Visual Informatics 5 (2021) 13–25

Contents lists available at ScienceDirect

Visual Informatics

journal homepage: www.elsevier.com/locate/visinf

Bridging cognitive gaps between user and model in interactive dimension reduction

Ming Wang [a], John Wenskovitch [a,b], Leanna House [c], Nicholas Polys [a], Chris North [a,*]

[a] Virginia Tech, Computer Science, United States of America
[b] Pacific Northwest National Laboratory (PNNL), United States of America
[c] Virginia Tech, Statistics, United States of America

ARTICLE INFO

Article history:
Available online 1 April 2021

Keywords:
Interactive machine learning
Visual analytics
Dimension reduction
Usability
Cognitive gaps

ABSTRACT

Interactive machine learning (ML) systems are difficult to design because of the "Two Black Boxes" problem that exists at the interface between human and machine. Many algorithms that are used in interactive ML systems are black boxes that are presented to users, while the human cognition represents a second black box that can be difficult for the algorithm to interpret. These black boxes create cognitive gaps between the user and the interactive ML model. In this paper, we identify several cognitive gaps that exist in a previously-developed interactive visual analytics (VA) system, Andromeda, but are also representative of common problems in other VA systems. Our goal with this work is to open both black boxes and bridge these cognitive gaps by making usability improvements to the original Andromeda system. These include designing new visual features to help people better understand how Andromeda processes and interacts with data, as well as improving the underlying algorithm so that the system can better implement the intent of the user during the data exploration process. We evaluate our designs through both qualitative and quantitative analysis, and the results confirm that the improved Andromeda system outperforms the original version in a series of high-dimensional data analysis tasks.

1. Introduction

High-dimensional data is a widely-used and valuable data form, but as this type of data is complex, analyzing and cognitively understanding it can pose a significant challenge for people who do not have mathematical knowledge. To analyze such data effectively, researchers require in-depth training on a variety of mathematical models and computational techniques. The complexity of these algorithms presents an obstacle to users who seek quick explorations to gain a fundamental understanding of their data. Additionally, educators require effective methods to promote students' interest in data analytics and to make it easier for them to learn and understand high-dimensional data (Ashaari et al., 2011). In response to these needs, the authors of this work have developed a variety of visual analytics (VA) systems to help untrained persons understand their data.

However, because of the "Two Black Boxes" problem (Wenskovitch and North, 2020) that underlies communication challenges between humans and algorithms, it is difficult to design a usable ML-based interactive VA system. In this problem, the first black box is the underlying algorithms within the system. Although these algorithms can process data and provide useful results, they often do not provide justification or rationale for their outputs, making it difficult for users to decide whether or not the results are acceptable. The Explainable AI (XAI) research agenda works to address this challenge. The second black box is human cognition. Users conduct a sequence of thought processes before interacting with the VA system, with the hope that their interactions will appropriately influence the results generated by the underlying algorithm. Unfortunately, the human mind is a black box that is closed to the algorithms, and thus the algorithms often fail to capture those user intentions. The existence of these two communication challenges leads to cognitive gaps between the user and model and produces many usability challenges. To solve these challenges, VA designers must open both black boxes and bridge the gaps between the user and model.

Our work presented here extends the interface of an interactive VA system called Andromeda (Self et al., 2015a,b, 2018), which reduces data from high-dimensional space into a 2D projection using weighted multidimensional scaling (WMDS). In the two-dimensional space, distance represents the similarity between data points and follows the "proximity ≈ similarity" principle: two points that are close to each other should be more similar in the high-dimensional space than two points that are distant

* Corresponding author.
E-mail address: north@vt.edu (C. North).

https://doi.org/10.1016/j.visinf.2021.03.002
2468-502X/

from each other. Andromeda permits users to explore their data using a variety of interactions, with observation-level interactions (OLI) being most relevant to human/model interaction challenges.

OLI represents a powerful method for exploring complex relationships in the underlying data by performing direct manipulation actions on the observations in the projection, demonstrating relationships that users hope to uncover in the data. For example, OLI allows users to demonstrate a similarity relationship in the projection by dragging a subset of observations into a new configuration, after which they can learn which dimensions are important to that relationship based on a computational update to the global layout. Andromeda hides the complexity of these underlying calculations and algorithms from users, allowing them to focus on exploring the data without having to first acquire statistical knowledge. Previous studies (Self et al., 2015a, 2018, 2016a; Chen et al., 2017) have clearly indicated the effectiveness of Andromeda in helping users to explore high-dimensional data, since users are able to gain complex insights and solve analytical tasks more efficiently. However, some cognitive gaps between user and model remain to be solved (Self et al., 2016b). Over the course of previous studies with novice users of Andromeda (Self et al., 2018; Zeitz et al., 2018), we identified three primary usability challenges in the interface that prevented fully clear communication between user and algorithm.

First, users did not understand they need to select and highlight all of the data points relevant to their demonstrated interaction. Interaction feedback in Andromeda is built around the idea of highlighting, whereby users can drag or select points to convey their interaction intention. The algorithms that underlie Andromeda only capture and consider moved or highlighted data points when calculating the updated projection. This demonstrates both an ambiguity in the interaction (Wenskovitch et al., 2020) and a cognitive gap between user's mental model and the mathematical model of the system. This interaction further violates both the "Don't Make Me Think" user experience principle (Krug, 2000), requiring a user to infer the way by which a system interprets their interaction, as well as the principles of semantic interaction (Endert et al., 2012a), permitting a user to stay focused on the data rather than considering the details of the model.

Second, when using Andromeda, users were confused about how to interpret the meanings and orientations of axes in the WMDS plot. They frequently wanted to map dimensions from the high-dimensional data to axes in the 2D WMDS plot, as seen in similar studies (Wenskovitch and North, 2021). Because of the separation of WMDS plot and dimensional sliders in the Andromeda interface, it was difficult for users to understand the effect and distribution of attributes in the projection.

Third, the previous Andromeda system could not correctly address and reflect user intent when constructing or de-constructing a single cluster, as there was no notion of a "cluster" encoded in the underlying models. The generated results hence did not match the user expectations.

These three issues cover both communication channels seen in the "Two Black Boxes" problem. The first two usability challenges, ensuring that the user knows how their behavior will change the model and understanding the projection, are caused by the fact that the users do not understand the Andromeda algorithm. The third challenge, constructing/deconstructing a cluster in accordance with the intent of the user, represents a cognitive black box challenge.

These communication challenges between the user and model are also common problems in other interactive VA systems, which also must (1) ensure that the human knows what actions to take to affect the model, (2) can properly interpret the presented data, and (3) verifies that the model updates in accordance with the human's intent. Our work presented here seeks to open the cognitive and algorithmic black boxes for Andromeda, addressing the usability challenges that we have identified. In summary, the contributions of this paper are:

1. We identified three cognitive gaps between user and algorithm in the interface of an interactive VA system.
2. We proposed, implemented, and evaluated solutions designed to resolve these usability issues.

2. Related work

2.1. Interactive visual analytics

Dimension reduction algorithms represent methods for interpreting high-dimensional data. These algorithms project data from high dimensions to a two- or three-dimensional projection, making it more accessible to users, but information is inevitably lost in the process. To retain and explore this hidden information, it is necessary to develop interactive tools that can adjust parameters and visualize the data from multiple perspectives.

Wong et al. (2004) develop a VA system called IN-SPIRE™, which can display the hidden relationships within corpus data using their Galaxy visualization and the ThemeView™ visualization. iVisClassifier was created by Choo et al. (2010) and is based on LDA. All of the reduced dimensions are represented by parallel coordinates, and users can interact with and explore individual dimensions of data. These tools incorporate surface-level interactions (SLI) so that users can explore the information they find most relevant, but understanding may be limited due to users having no control over the model parameters with these interactions.

More complex tools enable parametric interactions (PI), direct interactions with the values of model parameters, so that users can visualize data from multiple perspectives. Soo Yi et al. (2005) propose the Dust &Magnet system. Users can change the magnitude of the dimensional "magnet", with different magnet layouts leading to different visualizations. Other systems, such as STREAMIT (Alsakran et al., 2011) and DimStiller (Ingram et al., 2010) also allow support PI. However, PI requires users to have strong knowledge of the system model, which would be less accessible for novice users.

