EuroVis Workshop on Visual Analytics (2020)
K. Vrotsou and C. Turkay (Editors)

Enhanced Attribute-Based Explanations of Multidimensional Projections

D. van Driel[1] and X. Zhai[1] and Z. Tian[2] and A. Telea[2]

[1]Bernoulli Institute, University of Groningen, Netherlands
[2]Department of Information and Computing Sciences, Utrecht University, Netherlands

Abstract

Multidimensional projections (MPs) are established tools for exploring the structure of high-dimensional datasets to reveal groups of similar observations. For optimal usage, MPs can be augmented with mechanisms that explain what such points have in common that makes them similar. We extend the set of such explanatory instruments by two new techniques. First, we compute and encode the local dimensionality of the data in the projection, thereby showing areas where the MP can be well explained by a few latent variables. Secondly, we compute and display local attribute correlations, thereby helping the user to discover alternative explanations for the underlying phenomenon. We implement our explanatory tools using an image-based approach, which is efficient to compute, scales well visually for large and dense MP scatterplots, and can handle any projection technique. We demonstrate our approach using several datasets.

CCS Concepts

• ***Human-centered computing*** → ***Visualization design and evaluation methods;***

1. Introduction

Multidimensional Projections (MPs) are used to create scatterplot-like depictions of high-dimensional data. While quite successful in showing the data structure (in terms of point clusters and outliers), MPs do not explain which data aspects (*e.g.*, dimensions, dimension ranges, or interactions between dimensions) create such structures. As such, MP visualizations, typically created using 2D or 3D scatterplots, need to be enhanced by various explanations to enable the analyst to understand the observed data patterns.

Several such explanatory mechanisms exist. Color-coding points is the easiest and most widespread, but requires manual dimension selection and does not scale well for many dimensions. *Global* explanations, such as biplot axes [Gre10, GLR11] and axis legends [BBT13, CMN*16] are easy to use, but do not provide local explanations for different subsets of data points. Linked views and tooltips show *local* explanations, but require one to manually select structures of interest. *Image-based* techniques offer a middle way: They compute and display local explanations everywhere on the projection, not requiring the user to select specific point subsets. Image-based techniques have been also used to analyze projection errors [Aup07, SvLB10, MCMT14, PPM*15]. They scale well visually and computationally, are clutter-free, and can generically handle any (high-dimensional) dataset.

Closest to our aim, da Silva *et al.* [dSRM*15] proposed an image-based explanation of MPs that colors every scatterplot point by the dimension that contributes most to making data points in that neighborhood similar. In this paper, we extend the approach in [dSRM*15] with explanations that describe the data from additional viewpoints besides attribute similarity. First, we use principal component analysis (PCA) to analyze point neighborhoods to deduce and encode the local (intrinsic) dimensionality of the data. This allows users to separate regions of high intrinsic dimensionality in the projection (hard to explain by just a few dimensions) from low-dimensionality regions where such explanations are feasible. Secondly, we analyze point neighborhoods to detect and depict strong linear relationships (not captured by PCA) between dimensions. Our techniques complement existing mechanisms for projection explanation, can be computed efficiently on the GPU, and can be applied generically on any high-dimensional dataset visualized by any MP technique. We illustrate our work on both synthetic (ground-truth) and non-synthetic datasets.

2. Explanatory Mechanisms

Let $D = \{\mathbf{x}_i\} \subset \mathbb{R}^n$ be a n-dimensional dataset with points $\mathbf{x}_i = (x_i^1, \ldots, x_i^n)$. A projection P maps D to a scatterplot $P(D) \subset \mathbb{R}^m$, where typically $m \in \{2,3\}$. To allow analysts to recover the structure of D from its visualization $P(D)$, projections typically aim to place points that are similar in D close to each other in $P(D)$. Similarity can be computed in different ways, *e.g.* based on $\mathbb{R}^n$ distances [TDL00, DT04, JCC*11] or $\mathbb{R}^n$ neighborhoods [vH08, MHM18]. Recent surveys provide more details on the technicalities of MPs [NA18, EMK*19]. Image-based explanatory techniques exploit precisely this distance or neighborhood preservation property of MPs: Let $\nu_i \subset P(D)$, $\nu_i = \{\mathbf{y} \in P(D) | \|\mathbf{y} - \mathbf{y}_i\| \leq \rho\}$, be a neighborhood of size ρ of scatterplot point $\mathbf{y}_i \in P(D)$. Since points in ν_i are, by construction, close, and since P is expected to (reasonably) preserve similarities, the points $\mu_i \subset D$ that project to ν_i are expected to be similar. Hence, it makes sense to compute an *explanation* of μ_i and next visually encode this on all scatterplot points $\mathbf{y}_i$.

