Detecting Out-of-Distribution Objects Using Neuron Activation Patterns

Bartłomiej Olber^{a,b;*}, Krystian Radlak^{a,b}, Krystian Chachuła^b, Jakub Łyskawa^a and Piotr Frątczak^a

^aWarsaw University of Technology, Pl. Politechniki 1, 00-661 Warsaw ^bSilesian University of Technology, Akademicka 16, 44-100 Gliwice

Abstract. Object detection is essential to many perception algorithms used in modern robotics applications. Unfortunately, the existing models share a tendency to assign high confidence scores for outof-distribution (OOD) samples. Although OOD detection has been extensively studied in recent years by the computer vision (CV) community, most proposed solutions apply only to the image recognition task. Real-world applications such as perception in autonomous vehicles struggle with far more complex challenges than classification. In our work, we focus on the prevalent field of object detection, introducing Neuron Activation PaTteRns for out-of-distribution samples detection in Object detectioN (NAPTRON). Performed experiments show that our approach outperforms state-of-the-art methods, without the need to affect in-distribution (ID) performance. By evaluating the methods in two distinct OOD scenarios and three types of object detectors we have created the largest open-source benchmark for OOD object detection.

1 Introduction

OOD detection, as described in a thorough study by Yang et al. [31], is a problem that arises when machine learning models are applied to data derived from a distribution that is beyond the one they were trained on. This results in the model's poor performance since the model had not acquired knowledge that would enable correct predictions. This issue is not limited to classification tasks because object detectors may also encounter OOD samples in two situations: when the inference image contains unknown object classes or when the image scenery is significantly different from the training examples. Practitioners tackle the latter by covering the entire range of scenarios. For example, the most popular autonomous driving datasets include images of environmental, weather, and geographic diversity [1, 32, 29]. On the other hand, identifying an unknown object may be possible thanks to algorithms relating to two highly overlapping research fields - open-set (OS) detection [28, 6] and OOD detection [15, 16, 7]. Both OS and OOD methods design a way to quantify the model's uncertainty regarding encountered data so that high uncertainty scores are assigned to unknown objects while low scores are to those well-represented in the training data.

Contrary to the image classification problem, we observe a severe deficiency of OOD detection methods in object detectors. Detectors process images in a much more complex way than classifiers. Analogously hampered are the evaluation process and posterior analysis of predicted uncertainty scores. This issue became a substantial obstacle for researchers to invent novel OOD algorithms for object detection. Finally, almost every published OS or OOD algorithm assumes Faster R-CNN as the default architecture, making the proposed solution architecture-specific, and unfeasible to apply to any other model. We notice the need for a universal, simple, and yet efficacious OOD framework, which we address in this work.

In this work, we present NAPTRON - i.e., a neuron activation pattern (NAP) OOD method adapted to object detection. It was proved that NAP is a highly efficient technique for OOD detection in image recognition problems [25]. Our algorithm leverages object detectors' internal feature representation and enables an understanding of training distribution to estimate the uncertainty of predicted bounding boxes. ReLU-activated layers, being the foundation of most network architectures [14, 26, 19, 2], are the natural source of binary patterns, because they set every neuron (or convolution unit) in either positive i.e., on- or zeroed i.e., off-state. The NAPs of ReLU networks display a very convenient property for OOD detection; namely, ReLUactivated networks generate much fewer nonidentical NAPs than they are theoretically capable of generating [12, 13]. This finding fuels the intuition that if one memorized all known patterns, which are not very numerous, then encountering an unseen pattern during inference would be a reliable indicator of OOD data.

The main contributions of this paper are:

- We present a theoretically inspired NAPTRON that uses binary NAPs extracted from hidden layers of object detectors to tell ID from OOD predictions. This method is both computationally efficient and practically effective for OOD detection, making it simple to incorporate into existing object detection architectures.¹
- We perform comprehensive experiments involving two datasets and three network architectures, which prove that the proposed method outperforms state-of-the-art OOD detectors.
- We introduce a novel OOD detection evaluation protocol that analyzes scores of OOD bounding boxes and allows for a more objective comparison of the methods.