Endert et al. (2011) and Leman et al. (2013) develop a more natural method – OLI – for communicating with underlying models. Based on the principles of semantic interaction (Endert et al., 2012b), users can manipulate the observation points in the visualization directly to create a new visualization, which is interpreted quantitatively by the system, leading to changes to the model's parameters. ForceSPIRE (Endert et al., 2012a) and DisFunction (Brown et al., 2012) are OLI-enabled tools that allow users to directly manipulate observations. With OLI interaction, tools hide these underlying models from users so that users can focus on the data exploration without having to learn the details of statistical models.

2.2. Computational overview of Andromeda

Andromeda is a tool designed by Self et al. (2015a,b). This system relies on the WMDS algorithm (Cox and Cox, 2000) and permits users to interact with both the input and output of the WMDS model to convey their intentions and affect the underlying parameters. In the low-dimensional projection generated by WMDS, distance represents the relative similarity between data points. To enable the exploration of the high-dimensional space, WMDS introduces a weight vector parameter, $\omega = [\omega_1, \omega_2, \ldots, \omega_p]'$, to reflect the importance of each dimension in the underlying high-dimensional distance function. As shown in Eq. (1), d_i and d_j are high-dimensional points, and r_i and r_j are their low-dimensional representations. Given the weights ω, the points' coordinates on a low-dimensional projection are calculated by minimizing the stress function that represents the error between the low- and high-dimensional pairwise distances. Considering both distances and weights, users can deepen their interpretation of the projection; for example, in a projection with high weight in dimension A, two points that are close to each other are more similar in dimension A than those that are far apart, enabling users to explore hidden information by changing the weights for different dimensions and obtaining various low-dimensional visualizations.

$$r = \min_{r_1,\ldots,r_n} \sum_{i=1}^{n} \sum_{j>i}^{n} (dist_L(r_i, r_j) - dist_H(\omega, d_i, d_j))^2, \tag{1a}$$

$$dist_H(\omega, d_i, d_j) = \sqrt{\sum_{k=1}^{p} \omega_k (d_{ik} - d_{jk})^2} \tag{1b}$$

The Andromeda system combines SLI, PI, and OLI, allowing users to explore and obtain a more complete data analysis. Users can perform SLI by hovering the cursor over a point (blue point) or selecting points (maroon points) to view their raw data on the parameter sliders. PI permits the user to directly manipulate the weights ω by dragging the parameter sliders, providing feedback on dimensional importance and triggering the forward WMDS model to recalculate the 2D visualization. The resulting new projection allows users to explore those highly-weighted dimensions more fully.

Users can perform OLI to provide input to the weight-learning "inverse" WMDS algorithm by selecting and dragging points on the screen, thereby creating new 2D coordinates. After users select the "Update Layout" button, the Andromeda system optimizes the dimension weights based on Eq. (2) to best represent the coordinates of the set of moved points, r_i^* and r_j^*. The standard "forward" WMDS dimension reduction algorithm then runs again with the new weights to recalculate the projection. Users get feedback about the new weights that express their desired distance relationships, as well as the re-expression of the rest of the points using these weights. Fig. 1(b) visually represents this computational pipeline.

$$\omega = \min_{\omega_1,\ldots,\omega_p} \sum_{i=1}^{n} \sum_{j>i}^{n} (dist_L(r_i^*, r_j^*) - dist_H(\omega, d_i, d_j))^2 \tag{2}$$

2.3. Existing solutions for usability issues

This paper focuses on discussing and solving the three usability issues noted in the Introduction: displaying the observations relevant to an interaction, understanding WMDS dimensions, and constructing and deconstructing clusters. Here, we briefly summarize other works that have proposed solutions for these issues.

There are multiple approaches to solve the issue of "highlighting relevant data points". First, systems can introduce control points, which include moved points, highlighted points, or anchor points on projection boundaries. Moving control points will result in all other points moving in relation to the control points. These types of systems include StarSPIRE (Bradel et al., 2014), IVC (Desjardins et al., 2007), iLAMP (dos Santos Amorim et al., 2012), PLP (Paulovich et al., 2011) and VRV (Sharko et al., 2008). Similarly, systems such as Dust &Magnet (Soo Yi et al., 2005) and OCI-MDS (Broekens et al., 2006) have considered all the points in the visualization when updating the layout. However, considering all points in a projection might misinterpret the users intent of expressing only a subset of relationships (using the Animals dataset from this study for example, a user may intend to only reposition the Zebra data point with respect to the Rabbit data point to demonstrate 1 one-to-one similarity relationship, with their intent considering no other points in the projection), and also would require more expensive computations when executed on large datasets.

A second approach is to apply visual feedback, which indicates how the system will interpret the user interactions in an explainable fashion (Gunning, 2017), such as "label feedback", which combines both highlighting and clustering. EluciDebug (Kulesza et al., 2015) and iVisClustering (Choo et al., 2010) use this type of feedback. Users can apply a label to an item, and the underlying model will use that label as part of its training set. Because items should be labeled in a certain manner, users are able to realize that they need to label all data relevant to an interaction.

Clustering is another solution to the issue of highlighting and is also one of the usability issues we explore in our work. Wenskovitch et al. (2020) discuss this problem, and propose a cluster membership solution, Pollux, for this usability issue (Wenskovitch and North, 2019; Wenskovitch et al., 2019). In addition, Choo et al. (2009) propose a two-stage framework for cluster visualization. ASK-GraphView (Abello et al., 2006) supports the interactive visualization of large graphs by filtering, coloring, and labeling. Linesets (Alper et al., 2011) uses node colors to represent clustering. Clustering is a good solution to the highlighting usability gap, but it is not suitable for tackling problems related to relative similarity.

In addition, it is challenging for users to understand the dimensional information in dimension reduction plots, especially dimensional orientation and correlation. Smart-Stripes (May et al., 2011) allows users to select a subset of features to explore the dependencies and independencies between various features. Biplots (Udina, 2005; Frutos et al., 2014; La Grange et al., 2009) allow for the display of both data points and dimensions. Stahnke et al. (2015) develop a probing projection system that can display how each dimension contributes to the projection. Dowling et al. (2018a,b) implement the SIRIUS system to visualize similarities both between points and between dimensions. These works help users obtain dimensional information; however, most require domain knowledge or even training to understand and apply them. Therefore, we see the need for a more intuitive and easy-to-understand method for showing dimensional information and correlations.

3. Usability problems, solutions, results

As noted in the Introduction, gaps exist between users' mental model and the mathematical model of the Andromeda system. Fig. 1 shows that in some circumstances the Andromeda interface did not allow users to understand the model correctly (highlighting relevant data points and understanding the projection dimensions), and other circumstances resulted in the system generating a response that did not match the user expectations (constructing and deconstructing clusters). In this section, we describe modifications to both the WMDS Model and the visualization to bridge these cognitive gaps.

3.1. Highlighting relevant data points

3.1.1. Problem

When users drag points in the visualization, they have their intended system response to that interaction in mind. The user can move an observation close to others to represent similarities or drag a point away from other points to express dissimilarities. Fig. 2 illustrates a user trying to express that the Zebra point is more similar to the Rabbit than to the Giraffe. In this case,

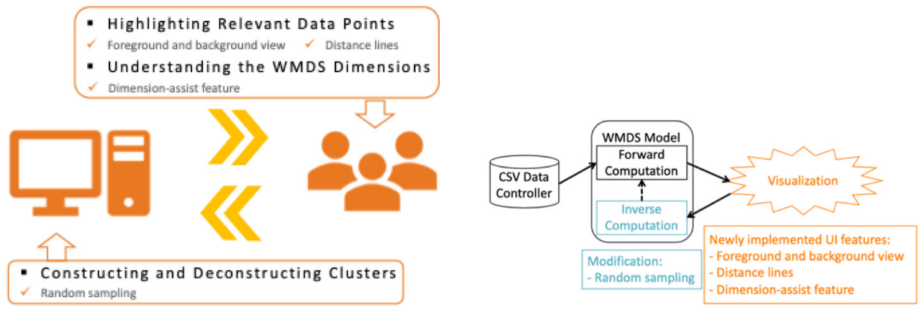

Fig. 1. (a) The identified usability problems and their solutions. (b) The semantic interaction pipeline of Andromeda, representing our changes to the inverse WMDS Model and Visualization.

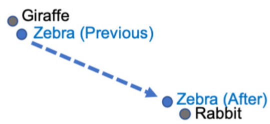

Fig. 2. The "highlighting relevant data points" usability issue. Moving the Zebra away from the Giraffe and toward the Rabbit.

the user moves the Zebra with respect to the Giraffe (near the Zebra's original position) and the Rabbit (near the Zebra's final location). However, the user does not wish to express any particular relationships to other points in the projection. So, how does the model know *with respect to which* points the user is moving the Zebra (Wenskovitch et al., 2020)? Andromeda requires the user to choose all points necessary to convey their interaction via interactive selection. This leads to a usability problem.