Da Silva *et al.* [dSRM*15] propose two such explanations. Let $\lambda^j_{\mathbf{x},\mathbf{x}'} = \|\mathbf{x}^j - \mathbf{x}'^j\|_1^2 / \|\mathbf{x} - \mathbf{x}'\|_n^2$ be the contribution of dimension j

DOI: 10.2312/eurova.20201084

to the distance between two points $\mathbf{x}$ and $\mathbf{x}'$ in D, with $\|\cdot\|_k$ being the k-dimensional Euclidean distance. This point-pair contribution is extended to neighborhoods μ_i by averaging the local contributions of $\mathbf{x}_i$ and all its neighbors, as $\overline{\lambda}_i^j = \sum_{\mathbf{x}\in\mu_i} \lambda^j_{\mathbf{x},\mathbf{x}_i}/|\mu_i|$, where $|\cdot|$ denotes set size. Finally, such distance contributions are normalized yielding

$$\lambda_i^j = \frac{\overline{\lambda}_i^j/\gamma^j}{\sum_{j=1}^n \left(\overline{\lambda}_i^j/\gamma^j\right)}, \tag{1}$$

where the normalization factor γ^j is the contribution $\overline{\lambda}^j$ of dimension j of the whole dataset D with respect to its centroid. Due to normalization, $\lambda_i^j \in [0,1]$, with lower values telling dimensions that contribute *little* to distances in μ_i, *i.e.*, explain well why points in μ_i are *similar*. An alternative to Eqn. 1 is to compute the relative variance v_i^j of dimension j over the neighborhood μ_i as

$$v_i^j = \frac{LV_i^j/GV^j}{\sum_{j=1}^n (LV_i^j/GV^j)}, \tag{2}$$

where LV_i^j is the variance of dimension j for all points in μ_i, normalized by the variance GV^j of the same dimension over all points in D. Just as λ_i^j, $v_i^j \in [0,1]$, with lower values telling dimensions that change little in a neighborhood. The scatterplot $P(D)$ is explained by color-coding its points by the C dimensions that have overall low values of λ_i^j (or v_i^j, depending on the user's choice) over all points. C is set to a low value, *e.g.* 8, since categorical colormaps should be small. Luminance is used to encode the *confidence* in the visual explanation: If j is the dimension picked to color point i, confidence κ is computed as the sum of λ_i^j (or v_i^j) values for all points in the neighborhood μ_i, normalized by the sum of the same terms over *all* dimensions over μ_i. If neighbors of point i are best explained by the same dimension j as i, the color will appear bright, and conversely.

Figure 1a shows a dataset of 3000 points spread over three faces of an axis-aligned cube (with added noise), projected with PCA to 2D, explained by variance (image from [dSRM*15]; explaining by distance creates a similar image). The image is simple to decode: Points on each cube face share very similar values of one of the dimensions, so are bright and colored by the respective dimension. Points along cube edges are dark, since two (or three, for the cube corner) dimensions are needed to explain their similarity with neighbors. It is important to see that these are *not* the original data dimensions, but dimensions synthesized by PCA (eigenvectors). Hence, we are not interested to showing their identities, but only their number.

2.1. Adding Dimensionality Explanation

Da Silva *et al.*'s explanations (Eqns. 1 and 2) cannot provide full insights into the structure of high-dimensional data. Take *e.g.* a non-axis-aligned cube like in Fig. 1a and embed it into a high-dimensional space. While the data structure stays the same, both distance contributions and variances cannot select a single dimension to explain the cube's faces, since all dimensions contribute to the data structure.

