Interactive Human-guided Dimensionality Reduction using Landmark Positioning

T. Cech[1], C. Raue[2], F. Sadrieh[2], W. Scheibel[2] and J. Döllner[1]

[1]Digital Engineering Faculty, University of Potsdam, Germany
[2]Hasso Plattner Institute, Digital Engineering Faculty, University of Potsdam, Germany

Abstract

Dimensionality Reduction Techniques (DRs) are used for projecting high-dimensional data onto a two-dimensional plane. One subclass of DRs are such techniques that utilize landmarks. Landmarks are a subset of the original data space that are projected by a slow and more precise technique. The other data points are then placed in relation to these landmarks with respect to their distance in the high-dimensional space. We propose a technique to refine the placement of the landmarks by a human user. We test two different techniques for unprojecting the movement of the low-dimensional landmarks into the high-dimensional data space. We showcase that such a movement can increase certain quality metrics while decreasing others. Therefore, users may use our technique to challenge their understanding of the high-dimensional data space.

CCS Concepts

• ***Human-centered computing*** → ***Visualization techniques;***

1. Introduction

Dimensionality Reduction Techniques (DRs) are a common way to visualize high-dimensional datasets in a scatter plot. Through this, a user can understand the high-dimensional data space by enabling the user a distorted view through the projection of a DR [TMW24]. One subclass of DRs are techniques that use landmarks, such as *Landmark Multidimensional Scaling (LMDS)* [ST02]. LMDS proposes that due to the high runtime of *Multidimensional Scaling (MDS)* only selected data points – the so-called landmarks – are projected down by the basic MDS technique first. Subsequently, the other data points are placed around those projected landmarks according to their distance from the landmarks in the high-dimensional data space. One important question for applying LMDS is how to choose the landmarks and how to position them on the low-dimensional plane. For example, Atzberger and Cech et al. have shown that the simple positioning of data points using MDS neither preserves local nor global properties of the high-dimensional data space well [ACT*24]. One way to overcome this shortcoming of MDS could be the usage of quali-quantitative methods that incorporate human expert feedback [BP14, vRMB*23]. For it, we argue that it is most appropriate to directly unproject the human feedback back into the data space instead of relying on another opaque technique to adapt the projection technique – here LMDS – itself [vRMB*23]. We propose a proof-of-concept of such a quali-quantitative technique for enhancing LMDS layouts.

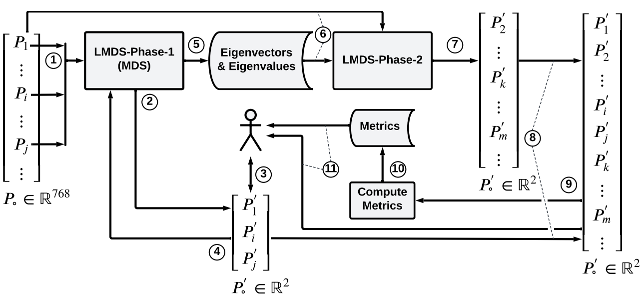

Figure 1: *The interactive Landmark MDS layout pipeline.*

2. Related Work

In terms of the taxonomy of von Rueden et al., we operate in the field of informing DRs with Human Feedback on the level of the Training Data [vRMB*23]. For this field, von Rueden et al., identified two previous papers by Brown et al. [BLBC12] and Heckerman et al. [HGC95]. Heckerman et al. proposed to combine statistical and expert knowledge to train more sophisticated Bayesian Networks [HGC95]. We are sympathetic to this notion and try to map onto DRs instead of Bayesian Networks. Like our system, Brown et al. proposed to interactively adapt an MDS projection by learning a distance function according to human expert feedback [BLBC12]. We differ from this work by not changing the LMDS algorithm itself by adapting the distance function. Instead, we propose to unproject the changes in the projection space back into the data space [EAS*23] and, subsequently, project this changed data space