The user often assumes that the underlying model will automatically consider the Giraffe and the Rabbit, and not other points; however, this is not true. The Andromeda algorithm only captures and considers *moved or highlighted* points in the visualization, which in this case is only the Zebra. They might also not realize the importance of blacking-out other points that they do not care about. The mismatch between the user intention and the mathematical model occurs here because of both the closed cognitive and algorithmic boxes. The model cannot read the users mind, and users are not provided with a means to understand how the algorithms interpret their interaction, process the data, and generate the results. To solve this usability problem, we open the black box of the underlying model using new interface features to help users understand the importance of highlighting those implicit but relevant points: in this case, the Giraffe and the Rabbit.

Our goal with this study is to understand the immediate usability of the system, determining whether or not a user would know how to correctly operate the system (i.e., highlighting the correct points to specify the desired interaction) without the need for training. As such, measuring correctness and evaluating accuracy are our primary means of demonstrating the effectiveness of this technique.

3.1.2. Solutions

Foreground and Background View: To clarify the boundary between observations that are and are not considered a part of the interaction, we introduce a "third dimension" in the two-dimensional observation view. As shown in Fig. 3(a), when no point is clicked, all the data points are colored in blue and projected according to their WMDS locations. After an observation is moved or highlighted (see Fig. 3(b)), the "third dimension" is added. The projection separates into a foreground and background. The interaction points pop out in the foreground and are colored in orange; these points will be considered in the calculation of the new layout. Any untouched points are dimmed to the background and grayed out, signifying that they will not count toward the interaction. These foreground points are also highlighted in the parameter sliders (Fig. 4b), demonstrating the precise attribute values for these touched points in each dimension and visually displaying their similarity. This design visually divides the considered points and unconsidered points into two layers, giving the user a cue that all the relevant points should be brought into the foreground layer.

Distance Lines: Along with the foreground and background view, we introduce another interface feature – distance lines – to help users recognize the relationships between points that they are specifying. Thus, all pairwise distances between the points in the foreground are represented by lines. However, to avoid clutter, the distance lines only display for the point currently hovered over. As shown in Fig. 3(b), lines between the actively dragging point (Zebra) and the other selected points (Giraffe and Rabbit) are shown. There are no distance lines to the unselected points, because these pairwise distances are not considered in the calculation of the new layout.

To reinforce the meaning of the user's interactions, distance lines also visually encode whether the user has increased or decreased the relative distance between each pair of points. Only relative distances are meaningful in the WMDS model. A distance line consists of two parts, a colored line and arrow heads, both of which encode the change in relative distance between the pair of points (the color scale can be seen in Fig. 4).

To determine changes in the relative distance between point i and point j, we calculate the pairwise distance ratio ($\phi_{i,j}$) and the average distance ratio ($\overline{\phi}$). The pairwise distance ratio between point i and point j ($\phi_{i,j}$) is calculated with Eq. (3), computing the ratio between the user-defined low-dimensional distance based on the interaction $dist_L(r_i^*, r_j^*)$ and the original low-dimensional distance $dist_L(r_i, r_j)$.

$$\phi_{i,j} = \frac{dist_L(r_i^*, r_j^*)}{dist_L(r_i, r_j)} \tag{3}$$

$$\overline{\phi} = \frac{\sum_{i=1}^{n} \sum_{j>i}^{n} \phi_{i,j}}{\binom{n}{2}} \tag{4}$$

$$w_{i,j} = \phi_{i,j} / \overline{\phi} \tag{5}$$

The average distance ratio ($\overline{\phi}$) of all the pairs in the foreground is calculated with Eq. (4), which acts as a normalization factor. The relative distance change $w_{i,j}$ is then calculated in Eq. (5).

If $w_{i,j} \approx 1$, meaning that the relative distance between the two points did not change much, then there are no arrow heads

(a) (b)

Fig. 3. The "highlighting relevant data points" solution: (a) Before OLI. (b) Dragging the Zebra away from the Giraffe and toward the Rabbit, showing the foreground and background view and displaying the distance lines between the three selected points.

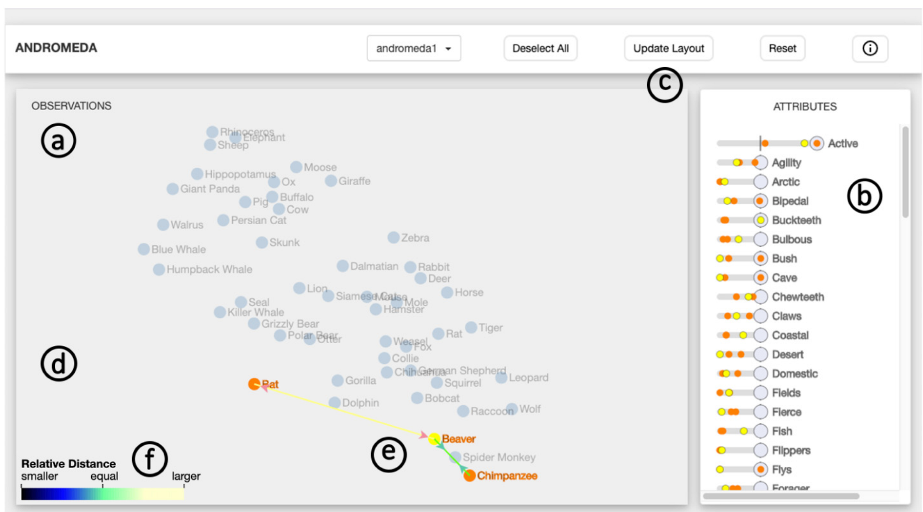

Fig. 4. Andromeda interface: (a) the observation view, (b) the parameter view, (c) the update layout button, (d) the foreground and background view, (e) distance lines, (f) the color scale for the distance lines.

on the distance line, and the color of the line is green. If $w_{i,j} \ll 1$, meaning that the relative distance between the two points was reduced by the user, then the distance line is colored darker blue or black and is capped by a pair of inward-pointing arrow heads to indicate compression. If $w_{i,j} \gg 1$, meaning that the relative distance between the two points was enlarged by the user, then the distance line is colored lighter yellow and is capped by a pair of outward-pointing arrow heads to indicate expansion.

3.1.3. Results

In a controlled usability study, we asked users to analyze a dataset by using two versions of Andromeda: Version A and Version B. Version A is the original version, whereas Version B contains the novel solutions of foreground/background and distance lines. We compare user performance of the two Andromeda versions when highlighting relevant data points for an interaction.

The participants in the usability study were sophomore and junior students in an undergraduate-level data science course who have learned about dimension reduction. There were 69 participants in total, divided into two groups for the two versions of Andromeda in a between-subjects study design. Prior to completing the study, the participants were given a brief introduction to WMDS and Andromeda Version A. The usability issues and solutions in Version B were never mentioned. In all, 33 students used Andromeda Version A, and the other 36 students used Andromeda Version B.

The participants were asked to explore a dataset about animals and their attributes (Lampert et al., 2009) using their assigned Andromeda version to answer identical sets of questions. This high-dimensional quantitative dataset contains 49 animals and 85 attributes, with the attributes expressing information about the color, habitat, diet, and behavior, among other properties. We have often used this dataset in previous studies because of its general knowledge applicability.

As shown in Table 1, we created three questions to evaluate whether the new foreground/background views and distance lines in Version B encouraged users to select and highlight all the relevant data points to complete the tasks correctly. The accuracy of participant answers was evaluated by determining if they performed the correct interactions and selected the correct observations to get the answers.

In Question 1, we asked users to create three specific clusters of two animals each. The correct answer involves highlighting all the six animals mentioned in the question. However, because of the usability issue previously mentioned, some users interacted with only one point in each cluster instead of two; for example, when a user sought to place the Deer and the Giant Panda close to each other, they may only drop the Deer next to the Giant Panda without also highlighting the Giant Panda as relevant to the interaction. In an incorrect manipulation, fewer (or more) than six points would be highlighted. In a correct manipulation, users would highlight all six relevant points, suggesting they benefited from the new visual cues in the visualization.

Table 1
The percentage of students in each group who answered correctly for questions 1–3.