We first propose to improve this by explaining the *local* (or intrinsic) *dimensionality*. For each neighborhood μ_i of each point $\mathbf{x}_i \in D$, we compute the n eigenvalues α_i of its covariance matrix, sorted decreasingly. From these, we explored computing the local dimensionality δ and its confidence κ in different ways (see also Tab. 1):

Total variance: We define dimensionality δ as the minimal number of largest eigenvalues $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_\delta$ needed to explain a user-set fraction θ of the data variance in μ_i. The confidence κ equals how much the sum of these largest δ eigenvalues deviates from the mean of all n eigenvalues.

Definition	Dimensionality δ	Confidence κ
Total variance	$\min \delta \left\| \frac{\sum_{i=1}^{\delta} \alpha_i}{\sum_{i=1}^{n} \alpha_i} \geq \theta \right.$	$1 - \frac{\sum_{i=1}^{\delta} \alpha_i - \overline{\alpha}}{\sum_{i=1}^{n} \alpha_i}$
Minimal variance	$\left\| \left\{ \frac{\alpha_i}{\sum_{j=1}^{n} \alpha_j} \geq \theta, 1 \leq i \leq n \right\} \right\|$	$\frac{\sum_{i=1}^{\delta} \alpha_i}{\sum_{i=1}^{n} \alpha_i}$

Table 1: *Definitions of local dimensionality and confidence.*

Minimal variance: The total variance model works well when eigenvalues significantly drop. However, consider the (limit) case where all eigenvalues are equal. The dimensionality δ computed by total variance equals then θ/n, even though locally the data is truly n-dimensional. To capture this, we define δ as the number of eigenvalues larger than a minimal user-set variance θ, and confidence κ as the sum of these divided by total variance, similar to Kaiser's criterion used in explanatory factor analysis [Cli88, Jol02].

Figure 1b,c show the total and minimal variance explanations for the noisy cube data. As in [dSRM*15], the explanations are color-coded on the projection points are color-coded, as detailed in a legend next to the projection. The legend bars' sizes indicate how many points are assigned a given explanation (dimensionality). The cube's faces are colored blue. The legend tells that these points are locally in $\delta = 2$-dimensional neighborhoods embedded in nD. Close and on the cube edges, the color is green, telling that $\delta = 3$ dimensions are needed to explain the data here.

2.2. Adding Correlation Explanation

High-dimensional data is often explained by how its dimensions *correlate*. Yet, assessing *global* correlation over an entire dataset is of limited value when the underlying phenomenon is a mix of local (linear) patterns. To address this, we compute and depict correlations over *local* neighborhoods. For each point neighborhood μ_i, we compute the $N = n(n+1)/2$ Pearson or Spearman correlations between all dimension-pairs $(j,k) \in [\![1,n]\!] \times [\![1,n]\!]$. We sort these pairs in descending correlation-strength order, and select the C top-ranked pairs that are most frequent over all points i. This resembles selecting the explaining dimensions in [dSRM*15], but now we select dimension-pairs rather than individual dimensions. We show these C pairs via a categorical colormap, using luminance to map the absolute correlation values. Figure 1c shows this for the noisy cube. The legend tells that the three faces map to strong correlations of the three dimensions x, y, and z, as expected. The edges orthogonal to faces show the same correlation. Indeed, for the face xy, for instance, the orthogonal edge has near-constant x and y, and strongly varying z, values, so x and y are correlated along it.

This visualization can only show the C top-ranked, most frequent, correlations from all possible N ones. However, users may want to examine the presence (or absence) of *specific* correlations. For this, we show the entire set of N dimension-pairs using a matrix view. Figure 2a shows this view, next to the projection, computed by t-SNE, for the *concrete* dataset [Yeh98, Lic13], that has 1030 8-dimensional samples measuring how 8 ingredients influence concrete strength. The matrix cells are colored by the same colormap as in the projection. Here, dark blue tells all dimension-pairs whose correlations have a frequency higher than zero but lower than the C top-ranked pairs.

