J Vis (2021) 24:1065–1081
https://doi.org/10.1007/s12650-020-00733-z

REGULAR PAPER

Zeyu Li · Changhong Zhang · Yi Zhang · Jiawan Zhang

SemanticAxis: exploring multi-attribute data by semantic construction and ranking analysis

Received: 1 July 2020 / Revised: 13 August 2020 / Accepted: 14 August 2020 / Published online: 10 March 2021

Abstract Mining the distribution of features and sorting items by combined attributes are 2 common tasks in exploring and understanding multi-attribute (or multivariate) data. Up to now, few have pointed out the possibility of merging these 2 tasks into a united exploration context and the potential benefits of doing so. In this paper, we present SemanticAxis, a technique that achieves this goal by enabling analysts to build a semantic vector in two-dimensional space interactively. Essentially, the semantic vector is a linear combination of the original attributes. It can be used to represent and explain abstract concepts implied in local (outliers, clusters) or global (general pattern) features of reduced space, as well as serving as a ranking metric for its defined concepts. In order to validate the significance of combining the above 2 tasks in multi-attribute data analysis, we design and implement a visual analysis system, in which several interactive components cooperate with SemanticAxis seamlessly and expand its capacity to handle complex scenarios. We prove the effectiveness of our system and the SemanticAxis technique via 2 practical cases.

Keywords Multivariable data · Multi-attribute rankings · Dimension reduction · Semantic modeling

1 Introduction

Multi-attribute (or multivariate) data are widely available in the real world and often rich in information. Hence, exploratory analysis toward understanding and interpreting multi-attribute data has always been a hot topic in visual analysis. Among existing work, dimension reduction (DR) is a commonly used technique. It can visually reveal the overall distribution pattern (e.g., a direction with certain semantics) and local characteristics of data (e.g., clusters and outliers). However, we believe that the following 2 points limit its application:

Z. Li (✉) · C. Zhang · Y. Zhang · J. Zhang
College of Intelligence and Computing, Tianjin University, Tianjin, China
E-mail: lzytianda@tju.edu.cn

C. Zhang
E-mail: ch_zhang@tju.edu.cn

Y. Zhang
E-mail: yizhang@tju.edu.cn

J. Zhang
E-mail: jwzhang@tju.edu.cn

J. Zhang
Tianjin Cultural Heritage Conservation and Inheritance Engineering Technology Center, Tianjin University, Tianjin, China

- Sometimes it is difficult to understand DR results. Especially when we are unfamiliar with the data, there is no significant clusters in reduced space[1], or data distribution does not match what it looks like in analysts' minds. Analysts often ask: What does this cluster mean? Why is there an outlier here? Is there a certain direction that may convey explicit semantics?
- It does not support data filtering and ranking based on single or combined attributes. However, these tasks are important and common in multi-attribute data analysis. For example, students expect to find universities that match their interests according to their performance in various research areas; teachers hope to ascertain students who are partial to several subjects according to their scores in each subject.

At present, few work intents to overcome these 2 limitations at the same time and points out the potential connections between them. In this paper, we present SemanticAxis, a technique that treats these 2 limitations as tasks and seamlessly merges them into a united exploration context. This is achieved by enabling analysts to interactively build semantic vectors that represent abstract concepts. To be specific, in reduced space, analysts can select the region of interest and another group of points as a target group and a control group, respectively. Then we compute the high-dimensional vector that connects the center of these 2 groups. Specially, if the vector has a distinct greater absolute value in some dimensions than others, and the combination of these dimensions can be interpreted as a reasonable concept by analysts, we call this vector a semantic axis. For instance, consider a multi-attribute data that describe scores of plenty of students across multiple subjects, and we have chosen 2 groups of students, who excel in natural sciences (e.g., mathematics and physics) and humanities (e.g., history and politics), respectively. If these selected students show no significant difference in other subjects, then the semantics of one end of the constructed semantic axis denotes "they perform better in natural sciences than humanities" and the other end shows the opposite.

On the one hand, based on the idea of contrastive analysis, SemanticAxis can explain the semantics of an arbitrary region, such as a cluster or some outliers, by simply selecting it as the target group and meanwhile selecting another region, such as the majority of the remaining points or another cluster, as the control group. The differences on each dimension between the 2 group of points highlight the semantics of the target group (corresponding to the first task); on the other hand, the semantics of the axis changes in concept or strength along its direction (Fig. 2), e.g., the semantics may transfer from "male" to "female" for words, or transfer from "excellent in natural sciences" to "advanced in humanities" for students. Hence, SemanticAxis can be viewed as a metric in which data points are sorted by their projected position on the axis (corresponding to the second task).

In addition, we designed and implemented a visual analysis system (Fig. 1), in which several visual components and interactions enhance the adaptability of our SemanticAxis facing complex real-world scenarios. For example, weight editor is used for attribute refinement in multi-attribute rankings as well as reduced space reconstruction; ranking rows supports combined filtering and detail inspection, making it possible for analysts to further validate the insights gained through the SemanticAxis.

We summarize our contributions as follows:

– We proposed SemanticAxis, a technique that merges the task of feature understanding and weighted ranking of multi-attribute data into a united exploration context;
– We designed a visual analysis system, in which easy-to-use visual components and concise interactions accommodate the SemanticAxis to complex analysis scenarios;
– We presented some interesting and valuable findings on academic strength distribution in computer science; for example, the differences between institutions at different levels are markedly different.

2 Related work

2.1 Semantic axis techniques

We refer to the technique of using vectors and their linear expressions to encode semantics and their transitions as semantic axis technique. Analogies in word embedding (Mikolov et al. 2013), such as man–woman, can be regarded as semantic axes. Heimerl and Gleicher (2018) helped understanding semantic differences between 2 corpora by checking word distribution in a Cartesian coordinate system that is spanned by 2 semantic axes which is trained from these 2 corpora and describe the same concept. Liu et al.

[1] Reduced space refers to the 2D plane created by DR algorithms.

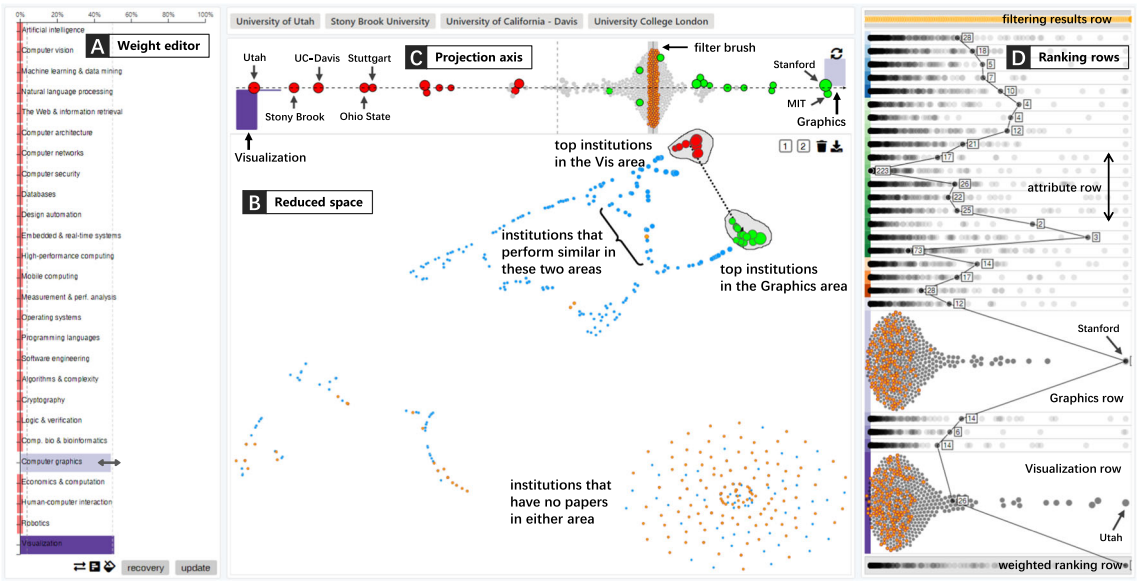

Fig. 1 Interface of our system. With weight editor, analysts are free to emphasize or ignore attributes by adjusting their weights (the light red rectangle reminds the analyst that the weight of the corresponding attribute has been or will soon be reduced to zero). Reduced space presents dimension reduction results and allows analysts to construct SemanticAxis by lassoing 2 groups of points. Projection axis is designed for interpreting a constructed axis and checking its semantic distribution in reduced space. Ranking rows are used to verify inferences by providing details of data points (space for those emphasized attributes is expanded to enable individual checking)

(2017) expanded the descriptors of a defined semantics by detecting words that fall on each side of its corresponding word vector. Explainers (Gleicher 2013) and InterAxis (Kim et al. 2015) learned linear functions from analysts' continuous decisions to model the relationship between data attributes and abstract concepts in their minds. Our SemanticAxis follows a similar technical roadmap to the InterAxis. However, the tasks, hypotheses, and motivation of the 2 are completely different. To be specific, the InterAxis emphasizes on creating an interpretable reduced space by semantic modeling, but we focus on understanding the generated reduced spaces and its combination with data ranking.

