

Are generative deep models for novelty detection truly better?

Vít Škvára
UTIA CAS CR
Pod Vodárenskou věží 4
Prague, Czech Republic
skvara@utia.cas.cz

Tomáš Pevný
FEL CTU Prague
Technická 2
Prague, Czech Republic
pevnytom@fel.cvut.cz

Václav Šmídl
UTIA CAS CR
Pod Vodárenskou věží 4
Prague, Czech Republic
smidl@utia.cas.cz

ABSTRACT

Many deep models have been recently proposed for anomaly detection. This paper presents comparison of selected generative deep models and classical anomaly detection methods on an extensive number of non-image benchmark datasets. We provide statistical comparison of the selected models, in many configurations, architectures and hyperparameters. We arrive to conclusion that performance of the generative models is determined by the process of selection of their hyperparameters. Specifically, performance of the deep generative models deteriorates with decreasing amount of anomalous samples used in hyperparameter selection. In practical scenarios of anomaly detection, none of the deep generative models systematically outperforms the kNN.

CCS Concepts

•Computing methodologies → Artificial intelligence;

Keywords

Anomaly detection, generative models, neural networks

1. INTRODUCTION

An anomaly is a data sample that is so different from the rest of the normal data that it was likely generated by a different underlying process. Anomalous samples are object of interest for various reasons and methods for novelty or anomaly detection try to identify them. These methods are used in a plethora of domains, including medical, computer security or sensor data collection. An overview of anomaly detection methods is presented in [15].

This paper asks the following question – how do anomaly detection methods based on deep neural-network generative models stand in comparison to the methods based on alternative paradigms? In the fashion of [13] or [5], we do not propose a new algorithm to prove it performs better than the existing methods. Instead, we aim to clarify whether the existing generative models actually bring a significant

improvement over classical approaches. To our knowledge, such comparison of anomaly detection methods based on neural networks to some simpler methods is missing. Sadly, most papers proposing new methods (especially those based on deep neural networks) focus on large image data from MNIST, CIFAR or some other publicly available database that pose an unsurmountable task for classical anomaly detection methods. However, image-based anomaly problems are a rather niche domain, therefore this paper focuses on evaluation on a varied range of datasets. When presenting a novel generative model, the authors usually compare to some baseline generative model against which they are trying to improve, a limited number of classical models such as PCA or OSVM, and at best they use the KDD dataset – see e.g. [1] or [22].

Furthermore, no one has ever done a thorough comparative study of generative deep models on a large number of different datasets. In the following text, a handful of selected generative models will be compared against each other and classical novelty detection methods in a statistical way on a large number of carefully crafted benchmark datasets. We do not claim to provide a complete overview of generative models applicable to anomaly detection task, but a general comparison that may simplify future assessment of novel methods. We admit that we have not tested or implemented state-of-the-art methods such as [24] or [23].

Additionally, we endeavor to create a standardized, publicly available implementation of different models. Apart from the actual model implementations, we create a framework for proper training, testing and comparing the models. The framework provides an interface for easy evaluation of new algorithms.

Finally, we propose a number of important questions that may lead to a better understanding of individual tested algorithms and their behavior on different datasets. Using experimental data and a thorough comparison methodology, we may be able to answer some of these questions or provide valuable insight.

2. BACKGROUND

Generative models are used to generate samples from some learned data distribution $p(x)$. In the case of anomaly detection, this is the distribution of normal data. Therefore, even though all following methods are unsupervised by default, we strive to train the generative model with (mostly) normal data. Even though it may be expensive to obtain labels, they are also useful for tuning of hyperparameters, as will be discussed later.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ODD workshop, SIGKDD '18 August 20, 2018, London, GB

© 2018 ACM. ISBN 978-1-4503-2138-9.

DOI: 10.1145/1235

When the model has learned the normal data distribution, it can be used to compute an anomaly score function $f : \mathcal{X} \rightarrow \mathbb{R}$ for a sample $x \in \mathcal{X}$ from the data space. The convention used here is that higher the anomaly score, the more likely it is that the sample is an anomaly. For generative models, reconstruction error (in case of autoencoders), discriminator score (for adversarial models) or their combination can be used.

This section contains a brief theoretical background of the used generative models based on neural nets and two classical methods. The description of cost functions, anomaly score functions and their parameters will be given.

2.1 Autoencoding models

2.1.1 Autoencoder

An autoencoder (AE) is a cornerstone of many multi layer perceptron (MLP) models. Although not a generative model by itself, it may be used as a baseline for the variational autoencoder generative model. It is easily used for the anomaly detection task (see [16], [20] or the comprehensive review of anomaly detection methods [15]).

It consists of two MLPs - an encoder and a decoder. The encoder represents the mapping $e_\phi : \mathcal{X} \rightarrow \mathcal{Z}$ that projects a sample x from the data space to code z in the latent space. The decoder reconstructs the code back to the data space via mapping $d_\theta : \mathcal{Z} \rightarrow \mathcal{X}$. From here on, θ and ϕ denote hidden parameters of neural nets. Both parts of the autoencoder are trained using backpropagation by minimizing the reconstruction error

$$\mathcal{L}_r(x, \phi, \theta) = \|x - d_\theta(e_\phi(x))\|_2^2. \quad (1)$$

Because the dimension of the code z is smaller than that of x , the autoencoder is forced to learn an efficient sparse representation of x while being robust to noise [21].

This roughly equals to learning to reconstruct samples coming from the distribution $p(x)$ of the normal data. When the trained autoencoder is shown an anomaly, it will likely not produce a good reconstruction as it has not been trained with similar samples. Therefore, the anomaly score is given by the reconstruction error of a sample

$$f_{\text{AE}}(x) = \mathcal{L}_r(x, \bar{\phi}, \bar{\theta}), \quad (2)$$

where $\bar{\phi}, \bar{\theta}$ are the learned parameters of the AE.

2.1.2 Variational autoencoder

The variational autoencoder (VAE) [11] and its modifications has seen a lot of success especially in generating realistic images. Its design is very similar to that of an ordinary autoencoder. To induce the generative property, we force the encoder to produce codes z that resemble samples from some easy-to-sample-from prior distribution $p(z)$ (e.g. $\mathcal{N}(0, 1)$). Afterwards, new samples that resemble the real data can be generated by inputting codes sampled from $p(z)$ to the decoder.

