A Concept for Consensus-based Ordering of Views

Wolfgang Jentner[1], Dominik Jäckle[1], Ulrich Engelke[2], Daniel Keim[1] and Tobias Schreck[3]

[1]Data Analysis and Visualization, University of Konstanz, Germany
[2]Commonwealth Scientific and Industrial Research Organization (CSIRO), Australia
[3]Institute of Computer Graphics and Knowledge Visualization, TU Graz, Austria

Abstract
High-dimensional data poses a significant challenge for analysis, as patterns typically exist only in subsets of dimensions or records. A common approach to reveal patterns, such as meaningful structures or relationships, is to split the data and then to create a visual representation (views) for each data subset. This introduces the problem of ordering the views effectively because patterns can depend on the presented sequence. Existing methods provide metrics and heuristics to achieve an ordering of views based on their data characteristics. However, an effective ordering of subspace views is expected to rely on task- and data-dependent properties. Hence, heuristic-based ordering methods can be highly objective and not relevant to the task at hand, which is why the user involvement is key to find a meaningful ordering. We introduce a concept for a consensus-based ordering of views *that learns to form sequences of subset views fitting the overall users' needs. This concept allows users to decide on the ordering freely and accumulates their preference into a global view that reflects the consensus. We showcase and discuss this concept based on ordering colored tiles from the controversially discussed rainbow color map.*

1. Introduction

There is no optimal view sequence. While this strong statement is incorrect for numerical data, time-dependent data, or any data that shows a natural ordering, it is held true for most sequences of views of high-dimensional data, as there is no universal criterion on how to order sets of views. High-dimensional data is special in a sense that both the vast amounts of records as well as dimensions hinder the effective identification of meaningful structures and relations, which can be traced to the curse of dimensionality [Bel61]. Therefore, the data is often split into subsets of records or dimensions and visualized using multiple connected or unconnected views on the data [Tuf91]. However, a key challenge in finding meaningful structures and relations across views is to find an appropriate sequence to lay out the views. We hereby focus on the order of one-dimensional sequences.

Recent research proposes various metrics and heuristics that can find a data-driven sequence for different classes of visualizations. Such approaches can be found in axis-based plots [CvW11], glyphs [War02, BKC*13], pixel-based visualizations [KAK95], and matrix visualizations [BBR*16]. The commonality between the aforementioned approaches is that they are based on certain data or image characteristics and do not take into account the users' subjective preference. However, a data- or image-driven ordering may go hand in hand with the users' expectations and preferences, raising the question: *What is an optimal ordering?* We argue that the notion of *optimal* is subjective and task-dependent. In general, we cannot assume that an optimal order exists, which all potential users may agree on. Consider, for example, a sequence that con-

sists of *n* different views on the data. There exist *n*! different sequential orderings. We aim to tackle the research question: *How to integrate the user into the ordering process to efficiently derive an optimal ordering of high-dimensional data representations?* Note the strong connection between the *representation* of the sequence and the order being *optimal*. Dependent on the visualization, users may prefer diverse sequences. Based on the assumption that any high-dimensional data can be brought into a meaningful sequence, we propose to integrate the user into the ordering process.

In this paper, we contribute a *concept for the consensus-based ordering of views on high-dimensional data*. Note that we consider high-dimensional data because it typically gives rise to many subset views for inspection. Our approach, in general, will work on any set of views, as long as a distance matrix can be defined on the views. Our concept accumulates multiple users' subjective preferences on the ordering and proposes the learned ordering to other users. To do so, we use distance as a quantitative measure between views and train a weighting scheme to influence these distances. This approach allows us to find a *consensus ordering* from multiple users through an averaging of weights. Additional statistical measures can be applied to sense disagreement and, further, propose multiple possible solutions reflecting different tasks or user groups. We are providing a concept and thus, leave a comprehensive evaluation open for future work.

2. Background and Challenges

One-dimensional ordering is relevant for a plethora of applications and visualization tasks. Liu et al. provide a categorization of

transformation steps within the information visualization pipeline, which also covers visual mapping techniques [LMW*17]. We consider three of these visual mapping techniques as relevant for the problem of 1D ordering and describe them in the following including the available approaches regarding the ordering as well as the problems that remain. In the following, we provide example approaches that motivate the need to incorporate the domain-specific knowledge in this problem in a feedback-based manner.

