Redesign and Extension of Source Code Package for Open-Set Classification on ImageNet

Team: Dean Heizmann, Nicolas Fazli Kohler and David Pierre Sébastien Lebrec

Supervisor: Prof. Dr. Manuel Günther

Intended Start: July 2023

1 Introduction

The automatic separation of images into several classes has gained a lot of interest in the last years. Starting with small-scale datasets such as MNIST and CIFAR-10, nowadays more realistic and large-scale datasets are used for this task. Particularly, the availability of the ImageNet dataset [Deng et al., 2009] and its usage in the International Large Scale Visual Recognition Challenge (ILSVRC) [Russakovsky et al., 2015] has fostered large improvements in this task. Especially, deep neural networks have shown great success in the ImageNet challenge [Krizhevsky et al., 2012, He et al., 2015].

However, most classifiers have not left academic areas since they have one important flaw: these classifiers can only classify samples of classes that they have seen during training. When presented with a sample from a different class that the network is not trained to predict, it has no other chance than predicting one of the known classes, which it oftentimes does with large confidence [Dhamija et al., 2018, Dhamija et al., 2020, Palechor et al., 2023].

Lately, researchers have understood this limitation and provided different ways of providing options for the classifier to reject such samples as unknown. There are two main directions of such approaches. The first approach is to take a pre-trained network that is trained on a closed-set task such as ImageNet, and provide means and options to reject a sample. For example [Hendrycks and Gimpel, 2017, Hendrycks et al., 2022] showed that thresholding the softmax probabilities, or the logits that feed into softmax, are good baselines for open-set classification. Other approaches try to approximate the probability of unknown by modeling deep feature representations [Bendale and Boult, 2016, Rudd et al., 2017, Lyu et al., 2023] of known classes, to provide a probability of sample exclusion.

The second main direction tries to train the networks such that they are better suited for open-set classification. Most approaches in this direction add another output for the unknown class to the network, while others include different loss functions to enable better thresholding of outputs [Dhamija et al., 2018]. To obtain training samples for the unknown class (which we term *negative* samples), different methods are developed. In the simplest way, negative samples are collected from classes that are not of interest [Dhamija et al., 2018, Palechor et al., 2023]. Other approaches artificially generate such negative samples in various different ways [Ge et al., 2017, Yu et al., 2017, Neal et al., 2018, Wilson et al., 2023, Geng et al., 2021]. For example, mixed representations of known classes are used as negatives [Zhou et al., 2021], with the hope that they are good predictors for unknown samples.

Most of the above algorithms are evaluated on small-scale datasets with particular properties, which do not reflect real-world applications well. Since there was a lack of large-scale evaluation protocols for open-set classification, we have developed our own evaluation [Palechor et al., 2023] based on the ImageNet dataset. In that paper, we compared simple approaches for open-set classification, for which we also have published an open-source code package.[1] Currently, we are extending this paper to include more algorithms to compare[2] – already implemented methods include [Dhamija et al., 2018, Zhou et al., 2021, Hendrycks and Gimpel, 2017, Hendrycks et al., 2022, Bendale and Boult, 2016, Rudd et al., 2017]. In this Master Project, the task is to re-design this source code package to make it easier to extend it; and then extend the package to include more open-source implementations of other methods. For this purpose, it is required to perform a literature review to understand the latest approaches for open-set classification, check the availability of open-source implementations for those algorithms and, finally, bind these implementations (or re-implementations) to the source code package.

[1] `https://github.com/AIML-IfI/openset-imagenet`

[2] `https://github.com/AIML-IfI/openset-imagenet-comparison`

At the end, the available documentation of the package needs to be adapted such that it reflects the new implementation. Examples on how to extend the current package with a novel algorithm should be provided. Possibly, test cases can be added to ensure the stability of the algorithms. Additionally, Jupyter notebooks can be created to showcase the usage of the package.

2 Schedule

Assuming 30 hours of work per week and a total of 15 ECTS with an average of 30 hours per ECTS, we arrive at a total workload of 15 weeks full-time. These should be distributed as follows.

Week 1-3 Setting up the work environment, installing all required tools, (partially) reproduce the results that are currently available in the source code package.

Week 4-5 Perform a literature review and search for available open-source implementations of other algorithms. Algorithms do not need to be restricted to image-based classification, for example, approaches from other domains, e.g., [Wu et al., 2020, Zhan et al., 2021] can be included if possible.

⇒ Milestone 1: A collection of open-source implementations is selected that should be included into the package.

Week 6-9 Designing the software package such that it is easily extensible and can include all methods collected in Milestone 1 in an easy way. This includes providing a training script to train a deep network using various loss functions and ways to create negative samples, as well as different post-processing techniques. Also, validation and evaluation metrics should be included.