In VAE, both the encoder and decoder model parameters of conditional distributions denoted as $q_\phi(z|x)$ and $p_\theta(x|z)$. We assume that the distributions are Gaussian, therefore at the output layer of the MLPs we obtain the mean and variance of the respective distributions. During training, the aim is to minimize the reconstruction loss of x while simultaneously minimizing the Kullback-Leibler divergence $D_{KL}(q_\phi(z|x)||p(z))$, which is zero if the two distributions

are equal. Also, an important part of the training is the *reparametrization trick*, which creates random decoder inputs in the following manner: $z = \mu_z + \sigma_z \epsilon$, where μ_z and σ_z are outputs of the encoder and ϵ is sampled from $p(z)$. The cost function for training a VAE is following

$$\mathcal{L}_v(x, \phi, \theta) = \mathbb{E}_{q(z|x)} [\|x - d_\theta(z)\|_2^2] + \lambda D_{KL}(q_\phi(z|x)||p(z)), \quad (3)$$

where $\lambda = \sigma_x^2$ is a tuning parameter equal to a known variance of data. Since both $q_\phi(z|x)$ and $p(z)$ are Gaussian, there is an analytical expression for their KL divergence.

There are numerous papers describing the use of VAE for anomaly detection - [19], [22], [6], but none make a more complete comparison with other generative and classical methods and only some use non-image datasets. In [1], the authors describe the advantages of VAE over AE - it generalizes more easily since it is working on probabilities. The anomaly score function of the VAE is the reconstruction error

$$f_{\text{VAE}}(x) = \mathbb{E}_{q(z|x)} [\|x - d_\theta(z)\|_2^2]. \quad (4)$$

Alternatively, the log-likelihood of the code z due to prior $p(z)$ can be used.

2.2 Adversarial models

2.2.1 GAN

The simplest adversarial generative model is the *generative adversarial network* - GAN [9]. It consists of two adversaries - a generator and a discriminator that are represented by MLPs. The generator creates samples that resemble the real data, while the discriminator is trying to recognize the fake samples from the real ones. During training, they both improve - generator creates more believable samples while the discriminator gets more proficient at recognizing fakes.

Inputs of the generator are codes $z \sim p(z)$, where $p(z)$ is e.g. standard or uniform distribution. The generator is a mapping into the data space $g_\phi : \mathcal{Z} \rightarrow \mathcal{X}$. Discriminator is a mapping $d_\theta : \mathcal{X} \rightarrow [0, 1]$, i.e. its output is a scalar representing the probability that a sample comes from the true data distribution $p(x)$. The training alternates between minimizing the logit cross-entropy for discriminator

$$\mathcal{L}_d(x, \theta) = -\mathbb{E}_{p(x)} [\log d_\theta(x)] - \mathbb{E}_{p(z)} [\log(1 - d_\theta(g_\phi(z)))] \quad (5)$$

and a simplified logit cross-entropy for the generator

$$\mathcal{L}_g(x, \phi) = -\mathbb{E}_{p(z)} [\log d_\theta(g_\phi(z))]. \quad (6)$$

After training, one can input samples z to the generator and it should be able to generate samples that resemble those from $p(x)$.

For anomaly detection application, $p(x)$ is the distribution of normal data. We do not need to know the true form of $p(x)$, we only need to be able to sample from it. The discriminator should ideally learn the however complicated form of $p(x)$ by backpropagation. Anomaly score of a sample x is computed as a weighted average of the discriminator output and a reconstruction error of the generated sample

$$f_{\text{GAN}}(x) = -(1 - \lambda) \log(d_\theta(x)) + \lambda \|x - g_\phi(z)\|_2, \quad (7)$$

where $z \sim p(z)$ and λ is a scalar scaling parameter.

2.2.2 Feature-matching GAN

While GAN enjoys success in generation of realistic images, it is famously difficult to train. The authors of [17] proposed a number of modifications to the original simple training process. One of them is the addition of the *feature-matching* loss to the cost function. It is based on the idea that backpropagation based on a loss computed somewhere else than the final scalar output of the discriminator may provide improved gradients for the generator, thus enabling a more stable training procedure.

The feature-matching GAN (fmGAN) uses an augmented generator cost function

$$\mathcal{L}_f(x, \phi) = \alpha \mathcal{L}_g(x, \phi) + \mathbb{E}_{p(x), p(z)} [\|h_\theta(x) - h_\theta(g_\phi(z))\|_2], \quad (8)$$

where h_θ is the output of some intermediate (e.g. the penultimate) layer of the discriminator and α is a scalar scaling parameter.

The fmGAN has been successfully used for anomaly detection in [18]. The anomaly score function is the same as in the case of the GAN model.

2.3 Classical outlier detection methods

2.3.1 kNN

The k -nearest neighbours anomaly detection algorithm [2] is a simple yet powerful model. It is based on the assumption that normal data are grouped in the data space and anomalies are distant from them and therefore can be detected by measuring their distance from the rest of the data. It is relatively easy to implement, well described in literature and quite well performing. A large study [5] concluded that no classical anomaly detection algorithm provides a comprehensive improvement over kNN.

For faster distance computation, the training data are encoded in a KDTree structure [3]. This significantly improves the prediction times on large datasets but requires additional overhead in construction of the tree. The only hyperparameter of the algorithm is k . The version of kNN used for experiments in this paper computes the anomaly score as the average distance to k -nearest neighbours in the training dataset.

2.3.2 Isolation forest

Isolation forest [12] is also one of the more widely used anomaly detection algorithms. During training, the algorithm randomly and recursively partitions the data and stores this partitioning in a tree structure. This partitioning is done numerous times – thus a forest is constructed. Anomalies should be ideally reached on a shorter path in the tree as they require a smaller number of partitionings to be separated from the rest of the data. The hyperparameter to be tuned is the number of trees n_t . The anomaly score is computed as the negative of the average path length of a sample over all trees in the forest.

3. EXPERIMENTS

This section describes the setup of experiment that was conducted on a large number of benchmark datasets. Datasets, algorithm and experiment implementation in the Julia language is made publicly available in a GitHub repository at <https://github.com/smidl/AnomalyDetection.jl>.