Axis-Based visualization techniques map dimensions of an element to different axes and, thus, allow a user to investigate relationships of the data across these dimensions. Well-known techniques include parallel coordinates, scatterplot matrices, and radial layouts [WGK10]. For each of these techniques, the order of the axes is crucial as it reveals different visual patterns. Claessen and van Wijk's approach on *Flexible Linked Axes* [CvW11] allows users to place axes in various ways for different visualization techniques. Their work highlights the high importance of ordering interaction for the user. While this system is capable of placing axes beyond a 1D ordering (e.g., 2D as for TimeWheels [TAS04]), it is not possible that multiple users agree in their ordering preferences. Similar to axis-based methods, also matrices require an ordering to reveal patterns and many ordering methods exist to date [BBR*16].

Glyphs are popular to visualize high-dimensional data within a small graphical entity. Ward [War02] and Borgo et al. [BKC*13] give an overview of different glyph placement strategies including data-, structure-, feature-, and user-driven. Chung et al. provide a so-called IMG plot where the user can actively influence the 2D axes that are used to place the glyphs [CLP*15]. The related work in this field either suggests automatic approaches [SFO*15, KDFB16] or interactive approaches [CLP*15] where features are selected by the user to find an order/placement for glyphs.

Pixel-Oriented visualizations provide the most compact encoding of information in a given display [LMW*17]. With pixel displays and appropriate layout techniques including aggregation and interaction, finding patterns and trends in large amounts of data is supported. Naturally, the way of mapping of data to pixel values is crucial. Keim et al. proposed recursive patterns [KAK95]. The general idea is to arrange the data records in such way that their spatial distance represents similarity of values to each other.

All of these techniques highlight the importance of ordering and placement, and there exist works to use features from the data (dimensions) to infer an order. Interactive approaches allow the user to explicitly select features or combinations to derive an order.

Furthermore, various approaches have been presented that include the user to explore *ranked* data. In contrast to orderings, ranking-based approaches help users to navigate the data content [GLG*13]. This is typically done by involving the user in attribute prioritization tasks or filter customization [dSSV15, dSSV16, WDC*18]. Other approaches, more related to the ordering issue, involve the domain knowledge of the user and enable the interactive selection of different ordering algorithms [FGS*17, JJ17]. There exists a rich variety of ordering algorithms, which we do not aim to survey here. Users perceive orderings differently and

have differing opinions as to what characterizes an optimal ordering. Typically, the ordering is determined based on features of interest in the data, depending on user and task. Hullman et al. [HKL17] evaluate the perception of sequences, finding that participants have different views on how a sequence should look like.

Recent work incorporates the idea of adapting distances in high-dimensional planar projections [YMSJ05, BLBC12, ECNZ15]. The analyst can make use of her domain knowledge and thus, change the relations between data points in the display. Each interaction impacts the corresponding visualization model, which is the parameterization of the projection. However, these works focus on the relationship between data points and do not consider the order.

In other works [WDC*18], the user is enabled to rank data items, and an algorithm finds the appropriate feature weighting. However, the order does not necessarily imply a ranking, and we do not weigh features but a given distance matrix.

We aim to incorporate the users' knowledge by providing an implicit solution, in which the user must not explicitly know about features and their importance, but in which this information can be extracted by analyzing the preferred order of one or more users.

3. Consensus-based ordering of views

Our general idea of consensus-based ordering of views on high-dimensional data is to incorporate the users' subjective ordering preferences on the ordering problem. This means that different users may not share the same ordering preferences, but prefer differing sequences. Note that we do not require users to issue a complete ordering, but that also partial orderings can be supplied. All the collected preferences of the users on the ordering problem are then incorporated as feedback into the generation of a new sequence that is learned based on a majority vote. We consider this learned sequence as the consensus-based ordering. In the following, we introduce and discuss a general model that implements this concept. We consider a view as a visual and application-dependent representation of a subset or a projection of the data. In our example case in the next Section, a view is merely a colored tile, but any form of small multiples, glyphs, or axes of a parallel coordinates plot are possible. Our approach learns a sequence from the way different users order the views and presents this learned sequence to the next user. The concept enables each user to bring in her knowledge and reorder the sequence, which again serves as input to the generation of the overall learned sequence.

A related approach, introduced by Stormo et al. [SSGE82], proposes Position Weight Matrices (PWM) that are trained to classify DNA/mRNA sequences. A PWM, hereby, encodes the probability for each element at each position, which is highly relevant to our ordering problem. In contrast to Stormo et al., we aim for an approach that is robust to shifts in sequences and enables users to sort-in new elements.