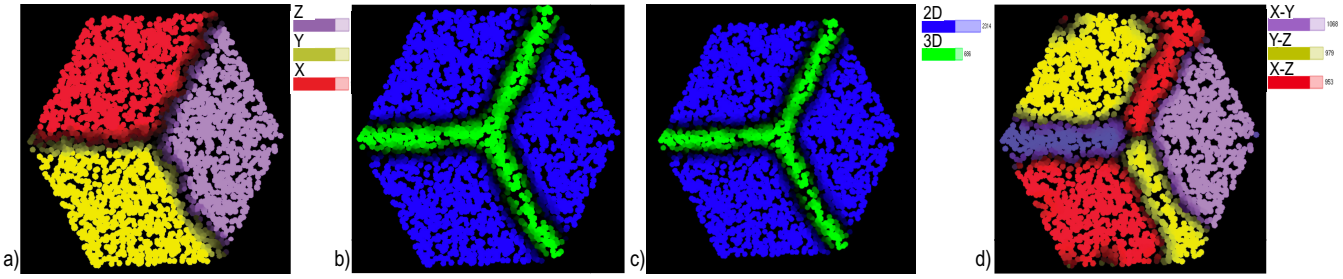

Figure 1: *Cube dataset explained by (a) variance; local dimensionality with total (b) and minimal (c) variance; and (d) dimensions correlation.*

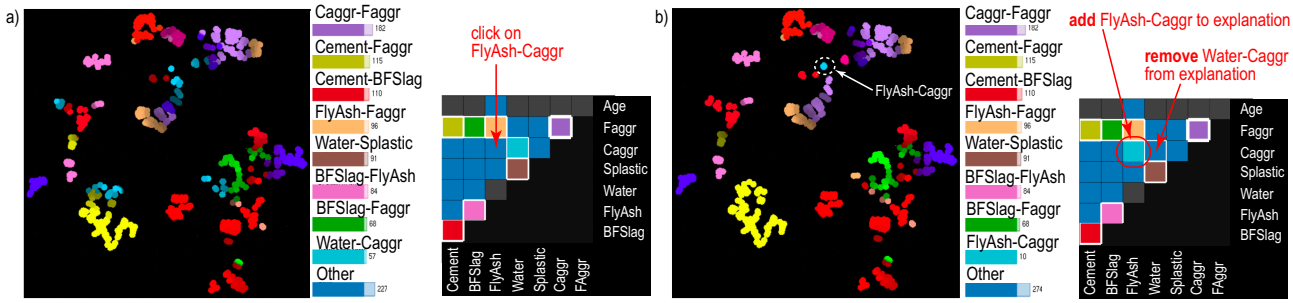

Figure 2: *Matrix view,* concrete *dataset. Clicking on the* FlyAsh-Caggr *cell (a) allocates a color to it, showing where in the projection view these two variables are strongly correlated. To make room for this, the weakest-correlated pair* Water-Caggr *is removed from the explanation.*

To see where, on the projection plot, such pairs correlate, the user clicks a dark blue cell, *e.g. FlyAsh-Caggr* dimension-pair in Fig. 2a. The color used for the C^{th} top dimension-pair, *i.e.* cyan used for *Water-Caggr*, is then allocated to the clicked pair and the C^{th} pair is made dark blue. The result shows a single cyan spot in the projection plot (Fig. 2b, dashed circle marker). This is the only place where the dimensions *FlyAsh* and *Caggr* strongly correlate.

The matrix view supports two other tasks. First, the cells of the top C (strongest correlated) dimension-pairs are outlined in white. This helps one to easily return to the original color allocation after having selected some other dimension-pairs for explanation. Secondly, rows and/or columns having many cells colored by the non-default (dark blue) color indicate *groups* of strongly correlated variables. For instance, the second top row in the matrix view in Fig. 2a, corresponding to the *Faggr* dimension, shows four such cells, that indicate its strong correlation with *Cement* (yellow), *BFSlag* (green), *FlyAsh* (orange), and *Caggr* (purple), respectively.

3. Applications

We consider the *wine* dataset, which has 6497 samples of Portuguese *vinho verde* [CCA*09], each with $n = 12$ physicochemical attributes such as acidity, residual sugar, and alcohol rate. Figure 3a shows the raw projection of this dataset using LAMP [JCC*11]. Besides a dense-point cluster bottom-right, there is not much else this image tells us. While other projection methods, *e.g.* t-SNE, may show better separated clusters, the question still remains how to explain these.

