analysis. In our analysis scenarios, no ground-truth is given. Thus, we employ and evaluate novel quality metrics derived from the distance distributions of our image feature vectors and the derived cluster model to guide the feature selection process. We visualize our results interactively, allowing the user to drill down from overview to detail into the pattern space and demonstrate our techniques in two case studies on Earth observation and biomedical genomic data.

Keywords

Pattern Analysis; Pattern-Driven Exploration; Quality Metrics; Visual Analytics; User Guidance

I. INTRODUCTION

Interactive pattern analysis is becoming an increasingly important topic in many research domains ranging from natural sciences with subjects of physics, chemistry, and biology, to social sciences such as economics, public health, and sociology. Its general goal is to extract interpretable, coherent chunks of information (patterns), and guide the user towards patterns which should be given priority in a subsequent in-depth analysis. In genomics, for example, biomedical researchers study pairwise physical interactions between regions on the genomes [14]. They are interested in areas that are frequently in close contact, forming re-occurring structures [29], to understand the functional consequence of the spatial organization of the genome. Sociologists study connection and diffusion patterns in social networks [8] to understand information spreading and human relationships. In the transportation domain movements are analyzed to reveal frequent traffic patterns between cities and suburbs and to optimize the infrastructure accordingly [59].

Thanks to today's sensor, recording, and storage technology, scientists can rely on large datasets that may comprise valuable insights into their object of research. Hence, pattern mining in databases becomes more and more important, but due to growing data sizes and complexity also increasingly challenging. Even when a suitable detection method is applied, patterns may be tough to interpret. Often hundreds or thousands of patterns are found, leaving it unclear which of them encode valuable information for the application domain, and how much support they have in the data.

To cope with these challenges, we contribute a stepwise analysis approach (see Figure 1), allowing the user to maintain overview in large and heterogeneous pattern-spaces. Given a set of found patterns and their visual representation we apply an image-based feature extraction to each pattern and cluster resulting vectors in a hierarchical manner. To evaluate the most suitable feature extractor, we calculate quality metrics on feature vector distribution and clustering results. The clustered pattern space can then be visually explored by interactively browsing the hierarchy, from overview to detail. While we believe the ideas of our approach to be generic, we focus on patterns in scatterplots and adjacency matrices, a common sub-problem in pattern mining.

The remainder of this paper is structured as follows: Section II discusses related work in the fields of feature extraction as well as automated and interactive visual pattern (space) mining. We present the concept of our analysis pipeline in Section III, and subsequently outline our prototypical implementation in Section IV. After that, in Section V, we exemplarily evaluate our concept by instantiating the process in two case studies on Earth observation and biomedical genomic data. We conclude by discussing limitations and future work in Section VI.

II. RELATED WORK

Our approach combines (A) feature extraction to cluster frequently appearing visual patterns and enables (B) pattern space exploration using advanced navigation techniques. To evaluate

feature extraction/ clustering we adapt quality metrics (C). Hence, related work discussed in the following is threefold.

A. Image-Based Feature Extraction

Image-based features allow characterizing data based on its visual representation. Commonly computed features comprise texture, color, and line/edge descriptors, shape, structure, and contour descriptors, interest point descriptors, as well as noise descriptors [42], [51]. As the extracted feature vectors are often similar to what the user visually inspects, they have been utilized to automatically suggest views of potential interest and to guide data exploration [13], [28], e.g., by feature-based aggregation and filtering. Doing so describes the data in a space that is different from the data itself and that instead is based on characteristics of visual patterns [13], [28], [44], [58].

Influential for this field is the work of Tukey [56] who formulates the problem that—as the number of plots to interactively inspect increase—exploratory data analysis becomes difficult and time-consuming. Tukey proposes to find the “interesting” plots automatically and to investigate those first. To that end, Wilkinson et al. [58] present a set of 14 measures for the quantification of distribution of points in scatter plots, called Scagnostics. Each measure describes a different characteristic of the data and helps, for example, to filter views with different Scagnostics measures than the majority. The underlying scatter plots are likely to exhibit informative relations between the two data dimensions.

Similar to these existing approaches we extract image based features. By contrast, we apply the extraction not on visualizations showing the whole data but on snippets containing detected patterns and use the resulting feature vectors as an input to hierarchically cluster the pattern space.

B. Semi-Automated Exploration

As datasets grow in size and complexity, analysts run the risk of overlooking interesting patterns in manual exploration scenarios. To this end, intelligent methods for compressing and filtering data for potential patterns of interest have recently become a research focus in the visual analytics community.

Overview-based approaches aim to generate effective layouts over many candidate data portions, to efficiently spot patterns of interest. Examples include the Value-and-Relation display [60], which lays out pixel-oriented views based on their data similarity and allows for further drill-down. Another similarity-based layout is proposed by Ward and Guo [57], where many time series are represented by small glyphs.

Besides overview approaches, automatic filtering of views for potential structures of interest has been proposed. As mentioned in the preceding section, the Scagnostics approach [58] automatically analyzes structures in scatter plots, which can be used to rank and filter. In case class information is given, scatter plots can be filtered for discriminative views by class consistency measures [50]. Also, projection pursuit approaches, such as initially presented by Friedman and Tukey [21], try to identify interesting 2D subspaces in high-dimensional

data (mostly depicted by scatter plot views). Further heuristic interestingness filters for Scatter- and Parallel Coordinate plots have been discussed [13], [54] and may narrow down the potentially large search space for high-dimensional data. While most approaches focus on global features, Shao et al. [48] propose means for determining frequent local scatter plot characteristics and summarize them in a motif-based dictionary.