3.1 Experiment setup

| algorithm | hyperparameters |
|------------------|---|
| kNN | $k \in \{1, \frac{1}{2}\sqrt{N}, \sqrt{N}, \frac{3}{2}\sqrt{N}, 2\sqrt{N}\}$ |
| Isolation Forest | $n_t \in \{100, 500\}$ |
| all deep models | $\dim(\mathcal{Z}) \in \{0.1\tilde{M}, 0.2\tilde{M}, \dots, 0.5\tilde{M}\},$ $\tilde{M} = \min(200, M)$ $n_h \in \{1, 2, 3\}$ |
| VAE | $\lambda \in \{10^{-4}, 10^{-3}, \dots, 1\}$ |
| GAN | $\lambda \in \{0, 0.2, \dots, 1.0\}$ |
| fmGAN | $\alpha \in \{10^{-4}, 10^{-2}, \dots, 10^4\}$ $\lambda \in \{0, 0.2, \dots, 1.0\}$ |

Table 1: Tested hyperparameter settings. M is the dimensionality of a dataset, N is the number of training samples, n_t is the number of trees, n_h is the number of hidden layers.

3.1.1 Benchmark data

In order to evaluate the selected models in the most realistic fashion, the methodology of creating appropriate anomaly detection benchmarks from [8] was adopted. A list of 35 pre-processed basic datasets from [14] was used. A more detailed overview of the basic datasets and of the preprocessing can be found in the same paper.

The basic datasets were originally created from multiclass classification or regression datasets that were split into normal and anomalous data, whitened and further processed. Afterwards, it is possible to sample normal and anomalous data according to different criteria. First, there are four levels of anomaly difficulty - *easy*, *medium*, *hard* and *very hard*. Second, a contamination rate (the percentage of anomalies in a sampled dataset) can be specified. Third, it is possible to sample anomalies that are clustered or unclustered.

For the purposes of our experiments, we have sampled from each basic dataset to obtain a non-contaminated (containing no anomalies) training dataset and a contaminated testing dataset. We have used 80% of normal data for training and 20% for testing. In each of the basic datasets, there is a different number of anomalies of a given anomaly difficulty level (see the summary table in [14]). Therefore, the anomaly difficulty levels were chosen (mostly *easy* or *medium*) so that there were enough anomalies to be sampled from with respect to the chosen contamination rate. The contamination rate of the testing dataset was 5% and we sampled for non-clustered anomalies. Random sampling was repeated 10 times for each basic dataset to obtain diverse training and testing datasets .

3.1.2 Algorithm setup

The table 1 contains an overview of tested hyperparameter values. An experiment consisting of training of a model and predicting anomaly scores was carried out for every possible combination of hyperparameters. Therefore, a different number of experiments per dataset sampling was realized for different algorithms – that is no more than 5 experiments for kNN but up to 450 experiments in the case of fmGAN. For training of the deep models, the Adam optimizer [10] was used with a learning rate of 0.001. We trained each generative model for a total of 10000 steps with a batchsize of 256 or lower in case of datasets that are not numerous enough. Relu activation functions and dense layers were used. Size of hidden layers was given by linear interpolation between the code dimension $\dim(\mathcal{Z})$ and a proper input/output dimen-

sion. The code dimension was proportionate to the input dimension $\dim(\mathcal{X})$ but never exceeded 100.

3.2 Evaluation methodology

3.2.1 Algorithm ranking

For comparison of different algorithms on multiple datasets, we use the statistical procedure described in [7]. For a given comparison metric, e.g. AUROC (area under ROC curve), we rank algorithms according to their performance on a dataset. If there is a tie, then average rank is given to the tied algorithms. Afterwards, an average rank over datasets is computed for each algorithm. See e.g. Table 5, where the evaluation metric is the maximum AUROC on the testing dataset (the first ranking criterion in the following text). Based on the average ranks, a number of statistical hypotheses may be tested. In the case of this study, we are interested in two: a) is there enough statistical evidence to reject the hypothesis that all the algorithms perform equally and b) do two individual algorithms perform differently on a statistically significant level? The answers to these questions are given by the Friedman and Nemenyi test, respectively. For our problem size (6 algorithms and 35 datasets), the critical value for the Nemenyi test at 5% confidence level is $CD_{0.05} \approx 1.2745$. If the average rank of two algorithms is larger than that, we can say that their performance is significantly different.

A concise way of visualization are the critical difference diagrams, see e.g. Fig 1. The average algorithm ranks are marked and the statistically equally performing algorithms are connected together with a wide black line.

3.2.2 Hyperparameter selection

After training of an algorithm on the training dataset, a vector of anomaly scores was computed for the testing dataset. Also, a small sample of anomalies was added to the training dataset and a vector of anomaly scores of the resulting ensemble was computed (reasoning behind this follows). Afterwards, testing and training AUROC were computed from these two vectors and known labels. This was done for each individual experiment specified by a combination of dataset, model, resampling iteration and model hyperparameter setting.

We assumed that tuning of hyperparameters is done during training. The rankings of models were produced using the testing AUROC of best-performant model in a resampling iteration, averaged over resampling iterations on a dataset. To simulate different real-world conditions in which hyperparameters are tuned, 3 methods of their selection were used.

1. Firstly, the hyperparameters were selected by the best AUROC on the testing dataset. This gives an idea of the maximum potential of the algorithm, which may be difficult to attain in reality. See Table 5.
2. Secondly, the hyperparameters were selected by best performance on training dataset with added anomalies (training AUROC). This simulates a more realistic scenario of hyperparameter tuning, where all the labels are known on the training dataset. See Table 6.
3. Thirdly, the hyperparameters were selected using the $p\%$ most anomalous samples of the training dataset

| | kNN | IForest | AE | VAE | GAN | fmGAN |
|-----------|------|---------|------|------|------|-------|
| test auc | 3.94 | 5.63 | 3.47 | 2.07 | 3.90 | 1.99 |
| train auc | 3.13 | 4.61 | 3.63 | 2.84 | 4.46 | 2.33 |
| top 5% | 2.57 | 4.07 | 3.24 | 2.73 | 4.90 | 3.49 |
| top 1% | 2.14 | 3.53 | 3.13 | 2.93 | 4.97 | 4.30 |

Table 2: Average ranks of algorithms for different hyperparameter selection criteria.

| | kNN | IForest | AE | VAE | GAN | fmGAN |
|-----------|------|---------|--------|---------|--------|--------|
| t_f [s] | 0.15 | 0.76 | 276.33 | 1088.71 | 844.83 | 879.94 |
| t_p [s] | 5.35 | 0.58 | 2.14 | 3.49 | 6.40 | 6.38 |

Table 3: Average fit t_f and predict t_p times.

with added anomalies. However, the performance was measured using the precision (rate of detected anomalies). This is the most realistic scenario in case when the training labels are expensive to obtain. The common practice is then to manually inspect the most $p\%$ anomalous samples in the training dataset and select the hyperparameters of the most precise model. For our experiments, we chose $p \in \{1, 5\}$. See Table 7 and 8.