Figures 3b-c show the contribution and variance explanations respectively. These are quite similar and split the projection roughly into four areas, explained by small variations of alcohol (purple), chlorides (yellow), sugar (red), and acidity (beige), respectively. The correlation view (Fig. 3d) brings additional insights: We see a large purple area bottom-right that matches well the area earlier explained by small variations of chlorides, alcohol, and acidity. Over this purple area, the legend of image (d) tells that sugar and density strongly correlate. Also, we see that the red area in Figs. 3b-c, where sugar has a low variation, is now roughly split in Fig. 3d into smaller areas – red (fixed acidity-citric acid correlation), yellow (fixed acidity-pH correlation), beige (fixed acidity-density correlation), and brown (chlorides-density correlation). Note that the contribution-variance and correlation explanations are *complementary*: They cannot, when taken separately, split the projection into fine-grained local explanations, but do so when *combined*. Indeed, the red area in Figs. 3b-c is further split (explained) by using correlation, as explained above; conversely, the purple area in Fig. 3d is further split (explained) by using contribution or variance.

At this point, the analyst may wonder which projection areas are *sufficiently* explained by the above views. The dimensionality view helps here. Figure 3e shows the local dimensionality of the projected data, computed by total variance (Sec. 2.1). We see how increasingly more dimensions are needed to capture increasing fractions $\theta \in [0.3, 0.9]$ of the total variance – in the limit, we need all $n = 12$ dimensions to explain $\theta = 100\%$ of the variance. More interestingly, we see in Fig. 3e a gradient of local dimensionality, from highest in the bottom-right area (red-purple colors for $\theta \geq 0.85$) to blue in the top-left area (blue for $\theta \leq 0.75$). Besides color hue, the local dimensionality gradient is also visible in the brightness, which tells the confidence κ that the color-coded number of dimensions locally explain θ percent of the variance. The effect is very similar to the enridged contour maps used to visualize scalar fields [vWT01]: The visual nesting of the 'cushions' created by varying brightness conveys the absolute value of the encoded signal, *i.e.*, the local dimensionality. The way we compute these cushions (Sec. 2.1) is, however, completely different to [vWT01].

The local dimensionality view helps interpreting the contribution-variance and correlation views as follows: As we have seen, local

