Autoencoder-based Semi-Supervised Dimensionality Reduction and Clustering for Scientific Ensembles

Hamid Gadirov, Steffen Frey

October 23, 2023

About the project

Prior studies have demonstrated the potential for enhanced dimensionality reduction (DR) and clustering performance using machine learning, specifically through the utilization of autoencoders for feature extraction, as compared to conventional DR techniques such as PCA, t-SNE, and UMAP [3]. Moreover, research in this domain has also revealed that methods incorporating labeled information tend to outperform their unsupervised counterparts, leading to improved clustering separation and other advantageous properties [2]. Notably, certain approaches, as presented in [8], permit the integration of unsupervised dissimilarity mixture functions into various types of autoencoders, enabling high-accuracy clustering and unsupervised classification.

This project aims to perform dimensionality reduction (DR) and clustering utilizing diverse types of autoencoders, including Vanilla, Sparse, Variational (VAE) [7], beta-Variational (β-VAE) [6], and Wasserstein (WAE) [12], within a self-supervised framework leveraging partially labeled ensemble data. These tasks can be achieved not only through convolutional neural networks (CNNs) but also by employing Transformer architectures [1]. The novelty of this project lies in the innovative approach of reformulating unsupervised reconstruction losses associated with the listed autoencoder variants into joint losses that incorporate available class information (labels). This strategic integration allows for the inclusion of dissimilarity measures and enhances clustering results by utilizing autoencoder-extracted features for scientific ensembles.

To assess the performance of the proposed self-supervised approach, clustering results, particularly 2D projections, will be compared with benchmarks like the supervised version of t-SNE which incorporates a dissimilarity measure [5], that can serve as a baseline. To evaluate the clustering quality, metrics such as neighborhood hit [13] and silhouette [10] will be employed. These metrics are commonly used for assessing the effectiveness of clustering algorithms.

Research steps

1. Obtaining partially labeled subsets of scientific ensembles.

2. Performing fine-tuning of a pre-trained EfficientNet classifier [11] to generate pseudo-labels for the unlabeled data subset.

3. Implementing various autoencoder variants, including Vanilla, Sparse, Variational, and Wasserstein, for the reconstruction task, following a similar approach to [3].

4. Experimenting with the addition of metric-based loss functions, such as neighborhood hit or silhouette scores, alongside the autoencoder loss.

5. Alternatively experiment in adding class information to the loss functions:

   - start by reproducing results similar to [2],
   - add classifier loss to autoencoder variants listed above,
   - experiment adding supervised dissimilarity measure from [5] to the autoencoder variants,
   - investigate unsupervised dissimilarity loss benefits from [8] for autoencoder based clustering,
   - finally integrate all loss components (reconstruction, regularization, classification, dissimilarity) for each autoencoder variant.

6. Training:

   - begin with benchmark datasets, such as MNIST,
   - progress to the application of the method to our scientific ensembles, including Kármán vortex street (KVS), Markov chain Monte Carlo (MCMC) [9], and Drop Dynamics [4].

7. Projection and clustering, either directly in 2D or after applying standard dimensionality reduction (DR) techniques, as described in [3].

8. Visualization of results in 2D, following a similar approach to [3].

9. Implementation of interactive visualization, supporting both 2D and 3D, accessible via web browsers.

Prerequisites

- Good knowledge of Python and experience with deep learning libraries, such as Keras, Tensorflow, or PyTorch.

- Completed courses such as Machine Learning or Deep Learning; Information Visualization, Scientific Visualization, or Computer Vision.

This project is suitable for a master's thesis or an internship and may also be adapted for a dedicated and enthusiastic bachelor's thesis student. Successful outcomes from this project hold the potential for publication.

References

[1] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

[2] M. Espadoto, N. S. Hirata, and A. C. Telea. Self-supervised dimensionality reduction with neural networks and pseudo-labeling. In *VISIGRAPP (3: IVAPP)*, pages 27–37, 2021.

[3] H. Gadirov, G. Tkachev, T. Ertl, and S. Frey. Evaluation and selection of autoencoders for expressive dimensionality reduction of spatial ensembles. In *International Symposium on Visual Computing*. Springer, 2021.

[4] A. Geppert, D. Chatzianagnostou, C. Meister, H. Gomaa, G. Lamanna, and B. Weigand. Classification of impact morphology and splashing/deposition limit for n-hexadecane. *Atomization and Sprays*, 26(10), 2016.

[5] L. Hajderanj, I. Weheliye, and D. Chen. A new supervised t-sne with dissimilarity measure for effective data visualization and classification. In *Proceedings of the 2019 8th International Conference on Software and Information Engineering*, pages 232–236, 2019.

[6] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.

[7] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[8] J. S. Lara and F. A. González. Dissimilarity mixture autoencoder for deep clustering. *arXiv preprint arXiv:2006.08177*, 2020.

[9] S. Reuschen, T. Xu, and W. Nowak. Bayesian inversion of hierarchical geostatistical models using a parallel-tempering sequential gibbs mcmc. *Advances in Water Resources*, 141:103614, 2020.

[10] P. J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 1987.

[11] M. Tan and Q. V. Le. Efficientnetv2: Smaller models and faster training. *arXiv preprint arXiv:2104.00298*, 2021.

[12] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf. Wasserstein autoencoders. *arXiv preprint arXiv:1711.01558*, 2017.

[13] E. F. Vernier, R. Garcia, I. da Silva, J. L. D. Comba, and A. C. Telea. Quantitative evaluation of time-dependent multidimensional projection techniques. *arXiv preprint arXiv:2002.07481*, 2020.