

CAVIR: Correspondence Analysis in Virtual Reality

Andrea Bönsch*, Frederik Graff[‡], Daniel Bündgens*, Torsten Kuhlen*

*Virtual Reality Group
RWTH Aachen University
52074 Aachen

[‡]RWTH School of Business & Economics
RWTH Aachen University
52074 Aachen

E-Mail: boensch@vr.rwth-aachen.de E-Mail: frederik.graff@org.rwth-aachen.de

Abstract: Correspondence Analysis (CA) is used to interpret correlations between categorical variables in the areas of social science and market research. To do so, coherences of variables are converted to a three-dimensional point cloud and plotted as several different 2D-mappings, each containing two axes. The major challenge is to correctly interpret these plottings. Due to a missing axis, distances can easily be under- or overestimated. This can lead to a misinterpretation and thus a misclustering of data.

To address this problem we present CAVIR, an approach for CA in Virtual Reality. It supports users with a three-dimensional representation of the point cloud and different options to show additional information, to measure Euclidean distances, and to cluster points. Besides, the motion parallax and a free rotation of the entire point cloud enable the CA expert to always have a correct view of the data.

Keywords: interaction techniques, user interfaces, exploratory data analysis, correspondence analysis, virtual reality, clustering

1 Introduction

Correspondence Analysis (CA) is a descriptive dimensionality reduction method of multi-variable statistics which allows a vivid graphical representation of complex correlations of two (or more) categorical variables as a point cloud in a (theoretically) three- (or more) dimensional space. Similarities between variables are converted into distances on three or more dimensions, and the positions of the variables, represented by points, are converted into coordinates. It is presumed that there is a dependency between rows and columns which can be explained by latent variables. Later on, the variables will be presented as a three-dimensional point cloud. The three axes can then intuitively be interpreted as these latent variables. CA is widely used in, i.a., social science, psychology, medicine, and, due to its relatively convenient graphical procedure, in the area of market research [BEPW06].

1.1 Key Concepts of Correspondence Analysis

The key concepts of CA, which are necessary for understanding the functionality and terminology of Correspondence Analysis in Virtual Reality (CAVIR), are briefly introduced. For more information, see [Mor12], [Ben77] or [BEPW06]; for a short introduction, see [NG07].

The database consists of frequencies n_{ij} in a contingency table K with I rows R_i and J columns C_j . With $N = \sum_i \sum_j n_{ij}$ and $p_{ij} = \frac{n_{ij}}{N}$, P is the matrix of relative frequencies.

$$P = \begin{pmatrix} p_{11} & \cdots & p_{1J} \\ \vdots & \ddots & \vdots \\ p_{I1} & \cdots & p_{IJ} \end{pmatrix}.$$

We define $r_i = \sum_{j=1}^J p_{ij}$ and $c_j = \sum_{i=1}^I p_{ij}$, as the masses of rows and columns. The frequencies in the cells are normalized with r_i and c_j respectively. As a result we get two matrices

$$R = \begin{pmatrix} \frac{p_{11}}{r_1} & \cdots & \frac{p_{1J}}{r_1} \\ \vdots & \ddots & \vdots \\ \frac{p_{I1}}{r_I} & \cdots & \frac{p_{IJ}}{r_I} \end{pmatrix} \text{ and } C = \begin{pmatrix} \frac{p_{11}}{c_1} & \cdots & \frac{p_{1J}}{c_J} \\ \vdots & \ddots & \vdots \\ \frac{p_{I1}}{c_1} & \cdots & \frac{p_{IJ}}{c_J} \end{pmatrix}.$$

Each row of R is called the i^{th} row profile of K . Each column of C is called the j^{th} column profile of K [Mor12]. The columns, interpreted as J axes, put up a J -dimensional space C_J , which the row profiles should be plotted in, analogously the rows. Obviously, the J -/ I -dimensionality causes problems. Plotting more than three-dimensional spaces makes an understanding of the mapping more difficult. Due to this, the number of dimensions of C_J and R_I is reduced to three (C^3, R^3) and both are integrated into one coordinate system, whereby three conditions have to hold:

1. The more similar the profiles are, the closer the points representing the respective row/column should lie to each other in the coordinate system. The distance is measured using the Euclidean distance.
2. The reduction of the overall variance of K over all p_{ij} , the so-called Total Inertia T , should be minimal, while T is given as $T = \sum_{i=1}^I \sum_{j=1}^J \frac{(p_{ij} - r_i * c_j)^2}{r_i * c_j}$.
3. The additive decomposition of T should be maximal [Mor12].