2.2 Understanding the DR results

For image or text, it is a simple and effective way to interpret their DR results by directly (Tenenbaum et al. 2000) or interactively (Heimerl et al. 2016; Kim et al. 2016; Li et al. 2019) drawing the original data as annotation in reduced space. For other types of high-dimensional data, Ji et al. (2019) applied parallel coordinates plots to identify hidden semantic features associated with recognized clusters. Cavallo and Demiralp (2018) explained the characteristics of a local region by plotting semantic curves of various dimensions centered on a certain data point. However, generally, their approach requires out-of-simple support of the DR algorithm they used. Axisketcher (Kwon et al. 2016) allowed analysts to draw a discretionary curve in reduced space and then helped them understanding the meaning of region that the curve passes through by expressing it as a combination of multi-segment linear functions. Stahnke et al. (2015), Liu et al. (2019), and we all uncovered the feature of local areas by analyzing the differences between the points inside and outside the areas. Compared to the former, we assign the differences a semantic explanation and apply it to the entire data set for rankings. Compared to the latter, we offer multiple interactions toward exploring the constructed semantics in depth. For example, we design a brush filter to support examining the distribution of semantics in reduced space with different granularity.

2.3 Interacting with DR model

We highlight 2 kinds of interactions in the context of interacting with the DR model in reduced space: parametric interaction (PI) and observation-level interaction (OLI) (Wenskovitch et al. 2017). PI refers to manipulating parameters directly in order to create a new projection. This presents a difficulty to novice or

non-mathematically inclined analysts. Typical examples include slider bars from Andromeda (PI view) (Self et al. 2016), Star Coordinates (Kandogan 2000), and SpinBox widgets from STREAMIT (Alsakran et al. 2011). OLI enables the analyst to directly manipulate the observations (data points), shielding the analysts from the complexity of the underlying mathematical models. Typical examples include StarSPIRE (Bradel et al. 2014), Paulovich et al. (2011), and Mamani et al. (2013). Recently, Self et al. (2018) have determined the differences, advantages, and drawbacks of PI and OLI, and drew the conclusion that these 2 serve different, but complementary. Semantic interaction is a similar concept to OLI whose interaction objects are also data points. It is just that the former puts more emphasis on the semantic interpretation of interaction. Semantic interaction follows the human-in-the-loop pipeline (Endert et al. 2014), in which analysts spatially interact with data models directly within the visual metaphor using expressive interactions, and then, system interface provides visual feedback of the updated model and learned parameters also within the visual metaphor (Endert et al. 2012). For example, Endert et al. (2011, 2012) enabled analysts to interactively generate a well-interpreted document space by semantic interactions, such as moving, highlighting, grouping, and annotating documents. These interactions are interpreted and mapped to the underlying parameters of a force-directed model (Endert et al. 2012) or a weighted MDS model (Endert et al. 2011). In our SemanticAxis system, analysts can rebuild projection by directly modifying attribute weights, which is clearly a PI. The interaction of creating a semantic axis by selecting 2 groups of points belongs to neither PI nor OLI, since it does not interact with the DR model.

2.4 Multi-attribute rankings

Gratzl et al. (2013) summarized some common visual designs for multi-attribute rankings, including spreadsheet (Few 2012), point-based, line-based (e.g., parallel coordinates plot (Inselberg 1985), slope graph (Tufte 2001, p.156), and bump chart (Tufte et al. 1990, p.110), and region-based (e.g., table with embedded bars Rao and Card 1994, multi-bar chart, and stacked bar Few 2012) techniques. We adopted a line-based technique: parallel coordinates plot, because it supports comparing the rankings of the same data point among various dimensions, multiple ranking criteria, and different time periods. Besides, dynamic weight adjustment is a widely used attribute refinement method (Gratzl et al. 2013; Weng et al. 2018; Carenini and Loyd 2004; Wall et al. 2017). We implemented it in our system to assist analysts in customizing desired ranking criteria.

3 Methodology

As introduced in Sect. 1, in order to understand the semantics of a target region in a reduced space, the analyst needs to construct a target group and a control group. We denote the vector of the ith and jth data point in the constructed target group and control group as $\boldsymbol{p}_i^t$, $\boldsymbol{p}_j^c \in R^{1\times N}$, respectively, and denote a weight vector that describes the importance of each attribute as $\boldsymbol{\omega} \in R^{1\times N}$ (sum up to 1). N represents the number of attributes. As shown in Fig. 2, we define the SemanticAxis $\boldsymbol{v} \in R^{1\times N}$ as the weighted vector that passes through the center of these 2 groups of points:

$$\boldsymbol{v} = \left(\frac{1}{T}\sum_{i=1}^{T} \boldsymbol{p}_i^t - \frac{1}{C}\sum_{j=1}^{C} \boldsymbol{p}_j^c \right) \odot \boldsymbol{\omega}$$

where $\odot$, L, and R represent the Hadamard product (also known as the element-wise product) and the number of points in the 2 groups, respectively. Therefore, in essence, our SemanticAxis is a linear combination of the original attributes.

In high-dimensional space, as we walk along an attribute axis, the semantic strength it expressed increases or decreases monotonically. This is also true for the axis formed by a linear combination of attributes, as long as the linear combination indeed conveys an interpretable semantics affirmed by analysts. For example, consider a multi-attribute data that record the academic performance of institutions (data points) in areas (attributes) of computer science and an axis $\boldsymbol{v}$ is a linear combination of all AI-related areas, then the semantics of this AI-axis can be interpreted as the comprehensive performance of institutions in AI. For an arbitrary institution $\boldsymbol{p}$ in high-dimensional space, its projection position $\text{proj}_{\boldsymbol{v}}^{\boldsymbol{p}} = \frac{\boldsymbol{p}\cdot\boldsymbol{v}}{\|\boldsymbol{v}\|}$ on the axis $\boldsymbol{v}$ can be used to measure its performance. In this way, SemanticAxis can be regarded as a ranking criterion

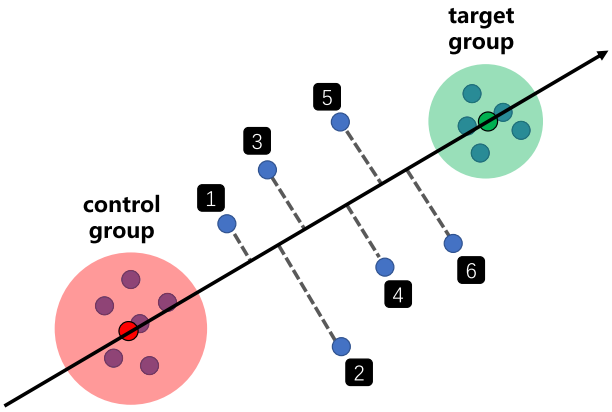

Fig. 2 Illustration for the SemanticAxis. SemanticAxis is a vector that passes through the center of 2 groups of high-dimensional data selected by analysts. It serves as a metric for the semantic strength of data points whose rankings according to their projected position on the axis

based on its represented abstract semantics, in which data points are sorted by their relative projection positions on the axis. It is exactly what the multi-attribute rankings required.

SemanticAxis utilizes the idea of contrastive analysis to interpret the semantics of any part of the reduced space. Attributes that have significant numerical differences between the target group and control group describe the characteristics (semantics) of the target cluster. Depending on the selection of the control group, semantic axis can be divided into 2 types:

- *Unipolar semantic axis* whose control group consists of the majority of the remaining data points (no need to be exactly precise). In this case, the semantic axis only takes a single semantics and its semantic strength is simply getting stronger or weaker along the axis. Take the above AI-axis as an example, institutions on the left/right always perform better in AI than those on their right/left. It is worth noting that being strong in semantic strength (or to say holding good semantic performance) does not mean being evenly strong in all involved attributes. On the contrary, it may only be strong in partial attributes.
- *Bipolar semantic axis* whose control group is another cluster. In this case, each end acts as a control group for the other end. As a consequence, each end holds a unique semantics that represents the characteristics of the corresponding cluster. The semantics gradually transfers along the axis. It is worth noting that, in this case, the projection position of data points reflects their *relative* instead of absolute differences on semantic performance between the 2 ends. Therefore, nodes toward one end are those whose semantic performance at the current end is far stronger than that at the other end, and nodes located at the middle of the axis are those whose semantic performance at both ends are similar. Imagine a "AI–Theory" axis, whose 2 ends consist of institutions that perform excellent in AI and Theory, respectively, then the semantics of the AI-end is "performance in AI is much better than that in Theory," while the Theory-end is the opposite. For those institutions located in the middle, their performances in AI and Theory are similar, either both excellent or both weak.