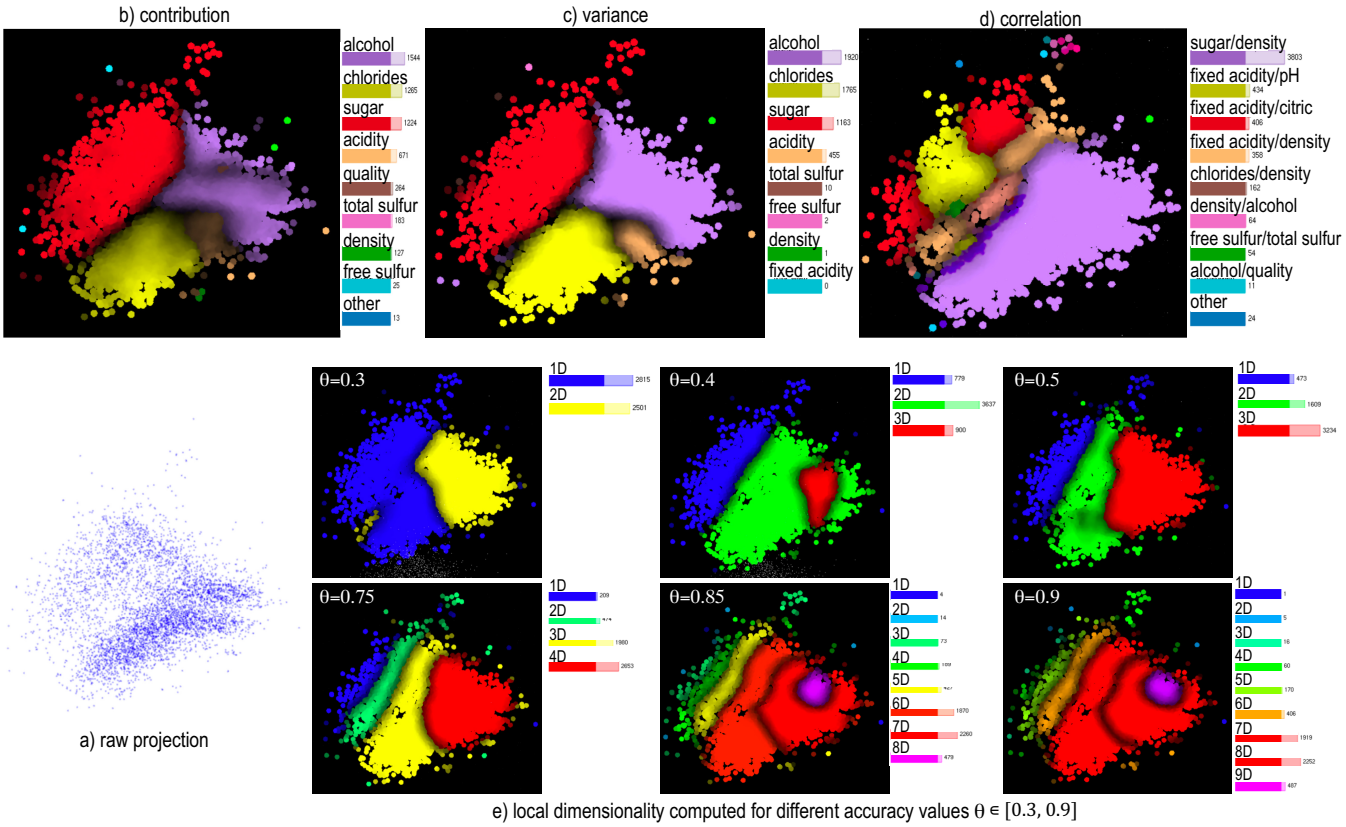

Figure 3: *Explanation of* wine *dataset by contribution (b), variance (c), correlation (d), and local dimensionality (e).*

dimensionality is high in the bottom-right (red-purple) area, where we need 7 to 9 dimensions to explain $\theta = 0.85$ of the data variance. In this area, the contribution-variance and correlation views jointly give us information about only *five* variables – alcohol, chlorides, acidity, sugar, and density. Hence, these two views do not *fully* explain this area, so we need to search for more explanations here. In contrast, the local dimensionality is low in the top-left (blue) area, where we can explain $\theta = 0.75$ of the data variance by a single dimension. From the contribution-variance views, we see that this area is well explained by a small variance of sugar. Hence, in this area, sugar's low variance is sufficient to explain the data.

4. Discussion and Conclusions

We detail several points on our method, as follows.

Genericity and scalability: Our method can handle any type of quantitative data projected by any MP technique. Correlations and PCA are computed with the Eigen library [eig20]. Since explanations are computed (and rendered) independently on local point neighborhoods, which is easily parallelizable. We did this using NVIDIA's CUDA, thereby generating all images in this paper in real time for datasets up to tens of thousands of points, tens of dimensions, on a commodity PC (2.4 GHz CPU, GeForce 320M GPU).

Limitations: Like its predecessor [dSRM*15], our method can *technically* handle datasets of any dimensionality n. Yet, as the local intrinsic dimensionality grows, one needs more variables for the explanation. Also, the correlation is $O(n^2)$ in computation and space needed for the dimension matrix (Fig. 2). Hence, our method works well up to a few tens of dimensions in practice. Separately, both our explanation and [dSRM*15] rely on the assumption that points close in $P(D)$ correspond to points close in D. The extent to which various MP techniques realize that varies [EMK*19]. One way to address this is to use projection error views [MCMT14] to exclude neighborhoods which do not respect this condition, or refine their computation by *e.g.* using larger radii ρ. Another option is to compute our explanations using the neighborhoods μ_i in D rather than ν_i in $P(D)$. Investigating all above options, and comparing the results with our current ones, is an interesting option that we consider for future work.

Several extensions are next possible. Adding more explanation types, such as inverse correlation, correlation of more than two dimensions, or the presence of specific nD data patterns, is a low hanging fruit. We aim to compute, in parallel, a wide range of local explanations based on a pattern library, and next show the most salient ones in the final view, thereby combining the so far separate contribution, variance, correlation, and dimensionality views. This would perform a scagnostics-like [WAG05] local analysis of the projection, but using patterns described by the high-dimensional data rather than by the scatterplot itself. Finally, computing a hierarchical explanation, where regions are recursively split based on additional explanations, is another direction we aim to pursue.

Acknowledgments: Z. Tian was supported by the China Scholarship Council under grant 201906080046.