2 Related Work

2.1 Out-of-Distribution Detection

Heyndrycks et al. [15] proposed a baseline OOD samples detection method. It only involves setting a threshold on the winning class softmax probability (maximum softmax probability, MSP). In object de-

^{*} Corresponding Author. Email: olberbartlomiej@gmail.com

tection, it is already necessary to set a softmax probability threshold for non-maximum-suppression (NMS) to filter redundant background predictions. Therefore, setting another one to separate OOD objects can not be particularly useful in practice. Nevertheless, it still provides a decent baseline.

Initially, Helmholtz **Energy** [21] had been successfully utilized in OOD detection in image classification. It is a very easy-to-use and universal way to estimate semantic uncertainty. It does not require any detector customization or loss handcrafting. Internally, object detectors perform the classification of proposals resulting in a classification vector for which one can compute the energy score.

Virtual outlier synthesis (VOS) [7] is the first work focusing on OOD identification in the object detection task. VOS allows OOD detection by synthesizing virtual outliers in feature (latent) space, thereby regularizing the model's decision boundary during training. VOS samples the virtual outliers from the low-likelihood region in the feature space and uses them as input for an unknown-aware training objective. The contrastive loss function shapes the uncertainty (Helmholtz energy) space between known data and synthesized outlier data. VOS is not architecture-agnostic because the outlier synthesis process occurs in feature space bound to the fully-connected layers of Faster R-CNN ROI head. The authors of VOS proposed an evaluation scheme that requires an OOD dataset i.e. a set of images that do not contain ID categories. Any output bounding box generated for these images is considered an "OOD object", while any bounding box generated for images from the ID test dataset is deemed an "ID object". We find this approach is over-simplified because model predictions are not evaluated, so any "OOD object" or "ID object" can actually be a background prediction.

2.2 Open-Set Detection

Gaussian mixture models (GMM) [23] approach also introduces a change in the default loss function, focusing on classification loss. The authors add an anchor loss term to facilitate learning a structured logit space. Next, they fit class-specific GMMs to the logit space with a validation dataset. For any test sample, uncertainty is estimated as a log-likelihood of belonging to any one of the known GMMs.

The authors of **OpenDet** [11] follow the intuition that known-class objects tend to be clustered to form high-density regions in the latent space, while unknown objects are distributed in low-density regions. Consequently, they propose identifying unknown objects by separating high- and low-density regions in the latent space using a contrastive loss. On top of this, they provide another loss function component responsible for learning predictive uncertainty directly as a softmax probability without the logit of ground-truth class. The authors of OpenDet provide the implementation for both Faster R-CNN and RetinaNet. However, the method does not apply to anchorfree architectures because IOU-based sampling of proposals for contrastive learning is not feasible.

The **OWOD** [17] approach was developed before OpenDet. Both ideas share the part of contrastive clustering in latent space. The authors of OWOD also propose an unknown-aware RPN head and the Helmholtz energy-based unknown identification. They model the energy distribution of the known and unknown energy values with a set of shifted Weibull distributions. Similarly to GMM, they fit the distributions using a validation dataset. Due to the RPN requirement, Faster R-CNN is the only compatible architecture.

2.3 Other uncertainty estimators in Object Detection

Previous work by Choi et. al. [5] introduced **Gaussian** YOLOv3. The authors redesign YOLO's loss function to model localization uncertainty directly. Each coordinate of a bounding box is represented by a pair of Gaussian parameters (mean and variance) instead of just one value. There is no ground truth for the variance of coordinates, but Gaussian negative-log-likelihood loss conveniently requires ground truth for the mean only, which allows one to learn variance without providing the ground truth. We managed to extend this approach to Faster R-CNN and RetinaNet architectures. In FCOS, however, a "localization centerness" branch is designed to perform a very similar role as Gaussian localization uncertainty; trying to combine them both led to poor results.

The popular **Monte-Carlo Dropout** (MCD) [9] method has already been adapted to the object detection task [22, 24]. The most straightforward adaptation technique was proposed in ref. [24], in which the authors suggested averaging both classification and regression vectors over the number of Monte-Carlo forward passes for each proposal separately. The variance of these vectors for each proposal is the final uncertainty measure.

