Out-of-Distribution Detection for Skin Lesion Images with Deep Isolation Forest

Xuan Li[1], Yuchen Lu[2], Christian Desrosiers[3], and Xue Liu[1]

xuan.li2@mcgill.ca, yuchen.lu@umontreal.ca, christian.desrosiers@etsmtl.ca, xue.liu@mcgill.ca

[1] McGill University
[2] Universite de Montrea
[3] ETS Montreal

Abstract. In this paper, we study the problem of out-of-distribution detection in skin disease images. Publicly available medical datasets normally have a limited number of lesion classes (e.g. HAM10000 has 8 lesion classes). However, there exists a few thousands of clinically identified diseases. Hence, it is important if lesions not in the training data can be differentiated. Toward this goal, we propose DeepIF, a non-parametric Isolation Forest based approach combined with deep convolutional networks. We conduct comprehensive experiments to compare our DeepIF with three baseline models. Results demonstrate state-of-the-art performance of our proposed approach on the task of detecting abnormal skin lesions.

1 Introduction

Deep learning models such as the convolution neural networks (CNN) have shown outstanding potential in dermatology for skin cancer classification [4,20,5]. However, the diversity of real life skin disease still hinder the application of automatic differential diagnosis to real life. E.g., the well-known HAM10000 dataset [18] contains eight different skin lesion classes in its training set. This is quite small compared to the actual number of known skin lesion types and subtypes, which can be in the thousands [4]. Hence, it is important to have methods that can make use of the limited amount of disease types in existing datasets to detect the unseen diseases. This is the problem of Out-Of-Distribution (OOD) detection, or abnormality detection. Recent work [9] proposes a simple but effective OOD detection framework. They model a class conditional Guassian distribution on the final feature of any pre-trained neural network, and they use Mahalanobis-distance-based metric to compute the abnormality score. However, skin lesions, even within the same class, are known to have huge intra-class difference. As a result, we argue that a uni-modal Gaussian distribution might not be expressive enough to capture the distribution of representation, which is shown in our paper.

To address this limitation, we propose to replace the simple Guassian estimation with a powerful non-parametric method Isolation Forest (IF) [11]. Unlike

traditional anomaly detection techniques, IF does not require normal profiling nor assuming a distribution family for normal samples. IF is designed based on the intuition that, abnormal samples are few and different, and as a result, they can be easily classified by a decision tree with fewer splits [11]. In this work, we propose to use IF on the features computed by a pre-trained deep CNN to detect OOD images of skin lesions, and hence the name *DeepIF*. Our contributions are as follows:

- We propose DeepIF as a modification to the existing OOD framework [9] to take into account the huge intra-class diversity of skin disease image.
- Our experiment on HAM10000 dataset [18] shows that DeepIF outperform the existing baselines on OOD detection, and it provides a 20% detection rate improvement compared to the metric based on simple Gaussian [9].
- We present a comprehensive analysis of hidden representations from different convolutional layers. Results show that the last convolutional layer has the most expressive representations among most of the diseases.

2   Related Works

In recent years, a broad range of approaches based on deep learning have been proposed for this problem. [7] introduce a simple heuristic by applying a threshold on the softmax probability of the predicted class. The ODIN approach, proposed by Liang et al. [10], uses softmax temperature scaling and adversarial input perturbation to make the softmax scores of in-distribution and out-of-distribution examples better separated. Based on the assumption that features computed by a pre-trained network follow a class-conditional Gaussian distribution, Lee et al. [9] use the Mahalanobis distance in the predicted class distribution to detect OOD and adversarial samples. Our method can be viewed as a non-parametric model extension on the above framework to take into account the high complexity of medical images like skin disease.