After further calculations we get XYZ-coordinates for each row and column. The axes of the integrated space can be interpreted as latent variables that explain the variance of the conditional frequencies.

1.2 Challenges of Plotting CA Point Clouds

The reason why CA is so widely used in social science and market research lies, i.a., in the way complex correlations between categorical variables can be intuitively interpreted using “plotting” (cf. fig. 1(a)). With plotting, sometimes called “mapping”, the researcher interprets the relative distances of the points from each other and their relative positions on the axes. Most often, one column point (■ in fig. 1(a)) makes up more or less the center of a sub-cloud consisting of several row points (● in fig. 1(a)). In market research, the column

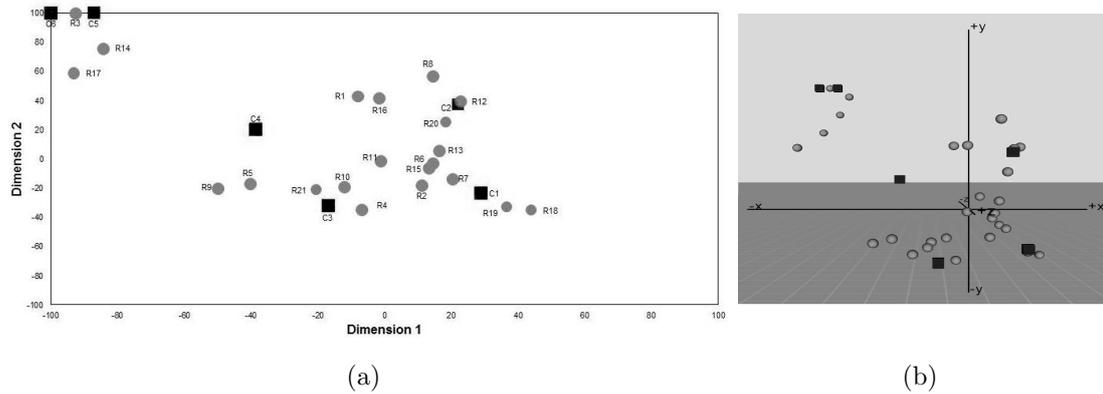


Figure 1: Comparison of the CA point cloud representation in a two-axis coordinate system with a simple 2D-mapping (layout similar to SPSS) (a) and in a three-axis coordinate system with our VR-based approach (b).

points usually represent variations of a product, brands, or firms, and the row points represent features, ratings, or (potential) purchasers [BEPW06]. Column points will henceforth be referred to as “trait carriers”, and row points as “traits”. Sometimes, researchers use the plotting to cluster elements graphically.

For a complete plotting of results, a coordinate system with three axes is required. This poses some difficulties, as a valid graphical representation may ultimately be only two dimensional. Plotting software solutions (e.g. SPSS, mapwise, XGobi) therefore usually offer the option of plotting three different mappings, where one of the three axes is omitted. Hereby, both an XY-, an XZ- as well as a YZ-coordinate system are produced [SCB98, NG07]. For the interpretation of results this may cause problems because only two axes can be considered simultaneously. This is why practitioners often choose the axes with the highest explained variance and ignore the third one and thereby neglect important distance information. To sum up, the main problems with non-three-dimensional plotting of a CA point cloud are:

1. Overestimation of distances
2. Underestimation of distances
3. Faulty clustering
4. Faulty interpretation

1.3 Approach

Several attempts to take account of all the three axes graphically have led to unsatisfactory results. Especially, the graphical summary of points into clusters is impossible when distance information is not realized validly as distances but rather as perspectives or as point characteristics. In the work presented, the authors have developed CAVIR, a tool that enables the graphical analysis and interpretation of CA point clouds in a virtual environment (VE),

providing a valid distortion-free impression of the spatial distances between points. To meet the challenges mentioned above, the key approaches to a solution are:

1. A virtual three-dimensional model of the CA point cloud
2. Free navigation through the VE
3. Arbitrary rotation opportunities of the point cloud to provide a valid impression of the spatial distances between points
4. Diverse options to display labels
5. Unhindered scaling of axes
6. Measuring of distances between points
7. Clustering of points

The rest of this paper is structured as follows: section 2 presents relevant related work. In section 3 we describe the technical realization of our approach. Section 4 covers the required handling of two different Virtual Reality displays. Finally, we give a brief conclusion and an outlook on future work in section 5.