Faster R-CNN-based **Object Localization Network** (OLN) [18] learns to detect objects focusing on localization instead of foreground-background classification. The authors argued that focusing on learning "objectness cues", for the price of obtaining class-agnostic outputs is the best way to achieve cross-dataset generalization. In practice, their method requires replacing the default classification losses in both Faster R-CNN stages with a pair of centerness and IOU losses.

As an alternative baseline focused on objectness for Faster R-CNN, we attempt to establish a threshold for the RPN objectness score of a bounding box proposal. This score is later converted into the final model output. Although the RPN objectness score is not used in the ROI head, we aim to determine its potential usefulness for the OOD detection tasks.

3 Problem Setup

Let us denote the set C of ID classes as $K_{ID} = \{1, \ldots, C\}$ and given object detection dataset $D = \{X, \Upsilon\}$ where X and Υ denote the input images and labels. The set of labeled images consists of Nsamples, $X = \{x_1, \ldots, x_N\}$ along with the labels associated with the sets of objects included in each of N images $\Upsilon = \{Y_1, \ldots, Y_N\}$. Next, for each image, we have $Y_i = \{y_1, \ldots, y_K\}$ that represents a set of labels for K object instances. Each instance y = [c, l] consists of class labels $c \in K_{ID}$ and locations l = [x, y, w, h], where x, y, w, and h denote the object's bounding box center coordinates and size.

An OOD dataset is defined identically, except for the fact that it is semantically disparate with the ID dataset, meaning there are no common classes between ID and OOD datasets - $K_{OOD} = \{C + 1, \dots\}$.

Let us also introduce the test dataset D_{test} consisting of ID and OOD objects and an object detector $\theta : x \mapsto Y$ trained on the train split (D_{train}) of D.

The OOD detection task is performed for θ -predicted ID and OOD $\hat{\Upsilon}_{\theta} = \{\hat{\Upsilon}_{ID}, \hat{\Upsilon}_{OOD}\}$ samples where ID samples are true positive predictions $\hat{\Upsilon}_{ID} := \{\hat{y}_i \in K_{ID} \times \mathbb{R}^4 : \exists y_t \in D_{test}, IOU(\hat{y}_i, y_t) > \lambda, \hat{c} = c_t\}$. The notation states that true positive predictions are all pairs of class labels and bounding box locations that sufficiently overlap with any ground truth test bounding box of the same class. Analogously, OOD samples are those predictions, that sufficiently overlap with any ground truth bounding box of unknown class $\hat{\Upsilon}_{OOD} := \{\hat{y}_o \in K_{ID} \times \mathbb{R}^4 : \exists y_t \in D_{test}, IOU(\hat{y}_o, y_t) > \lambda, c_t \in K_{OOD}\}$. The IOU threshold λ is typically set to 0.5. Now, the goal of an OOD detector is to perform binary classification of $\hat{\Upsilon}_{\theta}$, which relies on the uncertainty estimator $\Phi_{\theta} : \hat{y} \mapsto \phi$ that is intended to assign high scores to OOD samples and low ones to ID samples. Next, the classification is performed by comparing a predefined uncertainty threshold with the uncertainty score associated with each sample.

4 Our approach

4.1 Neuron Activation Patterns

NAPs have been previously used for uncertainty estimation in image classifiers [4, 25]. Image processing neural networks (NNs) use multiple layers that transform input samples sequentially into the desired form. Nonlinear activation functions present in each hidden layer of a NN significantly contribute to the network's ability to approximate complex, multidimensional relationships. Of all activation functions, ReLU is nowadays the most common choice for CV problems. All ReLU-activated layers produce matrices of either positive or zeroed values, which correspond respectively to active and inactive layers' units. An activation pattern is obtained by assigning a *true* to positive units and *false* to zeros. Thus, for a given network layer, NAP is a binary interpretation of which neurons or convolution units of the layer were activated during the processing of an image.