In [3], Devries et al. use an auxiliary loss function to generate a confidence score in another branch. The extra loss function encourages the network to identify examples for which its prediction is unsure. Vyas et al. [19] train an ensemble of classifiers in a self-supervised manner, considering a random subset of training examples as OOD data and the rest as in-distribution data. A margin-based loss is proposed to impose a given margin between the mean entropy of OOD and in-distribution samples. In [14], Masana et al. use metric learning to derive an embedding space where samples from the same in–distribution class form clusters that are separated from other in–distribution classes and OOD samples. [15] propose to use transfer learning as a general abnormality detection for medical images. [17] propose using the likelihood ratio between the output probability of two deep networks, the first one modeling in-distribution data and the second capturing background statistics, as measure of normality.

While all these approaches require modifying the original training algorithm of the model, our method is more flexible as it only needs a pre-trained network
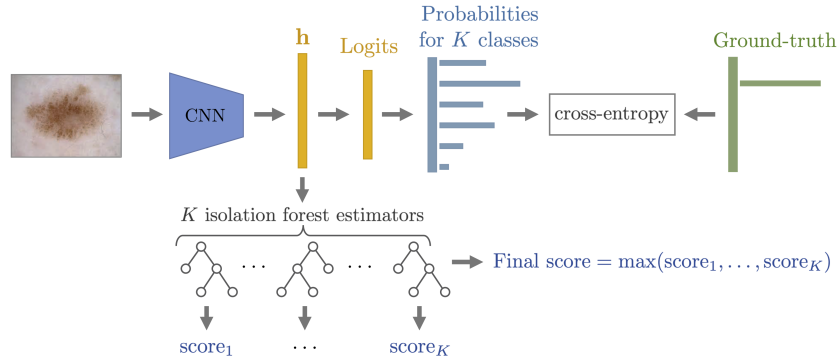
Fig. 1. Proposed DeepIF method for detecting OOD skin lesion images.

and can use a black-box algorithm for training. In addition, these studies focus on natural images and, as shown in our experiments, do not work well on skin lesion images which have less inter-class variability. So far, only a few works have investigated OOD detection for this type of image. Pacheco et al. [16] use the mean Shannon entropy of the softmax output for correctly classified and misclassified validation examples to detect outliers, yielding a 11.45% OOD detection rate for the ISIC 2019 dataset. In a different approach, Lu et al. [12] consider the likelihood of a variational autoencoder (VAE) to identify OOD skin lesion images. Different from these approaches, our method does not presume any distribution for the anomaly class. As we will empirically demonstrate, this makes our OOD method more robust.

3   Method

Isolation Forest Isolation Forest (IF) is an anomaly detection algorithm built on the idea of decision tree ensembling. Each decision tree is constructed by the data points in the training set. At each node of a tree, select a random feature from a subset of features (the proportion of the size of subset is N_f). A random value between the minimum and maximum values of that feature is chosen to make a split at that node. We construct a total of N_e decision trees.

For a given isolation forest IF and the test data $\mathbf{x}^{\text{test}}$, we calculate the normality as:

$$IF(\mathbf{x}^{\text{test}}) \;=\; -2^{-\frac{E[P_e]}{P_{\text{avg}}}} \;+\; 0.5. \tag{1}$$

where P_e is the number of tree nodes (i.e., path length) traversed by $\mathbf{x}^{\text{test}}$ from the root node to the terminal leaf node on the e-th decision tree, and we take its average across all trees in IF. P_{avg} is the average path length for training data. We refer to the original paper [11] for detailed information. The intuition is that anomaly data points have extreme values on certain features, such that they can be easily isolated and have shorter paths. Thus $IF(\mathbf{x}^{\text{test}})$ would be small if $\mathbf{x}^{\text{test}}$ is an OOD data.

OOD Detection Framework An arbitrary CNN $f(\cdot)$ is pre-trained to predict the K normal classes of the training data. The parameters of f are then fixed when training finishes. Afterwards, training examples $\mathbf{x}^{\mathrm{train}}$ are fed into f to obtain their hidden representation $\mathbf{h}^{\mathrm{train}}$ from the last convolutional layer. Lee et al [9] calculate the class mean and covariance as class-conditional Gaussian distributions based on the $\mathbf{h}^{\mathrm{train}}$. For OOD detection, they extract the $\mathbf{h}^{\mathrm{test}}$ from $\mathbf{x}^{\mathrm{test}}$ and calculate the Mahalanobis distance of each class, and assign $\mathbf{x}^{\mathrm{test}}$ the shortest distance as the final anomaly score.