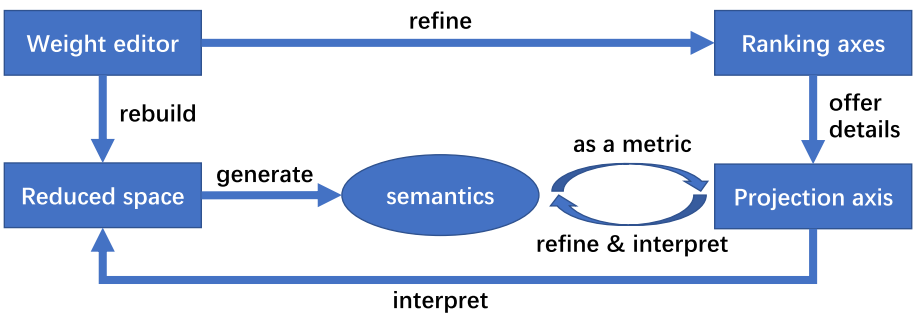

Fig. 3 Relationships between four main components of our system

4 Interface and interactions

To enhance the capabilities and applicability of our SemanticAxis, we designed and implemented a visual analysis system. In this section, we introduce the visual design and interaction design of each component and the cooperation between them. The relationships between the four main components are summarized in Fig. 3.

4.1 Reduced space

Reduced space uncovers the distribution of weighted high-dimensional data through embedding them into a two-dimensional plane. The DR algorithm we used in our paper is *t*-distributed stochastic neighbor embedding (t-SNE) (Maaten and Hinton 2008), since it can give a visually explicit DR result where the potential clusters are usually well separated without losing the details of individual cluster. In fact, any DR algorithm can be used here, we will talk about this in detail in Sect. 7. In reduced space, each circle represents a data point, and its radius is proportional to its weighted score given by $\sum_{i=1}^{N} \boldsymbol{p} \cdot \boldsymbol{\omega}$, where $\boldsymbol{p}$ and $\boldsymbol{\omega}$ are the data point vector and the weight vector, respectively. Analysts are free to lasso 2 sets of points in the reduced space as 2 ends of a semantic axis. The DR algorithm and the lasso interaction ensure the selected data group usually holds a stable and meaningful semantics, which is a precondition for the created semantic axis to be interpretable.

4.2 Weight editor

Analysts can adjust the weight of each attribute in weight editor by changing the length of its corresponding rectangle. With the constraint that the sum of the weights remains 1, as an analyst increases (decreases) the weight of one attribute, and the others decrease (increase) equally until one drops to zero. Compared to the design in which each attribute can be adjusted independently without the constraint and the linkage, our design greatly reduces the amount of operation needed for ignoring numerous undesired attributes. Weight adjustment is necessary for the following purposes:

- First, to avoid the situation where an analyst is intent to describe a concept in mind, but has difficulty finding its embodiment in the initial reduced space. For example, an analyst wants to construct a semantic axis that indicates how strong an institution is in computer theory. A conventional idea is to select institutions that excel in computer theory as one end and the rest as another end. However, the target institutions may be scattered throughout the current reduced space which prevents the analyst lassoing them. At this point, the analyst can change the weight of relevant attributes to reshape the reduced space so that the target institutions are clustered together.
- Second, to implement attribute refinement of multi-attribute rankings. On the one hand, analysts can build ranking criteria according to their own preferences, such as increasing the weight of the attributes that they care about; on the other hand, analysts can perceive the influence of a focused attribute on rankings by observing how the rankings change after adjusting their weights. The weighted ranking results are shown in weighted ranking row and we will mention it later.

4.3 Projection axis

In projection axis, circles denote data points, radiuses are proportional to their weighted scores, and x-positions are scaled according to their projection positions on the current semantic axis. A force-directed algorithm is utilized to prevent the overlap of data points. Rectangles represent all attributes of data, and we call them attribute rectangles. Their height is proportional to the absolute value of its corresponding element of the current semantic axis vector $\boldsymbol{v}$. They are evenly placed and sorted by the absolute value and the sign of the value they bind to, which means the rectangles with large absolute value are listed near the ends and the rectangles with positive/negative value are placed above/below the axis. Attribute rectangles are used to illuminate the semantics of the current axis. For example, in Fig. 1, the 2 rectangles at 2 ends reveal that the current axis describes the performance difference of institutions in visualization area and computer graphics area. Notice that $\boldsymbol{v}$ is scaled by the weight vector $\boldsymbol{\omega}$ element-wisely, as shown in Sect. 3. So there are 2 possible situations where several attribute rectangles are too short to be visible: 1. the 2 selected data groups

of the current axis show no difference on these attributes and 2. the weight of these attributes is set pretty low by analysts in the weight editor.

Analysts can fine-tune the semantics by slightly changing the height of rectangles. The benefits of this interaction are twofold: First, it allows analysts to eliminate the deviation (values of the irrelevant nonzero attributes) between the constructed semantics and the expected semantics, and second, post-adjustment makes it unnecessary for analysts to painstakingly picking the data points during constructing axis in the reduced space, which improves the efficiency of analysis. Projection axis will be updated automatically each time finishing lassoing 2 groups of points in reduced space.

4.4 Ranking rows

Ranking rows component has 2 functions: one for validating observations obtained from projection axis by further checking details of data points and the other for supporting several multi-attribute ranking tasks with simple designs. From top to bottom, all rows are divided into three groups: filtering results row, attribute rows and weighted ranking row. Attribute rows are similar to parallel coordinates plot where the position of a data item in each row indicates its performance in the corresponding attribute. The space for those emphasized attributes, like the visualization and computer graphics in Fig. 1, is expanded, while all data points are spread out, enabling individual checking. We design three linear scales to compute the position, and each of them focuses on an aspect of the data. The first scale is a local scale whose domain is the extent of attribute values of the current attribute. It is suitable for comparing the distribution of attribute values between attributes (see Fig. 4 a). The second scale is a global scale whose domain is the extent of attribute values of all attributes. It is suitable for comparing the extent of attribute values among attributes (see Fig. 4b). The third scale is a local scale. It places data points by their rankings in current attribute. It can reveal the distribution of the rankings (see Fig. 4c). The polyline crossing all attribute rows connects the same data item and uncovers the characteristics of the data item. In each row of attribute rows, a filter can be created by brushing. Created filters are combined by the "and" operator, and the final filtering results are presented in the filtering row. Weighted ranking row is used to present the weighted ranking results of data

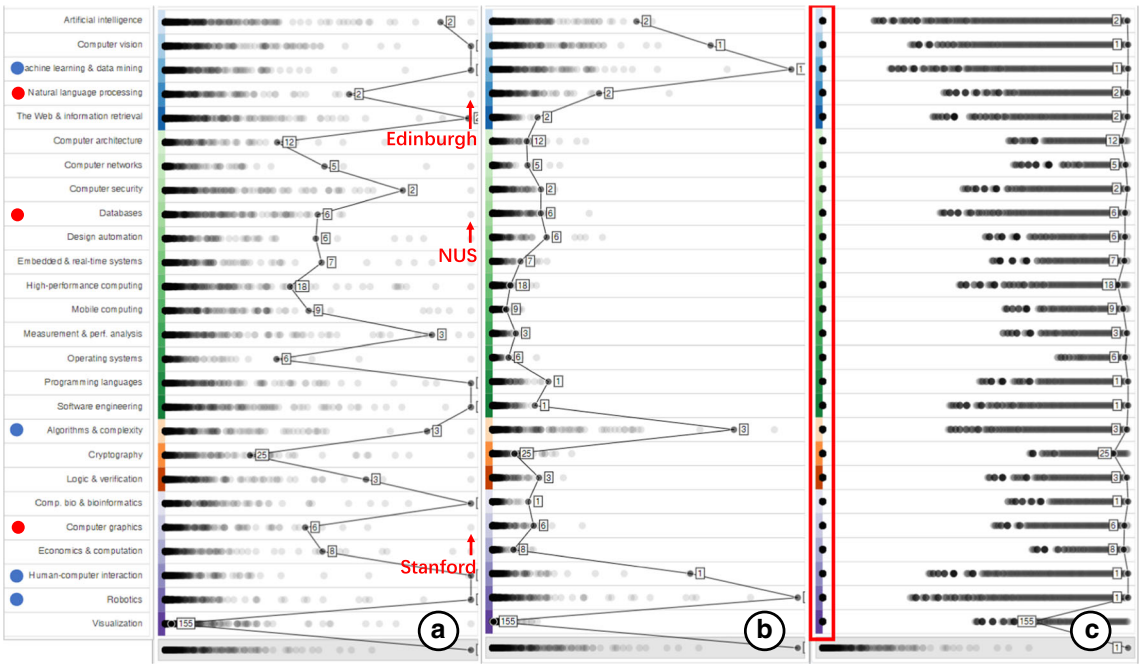

Fig. 4 Illustration of the three scale strategies in attribute rows: scale with values of individual attribute, values of all attributes, and rankings. ⓐ demonstrates that, for most areas of computer science (especially for the areas labeled with red node), there are a handful of institutions that perform noticeably well. ⓑ shows that the scale of the top conferences of blue node areas is much larger than that of other areas. The large blank on the left of ⓒ indicates that, for all areas, there are a significant number of institutions (in the red box) failed to publish in the corresponding top conferences

points based on their weighted scores. Analysts can split the row into time slices to track changes in the value or ranking of data items in different stages (see Fig. 5).