4.2 Uncertainty estimation

Algorithm 1 describes how NAPTRON estimates the uncertainty of predicted bounding boxes. First, for every training image, we perform an object detection inference, simultaneously extracting binary NAPs corresponding to output bounding boxes. The patterns are extracted from a pre-selected layer. Then, for every true positive prediction, we store the extracted pattern in a memory structure. Each object class has a dedicated memory instance. For a given test sample we perform the inference and NAP extraction again. Next, for each inferred bounding box, we find the Hamming distance between its NAP and the nearest pattern out of those stored in the memory structure corresponding to the predicted label. The Hamming distance to the nearest known NAP is the NAPTRON uncertainty estimate.

The authors of [25] provided an efficient implementation of finding the minimal Hamming distance to known NAPs, which makes real-time uncertainty estimation possible. In this implementation, the Hamming distances between the test NAP and all known NAPs are computed concurrently. The minimal distance is the output of the uncertainty estimation algorithm.

4.3 Pattern extraction for bounding boxes

In every state-of-the-art object detector, one can distinguish a detection backbone part (e.g. ResNet50) that extracts features of an input image and a detection head that performs classification and regression of proposal bounding boxes (priors) based on the extracted features. Typically, the classification and regression are computed in two separate subnetworks that work in parallel. The *extractNAPS* function extracts binary NAPs of the *l*-th ReLU-activated layer of the classification branch in the detection head of object detector θ . For Faster R-CNN, the ROI-head consists of fully connected (FC) layers. The operation is straightforward because there is no ambiguity in the

```
Data: \theta, D_{train}, x_{test}, Layer index l
Result: \phi_{test}
for c in K_{ID} do
     // initialize a data structure for
         each class
    M_c \leftarrow \emptyset
end
for (x, y) in D_{train} do
    \hat{Y} \leftarrow \theta(x);
    NAPS_x \leftarrow extractNAPS(\theta, x, l);
    for (\hat{y}, NAP) in (\hat{Y}, NAPS_x) do
         if \hat{y} is a true positive prediction then
             // Store NAP in the data
                  structure associated with
                  predicted class \hat{c}
             M_{\hat{c}} \leftarrow M_{\hat{c}} \parallel NAP;
         end
    end
end
// test phase
\hat{Y} \leftarrow \theta(x_{test});
NAPS_{test} \leftarrow extractNAPS(\theta, x_{test}, l);
for (\hat{y}, NAP) in (\hat{Y}, NAPS_{test}) do
    NAP_{nn} \leftarrow NearestNeighbour(M_{\hat{c}}, NAP);
    // Uncertainty of \hat{y}
```

Algorithm 1: NAPTRON uncertainty estimation

 $\phi \leftarrow HammingDistance(NAP, NAP_{nn});$

link of hidden layers' activations to the final output. Assuming P proposals generated for a single image by the RPN-head, each layer processes [P, In]-shaped feature maps into [P, Out]-shaped feature maps, where In is the number of neurons of the previous layer and Out is the number of neurons of the current layer. Thus, for each layer and each proposed bounding box, we can easily distinguish an activation pattern of length Out.

Nevertheless, detection heads of single-stage architectures such as RetinaNet and FCOS, use convolutional layers instead of FC layers. Each hidden layer takes as an input [W, H, C]-shaped feature maps and processes them without changing their dimensions. The final layer changes the number of channels of the matrix from C to K * A, where W and H are the width and height of processed feature maps, respectively, K is the number of classes, and A is the number of priors with centers in each of W * H locations of the feature maps. In our approach, all the A possible bounding boxes predicted in the (w, h) location are associated with a C-long activation vector located at (w, h) of the chosen hidden layer's output. In other words, any (w, h)-centered output bounding box is attributed to the activation values located in the same coordinates of the hidden feature maps.

5 Experiments

end

We propose an evaluation protocol to examine the effectiveness of the proposed NAPTRON detector in various aspects. These include the ability to identify OOD objects, recognize known from unknown samples, and gradually acquire knowledge about new categories when labels are available for certain unknown objects.