We introduce a general, linear concept that incorporates user-defined sequences by weigh adaption as depicted in Figure 1. In the first step, the data is visualized in the form of views, and basic interaction techniques such as drag and drop are enabled allowing to adapt the ordering. We hereby refer to the well-known visualization pipeline by Card et al. [CMS99]. The second step is key to

Figure 1: *General model for consensus-based ordering. Initial ordering of views is generated based on the well-known visualization pipeline by Card et al. [CMS99]. A learned weighting scheme is then applied to the sequence, causing the views to reorder about a common user-driven understanding of the sequence.*

Algorithm 1 A weight matrix (*wm*) is calculated for one user.

Input: *odm* =original distance matrix; *wm* =weight matrix
Output: *wm*, a weight matrix that, multiplied with *odm* reflects one user's optimal order
do
    $dm = odm \circ wm$ //Hadamard product
    $order = projectlinear(dm, 1)$ //e.g., MDS
    display order to user
    dm = switch rows and columns of *dm* according to the user's change
    $wm = dm \oslash odm$ //Hadamard division
while user not satisfied with order
return *wm*

our concept: Based on the visual representation of the views, we initialize the ordering of views taking into account a given distance matrix and learned distance weights. The notion of distance is a natural way of indicating how similar elements are and what order they follow. For high-dimensional data, the notion of distance becomes blurred due to the curse of dimensionality [Bel61], which is why different user groups may find different orderings to be correct.

Algorithm 1 shows the process of deriving a new weight matrix (*wm*) for *one* user. The algorithm takes a given distance matrix that can be based on any distance measure (data- or image-based). The second parameter is a weight matrix that either reflects the original distances (all elements are 1), or any consensus based on the input of previous users (in our case, obtained by averaging, see below). When multiplied (Hadamard product) with the original distance matrix, the resulting distance matrix represents the users' subjective feedback. The distance matrix can be transformed into a 1D ordered sequence based on well-known linear approaches such as the classical Multi-dimensional Scaling (MDS) [CC00] or force-directed drawing using Hooke's Law [Kob12]. Examples include [JHB*17, vdEHBvW16, BSH*16]. When the user changes the order, i.e., switches two views of the sequence, the distances (rows and columns) of the respective elements are switched. The weight matrix can be derived by a Hadamard-division of the current distance matrix with the original distance matrix. Therefore, the original distance matrix multiplied by the weight matrix and piped into the projection method always reflects the currently preferred sequence of the *current* user. The final weight matrix of each user is stored. The weight matrices can be averaged to reflect a consensus. This is, however, not robust to outliers but other statistical measures can be included in combining the various weight matrices. We consider this as future work.

The described approach uses a projection to one dimension to derive the order. This can be extended to additionally reflect distances by visually separating the views. This requires to adapt the distance matrix differently when the user swaps two views, but the rest of the procedure is not subject to any change.

The issue with storing pairwise distances is that the layout algorithms are unaware whether the sequence starts with the first or the last view. In other words, the sequence is flipped and still reflects the correct pairwise distances, which can be traced back to the rotation invariance; the distance matrix does not reflect any order using an ordering flipped by π. In fact, for a flipped ordered sequence the distance matrix is identical, and the layout approach cannot determine which view is the first. We, therefore, add two additional and artificial elements to our distance matrix, which also hold the information if they start or end the sequence [PZS*15]. These two elements serve as anchors. In a visual metaphor, these anchors are placed at the beginning and ends of the sequence. This implies that the distance between these two anchors is maximal and no other distance must be greater.

When new and previously unseen elements are added to the sequence, their position can be extrapolated. A generated consensus distance matrix can be compared to any other distance matrix that is being generated from a data- or image-based distance measure. A simple comparison can be achieved by calculating the mean squared error and choose the distance matrix with the lowest error. An extension could be to derive feature weights from a given consensus-based distance-matrix [WDC*18]. We consider this as future work. With this information, the position of the new element can be recommended to the user which may or may not result in additional iterations for refinements by the user.

In summary, our concept incorporates the user using interaction and learns an ordering based on the user consensus. The concept is not restricted to 1D sequences and can be extended to higher dimensions such as 2D and 3D; the notion of distance is not restricted to a number of dimensions.