	Task	Version A	Version B
Q1	Create three clusters	54.5%	80.5%
Q2	Identify the feedback points for the new layout	60.0%	85.3%
Q3	Create one cluster	56.3%	81.3%

Table 2
Results of the t-test difference between the users' total correctness scores on Version A and Version B.

	Version A	Version B
Mean	0.5827	0.7939
Standard deviation	0.3338	0.2805
Number of observations	33	36
df	63	
t-statistic	−2.8309	
$P(T \leq t)$ one-tail	0.0031	

Question 2 was an extension of Question 1, directly asking users which points would be considered by WMDS. If users understood the model from the interface features, they would provide the correct answers for all six points; however, if the new features failed to convey the concepts to users, we would obtain incorrect answers such as indicating that all points would be considered.

Question 3 asked users to create one cluster of three points. A correct interaction would involve dragging and highlighting all the three points mentioned in the question. If users failed to understand the visual cues to highlight all the relevant points, they might move/highlight fewer than three points. We asked this question after Question 1 because an action with two or fewer highlighted points triggers an error pop-up in Andromeda Version A, prompting that the learning algorithm requires at least three highlighted points to determine new relative distances.

The results in Table 1 demonstrate a significant improvement with Version B. The results indicate that the visual cues helped participants learn to distinguish between the points that would be considered by the algorithm and those that would not. This is the primary factor influencing user decisions about which points need to be highlighted to update the layout.

We calculated the mean correctness score of all questions for each user to represent the aggregated correctness for each user. As shown in Table 2, participants using Version B had a 21% higher score than those who used Version A. The t-test (p-value $<$ 0.01) result indicates that the difference is statistically significant. We conclude that the newly-implemented features helped users bridge the cognitive gap and understand the need to highlight all of the relevant data points.

3.2. Understanding the WMDS dimensions

3.2.1. Problem

In a WMDS projection, the coordinate axes have complex meanings, representing the combination of multiple high-dimensional features. However, people with no WMDS knowledge are confused about the meaning of the horizontal and vertical coordinates of WMDS plots, often trying to match the 2D coordinate axes to specific dimensions as if they were traditional scatter plots. Because of this understanding gap, it becomes a problem for some users to address common tasks (Amar et al., 2005) such as finding a dimensional extremum, dimensional correlations, and the data distribution of a dimension.

Andromeda previously required users to interact with the parameter sliders to align the data points in the observation view if they wish to explore a single dimension or find the extremum in a specific dimension. After performing these parameter-tuning operations, the users could then hover the cursor over points to view their raw data on each respective parameter slider. A similar approach was necessary to explore dimensional correlations and data distribution. However, this method of viewing the dimension information is time-consuming, and users must also change the projection in the observation view to find the result. We saw the need for a more convenient and straightforward way to help users quickly identify dimension-related information.

3.2.2. Solution

Dimension-Assist Feature: Taking inspiration from the size encoding found in the Dust &Magnet system (Soo Yi et al., 2005), we designed a dimension-assist feature that enables a quick overview of feature distributions. When a user hovers the cursor over a parameter slider, such as the Spots slider in Fig. 5, the size of each circle within the observation view is mapped according to its Spots value such that the data points with larger Spots values have bigger circle sizes. By hovering the cursor over the parameter slider, users can quickly determine the data distribution of Spots, learning in the case of Fig. 5 that the Deer has the largest Spots value. With this feature, users can glance at the one-dimensional distribution without interacting with the parameter sliders or affecting the underlying model parameters. Furthermore, we hoped that this feature can help users to understand that the coordinate axes in the observation view have complex meanings.

3.2.3. Results

We conducted a pilot study to evaluate the effectiveness of the dimension-assist feature. Again, we used two versions of Andromeda: Version A and Version B. Version B enabled the new dimension-assist feature. In Version A, however, the dimension-assist feature is disabled. We ask eight undergraduate and graduate participants to complete an online survey regarding dimensional information. The participants span multiple disciplines, including computer science, business, data analytics, and chemistry. Of the eight participants, none consider themselves WMDS experts: three have learned and used MDS; three have heard of it but never used it; and two have never heard of it. We separated the participants randomly into two equal groups, one group for each Andromeda version.

The purpose of this pilot study is to assess the impact of the dimension-assist feature on four tasks: finding an extremum, characterizing a distribution, describing attribute correlations, and understanding the meanings of the WMDS plot axes. To further explore our results, we perform follow-up interviews with study participants. For this pilot study, we focus mainly on qualitative results, observing how the dimension-assist feature could impact user interactions in dimension-related tasks. From our observations and follow-up interviews, we find obvious advantages for the dimension-assist feature.

First, the dimension-assist feature helped understanding the results of PI (as defined in Section 2.1, interactions that enable direct manipulation of individual parameters). When users tried to find a dimension extremum, they increased the weight on the dimension of interest with the parameter slider. However, because the projection rotation has no meaning in WMDS, they did not know how the dimensions are oriented. By monitoring the participants' actions, we found that Version A users obtained the results using SLI (interactions that have no effect on model parameters) by hovering the cursor on the projection boundary points one by one and viewing the raw data in the parameter view. This was not only time-consuming, but also sometimes overlooked the extremum. Version B users, by contrast, quickly

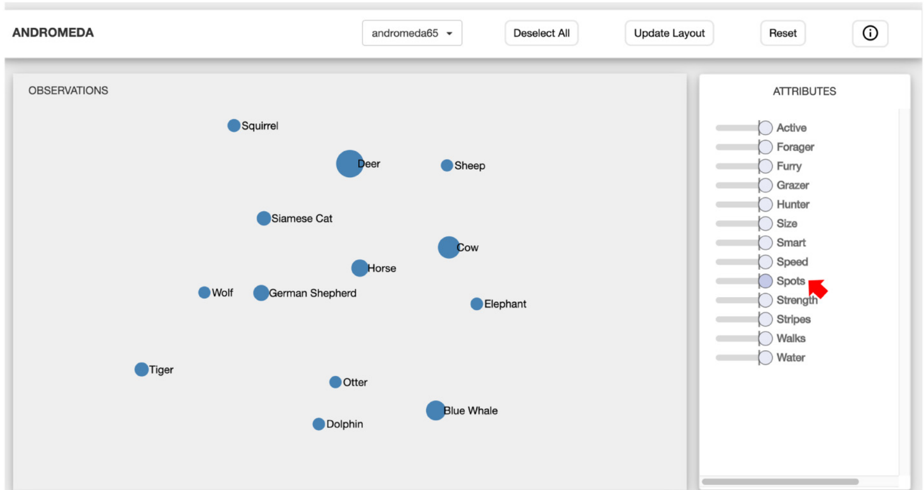

Fig. 5. When hovering on the Spots slider, the radius of each data point changes in proportion to its Spots values.

identified the dimensional orientation in the scatter plot generated by PI and identified the extremum with the dimension-assist feature.

Second, the dimension-assist feature had a shallower learning curve than PI. Two participants using Version A preferred to use SLI alone rather than PI to answer questions. One Version A participant even mentioned "I didn't realize how to use the dimension sliders until I almost finished the survey". All the participants using Version B obtained answers by using the dimension-assist feature solely or jointly with PI, except one participant who used PI instead of the dimension-assist feature for only one question. However, this participant quickly switched to using the dimension-assist feature for subsequent questions, which indicated the fast learning time of the dimension-assist feature.

Third, the dimension-assist feature was more efficient and accurate than SLI. When not using PI, Version A participants used SLI to obtain answers. They hovered on every point until they found the extremum point. However, it was easy to overlook the correct point using this strategy, which caused one Version A user to get the wrong answer. Whereas, with the dimension-assist feature, users quickly recognized the extremum as the point with the largest or smallest size.

Finally, the dimension-assist feature helped partially understand the dimensions within the projection. In the follow-up interview, participants used both Andromeda Version A and Version B, so that they could compare these two versions. All participants agreed that the dimension-assist feature made it easier to locate extrema and describe attribute correlations; however, one participant pointed out the shortcoming of the size-encoding, that "it is difficult to see the distribution without using [PI], especially with a large number of overlapping points in the visualization". Using dimension-assist, participants did also recognize that the data attributes were not straightforwardly mapped to the WMDS projection axes; however, none of the participants believed that the dimension-assist feature helped them understand the actual meanings of the WMDS axes. They stated that such understanding depended on background knowledge of the statistical methods.