Combining overview and filtering techniques, Tatu et al. [55] proposed an approach to analyze subspaces contained in high-dimensional data using hierarchical clustering for exploration on different levels of magnitude. Another approach, called ScagExplorer [12] clusters data based on Scagnostics [58] and allows for further drill-down and filtering. Also relying on clustering, Bach et al. [2] present an approach to aggregate adjacency matrices, so-called piles, indicating topological states in brain connectivity at different levels of magnitude.

Similarly, our approach extracts features from patterns in matrices and clusters them to explore the result space. A main contrast to the work by Bach et al. is that our approach utilizes and evaluates image-based features to characterize and cluster the data and does not rely on a temporal order.

C. Clustering Comparison

The quality of clustering results depends on a dataset's characteristics and distributions but also on multiple consecutive data processing steps such as extracted features, distance measures, and the clustering strategy. In such a pipeline interim results add up and can lead to significant different clusterings. Researchers thus developed several statistical methods to evaluate and compare clustering quality, as summarized by Rodriguez et al. [39]. One widely used metric is the silhouette coefficient [41] comparing intra cluster cohesion to inter cluster separation. Another class of clustering comparison approaches, e.g., often used in the bioinformatics domain, make use of labeled reference data, where cluster memberships are already known [22], or they rely on domain knowledge to judge the quality of the clustering [17].

Our approach considers situations where no ground truth is given. Similar to measures summarized by Rodriguez et al. [39], we apply quality metrics to rank clustering results. However, instead of calculating metrics for one partitioning, we measure and compare the quality of the whole clustering hierarchy by recursively computing coefficients for all levels.

III. VISUAL PATTERN-DRIVEN EXPLORATION CONCEPT

Our goal is to analyze large amounts of high-dimensional data in a pattern-driven fashion. While the collected data may be rich in information, the exploration challenge is to find data decompositions, e.g., subspaces or data intervals, which expose the underlying core dataset characteristics. Finding and understanding these *central patterns* requires not only automated data analysis, but also the analysts' understanding, background knowledge, and experience. Figure 1 depicts our conceptual pipeline for supporting analysts in a *pattern-driven exploration* of large high-dimensional (HD) datasets.

(1) Set of Representations

Large-scale HD data analysis can hardly be accomplished on the raw data objects, and visual representation can help to facilitate a mental model. Hence, we center our analysis pipeline on the basic idea to regard the visual representations as a *proxy* to the data of interest, and base similarity and relevance computation tasks on the visual data representation, instead of the original (raw) data. Our aim in doing so is to provide user-friendly, interpretable, and interactive assessment functions as a basis for search and analysis tasks.

In this work, we do not focus on the question “*Which visual representation is the best for the underlying dataset?*”; but rather start from the hypothesis that a data representation is established in the current analysis domain. In many domains, such as in the biomedical sector where results are frequently represented in heatmap displays [16], [26], [29], [30], or the environmental sciences, which oftentimes rely on scatterplot or line charts [9], [47], [49], analysts are oftentimes trained over years to see and retrieve patterns in a particular visualization type. While this assumption is not generally pertinent for all scientific domains, we claim that our exploration pipeline (Figure 1) is still applicable for the general task of understanding distributions in pattern spaces for a given data representation.

(2) Image Feature Analysis

Searching and analyzing are key tasks for retrieving, relating, and reusing complex data sets. The standard approach to similarity computation extracts feature representations from the raw data (see Figure 2 Top). We propose to extract feature representations from the visual transformation of input data (Figure 2 Bottom). This approach may provide several advantages. First, a visual transformation of data is naturally linked with the user interface: visualization of the data is used in many applications, and it can be intuitively shown why two data representations are considered similar, e.g., by representing corresponding visual features in the visual data representation. A second main advantage is that one can implement visual analysis interfaces that allow the user to explore visual similarity of many data instances. We will develop this idea further in the steps (3) *Clustering* and (4) *Quality Metrics for Feature Ranking*. Finally, the similarity notion can flexibly adapt to user needs: different visual abstractions give rise to different similarity notions and feature extraction approaches and thus a flexible means to adapt the exploration pipeline to the context of the user task.

One of the cornerstones of our pipeline is to employ image-based feature descriptors into the process. Many feature descriptors were proposed to automatically extract the most descriptive feature set from a given image or visual representation, among them novel visual descriptors that try to mimic human perception aspects. Applied in our context, these so called **-gnostics* approaches [4], [13], [28], [44], [58] can help the user to improve the interactive query specification and result interpretation stages of the visual analysis process.

(3) Pattern Space Clustering

Cluster-based navigation systems divide the exploration space into a range of distinctive clusters, such that each grouping corresponds to a meaningful data sub-selection. In the case

of visual pattern-driven exploration, a “good” feature encoding applied on a clustering approach will result in groupings that reflect pattern-cluster memberships.

In this work, the central analysis task is to partition the pattern space into interpretable, coherent chunks. These chunks should expose similar visual patterns and guide the analysts to beneficial exploration paths for a subsequent in-depth analysis. Hierarchical or density-based cluster approaches [40] are consequently the data analysis method of choice since it allows the analysts to perceive (dis-)similarity of patterns, assess the pattern-to-noise ratio, and the pattern-distribution without predetermining a k-fold pattern space partitioning as done by, e.g., k-means clustering [32].

(4) Quality Metrics for Feature Ranking

The extraction of relevant information from high-dimensional data is complex and time consuming. In that respect the notion *curse of dimensionality* represents a whole set of issues encountered in the analysis of these data sets: finding relevant data attributes, selecting meaningful and descriptive dimensions, removing noise are just a few of them.

Researchers have been trying to solve the aforementioned analysis problems through either automatic data analysis or interactive visualization approaches. However, we claim that integrated visual analytics approaches, where a machine searches automatically through a large number of potentially interesting data transformations and mappings based on quality metrics, and the user interactively steers the process and explores the output through visualizations will outperform isolated approaches.