4. Example Case: Color Ordering

We use the controversially discussed rainbow color map [BI07] as an example case as it nicely represents the subjectiveness of an optimal ordering. Furthermore, it allows a simple visual representation, but this could be extended to more complex views such. Our approach scales to any number of dimensions yet is only limited by the curse of dimensionality which is subject to current research. Even though the example case is not of a very high-dimensional nature, any derived distance matrix can be used in our approach.

The colors are selected from the HSV color space where the saturation and value are set to the maximum, and the hue is alternated to generate differently colored tiles (see Figure 2). While there exist some *short sections* that can be perceptually ordered, implying

Figure 2: *Different possible orderings of the rainbow color map. (1) and (6) represent an ordering according to the wavelength in an ascending and descending order.*

a total order is difficult or even impossible [War12, p.128]. More recent studies [CAB*16] are consistent with this claim. Figure 2 shows different possible orderings. (1) and (6) hereby reflect the sequence of the wavelength in an ascending and descending order. The other sequences diverge as the starting/ending point of the hue in the HSV color space is varied. Nevertheless, the *short sections* are still in order.

Four possible input orders can be presented to the user which originate only from the adaption of the input weight matrix. (i) the order based on original distance matrix (all weights are 1); (ii) a randomized order (random weights); (iii) other distance matrices from different distance measures reflected through the weight matrix; (iv) a consensus-based order of previously user-defined orderings. The second and third might be useful to prevent converging into a local optimum as a user might be tempted to agree on a given ordering. Especially, when the user knows about the fact that the input order is not randomized. In any case, the interactions of the user trying to find a new optimal ordering are captured, and the weights reflecting the majority vote are adjusted.

It is likely that, in this case, random users do not agree on one or a limited number of orderings. However, we can assume that specific user groups do. The rainbow color map originates from the field of physics. In this example user group, we could assume that the orderings converge to the ordering according to the wavelength as physicists know about this matter. While there is no correct ordering per se, such a convergence would reflect the background knowledge of this specific user group.

5. Discussion and Concluding Remarks

We propose a general concept for the consensus-based ordering of views. With our concept, it is possible to find a consensus of multiple users' opinions about the subjective preference of order. While we do not allow the explicit selection or ranking of features, we encourage the user to interactively change the order until she thinks it is optimal. The concept then allows us to derive the ranking/selection and use this information for additional data or similar tasks. Note that our concept is based on the assumption that a consensus can be found by averaging. This assumption may be violated in practice. Future work will include mechanisms to

detect disagreement and possible contradictory preferences of users, and visualization to show these to the user. In case of strong contradictions, users should be enabled to overrule influence by particularly contradictory other user preferences. We intend to use methods and best practices of related fields such as consensus clustering and consensus trees [SG02].

Our concept provides much control to the user which includes all of the user's possible cognitive biases.While some biases can be limited by the consensus of a user group others might be leveraged. However, we do not allow the user(s) to influence the underlying distance function directly but, moreover, change the weights which might impact the global order. The rapid feedback that is provided to the user allows her to reconsider the actions and she might come to a different solution.

The concept is flexible enough to be aware of different tasks, user groups, or even converging optimal solutions over time, for example, due to changing data or real-world knowledge. The modular weight adaption method can be adjusted or exchanged to cover all of these cases. A simple use case is to weight each weight matrix (scalar multiplication) according to the user's expertise. However, more complex scenarios are imaginable in practice.

Another interesting aspect is the combined weight matrix itself. Analyzing this data can shed light on what views of a sequence are mostly subject to change and also to what magnitude. More statistics such as the variance or distributions may help to identify various subgroups of users that have a diverging consensus of optimal order. Comparing the resulting distance matrix to other various matrices provides insight on what data- or image-based distance measure best reflects the users' optimal ordering – and also for which views it fails. Furthermore, it is possible to derive feature weights based on the resulting distance matrix, which then provides a similar use case to Wall et al. [WDC*18]

The initial distance matrix in combination with the learned weights can be combined and compared to various other distance metrics based on data or image features. This may be fruitful in a scenario where the user is presented with a (small) subset of the views, defines an order, and then additional views can be placed automatically based on the available feedback.

High-dimensional data and the curse of dimensionality impose the risk of distortions that can occur in projection techniques. This challenging problem conveys an active field of research. While we cannot solve this problem here, we argue that orderings impose a great difference in the weights and that such great differences are less impacted by distortions. Again, the rapid feedback supports the user to spot possible distortions and find countermeasures.