2 Related Work

Van Dam et al. already stated that the gap between the rate of data generation and the capacities to analyze this data is widening. As a possible solution they propose Immersive Virtual Reality (IVR) technology, which combines interactive visualization with immersive sensation [vDFL⁺00].

Arns et al. integrated 3D-scatterplots into a CAVE-like VE [ACNC99]. The authors performed a user study which contrasted their methods with the well-known XGobi system [SCB98]. Interestingly, users performed better with the VE but felt more comfortable when working with the classical desktop-based tool.

Another approach that uses a “grand tour” view embedded in a VE to show raw data as well as clustering results is presented by Yang [Yan99]. The same author specifically addresses the problem of overdrawing, and presents a solution based on direct volume visualization, namely a splatting approach for scatterplots and similar views [Yan03].

Very little research has yet been done on alternative ways of plotting the point cloud of a CA. [BEPW06] gives an introduction in the method itself and its mathematical procedures, pointing out that if more than two axes are needed to achieve a satisfactory explained variance, the researcher has to decide which two are the most important. These should be plotted. However, the authors admit that a two-axes-plotting is never possible without a loss of information.

Whitlark and Smith propose chi-square residuals to measure item distances (so-called attribute-brand relationships) and to thus escape the dimensionality problem. Previously

they had stated: “Relying on a two-dimensional map may be risky. In our experience, it is rare to see a two-dimensional map tell a complete or even an accurate story.” [WS01].

The statistical software package R, as described by [NG07], provides a non-immersive three-dimensional display of CA point clouds, which alleviates (but does not totally solve) the problem of perceptual distortion. Helpful analytical tools, such as brushing, labeling, distance measuring, and clustering, are not included in R.

The first application of Virtual Reality (VR) on CA was done by [MMG⁺03]. In the context of a dermatological study, they plotted point clouds along three dimensions of skin characteristics using a stereoscopic display. Point clustering was ex ante implemented by a hierarchical cluster analysis [MMG⁺03]. However, we want to enable the researcher to cluster the points manually within the model. In contrast to existing CA analysing and visualizing tools, the approach presented here proposes an immersive 3D-display providing the researcher with a distortion-free impression of spatial distances as well as with new interaction methods for valid interpretations and clustering.

3 Technical Realization

In this section we explain the set-up of our approach to enable CA.

CA is a complex process divided into several substeps, but in this paper we focus on one of the last steps, namely clustering and the supportive techniques. Due to the fact that the approach should be embedded into the workflow of CA experts, the format of our input and output files is determined by the notation of SPSS [ASDVV09], one of the standard programs in the field of CA. This enables a quick and uncomplicated switch between both working areas.

Our initial situation is an input file containing a list of data points described by three attributes

- one character **cMemshp** classifying the data point,
- one string **sDesc** giving a concrete point label and
- three numerical values defining the position of a point **P** in the 3D space.

The data points are divided into two sets, the traits **A_Trait** and the trait carriers **B_Carrier**. The data points’ membership of one of these two sets is encoded in *cMemshp*. A unique natural number indicates the inclusion in *A_Trait*; unique small letters, accordingly, in *B_Carrier*. To be able to cluster the data points meaningfully, users need to be able to easily distinguish both sets during exploration in the VE. For this purpose, different shapes, analogously to the representations in standard 2D-CA-programs, are used: a sphere is chosen for the traits and a box for the trait carriers (cf. fig. 1(b)). With this similarity, irritations on the user side will be avoided. Then, the representations are shown at the respective point *P*. A Cartesian coordinate system used as a reference frame (see 3.3) completes the basic set-up.

Due to the fact that the point cloud is relatively small compared to the complete VE, we added a tiled reference plane as a “ground floor” to give some spatial cues. This supports the user during exploration of the scene via navigation [BKLP04]. Here, we offer two techniques: a pointing-metaphor for free travel through the VE and a trackball metaphor for free rotation of the point cloud.