3.3. Constructing/deconstructing clusters

3.3.1. Problem

The learning model in Andromeda required users to make opposing changes in the relative distances between selected points, meaning that some points must be moved closer together, while others must be moved further apart. This is because WMDS is scale invariant. However, in some circumstances, one of these relationships may be implicit. For example, when using Andromeda, users frequently try to create a cluster of a few points. They drag several points closer together, without moving any other point away from the cluster. As shown in Fig. 6(a), the user creates a cluster by dragging the German Shepherd, Otter, and Dolphin together. While the absolute distances have decreased, the relative distances have not changed much. It suggests only a global change in scale of the projection. So, the question is: with respect to which other distances in the projection should these three distances be reduced? But the user typically may not have any specific other points in mind for comparison.

This leads to a mismatch between the user's understanding and the WMDS model. In the updated visualization that follows the interaction (Fig. 6(b)), the model fails to capture the intention of the user to create a cluster, only focusing on the minor changes in the relative distance relationships between the three point. Thus, while the user intends that the Otter, German Shepherd, and Dolphin should be close, they are still spread out in the updated projection.

The same problem occurs when users try to deconstruct a cluster by separating a group of points. Unless users specify other points that move closer as a result, Andromeda fails to capture the intention of the user to deconstruct the cluster. Therefore, we needed to find an approach to convey clustering intent to the model so that it can react correctly to these common user manipulations.

3.3.2. Solution

Random Sampling: One possible solution is to consider the relative distance changes of all points, which means the clustered/unclustered points are moved with respect to all the other untouched points. It could solve the "constructing/deconstructing clusters" issue to some extent, but other issues would arise. First, considering the pairwise distances of all points requires significantly more expensive computation, sacrificing the efficiency of the interaction. Additionally, changes would be very minimal when dragging a few points in relation to a large number of untouched points, and the effect of the relative distance changes would be canceled out by a large amount of unchanged data. To minimize these shortcomings, we randomly select only m untouched points to represent all points in the background. This strategy offers a compromise between the advantages and disadvantages brought by considering all untouched points.

Fig. 6. The "constructing a cluster" usability issue: (a) Create a cluster by dragging the German Shepherd, Otter, and Dolphin together. (b) The interaction result in the original Andromeda system.

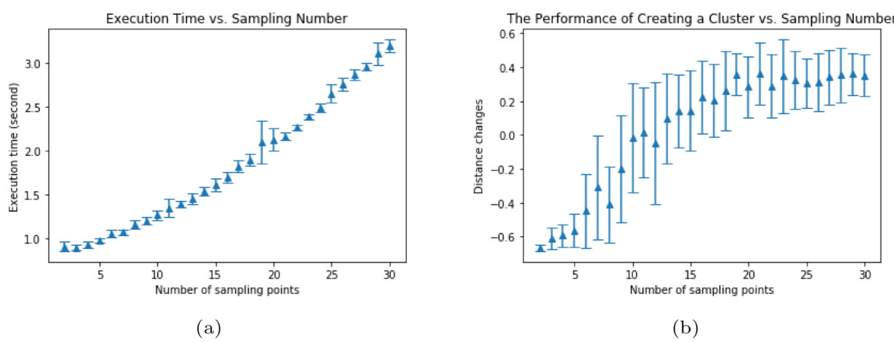

Fig. 7. (a) The execution times and (b) the absolute distance changes for the number of sampling points from 2 to 30.

The number of sampling points, m, significantly impacts the system's performance. To determine an optimal value of m, we experiment with the execution time and the clustering performance with different numbers of sampling points. As shown in Fig. 7(a), the execution time increases with the increasing number m due to the increasing demand for calculations of both the pairwise distances and the best-fit projection. Fig. 7(b) reveals the clustering performance for different numbers of sampling points m. A small value of distance change means a good clustering performance, whereas a large value of distance change means a bad clustering performance. As shown in the results, the distance values sharply increase along with the increasing number of sampling points, indicating that a larger m leads to worse clustering performance, with some instability beyond $m = 5$. Therefore, to achieve a better updated layout, we should select a small number of sampling points. To retain the relative distance between sampling points and achieve good time efficiency and clustering performance, we use three sampling points in the revised inverse WMDS computation.

As shown in Fig. 8, a user creates a cluster of the German Shepherd, the Otter, and the Dolphin. Once the user updates the layout, three points (the Sheep, the Elephant, and the Siamese cat) are randomly sampled and displayed with an orange border. Their relative distances will be considered along with those of the moved/highlighted points to project the new visualization with Eq. (6). In this equation, points i and j are the moved/highlighted points, r_i^* and r_j^* are their user-defined low-dimensional coordinates, and d_i and d_j are their high-dimensional coordinates. Points α and β are sampling points, r_α and r_β are their low-dimensional coordinates, and d_α and d_β are their high-dimensional coordinates. The $dist_L$ and $dist_H$ functions calculate the pairwise Euclidean distances in the low-dimensional and high-dimensional spaces. The pairwise distances between the moved/highlighted points and the sampling points are ignored. The generated new weight vectors, ω, will be used for the next forward WMDS pipeline to replot a new visualization.

$$\omega = \min_{\omega_1,\ldots,\omega_p} \Bigg[\sum_{i=1}^{n}\sum_{j>i}^{n}(dist_L(r_i^*, r_j^*) - dist_H(\omega, d_i, d_j))^2 + \sum_{\alpha=1}^{m}\sum_{\beta>\alpha}^{m}(dist_L(r_\alpha, r_\beta) - dist_H(\omega, d_\alpha, d_\beta))^2 \Bigg] \tag{6}$$

How does the system know when to apply this strategy? When the system detects the intent of a cluster construction, a cluster deconstruction, or a two-point-only manipulation, random sampling of other points is automatically applied. To detect these cases, before running WMDS, the system calculates all of the pairwise distance ratios, $\phi_{i,j}$. Points i and j are moved/highlighted points in the foreground. If all $\phi_{i,j} < 1$, all pairwise distances have decreased, then the system infers that the user is trying to create a cluster, and applies random sampling before the inverse WMDS computation. If all $\phi_{i,j} > 1$, all the pairwise distances increased, then the system infers that the user is trying to deconstruct a cluster, and applies random sampling before the inverse WMDS computation.

Likewise, in the original Andromeda, moving two points was not possible for the same reasons. Users were required to move at least three points, some closer and some further, to update the projections. This requirement prevented users from exploring the similarity or dissimilarity between only two observations. This problem is equivalent to the construct/deconstruct cluster problem with $n = 2$. With random sampling, two-point manipulation is enabled, making the Andromeda system more flexible for new exploration tasks.

3.3.3. Results

We performed a simulation analysis for our random sampling solution and compared the performance between Andromeda

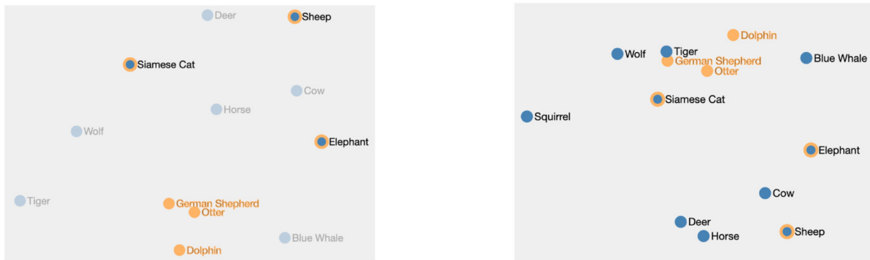

(a) (b)

Fig. 8. The random sampling solution: (a) Create a cluster by dragging the German Shepherd closer to the Otter and the Dolphin. After clicking the "Update Layout" button, three untouched points (the Sheep, Elephant, and Siamese Cat) are randomly sampled. (b) The model result showing the new cluster.

Versions A and B. Version B included the random-sampling features. Version A did not have the random sampling and the related model modifications. We tested combinations for constructing/deconstructing a cluster with varying numbers of moved points ranging from two to five. Because of the randomness of sampling points, the experiment is repeated 20 times for each combination to ensure reliability. The ability to move two points is not enabled in Andromeda Version A; therefore, we test this situation only in Version B.

In Fig. 9, the blue line represents the regression line for the new distances of cluster points in Version A, while the orange line is the regression line for the new distances of cluster points in Version B. All of the orange lines fall below the green original distance line, indicating that for all combinations in Version B, the mean values of their new distances are smaller than their original distances. Therefore, data points create a cluster in the updated layout in Version B. The orange regression lines for the new distances in Version B are lower than blue regression lines for the new distances in Version A, indicating that Version B performs better than Version A in creating clusters.