DOI: 10.2312/evp.20241085

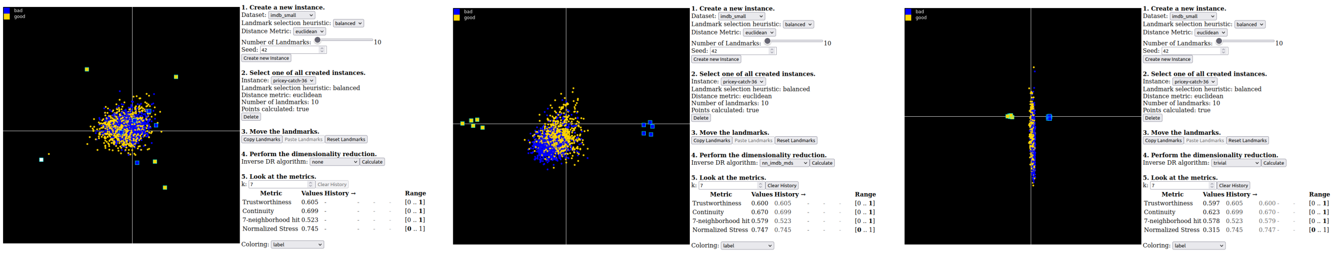

(a) *The original LMDS layout.* **(b)** *The modified layout with the NN approach.* **(c)** *The modified layout with the trivial approach.*

Figure 2: *An example for our approach on the IMDb dataset. In b and c the NN and the trivial approach are used for unprojection, respectively.*

back into a new projection. In addition, Iwata et al. also proposed to adapt a projection by letting human experts change the position of data points in the projection space based on an information criterion [IHG13]. In contrast to this work, we let the user decide which movements are most significant for them.

3. Approach

Our pipeline is showcased in Figure 1. First, we choose in our high-dimensional dataset our landmarks randomly (1) and project them with the classical LMDS technique by projecting some landmarks with MDS [ST02] (2). For it, we enable the user to choose from one of two distance functions – namely the Cosine Similarity and Euclidean Distance [ACT*24]. Afterward, we enable users to compare the landmarks (3) and freely move them along the two-dimensional plane (4). We expect users to include their qualitative feedback in that movement therefore enriching the positioning by expert knowledge [AAA*24, BLBC12]. To reflect the movement of the points in the high-dimensional data space we have to unproject the landmarks (5). For this, we provide the user with two techniques. First, we follow the argument by Kriegskorte and Mur who propose a technique which we call "trivial approach" [KM12]. They argue that the movement of the low-dimensional data points may best be reflected in the distance matrix used by MDS by simply setting the low-dimensional distances as the new high-dimensional distances. Second, we follow Espadoto et al. who propose the usage of a small *Neural Network (NN)* to unproject in their case hypothetical data points [EAS*23]. In our case, we adapt this approach by training the NN to recover the high-dimensional placement of the landmarks. We tested that already a relatively small NN achieves this target well as long as the distance function chosen in the high-dimensional and low-dimensional space is kept constant. For further details about our NN architecture, we refer to our supplemental material. After unprojecting the high-dimensional data points we perform the same second phase as the classical LMDS algorithm (6) to obtain a layout of all data points [ST02] (7). Finally, the so-generated layout of the other (8) data points together with the moved landmarks (9) is evaluated by a part of the metric suite which was proposed by Espadoto et al. [EMK*21] (10). We use the Trustworthiness, Continuity, 7-Neighborhood-Hit and Normalized Stress from their metric suite. The first three metrics are optimized when reaching one while the Normalized Stress is optimized when reaching zero. The user can reiterate this process by moving the landmarks again and choosing another distance function and unprojection technique until the layout satisfies their qualitative demands [BP14] (11).