3.1 Interaction

In addition to navigation, different interaction techniques are offered, based on IntenSelect [dHKP05] and simple ray intersection calculations. We scale the represented point clouds to enable an accurate selection with both techniques of a single point, even in a tight clouds. To give some additional help by all selection tasks, the position of the input device is displayed by either a beam or a small sphere, as will be explained in section 4.

To enable an intuitive and well-defined handling, each technique provided is bound to one application mode. A menu hierarchy is used to for unhindered switching between the modes and their submodes during runtime. We choose a so-called *pie menu* [CHWS88]. Callahan et al. compared *pie menus* to linear menus and discovered, that the circular arrangement of the entries around the menu center reduces the target seek-time and lowers the error rate. Thus, such a circular menu enables the user to work more efficiently.

Besides the layout, we try to achieve an effective working by using just a small menu hierarchy. The first menu consist of four entries, showing the four available modes, presented in the following. Each entry has a submenu for mode specific tasks. Only in urgent cases those entries have another submenus.

To always indicate which mode is currently active, we show a tooltip that provides all the necessary information. Additionally, a color coding of the floor tile borders and the scene background is used. These colors are also linked to the respective circular menu’s entries to quickly ensure a correct handling.

Users should always start with an exploration of the point cloud, so no application mode is chosen in the beginning. Besides the different shapes, we add a color coding for an easy and quick identification of the points’ membership. Traits are visualized in red and trait carriers in green. This color association is reused as often as possible in the modes.

3.2 Labeling

Besides the general set-up, additional information is required for an effective clustering process. The challenge is to position this information in an appropriate way. Superimpositions of the point cloud have to be avoided and quick access to the information has to be ensured. Annotations are a standard technique for solving these problems. Besides readability, more important aspects concerning the layout have to be taken into account. Those are described in [AHS05].

We use textual annotations, as shown in figure 2(a), that are based on the annotation library IDEA [PHW⁺10], which can be faded in or out for each point individually in the

“label” mode. Our layout refers to the column-wise “Flush-Left-Right Layout” from [AHS05]. We can strictly separate the scene and the labels to always enable an uncovered view of the pure data. In contrast to Ali et al., we use segmented connection lines between the anchor point of P and the respective label. To ensure a clear allocation, the anchor point coordinates are exactly the same as those of P . Due to both definitions, we reduce the space used for the labels and show a more compact setting. Nevertheless, it becomes difficult to find one particular label while displaying a huge amount of labels. Users have to trace the connection line from a data point to its label or they have to read the labels to find out whether they belong to a trait or a trait carrier. For simplification, we use the red-green coloring here, not only for the points but also for the labels.

We provide three label types, given in the initial input file, which are used to support the user during clustering.

1. The unique id *cMemshp* allows an explicit and a fast identification of data points.
2. Sometimes it may be useful to know what information a point represents, i.e. which product or key features it stands for. This information is contained in the attribute *sDesc*.
3. The explicit point position is also helpful during clustering to estimate point distances without using the distance mode, or to find outliers.

Users can switch between those label types via different entries in the circular menu.

3.3 Coordinate System

To quickly understand the distribution of the point cloud, we display a Cartesian coordinate system with labeled axes. Due to this, the origin of the point cloud can be clearly identified and this allows a first, rough classification of the distribution. To allow extraction of more information, the length of the individual positive and negative axis intercepts can be set to one of three predefined values using the mode “coordinate system”.

- To quickly understand the spatial distribution, the axes can be adapted to the **maximal absolute values** of the point cloud extended by a user specified value. Short axes imply small point distances, while large axes indicates an extensive distribution or outliers.
- In CA, points are often deemed to be outliers if they have at least one absolute value greater than a maximal bound. According to the concrete application, those values may be treated differently, e.g. they may be clamped to a special value. In our approach a **maximal bound** can be predefined by the user to enable the identification of outliers.
- In addition to the *maximal bound*, a **minimal bound** can be set. This value will be predefined by the user and used to quickly identify points in or over a special bound.