4.5 Interactions between components

After adjusting the weights in weight editor, analysts can click the update button to update reduced space and ranking rows. To examine the distribution of single or composite attributes in reduced space, analysts can check the corresponding attributes in weight editor, and accordingly, the circles in the reduced space are colored by their weighted scores (the case of composite attributes) or attribute values (the case of single attribute) on the checked attributes (see Fig. 6). To get rid of the limitation that analysts can only learn the data distribution in a single semantic axis, we designed a two-dimensional composite semantic space whose x-axis and y-axis are 2 created semantic axes, respectively. Its creation process is stated as follows: During the exploration, analysts are allowed to save the created semantic axis and check it by hovering its corresponding filled square at the top right corner of reduced space. Once 2 axes are saved, analysts can click the rightmost icon to show the custom composite semantic space in a pop-up window. The x and y coordinates of data points inside the space correspond to their projection positions on the 2 saved axes. For example, as shown in Fig. 7, analysts can identify institutions that excel in VIS area while having strong comprehensive strength. In projection axis, we allow analysts to brush some data points, and the selected points will be highlighted in reduced space and attribute rows. The observed granularity can be adjusted according to the width of brush. This interaction is simple but very useful, giving analysts the ability to scrutinize the distribution of captured semantics in reduced space. We will elaborate on this in Sect. 5.

We notice that it is hard for analysts to remember the semantics of all clusters they have ever explored. They may repeatedly construct similar semantic axes to examine the same cluster, which greatly reduces the efficiency, especially when there are numerous clusters and their boundaries are not clear. In order to alleviate the memory burden of analysts and prevent repetitive operations, we designed a storage mechanism called checkpoints. It allows analysts to save all information of the new created semantic axis, including a snapshot of the whole projection axis, the lassoed data points, and the lassoed regions. Meanwhile, a circle representing the checkpoint is pinned at the center of each lassoed region. Analysts can choose to hide the checkpoints representing control groups. When hovering over a checkpoint, its saved information emerges; if the checkpoint is clicked, the information would be embedded back in appropriate panels for further viewing. In addition to solidifying knowledge during exploration, checkpoints can serve as landmarks for reduced space, providing focus points and aiding to navigations (Ware et al. 2013, p.156), for example, guiding analysts to check unexplored regions (Han et al. 2019).

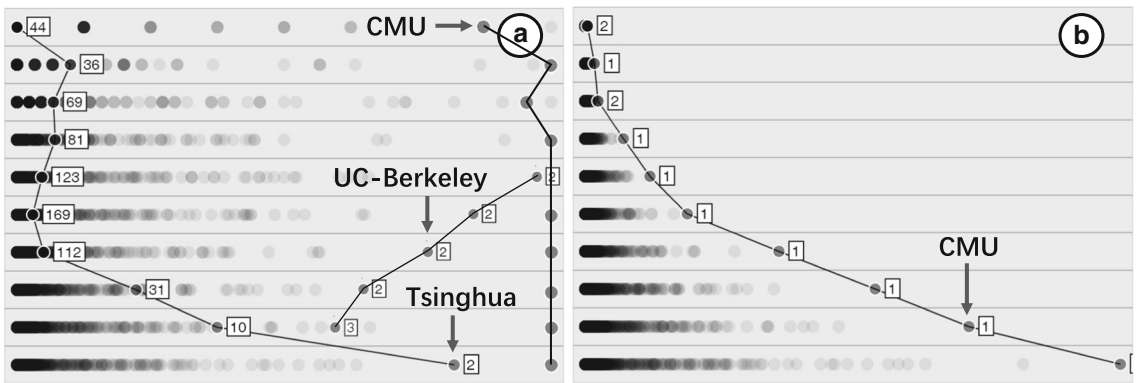

Fig. 5 Example of a weighted ranking row unfolded by time stages. We split the period 1970–2020 into 10 segments with a 5-year time interval. Each row reflects the comprehensive strength of institutions in the corresponding time period. Same as in the attribute rows, analysts can select local scale (ⓐ) or global scale (ⓑ). ⓐ indicates that the gap between the first (basically, Carnegie Mellon University (CMU)) and the second (basically, UC-Berkeley) institutions gradually widen since 1990, until Tsinghua University rushed to the second. ⓑ presents the rapid development of computer science in the past decades

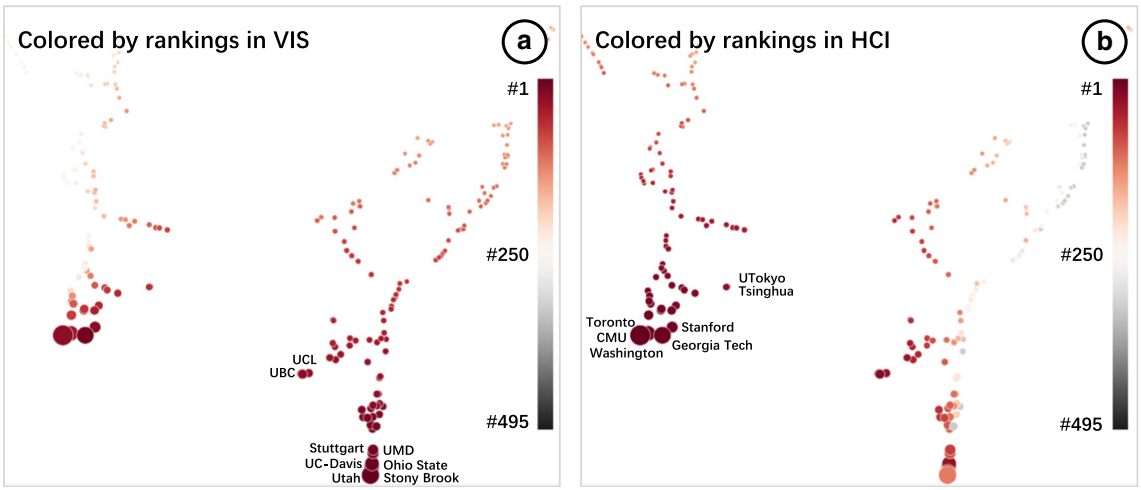

Fig. 6 Examples of coloring items by their values or rankings on selected weighted attributes. In ⓐ and ⓑ, institutions (points) are colored according to their rankings in the VIS and HCI areas, respectively. We can see that the bands on the left and right roughly represent institutions that perform well in the HCI and VIS, respectively

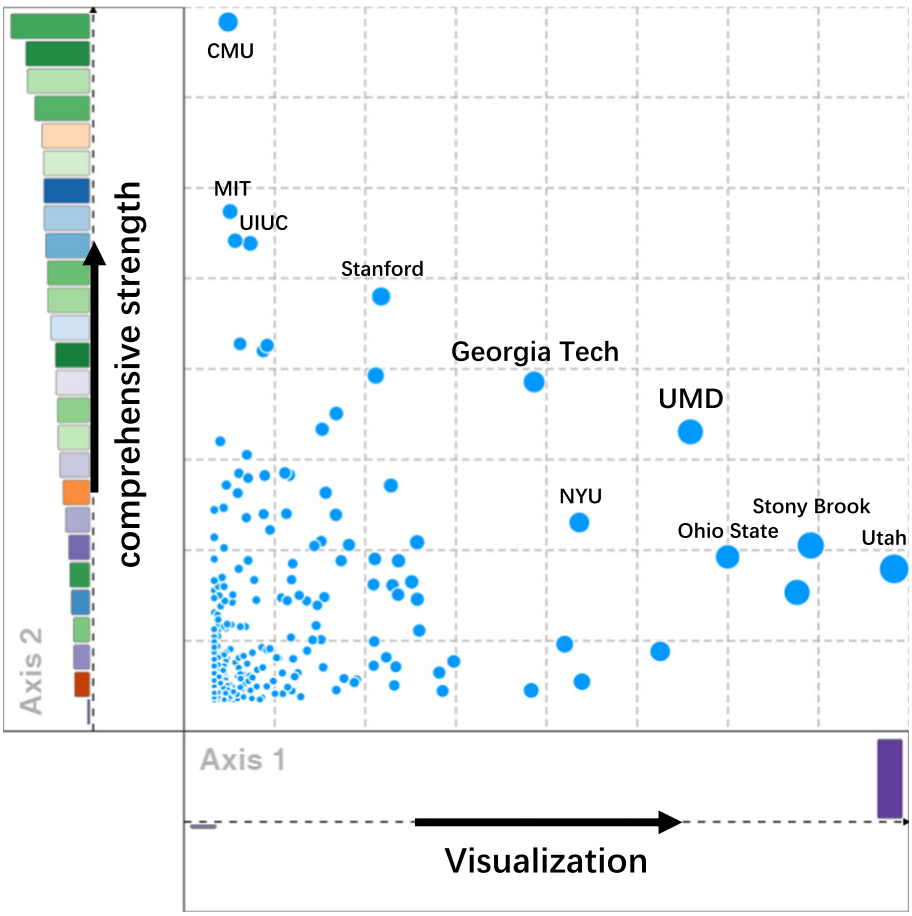

Fig. 7 A simple example of a composite custom semantic space. The distribution of institutions (points) indicates that the correlation between the performance in the VIS and the comprehensive strength is small. Georgia Tech and College Park are among the few options for institutions that have achieved great rankings in both aspects

5 Case study

5.1 Case 1: score data of students

In this case, we aim at finding biased students according to their scores in an examination. The biased students refer to the students who go overboard on partial subjects, but perform weakly in others. For example, some students may excel in natural sciences, but do badly in humanities. It is necessary for teachers to find out these biased students, because they should be given guidance toward balanced development or be encouraged to dive into their specialties. Our data record the scores of 494 students at a middle school in Ningbo in a final examination involving nine subjects. The original scores are converted to z-scores to ignore differences in the data distribution between subjects. In this case, we pay attention to the performance differences between natural sciences and humanities. Traditionally, we consider natural sciences to include mathematics, physics, chemistry, biology, and humanities to include Chinese, English, history, politics, and geography.