Datasets. The experiments were conducted in two domain shift scenarios. In the first scenario, the detector was trained on the train split

of the PASCAL VOC [8] dataset consisting of 20 classes. Evaluation protocol for known and unknown objects was performed on validation split of the COCO [20] dataset consisting of 80 classes, including the 20 known ones. The validation split includes 20631 known objects and 15704 unknown objects. In the second scenario, the detector was trained on images from the train split of the BDD100k [32] dataset that consists of only 4 known classes (i.e., "pedestrian", "bicycle", "car" and "traffic sign"). Evaluation protocol was performed on the full validation split of the BDD100k dataset consisting of all 10 classes, including the 4 known ones. The validation split includes 152025 known objects and 33920 unknown objects.

Architectures. For our experiments, we chose three well-established object detection architectures: Faster R-CNN (two-stage, anchorbased) [27], RetinaNet (single-stage, anchor-free) [19], and FCOS (single-stage, anchor-based) [30]. The selected architectures represent different properties that enable our findings to be universal across many state-of-the-art object detection architecture types. All models were trained with default configuration parameters provided in the MMDetection [3] framework.

5.1 Parameter sensitivity analysis

The NAPTRON algorithm has a couple of parameters that may affect OOD detection performance quality. We examined in detail the impact of those parameters, i.e.:

- layer index,
- distance reduction,
- binarization percentile threshold p,
- train samples softmax probability threshold s,
- NMS softmax score threshold.

Layer index. Choosing an optimal layer to extract binary NAPs from is very difficult when detecting OOD samples in the image classification task [25] since modern classifiers consist of a large number of layers. Detection heads of standard architectures - such as Faster R-CNN, RetinaNet, and FCOS - have only a few layers (i.e., 2-4), so choosing the correct one should be much easier. We also try extracting patterns from the [W, H, C]-shaped feature maps originating from the feature extracting backbone to the detection head - denoted Layer 0 in Tables 1, 2, and 3. For RetinaNet and FCOS, extracting NAPs from Layer 0 requires identical operations as extracting from other layers, as described in Section 4.3. However, for Faster R-CNN, we had to flatten the feature maps to match the dimensionality of the FC layers. To do so, we computed the C-long vector of the means of every channel and binarized it by zeroing p percent of the lowest values in the vector.

Distance reduction. The authors of [25] estimate uncertainty by finding the minimal Hamming distance between the test NAP and all known NAPs of the predicted class; we check whether computing the average distance to the known NAPs yields improved results.

Binarization percentile threshold *p*. Zeroing a certain percentage of units, which have the lowest magnitude across an activation pattern, was effective in the image classification setup. Our technique of extracting NAPs from detection heads makes the binarization step optional, so the default binarization threshold value equals 0.0.

Training samples IOU threshold. To construct the database of known patterns, one needs to perform object detection on the training images and choose only those activation patterns that correspond to the true positive predictions. Typically, a prediction is deemed correct if it overlaps any object with IOU greater than 0.5. We want to check if setting a higher IOU threshold (e.g., 0.9), and consequently

choosing only the patterns that correspond very accurately to an object might lead to more accurate OOD detection results.

Training samples softmax probability threshold *s***.** The reason for examining the effect of this parameter is the same as for the IOU threshold explained above. The default softmax threshold value equals 0.0, meaning no true positive (TP) sample is discarded.

NMS softmax score threshold. Every OOD detector is sensitive to this parameter. The impact of the NMS threshold on all OOD methods' performance is described and studied in Section 5.3.

Results. The experiments were performed on the first evaluation scenario (PASCAL-VOC \rightarrow COCO); AUROC is used as the performance metric. Note that the values of the first columns of each Table 1, 2, and 3 are identical - all were generated for the default setup of the parameters (IOU ≥ 0.5 , $s \geq 0.0$, p = 0.0).

 Table 1: Impact of the binarization percentile threshold on AUROC

 metric using Faster R-CNN and COCO evaluation protocol.

Layer index	Distance reduction	p = 0.0	p = 0.45	p = 0.9
Layer 0	minimal dist. avg. dist.	0.7104 0.7660	0.7083 0.7621	0.6946 0.7426
Layer 1	minimal dist. avg. dist.	$0.7302 \\ 0.7328$	$0.7272 \\ 0.7288$	0.7139 0.7168
Layer 2	minimal dist. avg. dist.	0.7712 0.7733	0.7671 0.7694	$0.7500 \\ 0.7565$

 Table 2: Impact of the NAP train patterns IOU threshold on AUROC

 metric using Faster R-CNN and COCO evaluation protocol.