Fig. 10 displays the results for deconstruction of a cluster of two to five points. Again, most of the mean values for the new distances of Version B (orange points and lines) are above the green original distance line, indicating that most pairs in Version B are separated after updating the layout. However, as the original distances increase, the separation becomes less obvious, likely because these points are already at their furthest possible positions in the original visualization. By comparing the results of the Andromeda Version A and B for deconstructing a cluster, the orange regression lines for the new distances of Version B are higher than the blue regression lines of Version A. This indicates that Version B could move the cluster points further apart in the updated layout than could Version A.

Overall, the random sampling improves the Andromeda performance on constructing and deconstructing clusters. In addition, the change in the relative distance is missing for two-point groups in Andromeda Version A. Andromeda Version B, with its random sampling, addresses this limitation and enables two-point manipulations. Figs. 9(a) and 10(a) present the results for constructing and deconstructing a cluster of two points, displaying all the two-point combinations. The results indicate that Andromeda Version B, with random sampling, performs well in creating a cluster of two points.

When running the usability study discussed in Section 3.1.3, we asked participants to compare similarities between two points so that we could explore how enabling two-point interactions would impact participants' data exploration. Two-point OLI is only enabled in Version B. As Table 3 makes evident, 27 of the 28 participants who use Version B choose OLI to compare the similarity between two points. This result indicates that the participants prefer to use OLI to interact with points directly rather than use SLI to check attribute values in the parameter view. For Version A users, we observed that six participants

Table 3
Comparing two-point interactions in Andromeda Version A and Version B.

	SLI	OLI	
		Two points	Select additional points
Version A	20	0	6
Version B	1	27	0

(23.1%) randomly selected additional points in order to bypass the Version A limitation of selecting at least three points, which resembles our random sampling strategy, as participants began to select additional points in an attempt to better communicate their intent to that version of the system. We randomly select three points to represent all the untouched points. This result confirms that our random sampling solution is consistent with user cognition and is unlikely to create a subsequent disconnect between users and the model.

4. Discussion

The previous section presents three usability challenges with the Andromeda visual analytics system, as well as presenting and evaluating our solutions to resolve these challenges. Our overarching goal with these solutions is to bridge the gaps that existed between the users' mental model and the mathematical model of the Andromeda system. While the solutions that we proposed may not be the optimal visual encoding or interaction to address the challenge, we were able to resolve the major issues that we identified from previous studies. In this section, we generalize our solutions to broader lessons learned for the visual analytics community, revisit the "Two Black Boxes" problem, and discuss the limitations of our approach with accompanying future work.

4.1. Lessons learned for visual analytics

Our experience with addressing each of the three usability problems in this work incorporates some lessons learned that can apply to bridging cognitive gaps between human and model in visual analytics tools.

Identifying target points that are relevant/irrelevant for an interaction (the "With Respect to What" Problem (Wenskovitch et al., 2020)) is a general problem in many interactive dimensionality reduction and clustering applications, particularly in those that include semantic interaction (Endert et al., 2012a), V2PI (Leman et al., 2013), and/or OLI (Endert et al., 2011). A key observation is that the data relationships within a projection should be expressed in both a clear and usable way. This must be true for both directions of communication (Wenskovitch and North, 2020); if the communication channel from the AI to the human shows pairwise distances, then the communication channel from the human to the AI should also make use of

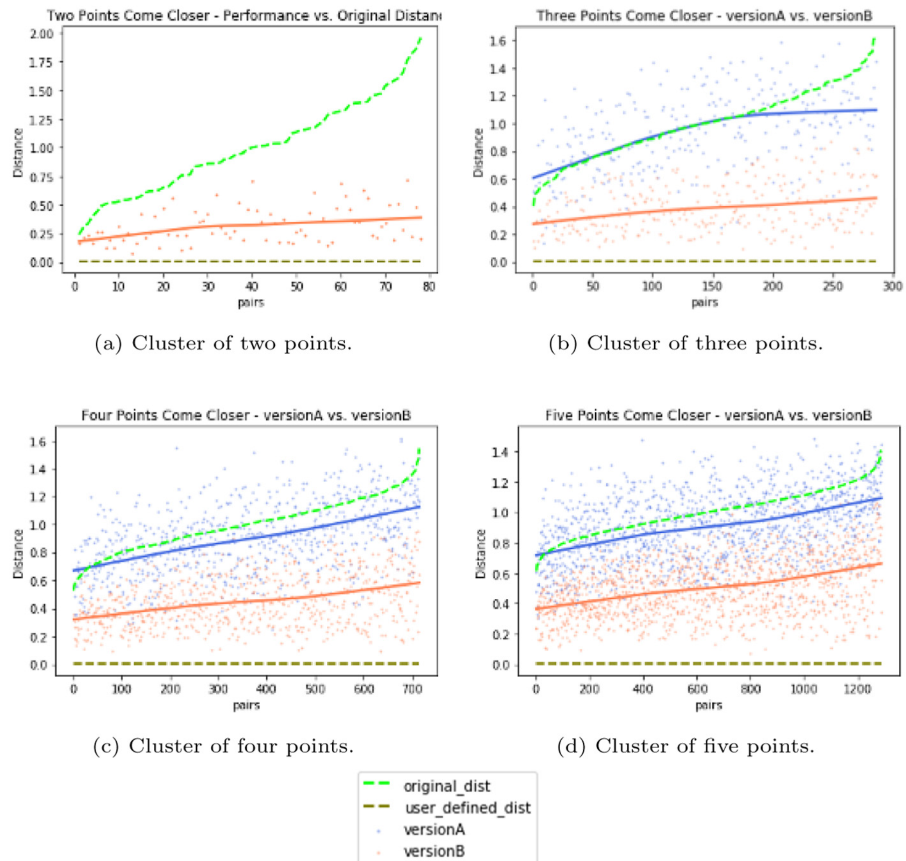

(a) Cluster of two points.

(b) Cluster of three points.

(c) Cluster of four points.

(d) Cluster of five points.

Fig. 9. Comparing the performances of Andromeda Version A and B in constructing a cluster of two to five points. The blue lines are the regression lines for mean values of the new distances of Version A (blue points). The orange lines are the regression lines for mean values of the new distances of Version B (orange points). The regression lines are calculated using a nonparametric lowess estimator.

pairwise distances. At times, the AI and the human may not be communicating the same information, such as when the user is expressing an interaction intent and the system responds with the relationships that it has inferred, so the data relationships must be made explicit. Our mapping of relevant/irrelevant points to the foreground/background visually represents one method that is immediately learnable to users without requiring training.

Presenting a means for users to understand the meaning of dimensionally-reduced spaces is also an active area of research in visual analytics (Wenskovitch and North, 2021). The sensitivity of nonlinear dimensionality reduction algorithms to parameters and hyperparameters is a known challenge for tuning and visually interpreting the resulting projection (Wattenberg et al., 2016). Tools such as DimReader (Faust et al., 2019) and CheckViz (Lespinats and Aupetit, 2011), among others (Aupetit, 2007), augment the visualization to add a level of explainability to the projection. In our case, simply showing one-dimensional sizes appeared to assist users in overcoming some common misunderstandings in the Andromeda data projections, but other methods such as biplots (Fry and Slifko, 2018) may still be needed in other applications to enhance the user's understanding of the visualization.

Underspecification is also a common problem in visual analytics, a fact that is especially true when the number of interactions is substantially more sparse than typical labeled data (Wenskovitch and North, 2017). Interaction machine learning applications can require complex labeling requirements that are not obvious, such as the need in the original Andromeda system to move three or more points while simultaneously expressing both similarly and dissimilarity relationships. However, these challenges often result from assumptions made in the machine learning algorithmic design, which may be different than the means by which a user thinks about the problem (Wenskovitch and North, 2021). Researchers should identify and study such common misunderstandings, developing new algorithms (such as the clustering case that we present in this work) to supplement and cover the non-obvious cases. Occasionally this may require systems to not rely on a single solution, but instead to integrate various edge case solutions.

4.2. "Two black boxes" revisited

The usability issues that we discuss in this work represent challenges both with communication from human to AI (the algorithm must understand the intent of the user's interaction) and from AI to human (the user must understand how their behavior will change the model). These two closed black boxes must be opened in order to maximize the performance of both actors when working together to explore a dataset. These usability issues are also representative of challenges that are common in visual analytics systems, including issues of transparency, explainability, interpretability, and verification of both interactions and models.