These ranking methods stimulate interesting questions.

- What is the reason of an algorithm performing well in the first method and not so well in the others? Insufficient tuning of hyperparameters? Overfitting on training? Or just very difficult applicability to real-world problems where labels are not present and such tuning is not possible?
- What is the robustness of different algorithms to variation of the training data? How robust is it to a possible presence of anomalies in the training data?
- How many samples need to be manually inspected to obtain robust optimal hyperparameter settings?
- Can we analyze on which datasets are some methods more successful than the others? Do these datasets have something in common?




Figure 1: Critical difference diagram for the first hyperparameter selection criterion.

3.3 Experimental results

Mean ranks of algorithms across datasets for all hyperparameter selection methods are summarized in Table 2 (the complete AUC scores are shown in Tables 5 – 8 in Appendix). Results from Table 2 shows that generative models, namely VAE and fmGAN, can produce excellent results if one has a large number of known anomalies to select right hyperparameters (methods *test auc* and *train auc*). Once the number of known anomalies decreases, as is the case of

the more realistic methods *top 1%* and *top 5%*, generative models becomes inferior to methods robust to hyperparameter selection, namely k-nearest neighbours.

The Friedman test rejects the hypothesis that the models perform equally well in all four ranking criteria. Critical difference diagrams shown in Figs. 1 – 4 reveal that unless one selects hyperparameters on testing set, which is obviously cheating, not a single generative model provides a statistically significant improvement over the naive kNN anomaly detection algorithm. VAE seems to be the most robust and therefore more promising deep model, as its ranks relatively well for all the hyperparameter selection criteria.

Although the deep models do not generally outperform kNN when tuned with a limited number of labels, they can still perform well on certain datasets. See Table 8, datasets *cardiotocography*, *libras* and *sonar* for the ranks of fmGAN. In Table 7, fmGAN again performs well on these (and additional) datasets. The same can be observed for VAE and the *ecoli*, *miniboone* and *multiple-features* datasets. This may mean that these datasets share some common characteristic which makes them suitable for application of deep models, rather than just a random fluctuation caused by the large number of experiments carried out for fmGAN and VAE. However we have not yet been able to find the connection. It is certainly not the difficulty of the anomalies, as kNN performs relatively well across all difficulty levels.

The only area in which the deep models outperform kNN is the mean prediction time on large datasets (in terms of number of training samples), see Table 3. Note that the largest benchmark dataset is *miniboone* with 93565 normal samples. In deep models, the prediction is independent on the training data size. This is however compensated for by the computational demands of their training.




Figure 2: Critical difference diagram for the second hyperparameter selection criterion.




Figure 3: Critical difference diagram for the third hyperparameter selection criterion at 5% most anomalous samples.




Figure 4: Critical difference diagram for the third hyperparameter selection criterion at 1% most anomalous samples.

4. CONCLUSIONS

In this paper, a selection of deep generative models adopted for anomaly detection were compared against traditional methods of k-nearest neighbours and Isolation Forests. This was done in a systematic way on a large number of benchmark datasets. Most authors of novel anomaly detection methods based on deep models do not compare to kNN, but instead compare to methods such as Local Outlier Factor [4], despite it has been shown in [14] and [5] that it does not systematically outperform kNN algorithm. We have shown that the robustness of kNN still holds even in comparison with some deep generative models.

From the conducted experiments, the main conclusion is that the major bottleneck for reliable employment of deep generative models for anomaly detection is the difficulty of hyperparameter tuning. Ideally, hyperparameter tuning should be automated and included in the training procedure, which is not the case. Especially adversarial models do not seem to be robust and their training is difficult. On the other hand, generative models showed promising potential and should be studied in a greater depth. We have studied performance of different models under a relatively simple, yet realistic, hyperparameter selection methods and criteria. An interesting question is the existence of a criterion that is more appropriate in the context of limited number of available labels.

This also leads us to questioning the appropriateness of the use of deep generative models in the unsupervised setting. Since we have shown that labels are necessary for hyperparameter tuning, maybe it is worth investigating a semi/fully-supervised deep generative models for anomaly detection.

5. ACKNOWLEDGMENTS

Research presented in this work has been supported by the Grant SGS18/188/OHK4/3T/14 provided by the Ministry of Education, Youth, and Sports of the Czech Republic (MŠMT ČR) and by the GACR project 18-21409S.

6. REFERENCES

- [1] J. An and S. Cho. Variational autoencoder based anomaly detection using reconstruction probability. *SNU Data Mining Center, Tech. Rep.*, 2015.
- [2] F. Angiulli and C. Pizzuti. Fast outlier detection in high dimensional spaces. In *European Conference on Principles of Data Mining and Knowledge Discovery*, pages 15–27. Springer, 2002.
- [3] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.
- [4] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander. Lof: identifying density-based local outliers. In *ACM sigmod record*, volume 29, pages 93–104. ACM, 2000.
- [5] G. O. Campos, A. Zimek, J. Sander, R. J. Campello, B. Micenková, E. Schubert, I. Assent, and M. E. Houle. On the evaluation of unsupervised outlier detection: measures, datasets, and an empirical study. *Data Mining and Knowledge Discovery*, 30(4):891–927, 2016.
- [6] S. Clachar. *Novelty detection and cluster analysis in time series data using variational autoencoder feature maps*. The University of North Dakota, 2016.

[7] J. Demšar. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7(Jan):1–30, 2006.

[8] A. F. Emmott, S. Das, T. Dietterich, A. Fern, and W.-K. Wong. Systematic construction of anomaly detection benchmarks from real data. In *Proceedings of the ACM SIGKDD workshop on outlier detection and description*, pages 16–21. ACM, 2013.

[9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[10] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

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

[12] F. T. Liu, K. M. Ting, and Z.-H. Zhou. Isolation forest. In *Data Mining, 2008. ICDM’08. Eighth IEEE International Conference on*, pages 413–422. IEEE, 2008.