References

[Aup07] AUPETIT M.: Visualizing distortions and recovering topology in continuous projection techniques. *Neurocomputing 10*, 7-9 (2007), 1304–1330. 1

[BBT13] BROEKSEMA B., BAUDEL T., TELEA A.: Visual analysis of multidimensional categorical datasets. *Computer Graphics Forum 32*, 8 (2013), 158–169. 1

[CCA*09] CORTEZ P., CERDEIRA A., ALMEIDA F., MATOS T., REIS J.: Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems 47*, 4 (2009), 547–553. 3

[Cli88] CLIFF N.: The eigenvalues-greater-than-one rule and the reliability of components. *Psychological Bulletin 103*, 2 (1988), 276–279. 2

[CMN*16] COIMBRA D., MARTINS R., NEVES T., TELEA A., PAULOVICH F.: Explaining three-dimensional dimensionality reduction plots. *Information Visualization 15*, 2 (2016), 154–172. 1

[dSRM*15] DA SILVA R., RAUBER P., MARTINS R., MINGHIM R., TELEA A.: Attribute-based visual explanation of multidimensional projections. In *Proc. EuroVA* (2015), pp. 97–101. 1, 2, 4

[DT04] DE SILVA V., TENENBAUM J. B.: *Sparse multidimensional scaling using landmark points*. Tech. rep., Stanford University, 2004. 1

[eig20] Eigen numerical library, 2020. http://eigen.tuxfamily.org. 4

[EMK*19] ESPADOTO M., MARTINS R., KERREN A., HIRATA N., TELEA A.: Towards a quantitative survey of dimension reduction techniques. *IEEE TVCG* (2019). doi:10.1109/TVCG.2019.2944182. 1, 4

[GLR11] GOWER J., LUBBE S., ROUX N.: *Understanding biplots*. Wiley, 2011. 1

[Gre10] GREENACRE M.: *Biplots in practice*. Fundacion BBVA, Bilbao, 2010. 1

[JCC*11] JOIA P., COIMBRA D., CUMINATO J. A., PAULOVICH F. V., NONATO L. G.: Local affine multidimensional projection. *IEEE TVCG 17*, 12 (2011), 2563–2571. 1, 3

[Jol02] JOLLIFFE I. T.: *Principal Component Analysis*. Springer, 2002. 2^{nd} edition. 2

[Lic13] LICHMAN M.: UCI machine learning repository, 2013. http://archive.ics.uci.edu/ml. 2

[MCMT14] MARTINS R., COIMBRA D., MINGHIM R., TELEA A. C.: Visual analysis of dimensionality reduction quality for parameterized projections. *Computers & Graphics 41* (2014), 26–42. 1, 4

[MHM18] MCINNES L., HEALY J., MELVILLE J.: UMAP: Uniform manifold approximation and projection for dimension reduction, 2018. arXiv:1802.03426v2 [stat.ML]. 1

[NA18] NONATO L. G., AUPETIT M.: Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE TVCG 25*, 8 (2018), 2650–2673. 1

[PPM*15] PAGLIOSA P., PAULOVICH F., MINGHIM R., LEVKOWITZ H., NONATO L.: Projection inspector: Assessment and synthesis of multidimensional projections. *Neurocomputing 150* (2015), 599–610. 1

[SvLB10] SCHRECK T., VON LANDESBERGER T., BREMM S.: Techniques for precision-based visual analysis of projected data. *Information Visualization 9*, 3 (2010), 181–193. 1

[TDL00] TENENBAUM J. B., DE SILVA V., LANGFORD J. C.: A global geometric framework for nonlinear dimensionality reduction. *Science 290*, 5500 (2000), 2319–2323. 1

[vH08] VAN DER MAATEN L., HINTON G. E.: Visualizing data using t-sne. *JMLR 9* (2008), 2579–2605. 1

[vWT01] VAN WIJK J. J., TELEA A.: Enridged contour maps. In *Proc. IEEE Visualization* (2001), pp. 69–74. 3

[WAG05] WILKINSON L., ARLAND A., GROSSMAN R.: Graph-theoretic scagnostics. In *Proc. InfoVis* (2005). 4

[Yeh98] YEH I.-C.: Modeling of strength of high performance concrete using artificial neural networks. *Cement and Concrete Research 28*, 12 (1998), 1797–1808. 2