As visual analytics systems become more advanced in inferring user intent and responding appropriately to interactions, the role of human–machine teaming research will become more applicable to visualization researchers in this space. One of the main principles of successful human–machine teams is having both shared knowledge and shared awareness (Lyons and Havig, 2014). Improving communication between human and machine by opening their respective black boxes will further enable these shared features when exploring a dataset in an interactive machine learning tool such as Andromeda.

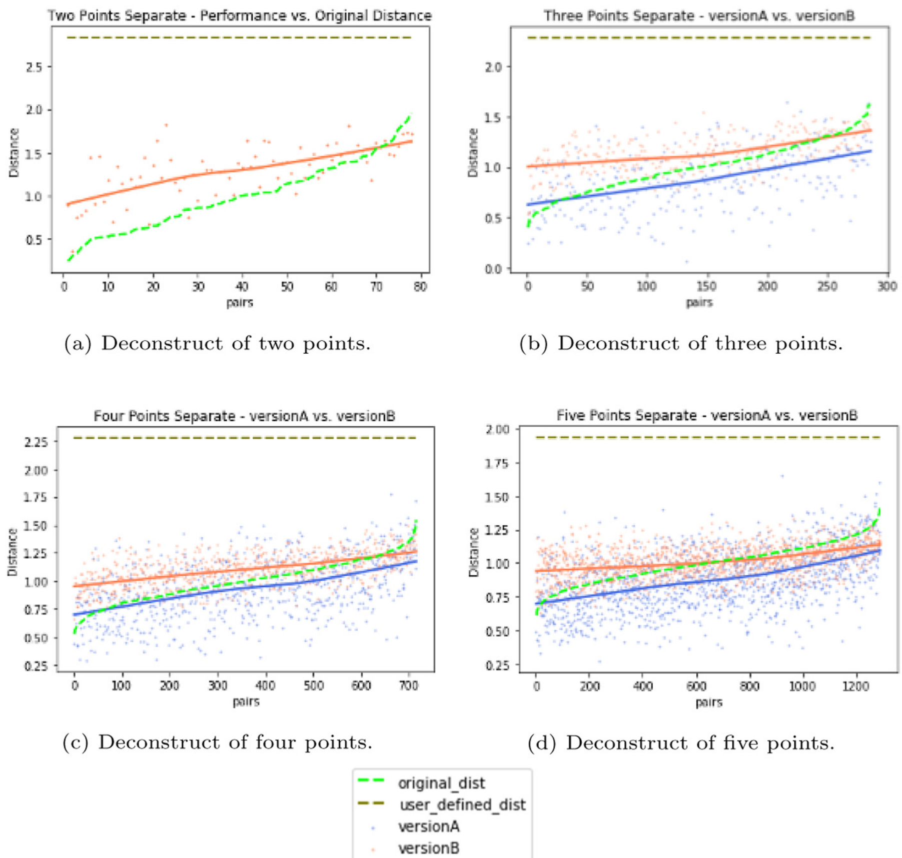

(a) Deconstruct of two points.

(b) Deconstruct of three points.

(c) Deconstruct of four points.

(d) Deconstruct of five points.

Fig. 10. Comparing the performances of Andromeda Version A and B in deconstructing a cluster of two to five points. The blue lines are the regression lines for mean values of the new distances of Version A (blue points). The orange lines are the regression lines for mean values of the new distances of Version B (orange points). The regression lines are calculated using a nonparametric lowess estimator.

4.3. Limitations and future work

In the pilot study for understanding the WMDS dimensions, we observe that the dimension-assist feature is somewhat helpful for characterizing the data distribution for one dimension, and only minimally improves the participants' understanding of the meanings of the WMDS plot axes. From these results, we conclude that the dimension-assist feature improves Andromeda's performance in some respects but does not meet all of our expectations. The main advantage of this dimension-assist feature was to aid users in understanding how the complex plot was transformed into simpler plots via PI, such as for finding extrema and identifying 2D correlations. We do not wish to overstate the value of this feature; it is simply one step toward making this tool more usable. In the future, we could introduce additional dimension related features to better minimize the understanding gap regarding WMDS dimensions.

Despite the positive results for the random sampling, it has limitations: its performance is unstable. For the same set of moved points, different sampling-point combinations heavily affect clustering performance. To investigate deeper, we evaluated the impacts of different sampling-point combinations, but were unable to discover any significant effects. Selecting sample points based on the original distances, new distances, or distance changes of sampling points did not significantly affect clustering performance in a consistent way. This suggests that it might be difficult to create heuristic methods for attempting to (non-randomly) select "good" sample points. Therefore, in this instance, the outcome still depends on the inverse WMDS model and whether it can find a best-fit projection that satisfies both the relative distances between highlighted/moved points and the relative distances between sampling points. An obvious partial solution would be to execute multiple runs of the sampling and inverse WMDS model, and then select the best clustering outcome. Future work could evaluate more impact factors to improve its performance. Additionally, the sudden introduction of the selected randomly-sampled points could be confusing for some users who are unfamiliar with the system. We consider it to be an acceptable tradeoff between introducing additional visual clutter and adding some explainability to the actions undertaken by the system, but other solutions are certainly possible.

5. Conclusion

This research focuses on designing and evaluating solutions for bridging gaps between user intent and model parameters. We study and discuss three major usability issues in the Andromeda system: highlighting relevant data points, understanding the WMDS dimensions, and constructing and deconstructing clusters. These three issues occur as the result of communication difficulties in both the algorithm misunderstanding the intent of the user and the user misunderstanding the actions of the algorithm.

To resolve these usability challenges, we implement and evaluate a number of modifications to the visualization and machine learning model in Andromeda. To resolve the uncertainty regarding which data points in the projection are used by the algorithm to learn and plot a new projection, we introduce two interface features – foreground and background view and distance lines – and evaluate these with a usability study in a classroom setting. To help users understand the dimensions of the WMDS plot, we introduce a dimension-assist feature that changes the size of data points according to their dimensional values. Finally, to improve Andromeda's performance when a user attempts to create or deconstruct a cluster, we introduce a random sampling strategy and modified learning objective function.

Our results indicate that the newly implemented features can help users realize that they need to select all the data points relevant to their intent. These features also help users in efficiently performing analysis tasks such as determining attribute extrema, identifying relationships between multiple attributes, and characterizing the distribution of attributes within the projection. Finally, our modifications to the model allow Andromeda to recognize and support the common user intent of clustering based on users' natural clustering interactions. In summary, each of our solutions performs well in addressing usability issues when compared to the original Andromeda interface, by aiding in bridging the cognitive gaps between user and model.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This work was supported in part by NSF grant CSSI-2003387 and NSF I/UCRC CNS-1822080 via the NSF Center for Space, High-performance, and Resilient Computing (SHREC).

References

Abello, J., Van Ham, F., Krishnan, N., 2006. Ask-graphview: A large scale graph visualization system. IEEE Trans. Vis. Comput. Graphics 12 (5), 669–676.

Alper, B., Riche, N., Ramos, G., Czerwinski, M., 2011. Design study of linesets, a novel set visualization technique. IEEE Trans. Vis. Comput. Graphics 17 (12), 2259–2267.

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. IEEE, pp. 131–138.

Amar, R., Eagan, J., Stasko, J., 2005. Low-level components of analytic activity in information visualization. In: IEEE Symposium on Information Visualization, 2005. INFOVIS 2005., pp. 111–117.

Ashaari, N.S., Judi, H.M., Mohamed, H., Wook, M.T., et al., 2011. Student's attitude towards statistics course. Procedia Social Behav. Sci. 18, 287–294.

Aupetit, M., 2007. Visualizing distortions and recovering topology in continuous projection techniques. Neurocomputing 70 (7–9), 1304–1330.

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). IEEE, pp. 163–172.

Broekens, J., Cocx, T., Kosters, W.A., 2006. Object-centered interactive multi-dimensional scaling: Ask the expert. In: Proceedings of the 18th Belgium-Netherlands Conference on Artificial Intelligence (BNAIC), pp. 59–66.

Brown, E.T., Liu, J., Brodley, C.E., Chang, R., 2012. Dis-function: Learning distance functions interactively. In: 2012 IEEE Conference on Visual Analytics Science and Technology (VAST). IEEE, pp. 83–92.

Chen, X., Self, J.Z., House, L., Wenskovitch, J., Sun, M., Wycoff, N., Evia, J.R., North, C., 2017. Be the data: embodied visual analytics. IEEE Trans. Learn. Technol. 11 (1), 81–95.