Deep Isolation Forest (DeepIF) Our DeepIF shares the same idea for extracting $\mathbf{h}$ from a pre-trained CNN (see Fig 1). Different from their distance-based approach, we construct models $IF_1, IF_2, ..., IF_K$ for each class. Then our final normality score is computed as $\max(IF_1(\mathbf{h}^{\mathrm{test}}), IF_2(\mathbf{h}^{\mathrm{test}}), ..., IF_K(\mathbf{h}^{\mathrm{test}}))$.

4 Experiments

Data and setup The data we use is from the HAM10000 [18] training set which contains 25,331 images with 8 classes: Melanoma (MEL), Melanocytic nevus (NV), Basal cell carcinoma (BCC), Actinic keratosis (AK), Benign keratosis (BKL), Dermatofibroma (DF), Vascular lesion (VASC), Squamous cell carcinoma (SCC). For each experiment, we hold out 1 class as an Anomaly Class, which we refer to as an *OOD set*. For each remaining class, a 90% - 10% split is made for the training and validation sets. We treat the validation set as *in-distribution set*. Since HAM10000 [18] contains 8 classes, we conduct 8 experiments with a single class being treated as the Anomaly class and the rest 7 are normal classes in each experiment.

Pre-trained CNN We train a skin lesion classification network with a standard approach: an image is feed into a ResNet152 [6] to get the predictions for each class. Cross-entropy loss is calculated and back-propagated to the network. SGD is adopted to optimize the network with a learning rate of 1e-4. We train the network 200 epochs with a batch size of 32. In the training stage, one class is held out to be treated as an anomaly class. Once the training procedure finishes, the parameters of the network is fixed through the rest of the procedures.

For constructing the IF models, we set N_e to be 200, and N_f to be 1.0. Final scores for in-distribution and OOD sets are stored separately for evaluation.

Baselines Our first baseline is to compare with the originally *Mahalanobis-distance* baseline using the implementation from [8]. We also compare to other strong baselines that beyond our framework. We compare to a *Confidence Score* baseline [3], which learns to predict the confidence score. We use the implementation from [2] but with the same network architecture as our DeepIF. Finally we compare with the *VAE* baseline [12] by measuring the negated reconstruction score.

Evaluation Metrics We adopt the same metrics as in other studies on OOD detection [3,9,10]: area under the ROC curve ($AUROC$); area under the precision recall curve where in-distribution is specified as the positive ($AUPR\,in$); area under the precision recall curve where OOD is specified as the positive ($AUPR\,out$); true negative rate (TNR) when the true positive rate is as high as 95% ($TNR95TPR$). In the latter, the TNR is computed as TN/(TN+FP), where TN is the number of true negative and FP the number of false positives. We also show the classification accuracy on the validation dataset.

5 Results

The results are shown in Table 1. We can first find that the confidence-based baseline would decrease the classification performance on validation data, with 4% mean accuracy drop than the other methods. We believe that learning to predict confidence would add extra requirement to the training process which might hurt the performance of the main task, and therefore an OOD framework that does not touch the training procedure like ours has the advantage to preserve the model performance.

DeepIF easily beat the Mahalanobis baseline, which confirms our hypothesis that medical images like skin lesion are too complex to be properly modelled by a uni-mode Gaussian even on the representation space. Our method also beat the VAE baseline, and VAE is known to be a very distribution modelling for high-dimensional data. We believe that this results show the potential of non-parametric OOD detection that does not depend on normal profiling [11]. The strongest baseline is the confidence score. DeepIF is better except in one metric (AUPR in), but DeepIF preserves the model accuracy.