We expect to construct a semantic axis to measure the performance differences. A reasonable design would be to set one end means "excel in the sciences" and the other represents "do well in humanities." Since the target semantics is clear, we can construct the semantic axis in projection axis by dragging attribute rectangles directly. As shown in Fig. 8, we set the value of the semantic axis on subjects in the 2 fields to a negative value and a positive value, respectively. Currently, the projection position of students on this axis should convey their degree of bias; that is, students at 2 ends may have a serious bias. To verify this, we select three representative students and check their rankings on each subject in ranking rows (Fig. 8). We find that the No. 14919 student ranked significantly better in natural sciences than humanities, the No. 14740 student did the opposite, and the No. 14760 student located in the middle of the semantic axis showed no obvious difference between the 2. This indicates the correctness of the axis we constructed.

Next, we want to check whether the current reduced space captures the "biased" semantics. We select 60 students at each end with brush filter and examine their distributions in the reduced space (Fig. 9). We find that these students are clearly clustered on the lower right side (ⓐ) and the upper left side (ⓑ), which suggests a positive answer to the previous question. Then we notice that the students prefer the sciences are distributed mainly in 2 distinct regions (ⓐ). Hence, we build 2 semantic axes to understand their differences. As shown in ⓒ and ⓓ, one region includes students who have a specific advantage in math alone, while the other region includes students who perform well in all subjects in the sciences, except for biology.

It is worth noting that there is no obvious cluster or other features (clues) in the original reduced space, which may cause analysts to have no idea where to start the exploration. We overcome this by allowing analysts to build an initial impression of the reduced space by exploring the semantics in their minds first.

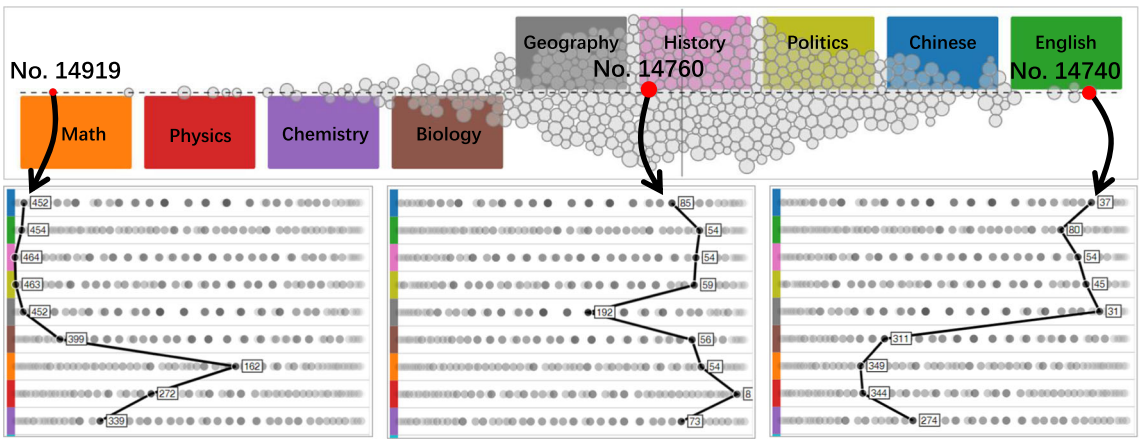

Fig. 8 Constructed SemanticAxis for detecting biased students. Students located at 2 ends of the axis hold a serious bias. It is verified by the ranking details offered by the ranking rows

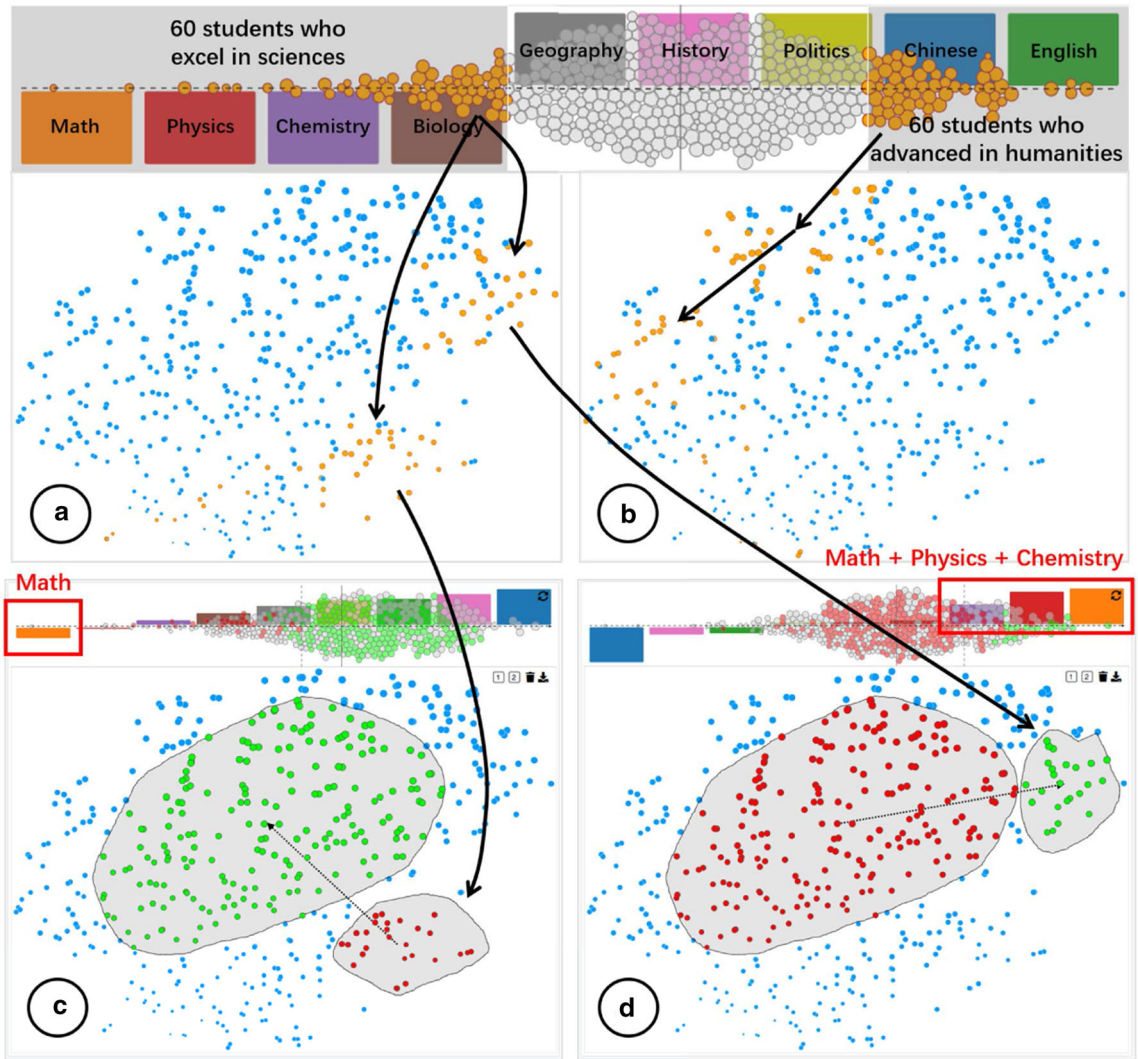

Fig. 9 Biased students found by SemanticAxis are distributed on the lower right (ⓐ, partial to the sciences) and upper left (ⓑ, partial to humanities) sides of the reduced space. The differences between the 2 groups of students in ⓐ lay in that the former only excel in math (ⓒ), while the latter is good at math, physics, and chemistry (ⓓ)

5.2 Case 2: academic performance data of institutions

In this case, we explore the academic performance and rankings of the whole world institutions in computer science. Potential users include decision-makers who manage and plan the subject of computer science and students who look forward to choosing an ideal school. Our data come from CSRankings (Berger 2020) and are updated to April 2020. CSRankings divides computer science into 4 categories and 26 sub-areas. It scores 495 institutions on all areas based on their papers published in corresponding top conferences. These conferences are carefully chosen by senior domain experts. We set an individual linear scale for each area to range all scores between 0 and 100. Then we assess a synthesis score for institutions by weighted arithmetic mean on areas.