Our work represents a specific example for the aforementioned quality metrics-driven exploration. We rely on a “good” image-feature extraction algorithm to derive an interpretable and useful pattern space clustering. One central assumption of our pipeline is that at least one image feature descriptor is capable to describe visual patterns. Earlier research has validated this hypothesis on relational data [4].

(5) Pattern Space Exploration

As mentioned before, pattern space exploration is a human-centered analysis approach in which the automatic analysis should support critical parameters, i.e., the feature descriptor selection. On the other hand, finding an appropriate partitioning of the pattern space is not enough. These approaches need to involve an analyst who can make sense of the results and give them meaning in their respective analysis domain.

A visual interface for navigating in stratified pattern space clusters is hence necessary. Ideally, such a system provides Overview+Detail functionality, i.e., it allows to perceive pattern distributions, feature descriptor, and cluster uncertainty, and allows an interactive exploration of the pattern space for unexpected findings. We showcase an exemplary interactive exploration prototype in this work as depicted in Figure 4.

IV. PROTOTYPICAL IMPLEMENTATION

We found that a range of implementation challenges need to be tackled to instantiate our analysis pipeline presented in Section III. The findings that we derived during the implementation of our prototype will structure this section and outline potentially beneficial research directions.

A. Quality Metrics Implementation

One general challenge of pattern space exploration is that no ground truth data is available, making standard external evaluation metrics, such as precision or recall, inapplicable. Also, the exploratory nature of our analysis lets the user develop experience-based heuristics earliest after getting intermediate results. In order to bootstrap the analysis for retrieving interpretable—intermediate—results we developed quality metrics to rank our set of 24 feature descriptors for their ability to (1) *differentiate visual patterns* and to (2) *produce (visually) coherent groupings* of patterns.

1) Pattern Discriminability: We approximate a feature descriptor's ability to differentiate visual patterns by calculating statistics on the normalized Euclidean distance scores:

$$\sigma^2(\text{FD}) = \frac{1}{n^2} \sum_{1 \leq i \leq n} \sum_{1 \leq j \leq n} (\text{ndist}(\text{FD}(i), \text{FD}(j)) - \bar{x})^2 \quad (1)$$

where $\text{FD}(i)$ calculates the feature vector of the i^{th} image in the dataset, $\text{ndist}()$ represents a normalized Euclidean distance and $\bar{x}$ corresponds to the average of all distance combinations. This variance calculation is a coarse-grained heuristic that solely allows the following conclusion: If there is a low distance score variance between the data feature vectors, the respective feature descriptor is not able to differentiate the inherent (visual) features; given the assumption that there are actually discriminative visual patterns.

Pattern discriminability for image feature descriptors has been initially studied, for example, in [4] in the context of matrix pattern research. However, more future research could be devoted for developing more fine-grained approximations. As an example, a correlation analysis of the image feature dimensions would be computationally more expensive but might result in potentially better retrieval of dimension subspaces that help in this context.

2) Clustering Structure Quality Metric: Pattern discriminability enables us to reject inappropriate feature descriptors early on, but it does not allow us to assess how well a feature descriptor can partition the feature space. We derive, therefore, a hierarchical clustering based on this feature descriptor and quantify its cluster separation and cohesion. A range of external measures quantifying the “accuracy” were presented in the past, such as the Rand index [36], the Fowlkes-Mallows index [19], and the Jaccard index [25], [34], which are only applicable to labeled datasets. In our case, however, we have to rely on

internal quality measures that base their quality understanding on (dis-)similarities of feature vectors.

While also here a range of quality measures exist, i.e., the Dunn index [15] or the Calinski-Harabasz score [10], we chose silhouette coefficient [24], because it represents an established and intuitive measure for both cohesion and separation of clusters. In our evaluation section (Section V-A) we present a comparative evaluation of the quality computations based on the Dunn index and the Calinski-Harabasz score. For our scenarios we found that the silhouette coefficient is well-suited. Its silhouette score bases on the mean intra-cluster distance and the mean nearest-cluster distance for each item in the dataset. A silhouette coefficient SC for a partition with k clusters is calculated by averaging the k individual silhouette scores:

$$SC = \frac{1}{n} \sum_{1 \leq i \leq n} \frac{b(FD(i)) - a(FD(i))}{\max(a(FD(i)), b(FD(i)))} \quad (2)$$

where $a(FD(i))$ is the average dissimilarity between $FD(i)$ and all other points in the same cluster and $b(FD(i))$ is the average dissimilarity between $FD(i)$ and the data points in the nearest neighbor cluster. Since the number of clusters k is unknown in our analysis, we introduce a cut-level balanced silhouette coefficient with the following formula:

$$\phi(SC) = \sum_{1 \leq k \leq h} \frac{1}{2^{(k-1)}} \times SC_k \quad (3)$$

where h is the height of the clustering dendrogram tree. The score calculates for all hierarchical clustering cut-levels the silhouette coefficients and aggregates them with a weighting score depending on the cut-level depth. The intuition for this calculation is the following: A clustering dendrogram that shows many well-separated, but highly coherent, clusters on the higher levels should be favored over degenerated trees that show coherence, but little separation. Since our score aggregates over all clustering levels and most dendrograms contain long, degenerated sub-trees we weight their influence depending on the cut-level depth.

Our quality score bases on considerations about hierarchical clustering results. While this quality metric could potentially also be used for density-based clustering methods, future research, and experiments should be devoted to proof its generalizability.

3) Compound Quality Score: Both scores above are equally weighted and aggregated to derive a final quality score.

$$QM(FD) = 0.5 \times \sigma^2(FD) + 0.5 \times \phi(SC) \quad (4)$$

In a future prototype, we are planning to let the user decide on the weighting and composition of the factors. With such a flexible understanding of quality, analysts could express—for example—their preference in scenarios where quality is sacrificed for faster computation times.