We are planning an user evaluation to seek answers to the following open questions: How (early) can a consensus be measured? Are users satisfied with the consensus? How scalable is our approach? We are confident that by answering these questions we can provide a practical solution to the initially stated problem of finding a subjective and task-dependent ordering strategy.

References

[BBR*16] BEHRISCH M., BACH B., RICHE N. H., SCHRECK T., FEKETE J.: Matrix reordering methods for table and network visualization. *Comput. Graph. Forum 35*, 3 (2016), 693–716. 1, 2

[Bel61] BELLMAN R.: *Adaptive Control Processes: A Guided Tour*. Princeton University Press, 1961. 1, 3

[BI07] BORLAND D., II R. M. T.: Rainbow color map (still) considered harmful. *IEEE Computer Graphics and Applications 27*, 2 (2007), 14–17. 3

[BKC*13] BORGO R., KEHRER J., CHUNG D. H. S., MAGUIRE E., LARAMEE R. S., HAUSER H., WARD M., CHEN M.: Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics 2013 - State of the Art Reports, Girona, Spain, May 6-10, 2013* (2013), Sbert M., Szirmay-Kalos L., (Eds.), Eurographics Association, pp. 39–63. 1, 2

[BLBC12] BROWN E. T., LIU J., BRODLEY C. E., CHANG R.: Disfunction: Learning distance functions interactively. In *2012 IEEE Conference on Visual Analytics Science and Technology, VAST 2012, Seattle, WA, USA, October 14-19, 2012* (2012), IEEE Computer Society, pp. 83–92. 2

[BSH*16] BACH B., SHI C., HEULOT N., MADHYASTHA T. M., GRABOWSKI T. J., DRAGICEVIC P.: Time curves: Folding time to visualize patterns of temporal evolution in data. *IEEE Trans. Vis. Comput. Graph. 22*, 1 (2016), 559–568. 3

[CAB*16] CHUNG D. H. S., ARCHAMBAULT D. W., BORGO R., EDWARDS D. J., LARAMEE R. S., CHEN M.: How ordered is it? on the perceptual orderability of visual channels. *Comput. Graph. Forum 35*, 3 (2016), 131–140. 4

[CC00] COX T., COX A.: *Multidimensional Scaling, Second Edition*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. CRC Press, 2000. 3

[CLP*15] CHUNG D. H. S., LEGG P. A., PARRY M. L., BOWN R., GRIFFITHS I. W., LARAMEE R. S., CHEN M.: Glyph sorting: Interactive visualization for multi-dimensional data. *Information Visualization 14*, 1 (2015), 76–90. 2

[CMS99] CARD S. K., MACKINLAY J. D., SHNEIDERMAN B.: *Readings in information visualization - using vision to think*. Academic Press, 1999. 2, 3

[CvW11] CLAESSEN J. H. T., VAN WIJK J. J.: Flexible linked axes for multivariate data visualization. *IEEE Trans. Vis. Comput. Graph. 17*, 12 (2011), 2310–2316. 1, 2

[dSSV15] DI SCIASCIO C., SABOL V., VEAS E. E.: *uRank*: Exploring document recommendations through an interactive user-driven approach. In *Proceedings of the Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, IntRS 2015, Vienna, Austria, September 19, 2015.* (2015), vol. 1438, CEUR-WS.org, pp. 29–36. 2

[dSSV16] DI SCIASCIO C., SABOL V., VEAS E. E.: Rank as you go: User-driven exploration of search results. In *Proceedings of the 21st International Conference on Intelligent User Interfaces, IUI 2016, Sonoma, CA, USA, March 7-10, 2016* (2016), ACM, pp. 118–129. 2

[ECNZ15] ENDERT A., CHANG R., NORTH C., ZHOU M. X.: Semantic interaction: Coupling cognition and computation through usable interactive analytics. *IEEE Computer Graphics and Applications 35*, 4 (2015), 94–99. 2

[FGS*17] FURMANOVA K., GRATZL S., STITZ H., ZICHNER T., JARESOVÁ M., ENNEMOSER M., LEX A., STREIT M.: Taggle: Scalable visualization of tabular data through aggregation. *CoRR abs/1712.05944* (2017). 2

[GLG*13] GRATZL S., LEX A., GEHLENBORG N., PFISTER H., STREIT M.: Lineup: Visual analysis of multi-attribute rankings. *IEEE Trans. Vis. Comput. Graph. 19*, 12 (2013), 2277–2286. 2