[13] A. Odena, A. Oliver, C. Raffel, E. D. Cubuk, and I. Goodfellow. Realistic evaluation of semi-supervised learning algorithms. 2018.

[14] T. Pevný. Loda: Lightweight on-line detector of anomalies. *Machine Learning*, 102(2):275–304, 2016.

[15] M. A. Pimentel, D. A. Clifton, L. Clifton, and L. Tarassenko. A review of novelty detection. *Signal Processing*, 99:215–249, 2014.

[16] M. Sakurada and T. Yairi. Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the MLSDA 2014 2nd Workshop on Machine Learning for Sensory Data Analysis*, page 4. ACM, 2014.

[17] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2234–2242, 2016.

[18] T. Schlegl, P. Seeböck, S. M. Waldstein, U. Schmidt-Erfurth, and G. Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International Conference on Information Processing in Medical Imaging*, pages 146–157. Springer, 2017.

[19] M. Sölch, J. Bayer, M. Ludersdorfer, and P. van der Smagt. Variational inference for on-line anomaly detection in high-dimensional time series. *arXiv preprint arXiv:1602.07109*, 2016.

[20] B. B. Thompson, R. J. Marks, J. J. Choi, M. A. El-Sharkawi, M.-Y. Huang, and C. Bunje. Implicit learning in autoencoder novelty assessment. In *Neural Networks, 2002. IJCNN’02. Proceedings of the 2002 International Joint Conference on*, volume 3, pages 2878–2883. IEEE, 2002.

[21] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11(Dec):3371–3408, 2010.

[22] H. Xu, W. Chen, N. Zhao, Z. Li, J. Bu, Z. Li, Y. Liu,

| | | | |
|-------------------------|-----|-------------------------|-----|
| abalone | aba | page-blocks | pag |
| blood-transfusion | blo | parkinsons | par |
| breast-cancer-wisconsin | brc | pendigits | pen |
| breast-tissue | brt | pima-indians | pim |
| cardiotocography | car | sonar | son |
| ecoli | eco | spect-heart | spe |
| glass | gla | statlog-satimage | ssa |
| haberman | hab | statlog-segment | sse |
| ionosphere | ion | statlog-shuttle | ssh |
| iris | iri | statlog-vehicle | sve |
| isolet | iso | synthetic-control-chart | syn |
| letter-recognition | let | vertebral-column | ver |
| libras | lib | wall-following-robot | wal |
| madelon | mad | waveform-1 | wa1 |
| magic-telescope | mag | waveform-2 | wa2 |
| miniboone | min | wine | win |
| multiple-features | mul | yeast | yea |
| musk-2 | mus | | |

Table 4: Overview of coded dataset names used in Tables 5 – 8.

Y. Zhao, D. Pei, Y. Feng, et al. Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications. *arXiv preprint arXiv:1802.03903*, 2018.

[23] S. Zhai, Y. Cheng, W. Lu, and Z. Zhang. Deep structured energy based models for anomaly detection. *arXiv preprint arXiv:1605.07717*, 2016.

[24] B. Zong, Q. Song, M. R. Min, W. Cheng, C. Lumezanu, D. Cho, and H. Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *International Conference on Learning Representations*, 2018.

APPENDIX

Here, selected results of scores and ranks for algorithms across all datasets are shown. For brevity, we use short codes for dataset names, see Table 4.

| dataset | kNN | IForest | AE | VAE | GAN | fmGAN |
|---------|------------|------------|------------|------------|-----------|------------|
| aba | 0.93(3.5) | 0.83(6.0) | 0.94(1.5) | 0.94(1.5) | 0.88(5.0) | 0.93(3.5) |
| blo | 0.99(4.0) | 0.98(6.0) | 0.99(4.0) | 1.00(1.5) | 0.99(4.0) | 1.00(1.5) |
| brc | 0.94(6.0) | 0.98(4.5) | 0.99(2.0) | 0.99(2.0) | 0.98(4.5) | 0.99(2.0) |
| brc | 0.99(5.5) | 0.99(5.5) | 1.00(2.5) | 1.00(2.5) | 1.00(2.5) | 1.00(2.5) |
| car | 0.66(6.0) | 0.70(4.0) | 0.67(5.0) | 0.83(3.0) | 0.85(2.0) | 0.91(1.0) |
| eco | 0.95(5.0) | 0.88(6.0) | 0.97(4.0) | 0.99(2.0) | 0.98(3.0) | 1.00(1.0) |
| gla | 0.90(5.0) | 0.86(6.0) | 0.98(4.0) | 1.00(2.0) | 1.00(2.0) | 1.00(2.0) |
| hab | 0.98(3.0) | 0.95(6.0) | 0.97(4.5) | 0.99(2.0) | 0.97(4.5) | 1.00(1.0) |
| ion | 0.96(5.0) | 0.94(6.0) | 0.99(3.5) | 1.00(1.5) | 0.99(3.5) | 1.00(1.5) |
| iri | 0.98(5.0) | 0.94(6.0) | 1.00(2.5) | 1.00(2.5) | 1.00(2.5) | 1.00(2.5) |
| iso | 0.78(3.0) | 0.59(6.0) | 0.76(4.0) | 0.81(2.0) | 0.64(5.0) | 0.83(1.0) |
| let | 0.80(3.0) | 0.57(6.0) | 0.79(4.0) | 0.83(1.0) | 0.62(5.0) | 0.81(2.0) |
| lib | 0.88(5.0) | 0.75(6.0) | 0.91(4.0) | 0.98(2.0) | 0.96(3.0) | 0.99(1.0) |
| mad | 0.59(6.0) | 0.60(5.0) | 0.65(4.0) | 0.71(2.5) | 0.71(2.5) | 0.79(1.0) |
| mag | 0.95(1.0) | 0.86(5.5) | 0.93(2.5) | 0.93(2.5) | 0.86(5.5) | 0.88(4.0) |
| min | 0.89(4.0) | 0.85(5.0) | 0.76(6.0) | 0.93(1.5) | 0.91(3.0) | 0.93(1.5) |
| mul | 0.99(2.0) | 0.81(6.0) | 0.99(2.0) | 0.99(2.0) | 0.85(5.0) | 0.97(4.0) |
| mus | 0.97(3.0) | 0.61(6.0) | 0.99(1.5) | 0.99(1.5) | 0.89(5.0) | 0.93(4.0) |
| pag | 0.98(3.5) | 0.96(5.0) | 0.99(1.5) | 0.99(1.5) | 0.95(6.0) | 0.98(3.5) |
| par | 0.94(5.0) | 0.87(6.0) | 0.98(4.0) | 0.99(2.5) | 0.99(2.5) | 1.00(1.0) |
| pen | 1.00(1.0) | 0.91(4.5) | 0.99(2.5) | 0.99(2.5) | 0.86(6.0) | 0.91(4.5) |
| pim | 0.91(4.5) | 0.88(6.0) | 0.92(3.0) | 0.94(2.0) | 0.91(4.5) | 0.98(1.0) |
| son | 0.84(5.0) | 0.81(6.0) | 0.95(4.0) | 1.00(2.0) | 1.00(2.0) | 1.00(2.0) |
| spe | 0.91(5.0) | 0.88(6.0) | 0.97(4.0) | 1.00(2.0) | 1.00(2.0) | 1.00(2.0) |
| ssa | 0.97(1.0) | 0.92(5.5) | 0.96(3.0) | 0.96(3.0) | 0.92(5.5) | 0.96(3.0) |
| sse | 0.96(3.0) | 0.85(6.0) | 0.95(4.0) | 0.97(1.5) | 0.92(5.0) | 0.97(1.5) |
| ssh | 1.00(2.0) | 0.96(5.0) | 1.00(2.0) | 1.00(2.0) | 0.94(6.0) | 0.98(4.0) |
| sve | 0.84(5.0) | 0.79(6.0) | 0.91(3.0) | 0.93(2.0) | 0.90(4.0) | 0.95(1.0) |
| syn | 0.98(3.5) | 0.89(6.0) | 0.98(3.5) | 1.00(1.5) | 0.96(5.0) | 1.00(1.5) |
| ver | 0.82(4.5) | 0.79(6.0) | 0.82(4.5) | 0.92(2.0) | 0.90(3.0) | 0.98(1.0) |
| wal | 0.86(2.0) | 0.73(5.0) | 0.79(4.0) | 0.80(3.0) | 0.72(6.0) | 0.87(1.0) |
| wa1 | 0.82(4.0) | 0.81(5.0) | 0.80(6.0) | 0.86(2.0) | 0.85(3.0) | 0.91(1.0) |
| wa2 | 0.82(4.0) | 0.80(5.5) | 0.80(5.5) | 0.86(3.0) | 0.87(2.0) | 0.94(1.0) |
| win | 0.98(5.0) | 0.96(6.0) | 1.00(2.5) | 1.00(2.5) | 1.00(2.5) | 1.00(2.5) |
| yea | 0.84(5.0) | 0.80(6.0) | 0.91(3.0) | 0.93(2.0) | 0.86(4.0) | 0.95(1.0) |
| avg | 0.90(3.94) | 0.84(5.63) | 0.91(3.47) | 0.94(2.07) | 0.90(3.9) | 0.95(1.99) |