4. Evaluation

For showcasing our approach we use a subsample of 2 000 examples out of the 25 000 labeled training examples from the IMDb movie review dataset [MDP*11]. In addition to the binary classification task of IMDb, we use the emotion dataset [SLH*18], which has six labels. The dataset contains English tweets and assigns each one of six emotions("sadness, disgust, anger, joy, surprise, and fear") [SLH*18]. We generate a sample 2 000 examples from the 16 000 split of the dataset. We only train on the IMDb and emotion datasets. As a hold-out set, we use the mnli dataset [WNB18] that contains premise and hypothesis sentences concatenated with a semicolon. The label indicates whether the premise entails the hypothesis, is neutral, or contradicts it [WNB18]. We have generated a sample of 2 000 examples out of the 393 000 train split of the dataset. To embed the textual input, we use the *all-mpnet-base-v2* [AER] Sentence-Transformer-model [RG19] from hugging face. This model is pre-trained on sentence similarity: It takes two sentences as input and outputs their similarity. The model generates a 768-dimensional embedding that includes the semantics of the text [AER]. For the training pipeline, we used parts of a template for fast ML training and iteration [DSZ23]. Figure 2 showcases one example. The left side is the original LMDS layout, the middle is a layout with moved landmarks and using the NN for unprojection, and the right is a layout with moved landmarks and using the trivial approach. We observe that by moving the landmarks with the NN technique the Trustworthiness drops, while the Continuity and Normalized Stress stay about constant and increase the 7-Neighborhood Hit. Furthermore, by moving the landmarks the NN technique the Trustworthiness and Continuity drop, while the 7-Neighborhood Hit and Normalized Stress are improved. In summary, the user can choose which metrics they want to optimize and choose techniques accordingly.

5. Conclusions

We proposed a technique for enhancing LMDS by using quali-quantitative human feedback loops. This technique showcases how users can insert their expert knowledge into the visualization process. By measuring this feedback in terms of quality metrics, we enable users to challenge their understanding of the high-dimensional data space. In the future, we want to extend our approach to more recent methods, especially t-Stochastic Neighbor Embedding (t-SNE) [VdMH08] or Uniform Manifold Projection and Approximation (UMAP) [MHM18] as well as investigating visual guidance techniques [IHG13].

Acknowledgements

This work was partially funded by the German Federal Ministry for Education and Research (BMBF) through grant 01IS22062 ("AI research group FFS-AI"). This work is part of project 16KN086467 ("DecodingFood") funded by the Federal Ministry for Economic Affairs and Climate Action of Germany.

References

[AAA*24] ANDRIENKO N., ANDRIENKO G., ARTIKIS A., MANTENOGLOU P., RINZIVILLO S.: Human-in-the-loop: Visual analytics for building models recognising behavioural patterns in time series. *IEEE Computer Graphics and Applications* (2024), 1–15. doi:10.1109/MCG.2024.3379851. 2

[ACT*24] ATZBERGER D., CECH T., TRAPP M., RICHTER R., SCHEIBEL W., DÖLLNER J., SCHRECK T.: Large-scale evaluation of topic models and dimensionality reduction methods for 2D text spatialization. *IEEE Transactions on Visualization and Computer Graphics 30*, 1 (2024), 902–912. doi:10.1109/TVCG.2023.3326569. 1, 2

[AER] AARSEN T., ESPEJEL O., REIMERS N.: all-mpnet-base-v2. URL: https://huggingface.co/sentence-transformers/all-mpnet-base-v2. 2

[BLBC12] BROWN E. T., LIU J., BRODLEY C. E., CHANG R.: Dis-function: Learning distance functions interactively. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2012), pp. 83–92. doi:10.1109/VAST.2012.6400486. 1, 2

[BP14] BLOK A., PEDERSEN M. A.: Complementary social science? Quali-quantitative experiments in a Big Data world. *Big Data & Society 1*, 2 (2014), 1–4. doi:10.1177/2053951714543908. 1, 2

[DSZ23] DOBLER K., SCHALL M., ZIMMERMANN O.: nlp-research-template, 2023. URL: https://github.com/konstantinjdobler/nlp-research-template. 2