[HKL17] HULLMAN J., KOSARA R., LAM H.: Finding a clear path: Structuring strategies for visualization sequences. *Comput. Graph. Forum 36*, 3 (2017), 365–375. 2

[JHB*17] JÄCKLE D., HUND M., BEHRISCH M., KEIM D. A., SCHRECK T.: Pattern Trails: Visual Analysis of Pattern Transitions in Subspaces. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2017). 3

[JJ17] JUGOVAC M., JANNACH D.: Interacting with recommenders - overview and research directions. *TiiS 7*, 3 (2017), 10:1–10:46. 2

[KAK95] KEIM D. A., ANKERST M., KRIEGEL H.: Recursive pattern: A technique for visualizing very large amounts of data. In *IEEE Visualization '95, Proceedings, Atlanta, Georgia, USA, October 29 - November 3, 1995.* (1995), IEEE Computer Society Press, pp. 279–286. 1, 2

[KDFB16] KRAUSE J., DASGUPTA A., FEKETE J., BERTINI E.: Seekaview: An intelligent dimensionality reduction strategy for navigating high-dimensional data spaces. In *LDAV* (2016), IEEE Computer Society, pp. 11–19. 2

[Kob12] KOBOUROV S. G.: Spring embedders and force directed graph drawing algorithms. *arXiv preprint arXiv:1201.3011* (2012). 3

[LMW*17] LIU S., MALJOVEC D., WANG B., BREMER P., PASCUCCI V.: Visualizing high-dimensional data: Advances in the past decade. *IEEE Trans. Vis. Comput. Graph. 23*, 3 (2017), 1249–1268. 2

[PZS*15] PÉREZ-LÓPEZ D., ZHANG L., SCHÄFER M., SCHRECK T., KEIM D. A., DÍAZ I.: Interactive feature space extension for multidimensional data projection. *Neurocomputing 150* (2015), 611–626. 3

[SFO*15] STOFFEL F., FLEKOVA L., OELKE D., GUREVYCH I., KEIM D. A.: Feature-Based Visual Exploration of Text Classification. In *Symposium on Visualization in Data Science (VDS) at IEEE VIS 2015* (2015). 2

[SG02] STREHL A., GHOSH J.: Cluster ensembles — A knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research 3* (2002), 583–617. 4

[SSGE82] STORMO G. D., SCHNEIDER T. D., GOLD L., EHRENFEUCHT A.: Use of the 'perceptron' algorithm to distinguish translational initiation sites in e. coli. *Nucleic acids research 10*, 9 (1982), 2997–3011. 2

[TAS04] TOMINSKI C., ABELLO J., SCHUMANN H.: Axes-based visualizations with radial layouts. In *Proceedings of the 2004 ACM Symposium on Applied Computing (SAC), Nicosia, Cyprus, March 14-17, 2004* (2004), Haddad H., Omicini A., Wainwright R. L., Liebrock L. M., (Eds.), ACM, pp. 1242–1247. 2

[Tuf91] TUFTE E. R.: Envisioning information. *Optometry & Vision Science 68*, 4 (1991), 322–324. 1

[vdEHBvW16] VAN DEN ELZEN S., HOLTEN D., BLAAS J., VAN WIJK J. J.: Reducing snapshots to points: A visual analytics approach to dynamic network exploration. *IEEE Trans. Vis. Comput. Graph. 22*, 1 (2016), 1–10. 3

[War02] WARD M. O.: A taxonomy of glyph placement strategies for multidimensional data visualization. *Information Visualization 1*, 3-4 (2002), 194–210. 1, 2

[War12] WARE C.: *Information visualization: perception for design*. Elsevier, 2012. 4

[WDC*18] WALL E., DAS S., CHAWLA R., KALIDINDI B., BROWN E. T., ENDERT A.: Podium: Ranking data using mixed-initiative visual analytics. *IEEE Trans. Vis. Comput. Graph. 24*, 1 (2018), 288–297. 2, 3, 4

[WGK10] WARD M. O., GRINSTEIN G. G., KEIM D. A.: *Interactive Data Visualization - Foundations, Techniques, and Applications*. A K Peters, 2010. 2

[YMSJ05] YI J. S., MELTON R., STASKO J. T., JACKO J. A.: Dust & magnet: multivariate information visualization using a magnet metaphor. *Information Visualization 4*, 3 (2005), 239–256. 2