B. Understanding Clustering Results

In order to be of practical use cluster results need to be represented by prototypes that aggregate the information of its underlying cluster. While many approaches try to give an example-based cluster prototype, such as most medoid entity [35], other approaches try to aggregate the cluster information, as done by Strobelt et al. [53]. For our purpose to understand clustering results on visualizations exposing structured visual patterns, we developed an image-based visual aggregation for visual views, as Figure 3 depicts.

As Figure 3 depicts, we construct a cluster representative from all entities contained in the cluster. We visually overlay each matrix image within one cluster such that their relative opaqueness value sums up to 1 (fully opaque). This opaqueness consideration shows the (un-)certainty of a hierarchical clustering with respect to the cut-level intuitively.

In our future research, we plan to investigate further visual aggregation methods, that focus on exposing the (un-)certainty of a hierarchical clustering. Since not all leaves are on the same height in a hierarchical clustering, one specific improvement will be to incorporate a leaf's path length to the root as a weighting factor.

C. Navigation in Stratified Pattern Spaces.

In this work, we showcase an interactive exploration system for pattern space exploration as depicted in Figure 4 on scatter plots. It follows the previously proposed framework for visual pattern exploration. Its purpose is to give an overview of pattern space clustering (Figure 4 1 and to enable the analysts to drill-down into the hierarchical—or stratified—cluster structure (Figure 4 2) to make sense of their certainty, quality, and cluster's membership composition.

As a proof of concept, we developed a hierarchical clustering visualization for general visualizations as an alternative to the standard dendrogram visualization for hierarchical clustering. In our experiments, we found that employing a radial layout is more space-efficient, thus allowing the user to render larger cluster prototypes. The user can switch between the two layout modes. In the radial layout, the cluster rings' sizes and positions reflect the hierarchy. The cluster's membership size is depicted additionally in the ring's outer appearance, as depicted in Figure 4 3: a pie chart metaphor shows the percentage of members in this cluster with respect to the overall number of data items; Small clusters (small arc angle) will be distinguishable from larger clusters (large arc angle). We render the cluster prototypes into the ring's center, such that all leaf nodes contribute equally to the prototype. As Figure 4 2 depicts, a “drill-down slider” can be used to examine a cluster's membership composition in the form of cut-levels in a dendrogram. This interaction allows retrieving meaningful dendrogram split levels and thus an appropriate number of visual pattern clusters in the dataset.

We found that our implemented interactive panning and zooming operations are helpful but may become tedious. On an abstract level, however, these interactions allow us to derive implicit and explicit interestingness considerations from the visualized pattern space. In a future work, we want to make use of these aspects by incorporating relevance feedback and learning, following the line of research suggested, e.g., in [6].

D. Scalability

Big Data analysis requires high-performance data processing schemes. In our presented approach, we combine image feature analysis, cluster evaluation, and interactive visualization with the goal to guide the user to interesting and interpretable (visual) patterns. Highly parallelized image feature extraction pipelines are already established in the biological domain, such as for cell classification with billions of cells [11], [33]. Cluster computation and analysis is still a computationally demanding, but distributed and parallelized hierarchical clustering approaches are already presented e.g., by Bandyopadhyay and Coyle [3]. Visual scalability was a basic requirement in our research project and is intrinsically manifested by the visual aggregation/ multi-scale overview-first approach, i.e., the user gets an (highly) aggregated overview first and analyzes specific pattern of interest by “zooming” into the pattern space. While our prototype is designed to demonstrate the general feasibility of the pattern exploration pipeline, we found that generating the cluster representatives (many reoccurring linear scans over the cluster leafs during the interactive exploration) and calculating the cluster quality metrics (for every cut-level in the hierarchy) are the most time consuming steps. Consequently, we developed a cluster representative caching for storing and retrieving the reoccurring visual aggregations. Also, the quality metric calculation be used for datasets with more than 10K elements, if we take advantage of the monotonically decreasing weighting score (inverse proportional to the cut-level depth). Simple cut-level thresholds can be used for approximating the cluster quality metrics.

V. EVALUATION

We evaluate our pipeline implementation with two distinct experiments. First, we focus on the parameter choices, outlined in the previous section, for constructing our quality metrics, second, we present two case studies on (a) Earth observation data and (b) on genome interaction data to showcase the applicability of our approach on a real-world dataset.

A. Quality Metric Evaluation

Figure 6 presents the impact of alternative parameter choices for our quality computation. We derive the data from the dataset showcased in Section V-B. For better interpretability, we depict the quality score ranking of our 24 feature descriptors in a scatterplot (Figure 5). The x-axis depicts the silhouette quality score SC (the higher, the better), and the y-axis shows the feature descriptor’s variance (the higher, the better). As one can see, the best performing feature descriptors are COLOR_OPPONENT_HISTOGRAM, its variation COLOR_FUZZY_COLOR_OPPONENT_HISTOGRAM, and the COLOR_AUTO_COLOR_CORRELOGRAM being present in the upper right quadrant.

The COLOR_OPPONENT_HISTOGRAM descriptors use normalized color differences between all primary colors. Although the COLOR_OPPONENT_HISTOGRAM descriptor has not the highest variance scores, indicating feature discriminability, it outperforms the other feature descriptors. Our silhouette quality score shows a gradual ranking improvement without big jumps. On the other hand, we can also see—and validate from other experiments—that the Calinski-Harabasz quality score shows some similarity with the ranking derived from the Silhouette score, while the Dunn index related score (needs to be considered inverse) shows a significant divergence. Generally, we found that the silhouette and Calinski-Harabasz score, often rank the same feature descriptor on the first place.