Week 10-13 Implementing the methods selected in Milestone 1, and those present in the current package, into the newly designed package. Running experiments with the new methods and compare/combine them with other methods.

⇒ Milestone 2: Old and new open-set methods are incorporated into the newly-designed package.

Week 14-15 Updating and writing the documentation of the package, as well as showcasing how to include new algorithms.

⇒ Milestone 3: The documentation is up-to-date includes all implemented methods.

If time allows Run experiments to optimize parameters of the implemented algorithms to get the best possible result on the protocols.

Optionally Create Jupyter notebooks show-casing how the algorithms can be used.

The project is designed for two to four students. For two students, the focus should be on designing the new interface such that it can be easily extended by other algorithms. When more students are present, more open-source methods should be collected and implemented.

Writing the project report is part of the Master project. As a template, the LaTeX thesis template from my webpage[3] should be used. I would recommend to start writing early and keep note of what was done when, and by whom. At the end of the project, there will be a joint presentation of the results in my research group.

3 References

[Bendale and Boult, 2016] Bendale, A. and Boult, T. E. (2016). Towards open set deep networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.

[Deng et al., 2009] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.

[Dhamija et al., 2020] Dhamija, A., Günther, M., Ventura, J., and Boult, T. E. (2020). The overlooked elephant of object detection: Open set. In *Winter Conference on Applications of Computer Vision (WACV)*, pages 1021–1030.

[3] `https://www.ifi.uzh.ch/en/aiml/theses.html`

[Dhamija et al., 2018] Dhamija, A. R., Günther, M., and Boult, T. E. (2018). Reducing network agnostophobia. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9157–9168.

[Ge et al., 2017] Ge, Z., Demyanov, S., and Garnavi, R. (2017). Generative OpenMax for multi-class open set classification. In *British Machine Vision Conference (BMVC)*.

[Geng et al., 2021] Geng, C., Huang, S.-J., and Chen, S. (2021). Recent advances in open set recognition: A survey. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 43(10):3614–3631.

[He et al., 2015] He, K., Zhang, X., Ren, S., and Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *International Conference on Computer Vision (ICCV)*, pages 1026–1034.

[Hendrycks et al., 2022] Hendrycks, D., Basart, S., Mazeika, M., Zou, A., Kwon, J., Mostajabi, M., Steinhardt, J., and Song, D. (2022). Scaling out-of-distribution detection for real-world settings. In *International Conference on Machine Learning (ICML)*. PMLR.

[Hendrycks and Gimpel, 2017] Hendrycks, D. and Gimpel, K. (2017). A baseline for detecting misclassified and out-of-distribution examples in neural networks. *International Conference on Learning Representations (ICLR)*.

[Krizhevsky et al., 2012] Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*.

[Lyu et al., 2023] Lyu, Z., Gutierrez, N. B., and Beksi, W. J. (2023). Metamax: Improved open-set deep neural networks via weibull calibration. In *Winter Conference on Applications of Computer Vision Workshops (WACVW)*.

[Neal et al., 2018] Neal, L., Olson, M., Fern, X., Wong, W.-K., and Li, F. (2018). Open set learning with counterfactual images. In *European Conference on Computer Vision (ECCV)*.

[Palechor et al., 2023] Palechor, A., Bhoumik, A., and Günther, M. (2023). Large-scale open-set classification protocols for imagenet. In *Winter Conference on Applications of Computer Vision (WACV)*, pages 42–51. CVF/IEEE.

[Rudd et al., 2017] Rudd, E. M., Jain, L. P., Scheirer, W. J., and Boult, T. E. (2017). The extreme value machine. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.

[Russakovsky et al., 2015] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. (2015). Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115(3).

[Wilson et al., 2023] Wilson, S., Fischer, T., Dayoub, F., Miller, D., and Sünderhauf, N. (2023). SAFE: Sensitivity-aware features for out-of-distribution object detection. *arXiv preprint arXiv:2208.13930*.

[Wu et al., 2020] Wu, M., Pan, S., and Zhu, X. (2020). OpenWGL: Open-world graph learning. In *International Conference on Data Mining (ICDM)*.

[Yu et al., 2017] Yu, Y., Qu, W.-Y., Li, N., and Guo, Z. (2017). Open-category classification by adversarial sample generation. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

[Zhan et al., 2021] Zhan, L., Liang, H., Liu, B., Fan, L., Wu, X., and Lam, A. Y. S. (2021). Out-of-scope intent detection with self-supervision and discriminative training. In *International Joint Conference on Natural Language Processing, (IJCNLP)*. Association for Computational Linguistics.

[Zhou et al., 2021] Zhou, D.-W., Ye, H.-J., and Zhan, D.-C. (2021). Learning placeholders for open-set recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.