First, let us take a general perspective to check the distribution of institutions in reduced space. In the initial settings, the weight of all areas is 1/26, and the value of each attribute of initial semantic vector is equal. This means that currently, the position of points in projection axis and the size of the nodes in reduced space indicate the comprehensive strength of its corresponding institution. We discover a probable general pattern in the reduced space: The stronger the comprehensive strength of an institution, the higher its position. To verify this observation, we construct a narrow filter brush in the projection axis and then move

it from right to left by inches, while paying close attention to the movement of highlighted (brushed) points in the reduced space. We find, generally, the highlighted points move from up to down (Fig. 10), which means that the vertical direction of the initial reduced space almost represents the comprehensive strength. Hence, the general pattern is approximately true.

Then we expect to understand the differences between different levels of institutions. As shown in the projection axis (Fig. 10), we divide all institutions into three classes: first (top 80), second (81st $\sim$ top 50%), and third class (last half). Then we build 2 semantic axes with these three classes as endpoints (Fig. 11). Unsurprisingly, institutions in upper class perform better than those in lower class in all areas. Nevertheless, the main differences between the first and second classes lay in areas under System category, while the main differences between the second and third classes lay in areas under AI category and Theory category.

We try to give a possible explanation for this finding. System category contains many old-line areas of computer science in which traditionally strong institutions have accumulated significant advantages. In contrast, most AI papers have been published over the past 10 years. The first class and second class almost stood on the same starting line. Besides, deep learning has swept through the whole academic, mobilizing the enthusiasm of almost all researchers in relevant fields, which makes the papers published in top conferences are no longer concentrated on a few top institutions (such as the first-class institutions). The second-class institutions have participated in the competition extensively. However, the third class is still unable to catch the deep learning fast train, resulting in a widening gap in AI with the second class. As for areas in Theory category, their leading institutions are in the second class, leading to a comparable status

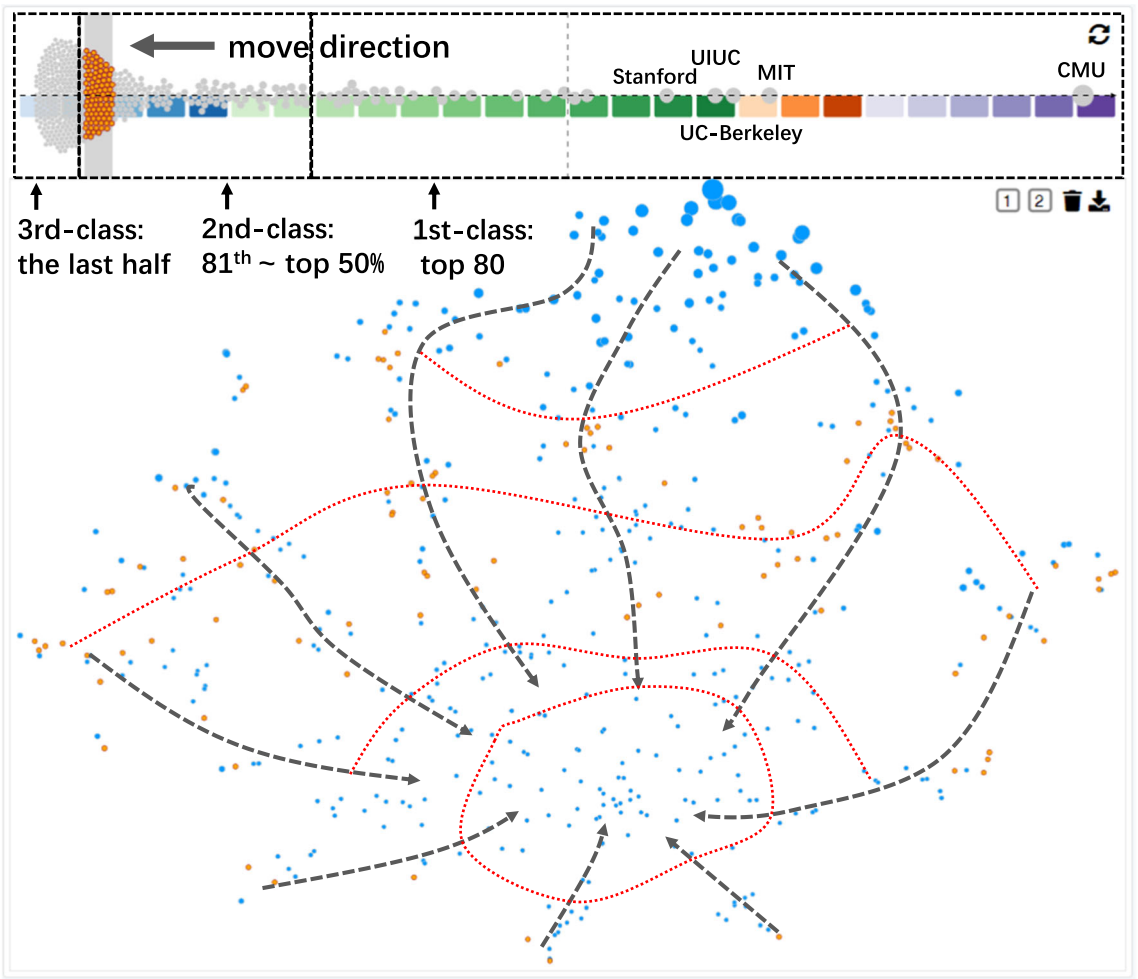

Fig. 10 Illustration of the movement of selected institutions while moving brush filter from the right to the left. Black lines represent moving trajectories, and red lines signify arrival moments

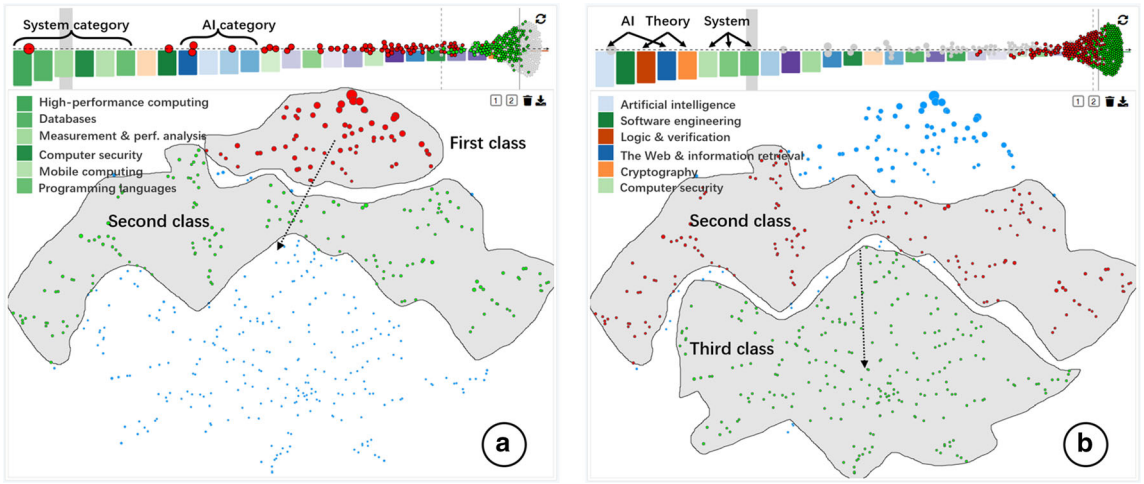

Fig. 11 Institutions in upper class perform better than those in lower class in all areas. Nevertheless, the differences between the first (top 80) and second classes (81st $\sim$ top 50%) are significantly different from that between the second and third classes (last 50%). It is mainly reflected in the conspicuous forward movement of areas under AI and Theory category in the later

between the first and second classes, while there is a relatively large gap between the second and third classes.

Next, we notice a cluster in the middle of the reduced space and intend to understand its meaning (Fig. 12). We lasso this cluster and almost all the other points as the 2 ends of a SemanticAxis. In projection axis, we see that all the other areas are close to zero compared to the visualization area, which means that the feature of the cluster is "excellent in visualization area" and institutions with the same feature should be placed near the left. Further, we notice that there are four institutions in the lasso, but away from the left, which means the current lasso is not accurate and these four institutions should not have been included in the cluster. Conversely, several institutions, such as Maryland-College Park, Stony Brook, Utah, and UC-Davis, are not in the lasso, but at the upper right corner of the reduced space. We infer that these institutions are excellent enough to follow the general pattern. This suggests another underlying semantics of the institutions in the focused cluster: They are mediocre in most other areas. All of these inferences are verified with the details provided by ranking rows. Now, we have shown that our system can help analysts to understand the precise semantics of clusters, correct inaccurate cluster boundaries, and detect and interpret outliers.

Finally, we expect to check how the rankings have changed in the field of AI. We set the weight of each of the five areas under AI category at 20% in weight editor and then divide the period from 2000 to 2020 into four segments at five-year intervals in weighted ranking row. We find that several Chinese institutions have made great progress (Fig. 13). For example, Tsinghua University rose from 90th to 1st, Peking University rose from 110th to 3rd, and the Chinese Academy of Sciences rose from 56th to 4th. The top universities in the USA, such as CMU, Stanford, and UC-Berkeley, have always been among the best. The competition between China and the US in AI has become increasingly strong.