We plot in Fig 2 the histograms of normality scores for in- and out-distribution data point between Mahalanobis baseline and DeepIF with MEL as the OOD set. It can be observed that DeepIF scores lead to a better separation of in-distribution and OOD examples, which explains our method's better ability in differentiating those two datasets. We also plot in Fig 3 the ROC curves with OOD set to be BKL and DF.

We analyze the effect of using the representation from different layers. Our default choice is to use the last layer f_{-1}. We evaluate the performance of DeepIF from f_{-2} to f_{-4} as well. The result is shown in Table. 2. We find that, with the exception of NV, the performance of DeepIF with shallower features is worse than using deep features. This highlights the importance of semantic information captured in deeper layers for OOD detection.

6 Discussion and Conclusion

In this paper, we studied the problem of OOD detection with a non-parametric approach on the HAM10000 [18] skin lesion dataset. We proposed a simple framework by adopting a pre-trained CNN and Isolation Forest models. Our

Table 1. Results for OOD Experiment on HAM10000. We take take one class of images out of dataset as OOD set and only train on the rest of them.

OOD Set	Method	AUROC	AUPR in	AUPR out	TNR at 95% TPR	Val. Acc %
MEL	DeepIF	**0.6918**	**0.6856**	**0.6909**	**0.1969**	**93.3**
	Mahalanobis	0.6108	0.5797	0.6073	0.1186	
	VAE	0.5653	0.5619	0.5301	0.0411	
	Confidence	0.6248	0.6536	0.5555	0.0249	89.5
NV	DeepIF	**0.6311**	**0.6513**	**0.5969**	**0.0894**	**90.7**
	Mahalanobis	0.5537	0.5564	0.5525	0.0807	
	VAE	0.5545	0.5606	0.5201	0.0362	
	Confidence	0.4375	0.5011	0.4301	0.0041	84.1
BCC	DeepIF	0.7539	0.6878	0.7503	0.2724	**89.0**
	Mahalanobis	0.5702	0.5785	0.5347	0.0464	
	VAE	0.5292	0.5324	0.5109	0.0453	
	Confidence	**0.8236**	**0.8325**	**0.7921**	**0.2996**	85.3
AK	DeepIF	0.6942	0.6271	0.6879	0.1693	**90.3**
	Mahalanobis	0.5509	0.5304	0.5398	0.0741	
	VAE	0.5151	0.5195	0.4938	0.0316	
	Confidence	**0.7908**	**0.8136**	**0.7416**	**0.1929**	86.4
BKL	DeepIF	0.6991	0.6743	**0.6847**	**0.1738**	**91.2**
	Mahalanobis	0.6126	0.6031	0.5790	0.0729	
	VAE	0.5151	0.5195	0.4938	0.0316	
	Confidence	**0.7384**	**0.7611**	0.6698	0.1032	87.2
DF	DeepIF	**0.7462**	0.7108	**0.7302**	**0.2676**	**88.9**
	Mahalanobis	0.5443	0.5409	0.5188	0.0584	
	VAE	0.5230	0.5468	0.5040	0.0375	
	Confidence	0.6972	**0.7389**	0.6279	0.0858	84.6
VASC	DeepIF	**0.7483**	**0.7480**	**0.7040**	**0.1635**	**89.6**
	Mahalanobis	0.5985	0.6229	0.5467	0.0466	
	VAE	0.5159	0.5490	0.4808	0.0221	
	Confidence	0.4813	0.5579	0.4489	0.0118	84.7
SCC	DeepIF	0.7612	0.7105	0.7523	0.2573	**89.2**
	Mahalanobis	0.5758	0.5705	0.5342	0.0397	
	VAE	0.5336	0.5444	0.5096	0.0404	
	Confidence	**0.8324**	**0.8505**	**0.7858**	**0.2682**	86.1
Mean	DeepIF	**0.7136**	0.6841	**0.6985**	**0.1979**	**90.3**
	Mahalanobis	0.5771	0.5728	0.5516	0.0672	
	VAE	0.5315	0.5418	0.5054	0.0357	
	Confidence	0.6783	**0.7137**	0.6315	0.1238	86.1

experiments showed our approach to achieve state-of-the-art performance for differentiating in-distribution and OOD data.