Note that this quality assessment is only valid and transferable between datasets if the underlying data representations expose similar visual structures. We also found that the feature ranking was entirely different for other datasets in our experiments. The appendix shows further experiments to validate this finding.

B. Case Study: Earth Observation Data

We apply our exploration pipeline in a case study on scientific data from Earth observation research. Our dataset contains a subset of the Baseline Surface Radiation Network (BSRN) repository, maintained by the World Climate Research Programme [43]. The repository hosts data on measurements of water, sediment, ice, and atmosphere, among others.

In our exploratory analysis, our primary goal is to develop an overview of the patterns contained in this dataset. Our quality metric evaluation, described in Section V-A, found that the COLOR_OPPONENT_HISTOGRAM differentiates patterns best. After exploring the first few Dendrogram Cutlevels (DC), we can assume that the descriptor differentiates mostly based on color density in the plot (Figure 4 shows DC4). After a closer inspection of the higher DC levels, however, we can also see that the feature descriptor seems to be able to group similar shapes while being rotation-invariant. Figure 7 shows some examples of the clustered patterns.

Generally, we found that scatter plots often share similar visual patterns if one of the axis dimensions remains static, such as “DIF” (Diffuse radiation) in Figure 7c or “LWD” (Long-wave downward radiation) in Figure 7d. While this finding seems obvious, we could use it as a starting point for finding deviations from this norm. As one example, we found that the pattern variability of scatter plots containing “DIR” (Direct radiation) seem to be worth examining.

C. Case Study: Genome Interaction Matrices

In a second case study, we are studying two collections of local patterns derived from two genome interaction matrices [37]. Genome interaction matrices capture pairwise interactions of up to 3 million regions on the genome and express several nested visual patterns, which act as a proxy to the spatial organization of the genome. Biologists study this spatial organization [14] as it has been shown to influence gene regulation [20], cell development [7], and pathogenic processes [23], [31], [45], [52]. In this case study we are focusing on two pattern types called loops and topologically associating domains (TADs) that ideally

exhibit an off-diagonal pronounced peak or dot and an on-diagonal block respectively. The patterns have been detected previously [37], but their quality is very diverse [18], [27], [29] and generally not measurable due to a lack of ground truth.

The focus of our study is two-fold. First, we want to evaluate the overall quality of the pattern collection. And second, we try to stratify the group of patterns into subgroup showing potentially distinct biological events. With our quality metric evaluation we found that the NOISE_STATISTICALSLIDING-WINDOW feature descriptor differentiates patterns best [1]. The feature descriptor derives for subsequent regions in the image statistical information about color intensities. These sliding window values can be interpreted as a time series of color differences. The final NOISE_STATISTICALSLIDINGWINDOW feature vector describes the time series with respect to its average, variance, and standard deviation. We show the performance results of all other feature descriptors in our quality metric performance evaluation in the appendix. We start by exploring loop patterns globally within the GM12878 dataset from Rao et al. [37]. An ideal loop pattern is shown in Figure 9a. After examining the first few DC of several chromosomes we realize that the diagonal is dominating the signal but starting with DC2 to DC4 we are able to identify clusters showing a pronounced loop pattern on average (Figure 9b highlighting that the NOISE_STATISTICALSLIDINGWINDOW descriptor is able to discriminate between noise and true dot-like pattern types, which we can use for to remove spurious patterns to increase the power of subsequent analyses. The TAD pattern type, which we extracted from the GM12878 and K562 data set of Rao et al. [37], is less uniform than the loop pattern type and often expresses distinct sub patterns as illustrated in Figure 9c. Using the NOISE_DISSIMILARITY feature descriptor [1], we are able to stratify the diverse collection of the extracted TAD patterns and find distinct cluster branches that group the two sub pattern types together. The NOISE_DISSIMILARITY feature descriptor calculates the dissimilarity between a given visualization—in our case snippets of the genome interaction matrix—and multiple random noise images. The feature descriptor follows a simple argumentation: Since the most distinctive feature of a noise image is the absence of structure, visualizations with a large NOISE_DISSIMILARITY are considered to have a higher potential structuring. We can use these identified sub groups to further analyze the differences between TAD patterns, to optimize our initial TAD pattern detector, or build curated sets of TAD sub patterns.

VI. DISCUSSION AND CONCLUSION

In this work, we present conceptually a pipeline for supporting a pattern-driven exploration of Big Data. As its foundation, we rely on an image feature analysis for given data representations and clustering to partition the pattern space into interpretable, coherent sub-clusters of patterns.

Yet, the underlying hypothesis is that at least one feature descriptor exists that is useful to quantify the existence of visual patterns in the dataset. The quantification of patterns in visualizations is an active research field with broadly two distinctive approaches: Either pattern measures are computed from the data space or the image space. Image-based quality

metrics have the advantage that a direct correspondence to the human perceptual system is imminent. Following this argumentation, we can also assess the limits of our approach: an image-based pattern analysis can only work if the pattern space is clearly defined and distinguishable; i.e., the patterns can be discerned computationally and perceptually (human). Moreover, image-based pattern analysis builds on the assumption that the applied visualization technique can express the pot. complex data patterns. For relational and bivariate data, several research papers [5], [46] have considered this question. For other visualization types a structured pattern space examination remains future research. Another limitation results from the choice of our data analysis machinery. Our quality scores rely on internal quality measures derived from feature vector distances. As, for example, shown by Reeb et al. [38], the choice of the dissimilarity calculation has an impact on the result interpretation. More research in this direction needs to be devoted to developing (more) dissimilarity score agnostic approaches.

In general, we found that a pattern-driven analysis is a viable approach to guide the analysis of big datasets towards patterns with high support but also draws attention to the unexpected and outlying structures.