Layer index	Distance reduction	$\rm{IOU} \geq 0.5$	$\rm{IOU} \geq 0.7$	$\rm{IOU} \geq 0.9$
Layer 0	minimal dist.	0.7104	0.7106	0.7066
	avg. dist.	0.7660	0.7658	0.7522
Layer 1	minimal dist. avg. dist.	0.7302 0.7328	0.7303 0.7328	$0.7371 \\ 0.7475$
Layer 2	minimal dist.	0.7712	0.7712	0.7665
	avg. dist.	0.7733	0.7733	0.7652

 Table 3: Impact of the softmax score threshold for training NAPs extraction on AUROC metric using Faster R-CNN and COCO evaluation protocol.

Layer index	Distance reduction	$s \geq 0.0$	$s \geq 0.3$	$s \geq 0.6$	$s \geq 0.9$
Layer 0	minimal dist.	0.7104	0.7105	0.7113	0.7158
	avg. dist.	0.7660	0.7659	0.7651	0.7623
Layer 1	minimal dist. avg. dist.	$0.7302 \\ 0.7328$	0.7295 0.7319	$0.7272 \\ 0.7309$	$0.7242 \\ 0.7283$
Layer 2	minimal dist.	0.7712	0.7692	0.7657	0.7576
	avg. dist.	0.7733	0.7716	0.7702	0.7666

Results in Tables 1, 2, and 3 suggest that tuning training samples' IOU and softmax threshold parameters is not worthwhile. Maintaining the standard 0.5 IOU threshold, the lowest (0.0) softmax threshold - thus letting as many training samples as possible be stored in the known patterns database - allows an obtaining of the optimal results. Moreover, shifts in these parameters do not introduce significant performance variation. As for the binarization percentile threshold, abstaining from binarization seems to be the best choice. We find that the higher the threshold, the worse the results we obtained.

On the other hand, introducing the averaging Hamming distance (instead of the previously proposed finding of the distance to the closest binary pattern) brings a slight improvement - especially if Layer 0 is chosen. The choice of the layer is the most important parameter to tune. We recommend extracting patterns from the penultimate layer of the detection head.

5.2 OOD detection

In the next experiment, we compared the performance of the proposed algorithm with the state-of-the-art OOD object detectors:

- confidence score from Faster R-CNN (Standard)
- objectness score from Region Proposal Network in Faster R-CNN (RPN)
- Energy [21]
- Virtual Outlier Synthesis (VOS) [7]
- Gaussian Mixture Models (GMM) [23]
- OpenDet [11]
- Open World Object Detection (OWOD) [17]
- Gaussian YOLOv3 (Gasussian) [5]
- Monte-Carlo Dropout (MCD) [24]
- Object Localization Network (OLN) [18]

All methods were compared using the officially published code of each algorithm with recommended configuration parameters, except the MCD method, which we implemented ourselves. It is important to notice that some of the methods were only designed for certain DNN architectures (e.g., Faster R-CNN); and we applied the methods only for the dedicated models.

Metrics. We measure OOD detection performance using two metrics: (1) the false positive rate at the true positive rate of 95% - FPR@95TPR and (2) the area under the receiver operating characteristic curve (AUROC). The metrics are computed separately for each class and then averaged.

Performance comparison. Table 4 shows the performance of the proposed algorithm and other methods. As can be observed, NAPTRON achieves the best FPR@95TPR for every evaluation scenario, detector architecture pair considered in this work, and the best AUROC for 3 out of 6 evaluation cases. No other method yields such consistently good results. VOS had been designed specifically for the Faster R-CNN architecture and detects OOD samples roughly just as well as NAPTRON. However, these methods could not be compared using other architectures since there is no straightforward way to apply VOS for anything but Faster R-CNN. We suspect that even if it could be applied to the single-stage architectures, it would perform similarly to Energy because both methods rely on the energy score.

Additionally, the results showed that the open-sets methods (OpenDet, GMM, and OWOD) are not the most effective OOD detectors. They fail to outperform the baseline detector (Standard) consistently. GMM manages to do so in 3 or 4 (depending on the metric) out of 6 cases, OpenDet in 1 out of 4, and OWOD in 1 or 2 out of 2.