Table 5: AUROC scores and ranks of algorithms using the first hyperparameter selection criterion. The last line is an average.

| dataset | kNN | IForest | AE | VAE | GAN | fmGAN |
|---------|------------|------------|------------|------------|------------|------------|
| aba | 0.93(1.0) | 0.83(6.0) | 0.92(2.5) | 0.92(2.5) | 0.84(5.0) | 0.86(4.0) |
| blo | 0.98(4.0) | 0.98(4.0) | 0.98(4.0) | 0.98(4.0) | 0.98(4.0) | 1.00(1.0) |
| brc | 0.93(6.0) | 0.98(1.0) | 0.95(3.0) | 0.94(4.5) | 0.94(4.5) | 0.96(2.0) |
| brc | 0.99(2.0) | 0.99(2.0) | 0.99(2.0) | 0.98(4.0) | 0.95(5.5) | 0.95(5.5) |
| car | 0.66(5.0) | 0.70(4.0) | 0.62(6.0) | 0.80(3.0) | 0.85(2.0) | 0.89(1.0) |
| eco | 0.93(3.0) | 0.87(5.5) | 0.94(1.5) | 0.94(1.5) | 0.87(5.5) | 0.89(4.0) |
| gla | 0.90(5.0) | 0.85(6.0) | 0.91(4.0) | 0.92(3.0) | 0.96(2.0) | 0.98(1.0) |
| hab | 0.97(1.0) | 0.94(3.5) | 0.86(5.0) | 0.94(3.5) | 0.83(6.0) | 0.95(2.0) |
| ion | 0.94(4.5) | 0.94(4.5) | 0.95(3.0) | 0.97(2.0) | 0.90(6.0) | 0.98(1.0) |
| iri | 0.97(5.0) | 0.92(6.0) | 0.98(4.0) | 0.99(2.5) | 0.99(2.5) | 1.00(1.0) |
| iso | 0.78(2.0) | 0.59(5.0) | 0.71(4.0) | 0.72(3.0) | 0.58(6.0) | 0.82(1.0) |
| let | 0.80(1.5) | 0.57(5.0) | 0.77(3.0) | 0.76(4.0) | 0.55(6.0) | 0.80(1.5) |
| lib | 0.87(2.5) | 0.75(6.0) | 0.82(5.0) | 0.83(4.0) | 0.92(1.0) | 0.87(2.5) |
| mad | 0.57(5.0) | 0.59(3.0) | 0.56(6.0) | 0.58(4.0) | 0.61(2.0) | 0.62(1.0) |
| mag | 0.95(1.0) | 0.86(5.0) | 0.93(2.5) | 0.93(2.5) | 0.85(6.0) | 0.88(4.0) |
| min | 0.75(6.0) | 0.85(4.0) | 0.76(5.0) | 0.93(1.5) | 0.91(3.0) | 0.93(1.5) |
| mul | 0.99(1.5) | 0.80(5.0) | 0.98(3.0) | 0.99(1.5) | 0.71(6.0) | 0.93(4.0) |
| mus | 0.97(3.0) | 0.61(6.0) | 0.98(2.0) | 0.99(1.0) | 0.89(5.0) | 0.93(4.0) |
| pag | 0.89(6.0) | 0.96(4.0) | 0.98(2.5) | 0.99(1.0) | 0.94(5.0) | 0.98(2.5) |
| par | 0.94(1.0) | 0.85(5.0) | 0.91(2.0) | 0.83(6.0) | 0.90(3.0) | 0.87(4.0) |
| pen | 1.00(1.0) | 0.91(4.0) | 0.98(3.0) | 0.99(2.0) | 0.86(6.0) | 0.88(5.0) |
| pim | 0.89(2.0) | 0.87(4.0) | 0.86(5.5) | 0.88(3.0) | 0.86(5.5) | 0.92(1.0) |
| son | 0.81(4.5) | 0.81(4.5) | 0.80(6.0) | 0.83(3.0) | 0.96(1.0) | 0.93(2.0) |
| spe | 0.89(1.5) | 0.86(5.0) | 0.89(1.5) | 0.88(3.0) | 0.85(6.0) | 0.87(4.0) |
| ssa | 0.97(1.0) | 0.92(5.0) | 0.95(3.5) | 0.95(3.5) | 0.90(6.0) | 0.96(2.0) |
| sse | 0.96(1.5) | 0.85(5.5) | 0.94(4.0) | 0.96(1.5) | 0.85(5.5) | 0.95(3.0) |
| ssh | 1.00(2.0) | 0.96(5.0) | 1.00(2.0) | 1.00(2.0) | 0.87(6.0) | 0.97(4.0) |
| sve | 0.83(4.0) | 0.78(6.0) | 0.84(3.0) | 0.85(2.0) | 0.82(5.0) | 0.93(1.0) |
| syn | 0.97(1.5) | 0.89(5.0) | 0.93(3.0) | 0.92(4.0) | 0.86(6.0) | 0.97(1.5) |
| ver | 0.68(5.5) | 0.78(4.0) | 0.68(5.5) | 0.84(2.0) | 0.82(3.0) | 0.89(1.0) |
| wal | 0.86(1.0) | 0.73(5.0) | 0.78(3.0) | 0.77(4.0) | 0.70(6.0) | 0.85(2.0) |
| wa1 | 0.77(6.0) | 0.81(4.0) | 0.79(5.0) | 0.83(3.0) | 0.84(2.0) | 0.90(1.0) |
| wa2 | 0.77(6.0) | 0.80(4.0) | 0.78(5.0) | 0.82(3.0) | 0.86(2.0) | 0.93(1.0) |
| win | 0.98(2.5) | 0.96(5.0) | 0.95(6.0) | 0.98(2.5) | 0.97(4.0) | 0.99(1.0) |
| yea | 0.82(3.5) | 0.79(5.0) | 0.89(1.0) | 0.88(2.0) | 0.72(6.0) | 0.82(3.5) |
| avg | 0.88(3.13) | 0.83(4.61) | 0.87(3.63) | 0.89(2.84) | 0.85(4.46) | 0.91(2.33) |