6 Expert interviews

In order to verify the effectiveness of our system in practice, we conducted an informal user study. We invited 3 frontline middle school teachers and three professors in charge of discipline construction at our college to take part in the analysis scenarios in Case 1 and Case 2, respectively. We spent 30 min explaining the goals, visual encodings, and interactions of our system and presenting the findings shown in Sect. 5, then asked them to operate it for 30 min, and finally performed a 20 min interview with each person. In the interview, they all agreed that our system is efficient, easy to use, and could help them get some valuable information that is difficult to grasp in their daily work. One teacher said, "I know the general learning situation for most students in my class, but I can not tell the specific characteristics of each student. Your interactive visualization shows many clear and diverse results that I think I should review many times in my future work." "... not only the individual level, the distribution shown in Projection Axis can reflect systematic biased among students," mentioned by another teacher. "We rarely get such macro knowledge

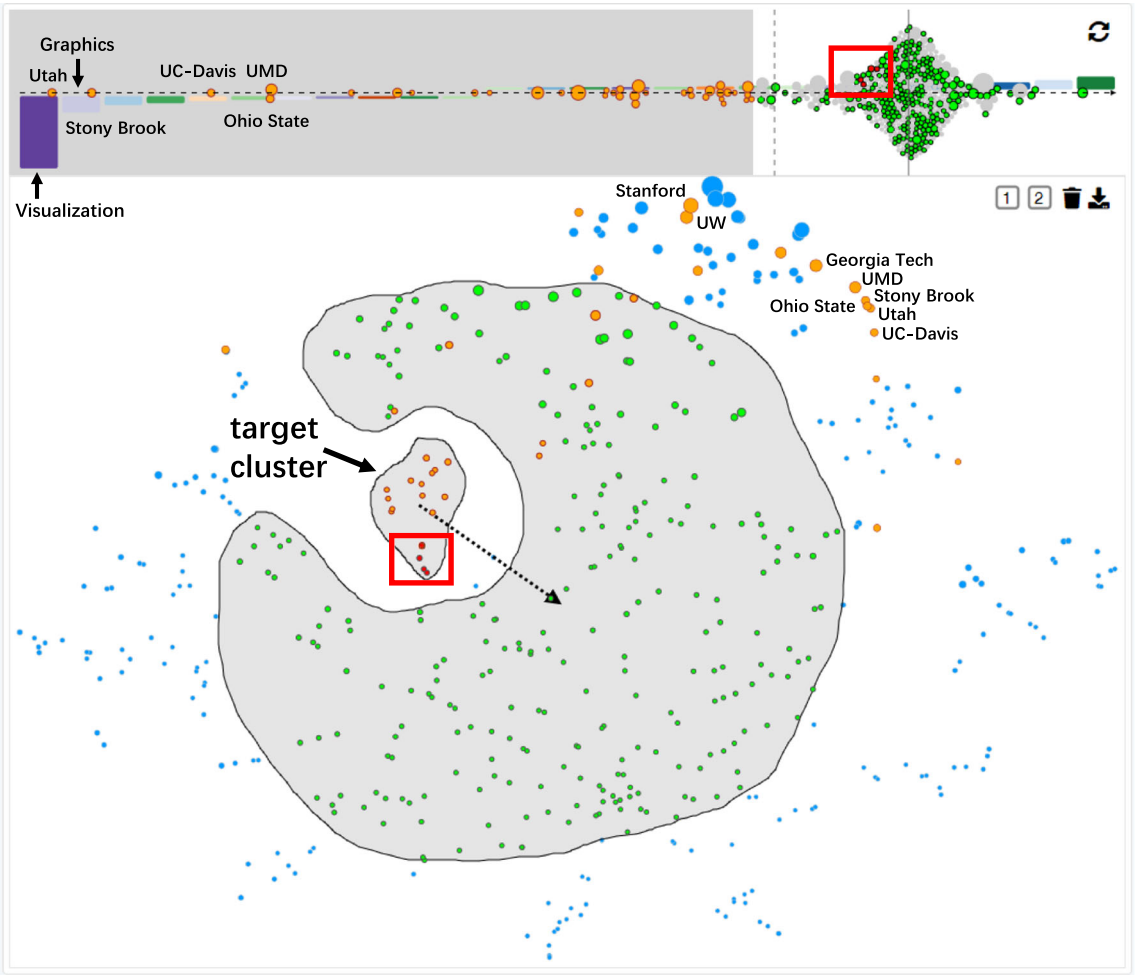

Fig. 12 Insights uncovered by checking the lassoed cluster: Four institutions (shown in the red box) should not have been selected; the exact semantics of the focused cluster is that it includes institutions that are nearly only excellent in the VIS area

[referring to the systematic differences between different levels of institutions]" said one professor, "It seems that focusing on hot areas is a feasible strategy for moving up the rankings quickly."

7 Discussion and future work

7.1 Nonlinear semantics

It is necessary to point out the caveats when analyzing nonlinear semantics using our linear SemanticAxis. We refer the semantics that analysts want to describe using the semantic axis the target semantics. It should be noted that only if the target semantics is linear, i.e., it can be described by a linear combination of the original attributes, the projection position of data points on our linear semantic axis is proportional to their strength on the target semantics. As the target semantics and the constructed semantic axes shown in Figs. 8 and 10, which are all linear, for nonlinear target semantics, such as the semantics implied in a cluster uncovered by nonlinear DR algorithms, the above proportional relation no longer holds. Constructed semantic axes even cannot be used to roughly sort data points by their strength on the target semantics when the target semantics is highly nonlinear. As in Case 2, the data points of the focused cluster are not at the left or right end of the constructed axis, i.e., the ranking function fails at this point. But this does not mean that our semantic axis technique is helpless against nonlinear semantics. We can see that the semantics of the constructed axis "performance in VIS area" is actually a linear approximation of the target semantics

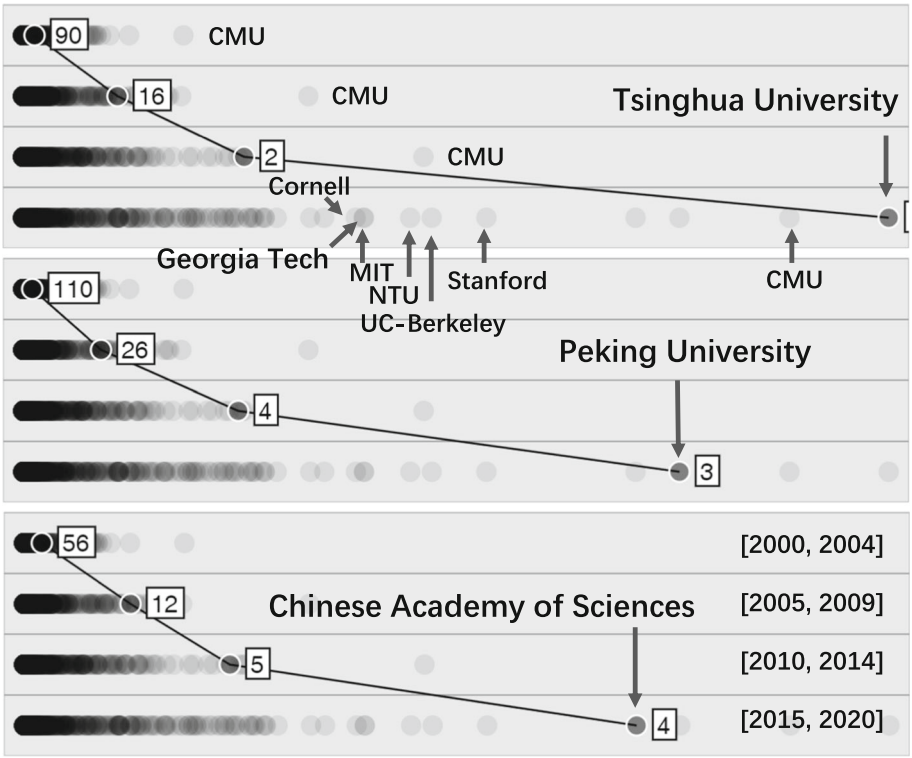

Fig. 13 In the field of AI, several institutions in China have made great progress over the last 20 years. The competition between China and the US has become increasingly strong

"excellent in VIS area only." This suggests that our linear semantic axis can give significative hints (the cluster is related to VIS area) about the precise semantics of the target cluster. Besides, as introduced in Case 2, interactions of our semantic axis with other components allow the analyst to infer and confirm the exact semantics of the target cluster.

7.2 Choice of DR algorithm

In fact, any DR algorithm can be applied in reduced space, its choice shares no relationship with the design of SemanticAxis and the linearity of target semantics, as long as it can uncover the latent global and/or local structure of the high-dimensional data of interest. We chose t-SNE (Maaten and Hinton 2008) because it tends to reveal more visually obvious local structures than other DR algorithms by emphasizing the different characteristics between latent clusters. For linear DR algorithms, like PCA (Wold et al. 1987), arbitrary direction in reduced space can be expressed as a linear combination of the original dimensions, which can be directly revealed by our SemanticAxis; for nonlinear DR algorithms, like t-SNE and UMAP (McInnes et al. 2018), our SemanticAxis can provide a linear approximation of nonlinear semantics revealed by them, as we mentioned in the previous subsection.