We demonstrated the usefulness of our proposed DeepIF, method on a skin lesion dataset. To further validate our method, we aim to cover a broader range of medical image datasets where there exists huge intra-class diversity, for instance,

Fig. 2. Comparison of normality score distribution between DeepIF and Mahalanobis baseline. The OOD set is MEL.
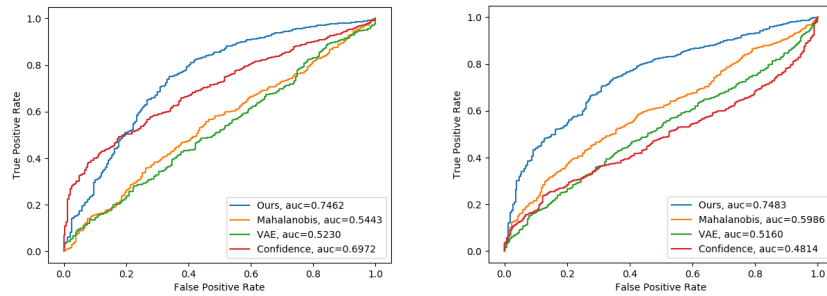

Fig. 3. ROC curves for OOD set to be BKL (left) and OOD set to be DF (right). DeepIF yields a better ROC curve compared with the other 3 approaches.

Diabetic Retinopathy, CT, and MRI datasets. Moreover, while our DeepIF focuses on image data, our method can be easily transferred to other non-image data, such as electric medical records data, or time sequence data including electroencephalogram (EEG) and electrocardiogram (ECG). In future work, we would also like to compare test method with more non-parametric algorithms such as Dirichlet Process Mixture Model (DPMM) [1] or a self-organizing network [13].

Table 2. Result of DeepIF using features from different layers of the pretrained network.

OOD set	Layer	AUROC	AUPR in	AUPR out	TNR at 95% TPR
MEL	f_{-1}	**0.6918**	**0.6856**	**0.6909**	**0.1969**
	f_{-2}	0.6240	0.6323	0.6081	0.1109
	f_{-3}	0.6001	0.5628	0.5921	0.1071
	f_{-4}	0.5600	0.5265	0.5653	0.0977
NV	f_{-1}	0.6311	**0.6513**	0.5969	0.0894
	f_{-2}	0.4891	0.4879	0.4941	0.0506
	f_{-3}	0.6599	0.6329	0.6461	0.1114
	f_{-4}	**0.6628**	0.6215	**0.6642**	**0.1908**
BCC	f_{-1}	**0.7539**	**0.6878**	**0.7503**	**0.2724**
	f_{-2}	0.6244	0.6655	0.5865	0.0718
	f_{-3}	0.5485	0.5484	0.5315	0.0596
	f_{-4}	0.5296	0.5150	0.5229	0.0465
AK	f_{-1}	**0.6942**	0.6271	**0.6879**	**0.1693**
	f_{-2}	0.6652	**0.7217**	0.6028	0.0734
	f_{-3}	0.5742	0.5992	0.5247	0.0352
	f_{-4}	0.5494	0.5809	0.5113	0.0415
BKL	f_{-2}	**0.6991**	**0.6743**	**0.6847**	**0.1738**
	f_{-2}	0.5494	0.5809	0.5113	0.0415
	f_{-3}	0.4920	0.4969	0.4891	0.0437
	f_{-4}	0.4600	0.4718	0.4754	0.0420
DF	f_{-1}	**0.7462**	**0.7108**	**0.7302**	**0.2676**
	f_{-2}	0.5116	0.5516	0.4886	0.0365
	f_{-3}	0.4600	0.4813	0.4532	0.0134
	f_{-4}	0.4521	0.4753	0.4518	0.0324
VASC	f_{-1}	**0.7483**	**0.7480**	**0.7040**	**0.1635**
	f_{-2}	0.4888	0.5338	0.4940	0.0607
	f_{-3}	0.5295	0.5394	0.5180	0.0731
	f_{-4}	0.5432	0.5249	0.5554	0.1182
SCC	f_{-1}	**0.7612**	**0.7105**	**0.7523**	**0.2573**
	f_{-2}	0.6575	0.6869	0.6159	0.0927
	f_{-3}	0.5518	0.5461	0.5267	0.0417
	f_{-4}	0.4774	0.4781	0.4904	0.0468