Table 6: AUROC scores and ranks of algorithms using the second hyperparameter selection criterion.

| dataset | kNN | IForest | AE | VAE | GAN | fmGAN |
|---------|------------|------------|------------|------------|-----------|------------|
| aba | 0.93(1.0) | 0.83(4.0) | 0.89(2.0) | 0.87(3.0) | 0.74(6.0) | 0.77(5.0) |
| blo | 0.99(1.0) | 0.98(4.0) | 0.98(4.0) | 0.98(4.0) | 0.98(4.0) | 0.98(4.0) |
| brc | 0.93(4.0) | 0.98(1.0) | 0.95(3.0) | 0.96(2.0) | 0.86(6.0) | 0.87(5.0) |
| brc | 0.99(2.0) | 0.99(2.0) | 0.99(2.0) | 0.98(4.0) | 0.95(5.5) | 0.95(5.5) |
| car | 0.66(4.0) | 0.70(3.0) | 0.63(5.0) | 0.61(6.0) | 0.71(2.0) | 0.77(1.0) |
| eco | 0.93(3.0) | 0.87(4.0) | 0.95(1.5) | 0.95(1.5) | 0.80(5.0) | 0.79(6.0) |
| gla | 0.90(5.0) | 0.84(6.0) | 0.91(4.0) | 0.92(3.0) | 0.93(2.0) | 0.94(1.0) |
| hab | 0.97(1.0) | 0.94(2.5) | 0.86(4.0) | 0.94(2.5) | 0.85(5.0) | 0.83(6.0) |
| ion | 0.94(4.0) | 0.92(5.0) | 0.95(3.0) | 0.97(1.0) | 0.82(6.0) | 0.96(2.0) |
| iri | 0.97(4.5) | 0.92(6.0) | 0.97(4.5) | 0.98(3.0) | 1.00(1.5) | 1.00(1.5) |
| iso | 0.78(1.0) | 0.58(5.0) | 0.74(2.5) | 0.74(2.5) | 0.57(6.0) | 0.69(4.0) |
| let | 0.80(1.0) | 0.56(5.0) | 0.78(2.5) | 0.78(2.5) | 0.55(6.0) | 0.68(4.0) |
| lib | 0.87(1.5) | 0.74(6.0) | 0.82(4.0) | 0.83(3.0) | 0.76(5.0) | 0.87(1.5) |
| mad | 0.57(5.0) | 0.58(3.5) | 0.56(6.0) | 0.58(3.5) | 0.59(2.0) | 0.64(1.0) |
| mag | 0.95(1.0) | 0.86(4.0) | 0.92(2.5) | 0.92(2.5) | 0.78(6.0) | 0.83(5.0) |
| min | 0.75(5.0) | 0.85(4.0) | 0.73(6.0) | 0.91(1.5) | 0.87(3.0) | 0.91(1.5) |
| mul | 0.99(1.5) | 0.80(5.0) | 0.98(3.0) | 0.99(1.5) | 0.72(6.0) | 0.92(4.0) |
| mus | 0.97(3.0) | 0.60(6.0) | 0.98(2.0) | 0.99(1.0) | 0.68(5.0) | 0.80(4.0) |
| pag | 0.89(6.0) | 0.96(3.0) | 0.98(1.5) | 0.98(1.5) | 0.91(5.0) | 0.95(4.0) |
| par | 0.94(1.0) | 0.86(4.0) | 0.91(2.0) | 0.83(5.0) | 0.78(6.0) | 0.87(3.0) |
| pen | 1.00(1.0) | 0.91(4.0) | 0.98(2.5) | 0.98(2.5) | 0.81(6.0) | 0.82(5.0) |
| pim | 0.89(1.0) | 0.86(2.0) | 0.85(3.0) | 0.80(4.0) | 0.79(5.5) | 0.79(5.5) |
| son | 0.81(4.5) | 0.81(4.5) | 0.80(6.0) | 0.83(3.0) | 0.95(1.0) | 0.93(2.0) |
| spe | 0.89(2.5) | 0.86(6.0) | 0.90(1.0) | 0.88(4.0) | 0.87(5.0) | 0.89(2.5) |
| ssa | 0.97(1.0) | 0.92(5.0) | 0.95(2.5) | 0.95(2.5) | 0.87(6.0) | 0.93(4.0) |
| sse | 0.96(1.0) | 0.84(6.0) | 0.92(2.5) | 0.92(2.5) | 0.85(5.0) | 0.88(4.0) |
| ssh | 1.00(1.0) | 0.96(4.0) | 0.99(2.0) | 0.98(3.0) | 0.81(6.0) | 0.92(5.0) |
| sve | 0.83(3.0) | 0.79(5.0) | 0.84(2.0) | 0.86(1.0) | 0.74(6.0) | 0.81(4.0) |
| syn | 0.97(1.5) | 0.86(5.0) | 0.94(3.0) | 0.92(4.0) | 0.85(6.0) | 0.97(1.5) |
| ver | 0.76(2.0) | 0.78(1.0) | 0.68(6.0) | 0.75(3.0) | 0.74(4.5) | 0.74(4.5) |
| wal | 0.86(1.0) | 0.73(5.0) | 0.75(3.5) | 0.75(3.5) | 0.62(6.0) | 0.79(2.0) |
| wa1 | 0.77(5.5) | 0.81(1.5) | 0.79(3.5) | 0.79(3.5) | 0.77(5.5) | 0.81(1.5) |
| wa2 | 0.77(4.5) | 0.80(1.5) | 0.78(3.0) | 0.80(1.5) | 0.75(6.0) | 0.77(4.5) |
| win | 0.98(2.0) | 0.96(5.0) | 0.95(6.0) | 0.98(2.0) | 0.97(4.0) | 0.98(2.0) |
| yea | 0.82(3.0) | 0.79(4.0) | 0.86(2.0) | 0.87(1.0) | 0.70(6.0) | 0.76(5.0) |
| avg | 0.88(2.57) | 0.83(4.07) | 0.86(3.24) | 0.86(2.73) | 0.76(4.9) | 0.80(3.49) |