The rest of the methods perform very poorly in our challenging experimental setup. Gaussian, OLN, RPN, and MCD are unable to beat the baseline detector (Standard).

NMS sensitivity. For the above experiments, we set a low NMS threshold (0.01) so that all the possible predicted objects are accounted for when computing the metrics. However, since varying NMS softmax confidence threshold makes a significant difference in regular object detection performance, we investigate the impact of the NMS threshold on OOD detection performance. Selected NMS sensitivity plots of all considered OOD methods are presented in Fig. 1, 2 and 3. We included 4 out of 12 plots generated for each architecture, dataset scenario, and metric combination.



Figure 1: Faster R-CNN-based OOD detectors NMS sensitivity evaluated on COCO. Metric - FPR@95TPR.



Figure 2: FCOS-based OOD detectors NMS sensitivity evaluated on COCO. Metric - AUROC.



Figure 3: Faster R-CNN-based OOD detectors NMS sensitivity evaluated on BDD100k. Metric - AUROC.

All the figures confirm that the quality of all the OOD methods is sensitive to underlying object detectors' NMS threshold variations. However, the ranking of OOD detectors does not change much for different NMS values since most of the methods' curves share a common trend. Typically, performance metrics reach their upper limit around 0.5 and decline sharply when the softmax probability threshold is between 0.8 and 0.99. These sharp shifts of metrics' levels for high thresholds occur because object detectors hardly ever predict OOD objects with such a high probability, so OOD metrics are being computed for scarce data. This issue does not matter much in practice because, in real-world applications of object detectors, the NMS threshold is usually set somewhere between 0.5 and 0.8 - thus filtering most of the FPs and letting through most of the TPs.

5.3 Object detection

Many of the state-of-the-art OOD detection methods require altering default components of models to obtain improved uncertainty

Network	Method	FPR@95TPR \downarrow	AUROC↑	FPR@95TPR↓	AUROC↑
		$\text{VOC} \rightarrow \text{COCO}$		$BDD100k \rightarrow BDD100k$	
Faster R-CNN	Energy Gaussian GMM MCD NAPTRON OLN OpenDet OWOD RPN Standard VOS	$\begin{array}{c} 0.6762\\ 0.8899\\ 0.7441\\ 0.8267\\ \textbf{0.5677}\\ 0.9023\\ 0.8783\\ 0.7547\\ 0.8784\\ 0.8426\\ 0.6551 \end{array}$	0.8058 0.7103 0.7881 0.6848 0.7786 0.5994 0.6965 0.6858 0.6761 0.7459 0.8217	0.7987 0.8458 0.8420 0.9111 0.7603 0.9532 0.9004 0.8232 0.8894 0.8358 0.7839	0.7165 0.6320 0.6742 0.5778 0.6990 0.5979 0.5663 0.6757 0.6220 0.6554 0.7249
RetinaNet	Energy Gaussian GMM MCD NAPTRON OpenDet Standard	0.7994 0.7984 0.8044 0.8947 0.6063 0.6198 0.7917	0.6865 0.7456 0.7340 0.5023 0.7750 0.7717 0.7426	0.9081 0.8906 0.9081 0.9482 0.8738 0.9584 0.8787	0.5438 0.5882 0.5925 0.5143 0.6468 0.5075 0.6041
FCOS	Energy GMM MCD NAPTRON Standard	0.8148 0.6985 0.7721 0.5984 0.7721	0.6762 0.7276 0.5783 0.7773 0.7256	0.8778 0.8295 0.9330 0.7791 0.8615	0.5528 0.7067 0.5095 0.6978 0.6334

Table 4: Comparison of OOD detectors' performance

scores. It is expected that these alterations should not impose any damage on regular object detection quality. Therefore, we performed additional experiments to evaluate the influence of the customization of standard object detectors. In other words, how modification of the original architecture affects the performance of the object detector.