[EAS*23] ESPADOTO M., APPLEBY G., SUH A., CASHMAN D., LI M., SCHEIDEGGER C., ANDERSON E. W., CHANG R., TELEA A. C.: UnProjection: Leveraging inverse-projections for visual analytics of high-dimensional data. *IEEE Transactions on Visualization and Computer Graphics 29*, 2 (2023), 1559–1572. doi:10.1109/TVCG.2021.3125576. 1, 2

[EMK*21] ESPADOTO M., MARTINS R. M., KERREN A., HIRATA N. S. T., TELEA A. C.: Toward a quantitative survey of dimension reduction techniques. *IEEE Transactions on Visualization and Computer Graphics 27*, 3 (2021), 2153–2173. doi:10.1109/TVCG.2019.2944182. 2

[HGC95] HECKERMAN D., GEIGER D., CHICKERING D. M.: Learning bayesian networks: The combination of knowledge and statistical data. *Machine learning 20* (1995), 197–243. 1

[IHG13] IWATA T., HOULSBY N., GHAHRAMANI Z.: Active learning for interactive visualization. In *Artificial Intelligence and Statistics* (2013), PMLR, pp. 342–350. 2, 3

[KM12] KRIEGESKORTE N., MUR M.: Inverse MDS: Inferring dissimilarity structure from multiple item arrangements. *Frontiers in Psychology 3*, 245 (2012). doi:10.3389/fpsyg.2012.00245. 2

[MDP*11] MAAS A. L., DALY R. E., PHAM P. T., HUANG D., NG A. Y., POTTS C.: Learning word vectors for sentiment analysis. In *Proc. 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies* (2011), ACL, pp. 142–150. 2

[MHM18] MCINNES L., HEALY J., MELVILLE J.: UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv CoRR stat.ML*, 1802.03426 (2018). preprint. doi:10.48550/arXiv.1802.03426. 3

[RG19] REIMERS N., GUREVYCH I.: Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *Proc. Conference on Empirical Methods in Natural Language Processing* (2019), EMNLP '19, ACL, pp. 3982–3992. doi:10.18653/v1/D19-1410. 2

[SLH*18] SARAVIA E., LIU H.-C. T., HUANG Y.-H., WU J., CHEN Y.-S.: CARER: Contextualized affect representations for emotion recognition. In *Proc. Conference on Empirical Methods in Natural Language Processing* (2018), EMNLP '18, ACL, pp. 3687–3697. doi:10.18653/v1/D18-1404. 2

[ST02] SILVA V., TENENBAUM J.: Global versus local methods in nonlinear dimensionality reduction. In *Advances in Neural Information Processing Systems* (2002), NIPS '02, MIT Press. 1, 2

[TMW24] TELEA A., MACHADO A., WANG Y.: Seeing is learning in high dimensions: The synergy between dimensionality reduction and machine learning. *Springer Nature Computer Science 5*, 3 (2024), 279:1–25. doi:10.1007/s42979-024-02604-y. 1

[VdMH08] VAN DER MAATEN L., HINTON G.: Visualizing data using t-SNE. *Journal of Machine Learning Research 9*, 86 (2008), 2579–2605. 3

[vRMB*23] VON RUEDEN L., MAYER S., BECKH K., GEORGIEV B., GIESSELBACH S., HEESE R., KIRSCH B., PFROMMER J., PICK A., RAMAMURTHY R., WALCZAK M., GARCKE J., BAUCKHAGE C., SCHUECKER J.: Informed machine learning – a taxonomy and survey of integrating prior knowledge into learning systems. *IEEE Transactions on Knowledge and Data Engineering 35*, 1 (2023), 614–633. doi:10.1109/TKDE.2021.3079836. 1

[WNB18] WILLIAMS A., NANGIA N., BOWMAN S.: A broad-coverage challenge corpus for sentence understanding through inference. In *Proc. Conference North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (2018), vol. 1 (Long Papers), ACL, pp. 1112–1122. doi:10.18653/v1/N18-1101. 2