Anomaly Detection through Density Matrices and Kernel Density Estimation (AD-DMKDE)

Joseph A. Gallego - jagallegom@unal.edu.co
Oscar A. Bustos-Brinez - oabustosb@unal.edu.co
Fabio A. González - fagonzalezo@unal.edu.co

Abstract

- AD-DMKDE is a novel anomaly detection method that combines density matrices (a mathematical formalism from quantum mechanics) and Fourier features. The method can be seen as an efficient approximation of Kernel Density Estimation (KDE).
- AD-DMKDE was systematically compared against eleven state-of-the-art anomaly detection methods on a variety of benchmark data sets, showing competitive performance.
- The method uses optimization in order to find the parameters of data embedding. Its architecture can be easily implemented on GPU/TPU hardware.
- The prediction stage complexity of AD-DMKDE is constant relative to the training data size, in contrast with KDE, and it performs well in data sets with different anomaly rates.

Neural Density Estimation

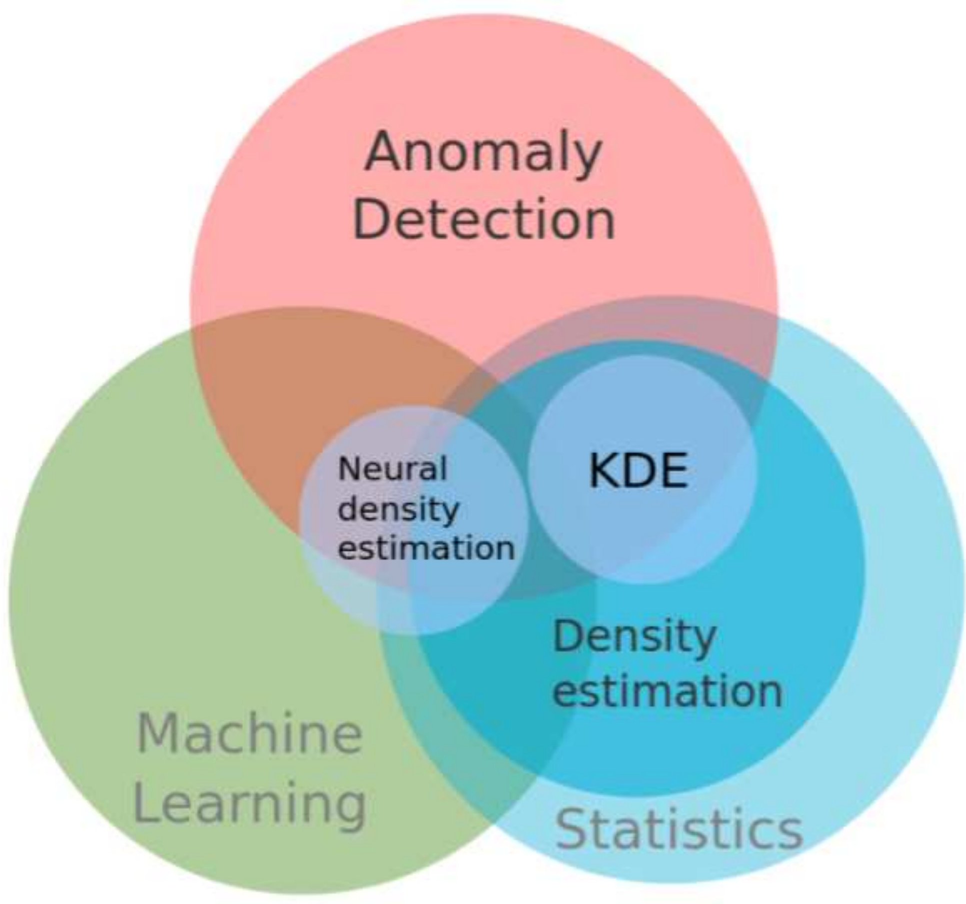

Anomaly Detection

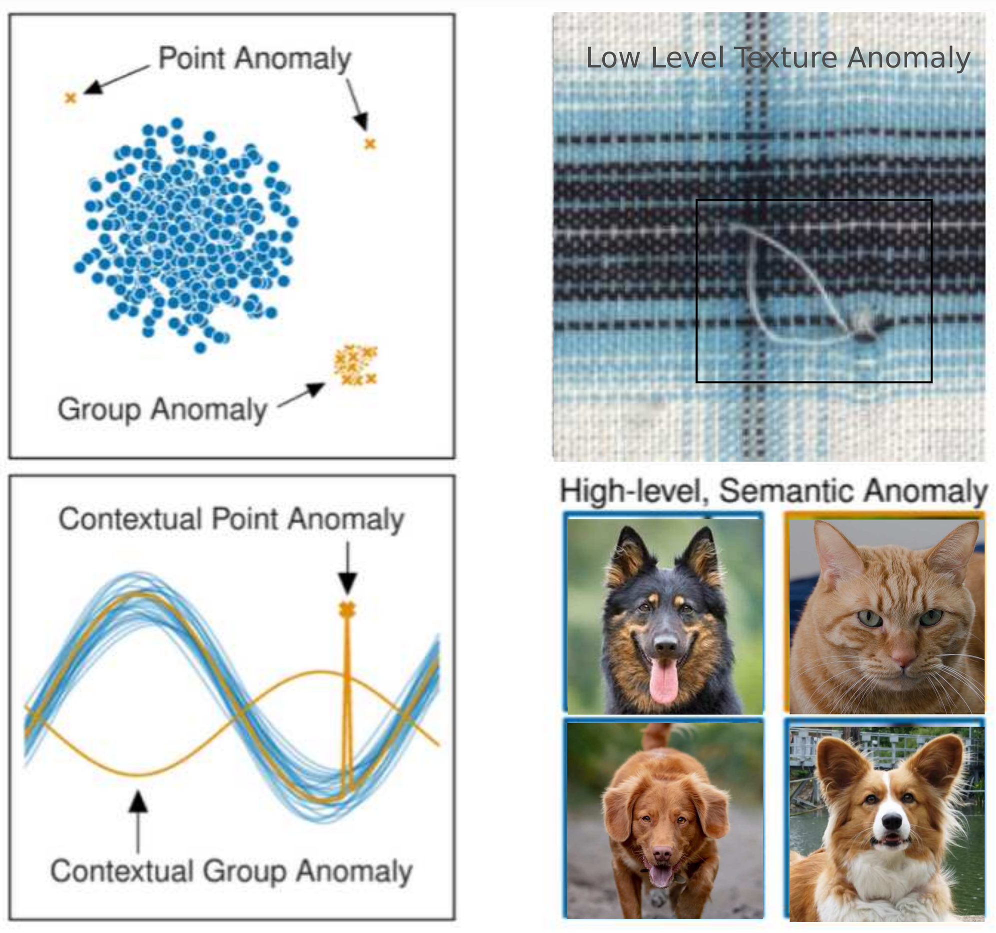

AD-DMKDE Architecture

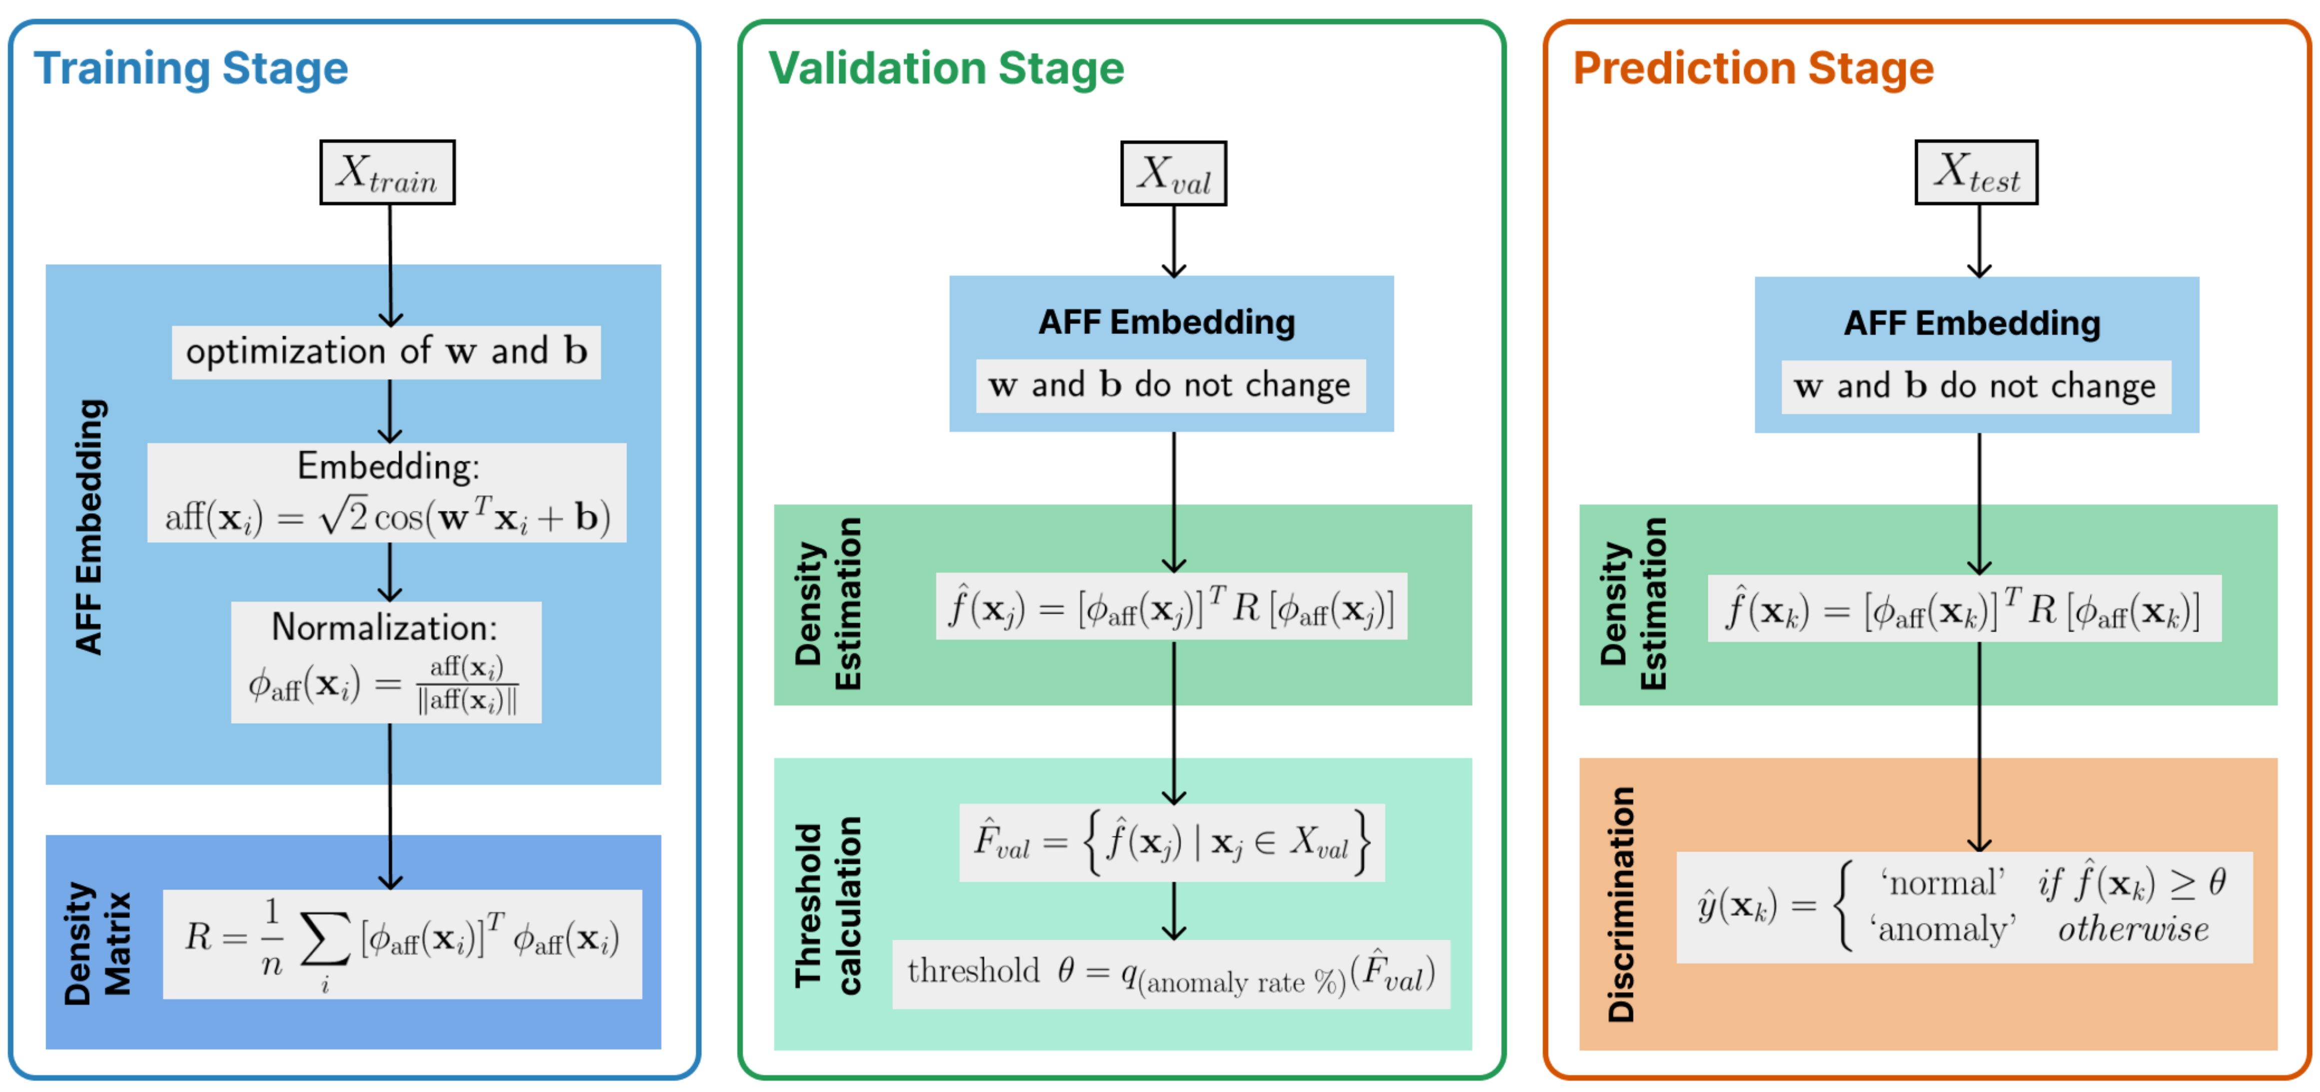

Datasets to work with

Data Set	Instances	Dimensions	Outlier Rate
Arrhythmia	452	274	0,146
Cardio	2060	22	0,2
SpamBase	3485	58	0,2
Thyroid	3772	36	0,0247
KDDCUP	10000	118	0,1934
Glass	214	9	0,042
Lympho	148	18	0,04
Ionosphere	351	33	0,359
Letter	1600	32	0,0625
MNIST	7603	100	0,092
Musk	3062	166	0,0317
OptDigits	5216	64	0,0288
PenDigits	6870	16	0,0227
Pima	768	8	0,349
Satellite	6435	36	0,3164
SatImage	5803	36	0,0122
Shuttle	10000	9	0,0715
Vertebral	240	6	0,125
Vowels	1456	12	0,03434
WBC	378	30	0,0556

Results

Average F1 Score over all datasets:

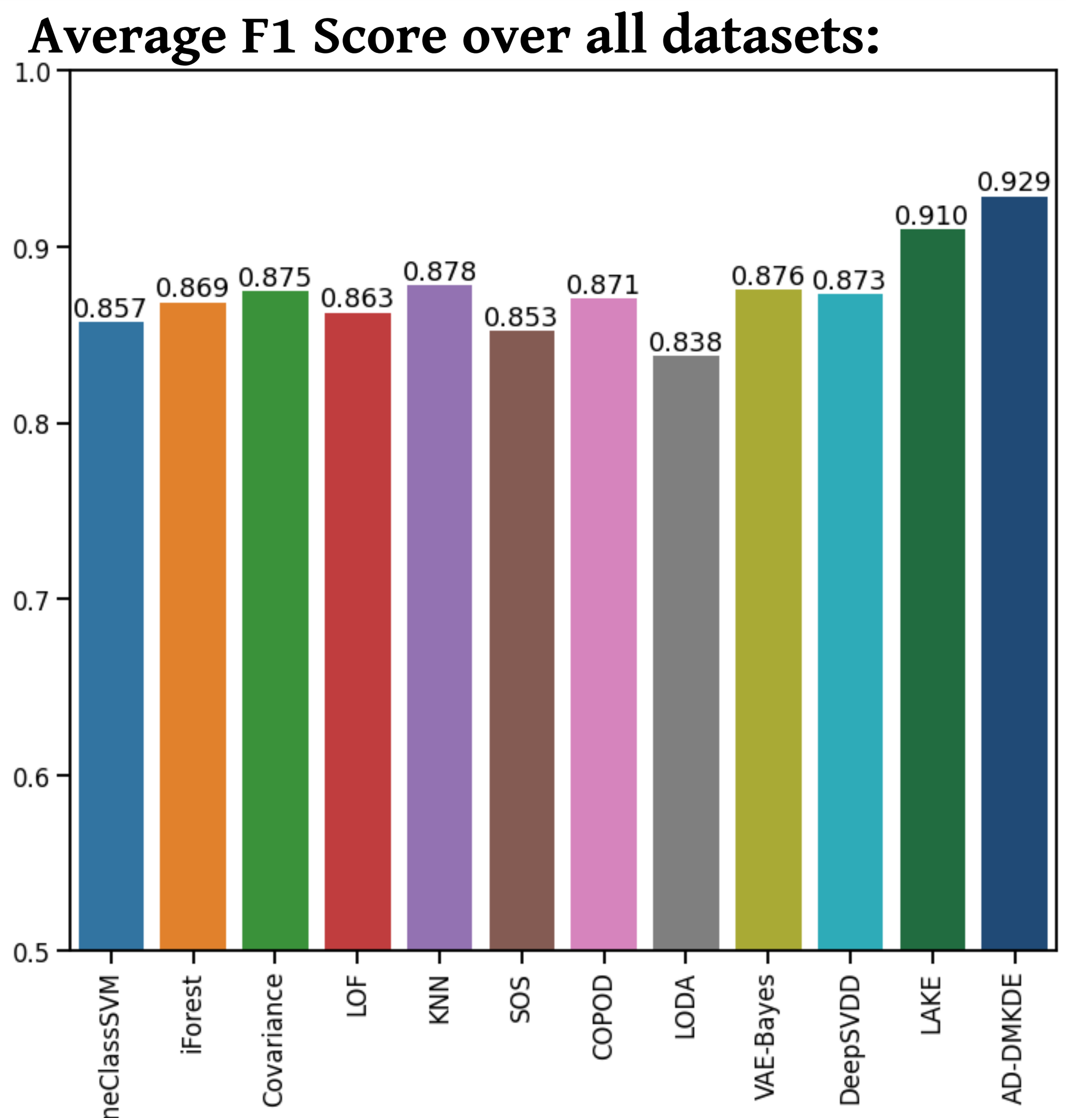

Boxplot of F1 Scores over all datasets:

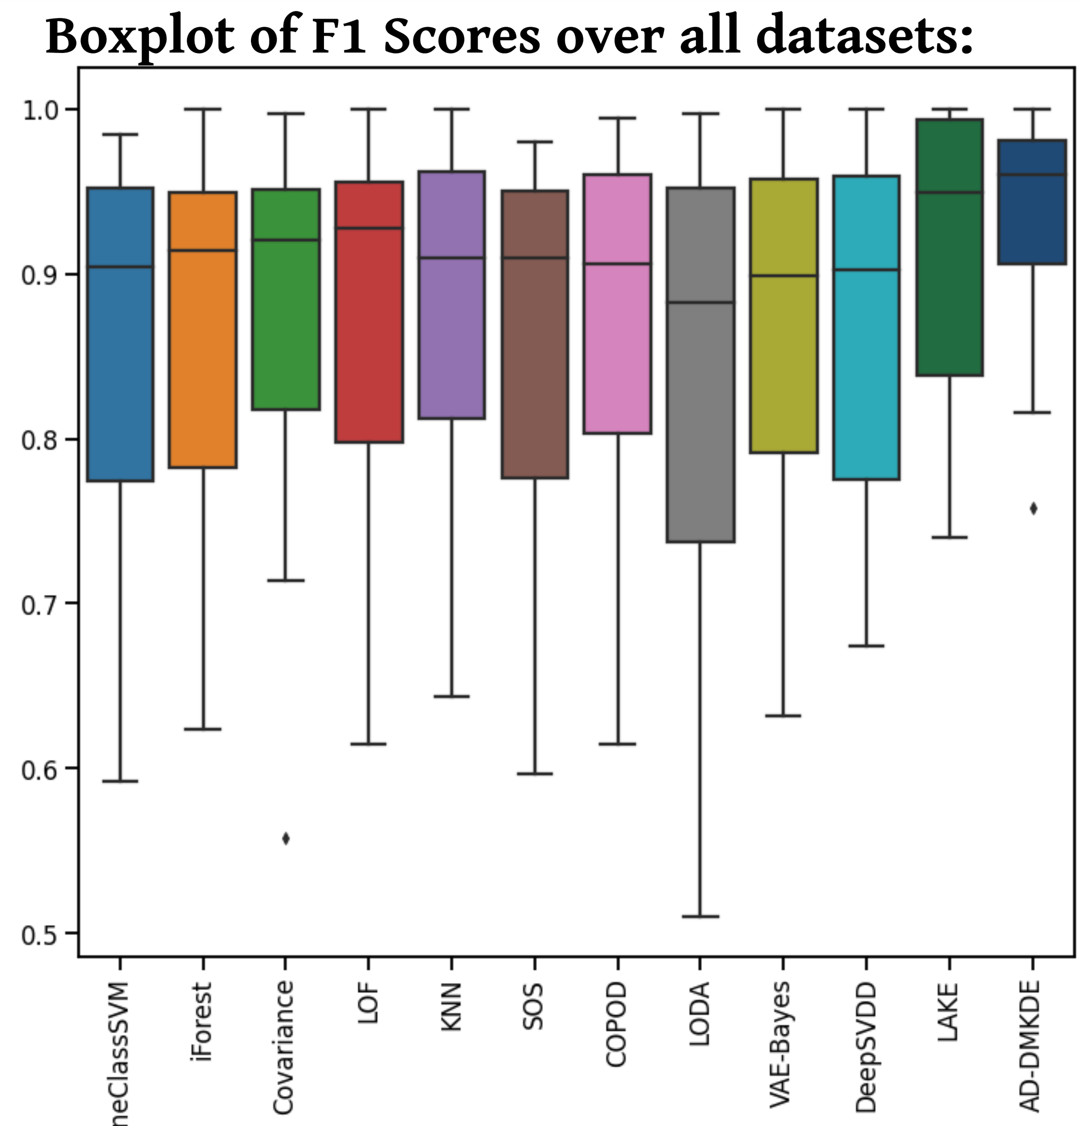

Conclusions / Future Work

- AD-DMKDE show better than state-of-the-art performance over twenty anomaly detection data sets, being superior than classic algorithms and comparable to deep learning-based methods.
- The performance of the method does not seem to be affected by the anomaly rate of the data sets or the size of the data sets, but it seems to perform better for low dimensionality data sets.
- The method does not have huge memory requirements due to that it constructs a single density matrix in all training phase whose size is defined by the embedding.
- As future work, we will continue to further develop the main concepts behind AD-DMKDE, building algorithms based on the combination of Fourier features and density matrices with deeper neural networks, such as autoencoders and variational autoencoders, whose reduction power can be coupled to AD-DMKDE in order to handle data sets with higher dimensionalities.

To read the full paper, please scan the following QR code:

Bibliography

- J. A. Gallego and F. A. González. Quantum adaptive fourier features for neural density estimation, 2022.
- F. A. González, A. Gallego, S. Toledo-Cortés, and V. Vargas-Calderón. Learning with density matrices and random features. *Quantum Machine Intelligence*, 4(2), 2022.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Proceedings of the 20th International Conference on Neural Information Processing Systems*, NIPS'07, page 1177–1184. Curran Associates Inc., 2007.
- L. Ruff, J. R. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, and K. R. Muller. A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109:756–795, 2021.
- Y. Zhao, Z. Nasrullah, and Z. Li. Pyod: A python toolbox for scalable outlier detection. *Journal of Machine Learning Research*, 20, 2019.
- P. Lv, Y. Yu, Y. Fan, X. Tang, and X. Tong. Layer-constrained variational autoencoding kernel density estimation model for anomaly detection. *Knowledge-Based Systems*, 196, 2020.
- S. Rayana. ODDS Library, 2016. URL: `http://odds.cs.stonybrook.edu`