ACKNOWLEDGMENT

This work was supported in part by the National Institutes of Health (U01 CA200059 and R00 HG007583). We also thank Lin Shao for his valuable help in preparing the BSRN dataset.

REFERENCES

- [1]. Albuquerque G, Eisemann M, Lehmann DJ, Theisel H, and Magnor M. Improving the visual analysis of high-dimensional datasets using quality measures In IEEE Conference on Visual Analytics Science and Technology, pages 19–26. IEEE, 2010.
- [2]. Bach B, Henry-Riche N, Dwyer T, Madhyastha T, Fekete J-D, and Grabowski T. Small multiples: Piling time to explore temporal patterns in dynamic networks In Computer Graphics Forum, volume 34, pages 31–40. Wiley Online Library, 2015.
- [3]. Bandyopadhyay S and Coyle EJ. An energy efficient hierarchical clustering algorithm for wireless sensor networks In INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies, volume 3, pages 1713–1723. IEEE, 2003.
- [4]. Behrisch M, Bach B, Hund M, Delz M, von Ruden L, Fekete J-D, and Scheck T. Magnostics: Image-based Search of Interesting Matrix Views for Guided Network Exploration. IEEE Transactions on Visualization and Computer Graphics, 23(1):31–40, 10 2017. [PubMed: 27514053]
- [5]. Behrisch M, Bach B, Riche NH, Schreck T, and Fekete J-D. Matrix Reordering Methods for Table and Network Visualization. Computer Graphics Forum, 35(3):693–716, 6 2016.
- [6]. Behrisch M, Korkmaz F, Shao L, and Schreck T. Feedback-Driven Interactive Exploration of Large Multidimensional Data Supported by Visual Classifier In IEEE Conference on Visual Analytics Science and Technology, pages 43–52. IEEE CS Press, 10 2014.
- [7]. Bonev B, Cohen NM, Szabo Q, Fritsch L, Papadopoulos GL, Lubling Y, Xu X, Lv X, Hugnot J-P, Tanay A, et al. Multiscale 3D genome rewiring during mouse neural development. Cell, 171(3): 557–572, 2017. [PubMed: 29053968]
- [8]. Borgatti SP, Mehra A, Brass DJ, and Labianca G. Network analysis in the social sciences. science, 323(5916):892–895, 2009. [PubMed: 19213908]
- [9]. Bremm S, von Landesberger T, Bernard J, and Schreck T. Assisted Descriptor Selection Based on Visual Comparative Data Analysis. Computer Graphics Forum, 30(3):891–900, 6 2011.

- [10]. Cali ski T and Harabasz J. A dendrite method for cluster analysis. *Communications in Statistics - Theory and Methods*, 3(1):1–27, 1974.
- [11]. Carpenter AE, Jones TR, Lamprecht MR, Clarke C, Kang IH, Friman O, Guertin DA, Chang JH, Lindquist RA, Moffat J, Golland P, and Sabatini DM. Cellprofiler: image analysis software for identifying and quantifying cell phenotypes. *Genome Biology*, 7(10):R100, 10 2006. [PubMed: 17076895]
- [12]. Dang TN and Wilkinson L. Scagexplorer: Exploring scatterplots by their scagnostics In *IEEE Pacific Visualization Symposium (PacificVis)*, pages 73–80. IEEE, 3 2014.
- [13]. Dasgupta A and Kosara R. Pargnostics: Screen-space metrics for parallel coordinates. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1017–1026, 11 2010. [PubMed: 20975139]
- [14]. Dekker J, Belmont AS, Guttman M, Leshyk VO, Lis JT, Lomvardas S, Mirny LA, Oshea CC, Park PJ, Ren B, et al. The 4d nucleome project. *Nature*, 549(7671):219, 2017. [PubMed: 28905911]
- [15]. Dunn OJ and Clark VA. *Applied statistics: analysis of variance and regression*. John Wiley & Sons, Inc., 1986.
- [16]. Eisen MB, Spellman PT, Brown PO, and Botstein D. Cluster analysis and display of genome wide expression patterns. *Proceedings of the National Academy of Sciences*, 95(25):14863–14868, 1998.
- [17]. Ferreira L and Hitchcock DB. A comparison of hierarchical methods for clustering functional data. *Communications in Statistics-Simulation and Computation*, 38(9):1925–1949, 2009.
- [18]. Forcato M, Nicoletti C, Pal K, Livi CM, Ferrari F, and Bicciato S. Comparison of computational methods for hi-c data analysis. *Nature methods*, 14(7):679, 2017. [PubMed: 28604721]
- [19]. Fowlkes EB and Mallows CL. A method for comparing two hierarchical clusterings. *Journal of the American statistical association*, 78(383):553–569, 1983.
- [20]. Fraser P and Bickmore W. Nuclear organization of the genome and the potential for gene regulation. *Nature*, 447(7143):413–417, 2007. [PubMed: 17522674]
- [21]. Friedman J and Tukey J. A projection pursuit algorithm for exploratory data analysis. *IEEE Transactions on Computers*, C-23(9):881–890, 9 1974.
- [22]. Gibbons FD and Roth FP. Judging the quality of gene expression-based clustering methods using gene annotation. *Genome research*, 12(10):1574–1581, 2002. [PubMed: 12368250]
- [23]. Hnisz D, Weintraub AS, Day DS, Valton A-L, Bak RO, Li CH, Goldmann J, Lajoie BR, Fan ZP, Sigova AA, Reddy J, Borges-Rivera D, Lee TI, Jaenisch R, Porteus MH, Dekker J, and Young RA. Activation of proto-oncogenes by disruption of chromosome neighborhoods. *Science*, 351(6280):1454–1458, 2016. [PubMed: 26940867]
- [24]. Hopke PK and Kaufman L. The use of sampling to cluster large data sets. *Chemometrics and Intelligent Laboratory Systems*, 8(2):195–204, 1990.
- [25]. Hubert L and Arabie P. Comparing partitions. *Journal of Classification*, 2(1):193–218, 12 1985.
- [26]. Kerpedjiev P, Abdennur N, Lekschas F, McCallum C, Dinkla K, Strobelt H, Luber JM, Ouellette SB, Azhir A, Kumar N, Hwang J, Alver BH, Pfister H, Mirny LA, Park PJ, and Gehlenborg N. Higlass: Web-based visual comparison and exploration of genome interaction maps. *bioRxiv*, 2017.
- [27]. Kerpedjiev P, Abdennur N, Lekschas F, McCallum C, Dinkla K, Strobelt H, Luber JM, Ouellette SB, Azhir A, Kumar N, et al. Higlass: Web-based visual exploration and analysis of genome interaction maps. *bioRxiv*, page 121889, 2018.
- [28]. Lehmann DJ, Kemmler F, Zhyhalava T, Kirschke M, and Theisel H. Visualnostics: Visual guidance pictograms for analyzing projections of high-dimensional data. *Computer Graphics Forum*, 34(3):291–300, 2015.
- [29]. Lekschas F, Bach B, Kerpedjiev P, Gehlenborg N, and Pfister H. Hipiler: Visual exploration of large genome interaction matrices with interactive small multiples. *IEEE transactions on visualization and computer graphics*, 24(1):522–531, 2018. [PubMed: 28866592]
- [30]. Lex A, Streit M, Schulz H-J, Partl C, Schmalstieg D, Park PJ, and Gehlenborg N. Stratomex: Visual analysis of large-scale heterogeneous genomics data for cancer subtype characterization In