7.3 Scalability

One of the advantages of the linearity of our SemanticAxis is that the computation it involved is simple. Hence, SemanticAxis technique does not suffer from scalability problems. However, force-directed algorithm and dimensionality reduction algorithm, which are frequently used in our system, take on high complexity. For the former, it is an alternative to discretize the continuous position of data points by bins (Rodrigues and Weiskopf 2017); for the latter, descending sampling and applying a more efficient dimensionality reduction algorithm (e.g., UMAP or PCA) are 2 mitigatory methods.

7.4 Future work

Future work involves three aspects: First, enhance the current SemanticAxis in analyzing nonlinear semantics; second, promote the efficiency in checking the semantics of clusters by breaking the limitation that the current semantic axis only allows analysts to inspect clusters one by one (unipolar semantic axis) or two by two (bipolar semantic axis); and third, embed our SemanticAxis into the human-in-the-loop analysis process, helping analysts to understand the model and add prior knowledge to the model in a positive feedback loop (Endert et al. 2011).

8 Conclusion

In this paper, we propose SemanticAxis, a technique toward exploratory analysis of multi-attribute (or multivariate) data, and then, we present a visual analysis system with the SemanticAxis at its core. SemanticAxis characterizes abstract semantics by a linear combination of original attributes, through which it can merge the tasks of understanding the distribution and semantics of features (e.g., clusters, outliers, general patterns) and sorting/filtering data into a unified exploration context. The visual analysis system complements this context by providing supporting components and rich interactions between them. The semantic axis is computationally efficient and can be used for large-scale data. However, the inherent linearity of our SemanticAxis may hinder its application in highly nonlinear semantic analysis.

Acknowledgements The authors wish to thank all anonymous reviewers. This work was supported by National NSF of China (No. 61702359).

References

Alsakran J, Chen Y, Zhao Y, Yang J, Luo D (2011) Streamit: dynamic visualization and interactive exploration of text streams. In: 2011 IEEE Pacific Visualization Symposium, pp. 131–138. IEEE
Berger E (2020) Csrankings. http://csrankings.org/. Accessed 25 April
Bradel L, North C, House L (2014) Multi-model semantic interaction for text analytics. In: 2014 IEEE Conference on Visual Analytics Science and Technology (VAST), pp. 163–172. IEEE
Carenini G, Loyd J (2004) Valuecharts: analyzing linear models expressing preferences and evaluations. In: Proceedings of the working conference on Advanced visual interfaces, pp. 150–157. ACM
Cavallo M, Demiralp Ç (2018) A visual interaction framework for dimensionality reduction based data exploration. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, p. 635. ACM
Endert A, Fiaux P, North C (2012) Semantic interaction for visual text analytics. In: Proceedings of the SIGCHI conference on Human factors in computing systems, pp. 473–482
Endert A, Han C, Maiti D, House L, North C (2011) Observation-level interaction with statistical models for visual analytics. In: 2011 IEEE conference on visual analytics science and technology (VAST), pp. 121–130. IEEE
Endert A, Hossain MS, Ramakrishnan N, North C, Fiaux P, Andrews C (2014) The human is the loop: new directions for visual analytics. J Intell Inf Syst 43(3):411–435
Few S (2012) Show Me the Numbers: Designing Tables and Graphs to Enlighten, 2nd edn. Analytics Press, Oakland, CA, USA
Gleicher M (2013) Explainers: expert explorations with crafted projections. IEEE Trans Vis Comput Graph 19(12):2042–2051
Gratzl S, Lex A, Gehlenborg N, Pfister H, Streit M (2013) Lineup: visual analysis of multi-attribute rankings. IEEE Trans Vis Comput Graph 19(12):2277–2286
Han Q, Thom D, John M, Koch S, Heimerl F, Ertl T (2019) Visual quality guidance for document exploration with focus+ context techniques. IEEE Trans Vis Comput Graph 26:2715
Heimerl F, Gleicher M (2018) Interactive analysis of word vector embeddings. Comput Graph Forum 37:253–265
Heimerl F, John M, Han Q, Koch S, Ertl T (2016) Docucompass: Effective exploration of document landscapes. In: 2016 IEEE Conference on Visual Analytics Science and Technology (VAST), pp. 11–20. IEEE
Inselberg A (1985) The plane with parallel coordinates. Vis Comput 1(2):69–91
Ji X, Shen H-W, Ritter A, Machiraju R, Yen P-Y (2019) Visual exploration of neural document embedding in information retrieval: Semantics and feature selection. IEEE Trans Vis Comput Graph 25(6):2181–2192
Kandogan E (2000) Star coordinates: a multi-dimensional visualization technique with uniform treatment of dimensions. In: Proceedings of the IEEE Information Visualization Symposium, vol. 650, p. 22. Citeseer
Kim H, Choo J, Park H, Endert A (2015) Interaxis: steering scatterplot axes via observation-level interaction. IEEE Trans Vis Comput Graph 22(1):131–140
Kim M, Kang K, Park D, Choo J, Elmqvist N (2016) Topiclens: efficient multi-level visual topic exploration of large-scale document collections. IEEE Trans Vis Comput Graph 23(1):151–160
Kwon BC, Kim H, Wall E, Choo J, Park H, Endert A (2016) Axisketcher: interactive nonlinear axis mapping of visualizations through user drawings. IEEE Trans Vis Comput Graph 23(1):221–230
Li Z, Zhang C, Jia S, Zhang J (2019) Galex: exploring the evolution and intersection of disciplines. IEEE Trans Vis Comput Graph 26(1):1182–1192

Liu S, Bremer P-T, Thiagarajan JJ, Srikumar V, Wang B, Livnat Y, Pascucci V (2017) Visual exploration of semantic relationships in neural word embeddings. IEEE Trans Vis Comput Graph 24(1):553–562
Liu Y, Jun E, Li Q, Heer J (2019) Latent space cartography: visual analysis of vector space embeddings. Comput Graph Forum 38:67–78
Lvd Maaten, Hinton G (2008) Visualizing data using t-sne. J Mach Learn Res 9:2579–2605
Mamani GM, Fatore FM, Nonato LG, Paulovich FV (2013) User-driven feature space transformation. Comput Graph Forum 32:291–299
McInnes L, Healy J, Melville J (2018) Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv1802.03426
Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J (2013) Distributed representations of words and phrases and their compositionality. Adv Neural Inf Process Syst 26:3111–3119
Paulovich FV, Eler DM, Poco J, Botha CP, Minghim R, Nonato LG (2011) Piece wise laplacian-based projection for interactive data exploration and organization. Comput Graph Forum 30:1091–1100
Rao R, Card SK (1994) The table lens: merging graphical and symbolic representations in an interactive focus+ context visualization for tabular information. In: Proceedings of the SIGCHI conference on Human factors in computing systems, pp. 318–322
Rodrigues N, Weiskopf D (2017) Nonlinear dot plots. IEEE Trans Vis Comput Graph 24(1):616–625
Self JZ, Dowling M, Wenskovitch J, Crandell I, Wang M, House L, Leman S, North C (2018) Observation-level and parametric interaction for high-dimensional data analysis. ACM Trans Interact Intell Syst 8(2):1–36
Self JZ, Vinayagam RK, Fry J, North C (2016) Bridging the gap between user intention and model parameters for human-in-the-loop data analytics. In: Proceedings of the Workshop on Human-In-the-Loop Data Analytics, pp. 1–6
Stahnke J, Dörk M, Müller B, Thom A (2015) Probing projections: interaction techniques for interpreting arrangements and errors of dimensionality reductions. IEEE Tran Vis Comput Graph 22(1):629–638
Tenenbaum J B, De Silva V, Langford J C (2000) A global geometric framework for nonlinear dimensionality reduction. Science 290(5500):2319–2323
Tufte ER (2001) The visual display of quantitative information, vol 2. Graphics press Cheshire, Connecticut
Tufte ER, Goeler NH, Benson R (1990) Envis Inf, vol 126. Graphics press Cheshire, Connecticut
Wall E, Das S, Chawla R, Kalidindi B, Brown ET, Endert A (2017) Podium: ranking data using mixed-initiative visual analytics. IEEE Trans Vis Comput Graph 24(1):288–297
Ware C (ed) (2013) Chapter Five - Visual Salience and Finding Information. Information Visualization, 3rd edn. Morgan Kaufmann, Boston, pp 139–177. https://doi.org/10.1016/B978-0-12-381464-7.00005-3.
Weng D, Chen R, Deng Z, Wu F, Chen J, Wu Y (2018) Srvis: towards better spatial integration in ranking visualization. IEEE Trans Vis Comput Graph 25(1):459–469
Wenskovitch J, Crandell I, Ramakrishnan N, House L, North C (2017) Towards a systematic combination of dimension reduction and clustering in visual analytics. IEEE Trans Vis Comput Graph 24(1):131–141
Wold S, Esbensen K, Geladi P (1987) Principal component analysis. Chemom Intell Lab Syst 2(1–3):37–52