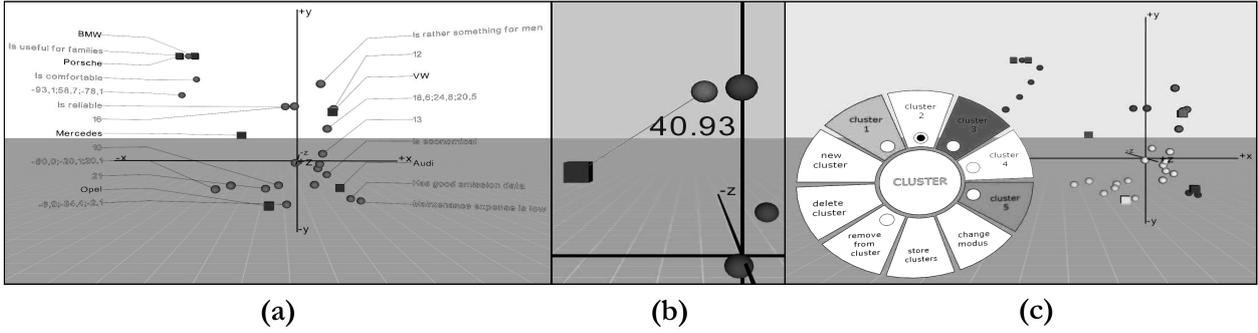


Figure 2: (a) The label layout consists of two columns next to the point cloud and different labels are offered, which can be selected individually for each data point. (b) A measured distance is visualized via a connecting line and a centered text label. (c) One example to divide a given point cloud into four clusters, currently adding points to cluster 2.

Besides scaling the axes, users can add an object-oriented bounding box to identify inliers and outliers more clearly.

3.4 Distance Measurement

The main aim of our approach is to facilitate clustering and to reduce errors during this process. One criterion used to find the correct cluster for a data point is that of comparing the distances between several points. This can be done in the mode “distance”. Users are asked to select two different data points P_i and P_j , between which the Euclidean distance will be computed. The Euclidean distance itself is calculated based on the exact coordinates given in the input file.

Figure 2(b) illustrates the distance visualization: a connection line between both selected points, and a centered label with the computed distance. To maximize the clarity of which points are selected, all points are inked in gray in the beginning. After selecting a data point, it is recolored to red or green according to the set it belongs to.

Points can be deselected and inked in gray by the user or the system. Users can select a marked point again to deselect it, or just choose a third point P_k . In this case, the system will automatically deselect the first two points P_i and P_j . Owing to this, just one distance is shown at the same time.

3.5 Clustering

In the “clustering” mode, the whole interaction is managed by the circular submenu, shown in figure 2(c). Thereby, the compliance of two rules is always ensured:

R1: The total number of clusters is greater than zero and less or equal to $|B_{carrier}|$.

R2: In each cluster there has to be at least one trait carrier.

Three interaction types are offered in this mode:

- a) Users can add and delete clusters, complying with R1. For each cluster an individual menu entry is provided, containing a radio button.
- b) Via brushing [WR05], points can be linked to an active cluster according to R2 or they can be removed from any cluster.
- c) The clustering can be stored to the filesystem. The format of the output file is determined as dummy codes, to enable the analyst to go on working in her used surroundings after using the VE to cluster the point cloud.

If the user selects a cluster entry in the circular menu, linking of data points to the according cluster is activated. This is shown by selection of the respective radio button and by the tooltip. Adding points to the cluster is done by brushing: first, the points are focused via *IntenSelect* [dHKP05] and then added by clicking. To be able to comply to R2, the first point added to a new cluster has to be a trait carrier. We again use colors to show the set membership. All points that are not yet linked to a cluster are gray. Each cluster has a unique color, which is used for the cluster entry in the circular menu and the linked points. Additionally, it is used for the representation of the input device to give the user some additional help besides the tooltip and the radio buttons. If the user adds a point to a wrong cluster, the circular menu provides an entry for deleting individual points from clusters. With this option, the points are recolored gray. If a single trait carrier in a cluster is deselected, all points of this cluster are deselected, too. With this automatic deselection, R2 is always ensured to hold.

4 Challenges of Different VR Displays

The aim of CAVIR is to accelerate the process of interpreting and clustering and thereby reduce the potential clustering errors described in section 1. [ACNC99] already showed in a user study that VEs support this requirement, because a third dimension is available. Users are able to make better decisions, due to a better and not misleading view of the data [WS01].

Our approach should be available for two different display system classes. On the one hand, we want to use a five-sided CAVE-like environment to enable a collaboration with several people during the cluster process and the possibility to freely walk around the point cloud. On the other hand, we want to provide a semi-immersive low-cost and quick possibility to cluster the point cloud in the office in order to embed its use into the normal workflow. Therefore, we decided to use a 3D monitor, as shown in 3(a).