Computer graphics forum, volume 31, pages 1175–1184. Wiley Online Library, 2012. [PubMed: 27942089]

- [31]. Lupiáñez DG, Kraft K, Heinrich V, Krawitz P, Brancati F, Klopocki E, Horn D, Kayserili H, Opitz JM, Laxova R, et al. Disruptions of topological chromatin domains cause pathogenic rewiring of gene-enhancer interactions. *Cell*, 161(5):1012–1025, 2015. [PubMed: 25959774]
- [32]. Mackinlay J. Automating the design of graphical presentations of relational information. *ACM Transactions On Graphics (TOG)*, 5(2):110–141, 1986.
- [33]. McQuin C, Goodman A, Chernyshev V, Kametsky L, Cimini BA, Karhohs KW, Doan M, Ding L, Rafelski SM, Thirstrup D, Wiegand W, Singh S, Becker T, Caicedo JC, and Carpenter AE. Cellprofiler 3.0: Next-generation image processing for biology. *PLOS Biology*, 16(7):1–17, 07 2018.
- [34]. Milligan GW and Cooper MC. A study of the comparability of external criteria for hierarchical cluster analysis. *Multivariate behavioral research*, 21 4:441–58, 1986. [PubMed: 26828221]
- [35]. Paley SM and Karp PD. The pathway tools cellular overview diagram and omics viewer. *Nucleic Acids Research*, 34(13):3771–3778, 2006. [PubMed: 16893960]
- [36]. Rand WM. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical association*, 66(336):846–850, 1971.
- [37]. Rao SS, Huntley MH, Durand NC, Stamenova EK, Bochkov ID, Robinson JT, Sanborn AL, Machol I, Omer AD, Lander ES, et al. A 3D map of the human genome at kilobase resolution reveals principles of chromatin looping. *Cell*, 159(7):1665–1680, 2014. [PubMed: 25497547]
- [38]. Reeb PD, Bramardi SJ, and Steibel JP. Assessing dissimilarity measures for sample-based hierarchical clustering of RNA sequencing data using plasmode datasets. *PLOS ONE*, 10(7):1–18, 07 2015.
- [39]. Rodriguez MZ, Comin CH, Casanova D, Bruno OM, Amancio DR, Rodrigues FA, and Costa L. d. F. Clustering algorithms: A comparative approach. *arXiv preprint arXiv:161208388*, 2016.
- [40]. Rokach L and Maimon O. Clustering methods. In *Data mining and knowledge discovery handbook*, pages 321–352. Springer, 2005.
- [41]. Rousseeuw PJ. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- [42]. Rui Y, Huang TS, and Chang S-F. Image retrieval: Current techniques, promising directions, and open issues. *Journal of visual communication and image representation*, 10(1):39–62, 1999.
- [43]. Scherer M, Landesberger T. v., and Schreck T. A benchmark for content-based retrieval in bivariate data collections. In *Conference on Theory and Practice of Digital Libraries*, pages 286–297, 2012.
- [44]. Schneidewind J, Sips M, and Keim DA. Pixnostics: Towards measuring the value of visualization In *IEEE Symposium on Visual Analytics Science and Technology*, pages 199–206. IEEE, 10 2006.
- [45]. Seaman L, Chen H, Brown M, Wangsa D, Patterson G, Camps J, Omenn GS, Ried T, and Rajapakse I. Nucleome analysis reveals structure–function relationships for colon cancer. *Molecular Cancer Research*, 15(7):821–830, 2017. [PubMed: 28258094]
- [46]. Shao L, Behrisch M, Schreck T, Sipiran I, Kwon BC, and Keim DA. Identifying Locally Interesting Motifs for Exploration of Scatter Plot Matrices Poster Presentation at GI Workshop Big Data Visual Computing Quantitative Perspectives for Visual Computing, September 22, 2014 Stuttgart, Germany, 9 2014.
- [47]. Shao L, Schleicher T, Behrisch M, Schreck T, Sipiran I, and Keim D. Guiding the exploration of scatter plot data using motif-based interest measures. In *IEEE Symposium on Big Data Visual Analytics*, 2015.
- [48]. Shao L, Schleicher T, Behrisch M, Schreck T, Sipiran I, and Keim DA. Guiding the exploration of scatter plot data using motif-based interest measures. *J. Vis. Lang. Comput*, 36:1–12, 2016.
- [49]. Sips M, Kthur P, Unger A, Hege H-C, and Dransch D. A Visual Analytics Approach to Multiscale Exploration of Environmental Time Series. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2899–2907, 12 2012. [PubMed: 26357199]