Table 7: AUROC scores and ranks of algorithms using the third hyperparameter selection criterion, using top 5% of samples.

| dataset | kNN | IForest | AE | VAE | GAN | fmGAN |
|---------|-----------|-----------|-----------|-----------|-----------|-----------|
| aba | 0.93(1.0) | 0.82(4.0) | 0.86(2.0) | 0.84(3.0) | 0.67(6.0) | 0.73(5.0) |
| blo | 0.98(1.5) | 0.98(1.5) | 0.95(5.0) | 0.96(3.5) | 0.96(3.5) | 0.92(6.0) |
| brc | 0.93(3.0) | 0.98(1.0) | 0.92(4.0) | 0.94(2.0) | 0.82(5.0) | 0.80(6.0) |
| brc | 0.99(2.0) | 0.99(2.0) | 0.99(2.0) | 0.96(4.0) | 0.93(6.0) | 0.94(5.0) |
| car | 0.66(3.5) | 0.70(2.0) | 0.64(5.0) | 0.62(6.0) | 0.66(3.5) | 0.75(1.0) |
| eco | 0.93(1.5) | 0.88(3.5) | 0.88(3.5) | 0.93(1.5) | 0.86(5.0) | 0.83(6.0) |
| gla | 0.90(2.5) | 0.84(6.0) | 0.90(2.5) | 0.88(4.0) | 0.87(5.0) | 0.92(1.0) |
| hab | 0.97(1.0) | 0.94(2.0) | 0.86(4.0) | 0.91(3.0) | 0.82(5.0) | 0.79(6.0) |
| ion | 0.96(1.0) | 0.93(3.5) | 0.95(2.0) | 0.93(3.5) | 0.87(5.0) | 0.86(6.0) |
| iri | 0.96(1.0) | 0.92(5.5) | 0.94(3.5) | 0.94(3.5) | 0.95(2.0) | 0.92(5.5) |
| iso | 0.78(1.0) | 0.58(5.5) | 0.74(2.5) | 0.74(2.5) | 0.58(5.5) | 0.69(4.0) |
| let | 0.80(1.0) | 0.56(5.0) | 0.78(2.5) | 0.78(2.5) | 0.55(6.0) | 0.65(4.0) |
| lib | 0.87(1.5) | 0.74(6.0) | 0.82(4.0) | 0.83(3.0) | 0.80(5.0) | 0.87(1.5) |
| mad | 0.57(5.0) | 0.58(3.5) | 0.56(6.0) | 0.58(3.5) | 0.59(2.0) | 0.63(1.0) |
| mag | 0.95(1.0) | 0.86(4.0) | 0.90(2.0) | 0.89(3.0) | 0.77(6.0) | 0.79(5.0) |
| min | 0.75(5.0) | 0.84(2.0) | 0.72(6.0) | 0.92(1.0) | 0.83(3.0) | 0.81(4.0) |
| mul | 0.99(1.5) | 0.80(4.0) | 0.98(3.0) | 0.99(1.5) | 0.73(5.5) | 0.73(5.5) |
| mus | 0.97(3.0) | 0.58(6.0) | 0.98(1.5) | 0.98(1.5) | 0.65(5.0) | 0.69(4.0) |
| pag | 0.89(4.0) | 0.96(3.0) | 0.98(1.0) | 0.97(2.0) | 0.72(6.0) | 0.73(5.0) |
| par | 0.94(1.0) | 0.86(3.0) | 0.89(2.0) | 0.82(5.0) | 0.73(6.0) | 0.84(4.0) |
| pen | 1.00(1.0) | 0.90(4.0) | 0.97(2.5) | 0.97(2.5) | 0.76(5.0) | 0.72(6.0) |
| pim | 0.89(1.0) | 0.87(2.0) | 0.82(4.0) | 0.85(3.0) | 0.74(5.0) | 0.70(6.0) |
| son | 0.81(4.5) | 0.81(4.5) | 0.80(6.0) | 0.83(3.0) | 0.87(2.0) | 0.93(1.0) |
| spe | 0.89(3.0) | 0.86(6.0) | | | | |