The low-cost solution has multiple advantages. First, it is easy to handle for novice users, who are often not familiar with VR techniques. Second, all the input devices known from standard desktop systems can be used and thus the interaction can be maintained similarly to the usual work environment. Due to this, the learning phase before users can handle the application correctly is shortened. Another advantage is that we do not need head-tracking.

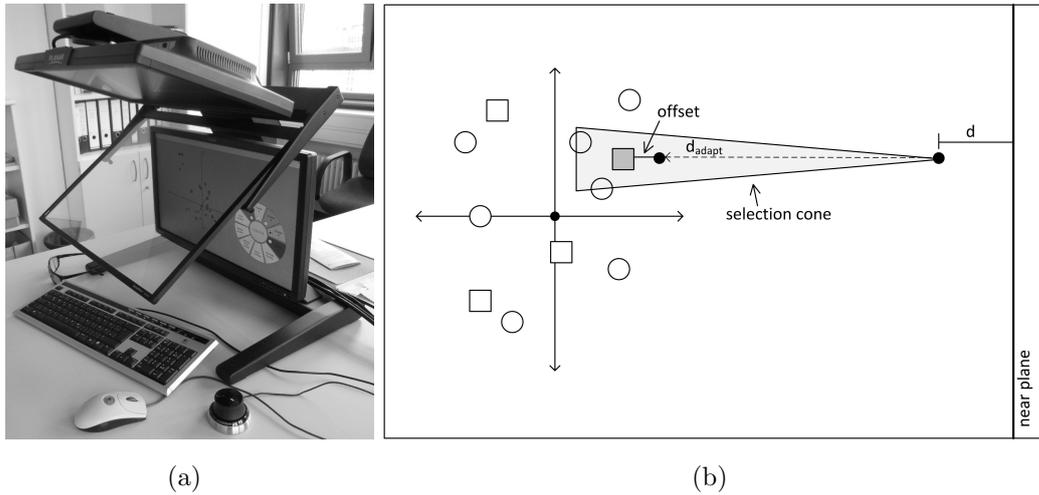


Figure 3: Using CAVIR on a 3D monitor without head-tracking (a) requires a technique to adapt the input device representation depth, as shown in the schematic drawing (b).

The motion parallax in combination with the stereo projection is enough to have the required correct view of the data. Nevertheless, this causes a difficulty. While we use a beam pointer in the CAVE-like environment for input device representation, the beam seems to directly penetrate the user’s head between her eyes at the low-cost device. This is disconcerting and leads to a headache after a short time of usage.

Our solution is a depth-adaptation into the viewing direction of the input device representation, influenced by the focused part of the scene. This is schematically shown in figure 3(b). Normally, the input device representation would be shown in distance d behind the near plane. Now, we figure out which data point is in the focus of our IntenSelect cone and we determine its exact position. Afterwards, we translate the representation in viewing direction to the focused data point. This distance d_{adapt} is computed by taking the difference between the old position and the data point position and subtracting a small offset. In addition to choosing a sphere representation rather than of the beam, our approach avoids any unwell feeling.

During movements of the input device, this technique leads to small jumps of the sphere if the focused data points lie in different depths. In a pre-test, this did not turn out to be disturbing. If no data point is focused, the depth is not adapted and the circular menu entries can be selected without problems.

5 Discussion and Future Work

We have presented a VR-based technique to ease valid interpretation and to realize clustering - one of the last steps in correspondence analysis (CA). The standard CA programs mostly provide 2D-views of the data, leading easily to misinterpretations and misclustering. None offers a clustering option. In the virtual environment (VE) we have the advantage of showing

the data three-dimensionally as a point cloud in a Cartesian coordinate system. Thus, a correct view is always provided and perceptual distortion avoided, which accelerates the clustering process and at the same time reduces the risk of misinterpretations and invalid clustering. Different navigation and interaction techniques are used to support the clustering, which is done by linking points to certain clusters via brushing.

During the implementation process we had close contact to one CA expert, who pretested all features. However, we plan to evaluate our approach in a user study to identify issues which need to be addressed for improvement, especially concerning usability. Besides, more features will be added, such as techniques to identify and handle outliers correctly, according to the concrete task; a way of dealing with data points, which should be assigned to at least two different clusters; and a bounding volume for each cluster in order to have an additional membership cue on top of colors.