- [50]. Sips M, Neubert B, Lewis JP, and Hanrahan P. Selecting good views of high-dimensional data using class consistency. *Computer Graphics Forum*, 28(3):831–838, 2009.
- [51]. Smeulders AW, Worring M, Santini S, Gupta A, and Jain R. Content-based image retrieval at the end of the early years. *IEEE Transactions on pattern analysis and machine intelligence*, 22(12): 1349–1380, 2000.
- [52]. Spielmann M, Lupiáñez DG, and Mundlos S. Structural variation in the 3D genome. *Nature Reviews Genetics*, 2018.
- [53]. Strobelt H, Bertini E, Braun J, Deussen O, Groth U, Mayer T, and Merhof D. HiTSEE KNIME: a visualization tool for hit selection and analysis in high-throughput screening experiments for the KNIME platform. *BMC Bioinformatics*, 13(Suppl 8):S4, 12 2012.
- [54]. Tatu A, Albuquerque G, Eisemann M, Bak P, Theisel H, Magnor MA, and Keim DA. Automated analytical methods to support visual exploration of high-dimensional data. *IEEE Transactions on Visualization and Computer Graphics*, 17(5):584–597, 5 2011. [PubMed: 21041874]
- [55]. Tatu A, Maass F, Faerber I, Bertini E, Schreck T, Seidl T, and Keim DA. Subspace Search and Visualization to Make Sense of Alternative Clusterings in High-Dimensional Data In *IEEE Conference on Visual Analytics Science and Technology*, pages 63–72. IEEE, IEEE CS Press, 10 2012.
- [56]. Tukey JW and Tukey PA. Computer graphics and exploratory data analysis: An introduction. *The Collected Works of John W. Tukey: Graphics: 1965–1985*, 5:419, 1988.
- [57]. Ward MO and Guo Z. Visual exploration of time-series data with shape space projections. *Computer Graphics Forum*, 30(3):701–710, 2011.
- [58]. Wilkinson L, Anand A, and Grossman R. Graph-theoretic scagnostics In *IEEE Symposium on Information Visualization*, volume 5, pages 157–164. IEEE Computer Society, 10 2005.
- [59]. Wood J, Slingsby A, and Dykes J. Visualizing the dynamics of London’s bicycle-hire scheme. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 46(4):239–251, 2011.
- [60]. Yang J, Hubball D, Ward MO, Rundensteiner EA, and Ribarsky W. Value and relation display: Interactive visual exploration of large data sets with hundreds of dimensions. *IEEE Transactions on Visualization and Computer Graphics*, 13(3):494–507, 2007. [PubMed: 17356216]

Fig. 1:

In our pattern-driven exploration pipeline visual representations (1) serve as a proxy to understand the underlying data of interest. We extract image features (2) of these data views with the goal to partition the (visual) pattern space into interpretable, coherent chunks. To select an appropriate feature descriptor, we calculate heuristic quality scores (4) assessing each descriptor's ability to discriminate the visualizations of interest. The clustering results (3) of the best performing feature descriptor, is shown in an interactive visualization that lets the user explore pattern clusters and distributions (5).

Fig. 2:

Image Feature Analysis. Top: The standard feature extraction operates on the raw data and is typically defined in a static and heuristic way. Bottom: Our approach extracts features from a visual representation. The approach is able to visually represent why objects are similar and provides a starting point for navigation and visual pattern space exploration.

Fig. 3: Visual depiction of a cluster prototype in a hierarchical clustering. We overlay all cluster entities visually, such that their relative opaqueness value aggregates to 1.

(a) Overview of the scatter plot pattern space. The screenshot in ① is taken at the Dendrogram Cutlevel (DC) 4. The left cluster contains significantly more locally dense elements (1,200) than the right cluster with globally dense scatter plots (76). The drill-down slider ② allows modifying the dendrogram cutlevels and controls the cluster prototype resolution. The pie chart metaphor in ③ depicts the absolute cluster size.

Fig. 4:

Pattern space visualization for 1,276 scatter plots representing a BSRN Earth observation dataset; The best performing feature descriptor COLOR_OPPONENT_HISTOGRAM differentiates globally dense from locally dense scatter plots on the first split levels and later differentiates the density aspects according to its shape. The adapted dendrogram visualization shows a cluster representative image and encodes the size of the respective cluster in the outer arcs. The orange border highlights the current dendrogram split level, which can be interactively modified with the slider below.

Fig. 5:

The scatter plot of the feature descriptors (FD) performance. The y-axis shows the variance of our quality score (more is better); the x-axis depicts the silhouette score (more is better). Nine FDs result in similar quality metric scores. The texture descriptor “Haralick” scores first with respect to the variance, but only fifth for the silhouette index, thus leading to an aggregate fourth QM score ranking place.

Fig. 6:
Quality metrics for the BSRN dataset. The Calinski-Harabasz quality index shows similar results to the Silhouette index.

(a) BSRN patterns 1

(b) BSRN patterns 2

(c) BSRN patterns 3

(d) BSRN patterns 4

Fig. 7:
Exemplary patterns of the BSRN dataset.

Fig. 8:
Pattern variability of measure “DIR” (Direct radiation).

Fig. 9:
Quality control of loop patterns (a & b). Stratification of two TAD sub pattern types (c & d).