Choo, J., Bohn, S., Park, H., 2009. Two-stage framework for visualization of clustered high dimensional data. In: 2009 IEEE Symposium on Visual Analytics Science and Technology. IEEE, pp. 67–74.

Choo, J., Lee, H., Kihm, J., Park, H., 2010. Ivisclassifier: An interactive visual analytics system for classification based on supervised dimension reduction. In: 2010 IEEE Symposium on Visual Analytics Science and Technology. IEEE, pp. 27–34.

Cox, T.F., Cox, M.A., 2000. Multidimensional scaling. Springer.

Desjardins, M., MacGlashan, J., Ferraioli, J., 2007. Interactive visual clustering. In: Proceedings of the 12th International Conference on Intelligent User Interfaces, pp. 361–364.

Dowling, M., Wenskovitch, J., Fry, J.T., House, L., North, C., 2018b. SIRIUS: Dual, symmetric, interactive dimension reductions. IEEE Trans. Vis. Comput. Graphics 25 (1), 172–182.

Dowling, M., Wenskovitch, J., Hauck, P., Binford, A., Polys, N., North, C., 2018. A bidirectional pipeline for semantic interaction. In: Proc. Workshop on Machine Learning from User Interaction for Visualization and Analytics (At IEEE VIS 2018), Vol. 11, pp. 74–85.

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., Fiaux, P., North, C., 2012b. Semantic interaction for sensemaking: inferring analytical reasoning for model steering. IEEE Trans. Vis. Comput. Graphics 18 (12), 2879–2888.

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). IEEE, pp. 121–130.

Faust, R., Glickenstein, D., Scheidegger, C., 2019. Dimreader: Axis lines that explain non-linear projections. IEEE Trans. Vis. Comput. Graphics 25 (1), 481–490. http://dx.doi.org/10.1109/TVCG.2018.2865194.

Frutos, E., Galindo, M.P., Leiva, V., 2014. An interactive biplot implementation in r for modeling genotype-by-environment interaction. Stoch. Environ. Res. Risk Assess. 28 (7), 1629–1641.

Fry, J., Slifko, M., 2018. Generalized biplots for stress-based multidimensionally scaled projections. Comput. Statist. Data Anal. 128, 340–353.

Gunning, D., 2017. Explainable artificial intelligence (xai). Def. Adv. Res. Proj. Agency (DARPA), nd Web 2.

Ingram, S., Munzner, T., Irvine, V., Tory, M., Bergner, S., Möller, T., 2010. DimStiller: Workflows for dimensional analysis and reduction. In: 2010 IEEE Symposium on Visual Analytics Science and Technology. IEEE, pp. 3–10.

Krug, S., 2000. Don't make me think!: a common sense approach to Web usability. Pearson Education India.

Kulesza, T., Burnett, M., Wong, W., Stumpf, S., 2015. Principles of explanatory debugging to personalize interactive machine learning. In: Proceedings of the 20th International Conference on Intelligent User Interfaces, pp. 126–137.

La Grange, A., Le Roux, N., Gardner-Lubbe, S., 2009. BiplotGUI: interactive biplots in r. J. Stat. Softw. 30 (12), 1–37.

Lampert, C.H., Nickisch, H., Harmeling, S., Weidmann, J., 2009. Animals with attributes: A dataset for attribute based classification.

Leman, S., House, L., Maiti, D., Endert, A., North, C., 2013. Visual to parametric interaction (v2pi). PLoS One 8 (3).

Lespinats, S., Aupetit, M., 2011. Checkviz: Sanity check and topological clues for linear and non-linear mappings. In: Computer Graphics Forum. 30, Wiley Online Library, pp. 113–125.

Lyons, J.B., Havig, P.R., 2014. Transparency in a human-machine context: Approaches for fostering shared awareness/intent. In: International Conference on Virtual, Augmented and Mixed Reality. Springer, pp. 181–190.

May, T., Bannach, A., Davey, J., Ruppert, T., Kohlhammer, J., 2011. Guiding feature subset selection with an interactive visualization. In: 2011 IEEE Conference on Visual Analytics Science and Technology (VAST). IEEE, pp. 111–120.

Paulovich, F.V., Eler, D.M., Poco, J., Botha, C.P., Minghim, R., Nonato, L.G., 2011. Piece wise laplacian-based projection for interactive data exploration and organization. In: Computer Graphics Forum, Vol. 30. Wiley Online Library, pp. 1091–1100.

dos Santos Amorim, E.P., Brazil, E.V., Daniels, J., Joia, P., Nonato, L.G., Sousa, M.C., 2012. ILAMP: Exploring high-dimensional spacing through backward multidimensional projection. In: 2012 IEEE Conference on Visual Analytics Science and Technology (VAST). IEEE, pp. 53–62.

Self, J.Z., 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), http://dx.doi.org/10.1145/3158230.

Self, J.Z., House, L., Leman, S., North, C., 2015a. Andromeda: Observation-level and parametric interaction for exploratory data analysis. Tech. rep., Department of Computer Science, Virginia Tech, Blacksburg.

Self, J.Z., House, L., Leman, S., North, C., 2016. Designing usable interactive visual analytics tools for dimension reduction. In: CHI 2016 Workshop on Human-Centered Machine Learning (HCML), p. 7.

Self, J.Z., Hu, X., House, L., North, C., 2015b. Designing for interactive dimension reduction visual analytics tools to explore high-dimensional data. Technical Report, Department of Computer Science, Virginia Tech, Blacksburg.

Self, J.Z., Vinayagam, R.K., Fry, J.T., 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.

Sharko, J., Grinstein, G., Marx, K.A., 2008. Vectorized radviz and its application to multiple cluster datasets. IEEE Trans. Vis. Comput. Graphics 14 (6), 1427–1444.

Soo Yi, J., Melton, R., Stasko, J., Jacko, J.A., 2005. Dust & magnet: Multivariate information visualization using a magnet metaphor. Inf. Vis. 4, 239–256.

Stahnke, J., Dörk, M., Müller, B., Thom, A., 2015. Probing projections: Interaction techniques for interpreting arrangements and errors of dimensionality reductions. IEEE Trans. Vis. Comput. Graphics 22 (1), 629–638.

Udina, F., 2005. Interactive biplot construction. J. Stat. Softw. 13 (5), 1–16.

Wattenberg, M., Viégas, F., Johnson, I., 2016. How to use t-SNE effectively. Distill 1 (10), e2.

Wenskovitch, J., Dowling, M., North, C., 2019. Simultaneous interaction with dimension reduction and clustering projections. In: Proceedings of the 24th International Conference on Intelligent User Interfaces: Companion, pp. 89–90.

Wenskovitch, J., Dowling, M., North, C., 2020. With respect to what? Simultaneous interaction with dimension reduction and clustering projections. In: Proceedings of the 25th International Conference on Intelligent User Interfaces. In: IUI '20, Association for Computing Machinery, New York, NY, USA, pp. 177–188. http://dx.doi.org/10.1145/3377325.3377516.

Wenskovitch, J., North, C., 2017. Observation-level interaction with clustering and dimension reduction algorithms. In: Proceedings of the 2nd Workshop on Human-in-the-Loop Data Analytics. In: HILDA'17, Association for Computing Machinery, New York, NY, USA, pp. 1–6. http://dx.doi.org/10.1145/3077257.3077259, https://doi.org/10.1145/3077257.3077259.

Wenskovitch, J., North, C., 2019. Pollux: Interactive cluster-first projections of high-dimensional data. In: 2019 IEEE Visualization in Data Science (VDS). IEEE, pp. 38–47.

Wenskovitch, J., North, C., 2020. Interactive artificial intelligence: Designing for the 'two black boxes' problem. Computer 53 (8), 29–39. http://dx.doi.org/10.1109/MC.2020.2996416.

Wenskovitch, J., North, C., 2021. An examination of grouping and spatial organization tasks for high-dimensional data exploration. IEEE Trans. Vis. Comput. Graphics 27 (2), http://dx.doi.org/10.1109/TVCG.2020.3028890.

Wong, P.C., Hetzler, B., Posse, C., Whiting, M., Havre, S., Cramer, N., Shah, A., Singhal, M., Turner, A., Thomas, J., 2004. IN-SPIRE InfoVis 2004 contest entry. In: IEEE Symposium on Information Visualization. IEEE, p. r2.

Zeitz, J., Self, N., House, L., Evia, J.R., Leman, S., North, C., 2018. Bringing interactive visual analytics to the classroom for developing EDA skills. J. Comput. Sci. Colleges 33 (3), 115–125.