References

1. Blei, D.M., Jordan, M.I., et al.: Variational inference for dirichlet process mixtures. Bayesian analysis **1**(1), 121–143 (2006)
2. DeVries, T.: Learning confidence for out-of-distribution detection in neural networks. `https://github.com/uoguelph-mlrg/confidence_estimation` (2018)
3. DeVries, T., Taylor, G.W.: Learning confidence for out-of-distribution detection in neural networks. arXiv preprint arXiv:1802.04865 (2018)
4. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. Nature **542**(7639), 115–118 (2017)
5. Han, S.S., Park, G.H., Lim, W., Kim, M.S., Im Na, J., Park, I., Chang, S.E.: Deep neural networks show an equivalent and often superior performance to der-

matologists in onychomycosis diagnosis: Automatic construction of onychomycosis datasets by region-based convolutional deep neural network. PloS one **13**(1) (2018)

6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)

7. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. arXiv preprint arXiv:1610.02136 (2016)

8. Lee, K.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. `https://github.com/pokaxpoka/deep_Mahalanobis_detector` (2019)

9. Lee, K., Lee, K., Lee, H., Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In: Advances in Neural Information Processing Systems. pp. 7167–7177 (2018)

10. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. In: 6th International Conference on Learning Representations, ICLR 2018 (2018)

11. Liu, F.T., Ting, K.M., Zhou, Z.H.: Isolation forest. In: 2008 Eighth IEEE International Conference on Data Mining. pp. 413–422. IEEE (2008)

12. Lu, Y., Xu, P.: Anomaly detection for skin disease images using variational autoencoder. arXiv preprint arXiv:1807.01349 (2018)

13. Marsland, S., Shapiro, J., Nehmzow, U.: A self-organising network that grows when required. Neural networks **15**(8-9), 1041–1058 (2002)

14. Masana, M., Ruiz, I., Serrat, J., van de Weijer, J., Lopez, A.M.: Metric learning for novelty and anomaly detection. arXiv preprint arXiv:1808.05492 (2018)

15. Ouardini, K., Yang, H., Unnikrishnan, B., Romain, M., Garcin, C., Zenati, H., Campbell, J.P., Chiang, M.F., Kalpathy-Cramer, J., Chandrasekhar, V., et al.: Towards practical unsupervised anomaly detection on retinal images. In: Domain Adaptation and Representation Transfer and Medical Image Learning with Less Labels and Imperfect Data, pp. 225–234. Springer (2019)

16. Pacheco, A.G., Ali, A.R., Trappenberg, T.: Skin cancer detection based on deep learning and entropy to detect outlier samples. arXiv preprint arXiv:1909.04525 (2019)

17. Ren, J., Liu, P.J., Fertig, E., Snoek, J., Poplin, R., Depristo, M., Dillon, J., Lakshminarayanan, B.: Likelihood ratios for out-of-distribution detection. In: Advances in Neural Information Processing Systems. pp. 14680–14691 (2019)

18. Tschandl, P., Rosendahl, C., Kittler, H.: The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. Scientific data **5**, 180161 (2018)

19. Vyas, A., Jammalamadaka, N., Zhu, X., Das, D., Kaul, B., Willke, T.L.: Out-of-distribution detection using an ensemble of self supervised leave-out classifiers. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 550–564 (2018)

20. Zhang, X., Wang, S., Liu, J., Tao, C.: Towards improving diagnosis of skin diseases by combining deep neural network and human knowledge. BMC medical informatics and decision making **18**(2), 59 (2018)