Acknowledgements

The approach described in this article was funded by the Interdisciplinary Management Practice (IMP) Pathfinder Funding.

References

- [ACNC99] Laura Arns, Carolina Cruz-Neira, and Dianne Cook. The Benefits of Statistical Visualization in an Immersive Environment. In *Proceedings of the IEEE Virtual Reality, VR '99*, pages 88 – 95, 1999.
- [AHS05] Kamran Ali, Knut Hartmann, and Thomas Strothotte. Label Layout for Interactive 3D Illustrations. *Journal of the WSCG*, 13:1–8, 2005.
- [ASDVV09] R. Akin, N.R. Sharpe, R. De Veaux, and P. Velleman. *SPSS Manual for Business Statistics*. Pearson Education, 2009.
- [Ben77] J.P. Bensécri. Histoire et Préhistoire de l'Analyse des Données. In *l'Analyse des Correspondances. Les Cahiers de l'Analyse des Données*, volume 5, pages 9–40, 1977.
- [BEPW06] Klaus Backhaus, Bernd Erichson, Wulff Plinke, and Rolf Weiber. *Multivariate Analysemethoden. Eine Anwendungsorientierte Einführung*. Springer-Lehrbuch. Springer, 11., überarb. Aufl. edition, 2006.
- [BKLP04] Doug A. Bowman, Ernst Kruijff, Joseph J. LaViola, and Ivan Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, 2004.
- [CHWS88] J. Callahan, D. Hopkins, M. Weiser, and B. Shneiderman. An Empirical Comparison of Pie vs. Linear Menus. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 95–100, 1988.

- [dHKP05] G. de Haan, M. Koutek, and F. Post. IntenSelect: Using Dynamic Object Rating for Assisting 3D Object Selection. In *Proceedings of the 9th International Immersive Projection Technology Workshop and 11th Eurographics Workshop on Virtual Environment*, pages 201–209, 2005.
- [MMG⁺03] N. Monmarche, H. Marteau, J.-P. Gerard, C. Guinot, and G. Venturini. Interactive Mining of Multimedia Databases with Virtual Reality. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4756, pages 367–374, 2003.
- [Mor12] U. Mortensen. Einführung in die Korrespondenzanalyse. <http://uwmortensen.de/caneu1c.pdf>, 11.06.2012.
- [NG07] Oleg Nenadic and Michael Greenacre. Correspondence Analysis in R, with Two- and Three-dimensional Graphics: The CA Package. *Journal of Statistical Software*, 20(3):1–13, 2007.
- [PHW⁺10] Sebastian Pick, Bernd Hentschel, Marc Wolter, Irene Tedjo-Palczynski, and Torsten Kuhlen. Automated Positioning of Annotations in Immersive Virtual Environments. In *the Joint Virtual Reality Conference of EuroVR - EGVE - VEC*, pages 1–8, 2010.
- [SCB98] Deborah F. Swayne, Dianne Cook, and Andreas Buja. Interactive Dynamic Data Visualization in the X Window System. *Journal of Computational and Graphical Statistics*, 7:pp. 113–130, 1998.
- [vDFL⁺00] A. van Dam, A.S. Forsberg, D.H. Laidlaw, Jr. LaViola, J.J., and R.M. Simpson. Immersive VR for scientific Visualization: a Progress Report. *Computer Graphics and Applications, IEEE*, 20(6):26–52, nov/dec 2000.
- [WR05] Michael A. E. Wright and Jonathan C. Roberts. Click and Brush: A Novel Way of Finding Correlations and Relationships in Visualizations. In *Proceedings Theory and Practice of Computer Graphics*, pages 179–186, 2005.
- [WS01] David B. Whitlark and Scott M. Smith. Using Correspondence Analysis to Map Relationships. *Marketing Research*, 13(3):22–27, 2001.
- [Yan99] Li Yang. 3D Grand Tour for Multidimensional Data and Clusters. In *Advances in Intelligent Data Analysis*, volume 1642 of *Lecture Notes in Computer Science*, pages 173–184. Springer Berlin / Heidelberg, 1999.
- [Yan03] Li Yang. Visual Exploration of Large Relational Data Sets Through 3d Projections and Footprint Splatting. *Knowledge and Data Engineering, IEEE Transactions on*, 15(6):1460–1471, nov.-dec. 2003.