Object detector performance heavily depends on a predefined, application-dependent confidence score threshold. The authors of [10] observed that decreasing the NMS confidence score threshold and consequently massively increasing the number of false positive detections increases mAP. This phenomenon undermines the validity of using mAP as the primary quality metric. Therefore, as a way to gauge the impact of the OOD methods on object detection quality, we conduct a visual collation of TPR vs FP curves (see Figures 4 and 5). All the plots were horizontally limited to 100 000 FPs.

Drawing conclusions about the performance can hardly be completed merely by looking at the curves, so we compute the area under each curve (AUC), limiting the FP number to 2N where N is the number of all known ground truth objects in the test dataset. Next, we compare each method's AUC with the standard detector's AUC - Table 5 shows the results. Positive Δ AUC values signify a higher (better) curve, while negative Δ AUC indicates inferior performance.

OLN is class-agnostic (unlike all other models) and, for that reason, is subject to a more relaxed evaluation regime. The classagnostic true positive prediction required to the correct localization of an object, without the need for accurate label assignment.



Figure 4: Faster R-CNN-based detectors evaluated on COCO.

Figures 4 and 5 and Table 5 show that altering the default config-



Figure 5: RetinaNet-based detectors evaluated on BDD100k.

 Table 5: Object detection quality - area under the curve compared to baseline.

Architecture	Method	ΔAUC	ΔAUC
		$\text{VOC} \rightarrow \text{COCO}$	$BDD100k \rightarrow BDD100k$
Faster R-CNN	Gaussian	-0.0214	-0.0041
	GMM	-0.0180	-0.0015
	MCD	-0.0044	0.0001
	OLN	-0.1524	-0.1584
	OpenDet	0.0073	-0.0002
	OWOD	-0.0054	0.0024
	VOS	-0.0098	0.0007
RetinaNet	Gaussian	0.0024	-0.0315
	GMM	-0.0631	-0.0341
	MCD	-0.0128	-0.0038
	OpenDet	-0.0168	0.0019
FCOS	GMM	0.0089	-0.0111
	MCD	0.0075	-0.0032

urations of object detectors usually harms their performance. Except for the FCOS PASCAL-VOC \rightarrow COCO that provides relatively poor results, standard detectors are the most effective. Despite more relaxed rules, OLN fails to reach the performance levels of other detectors. GMM's anchor loss proves to be detrimental, especially for RetinaNet architecture. Gaussian negative log-likelihood loss bears a negative effect on object detectors too. For a fixed number of FPs, Gaussian Faster R-CNN generates hundreds of TPs less than the standard Faster R-CNN. RetinaNet OpenDet obtains significantly worse precision in the PASCAL-VOC \rightarrow COCO scenario, but in the remaining scenarios works very well. The AUC of VOS, OWOD, and MC-Dropout detectors is, on average, lower than the baseline, but the differences are smaller than for other methods.

5.4 Visual evaluation

Ideally, an OOD detector provides an independent uncertainty score, yet complementary to the objects detector softmax score. By design, the softmax score is solely responsible for distinguishing background from foreground objects, and the uncertainty score separates known objects from unknown ones.

The relationship between both the scores is nontrivial and simple statistic coefficients, such as the Pearson correlation coefficient, do not explain it sufficiently. Therefore, we provide visual 2D characteristics for the best-performing OOD detectors. A perfect characteristic would depict a cloud of blue points (OOD objects) vertically separable from a cloud of green points (TPs), and a cloud of red points (FPs) horizontally separable from the TPs. Figure 6 shows the NAPTRON uncertainty and softmax score relationship. We can observe that the blue OOD triangles are placed, in general, higher than the green TPs, especially for the most certain samples on the right side of the plot.

Overconfident OOD and FPs are the users' nightmare but our approach enables us to filter at least some of them. Examples of the samples identified as OOD by the proposed NAPTRON are in Figures 7 and 8. For all the examined algorithms, both scores are imperfect. In many cases, FPs have a high softmax score, whereas TPs have a lower score. Analogously, a low uncertainty score may be attributed to an OOD object. We observe that every type of uncertainty score is correlated with the softmax confidence score of an underlying object detector. This outcome is intuitive because we expect the TPs to have low uncertainty and high softmax probability. Thankfully, object detectors tend to assign lower probabilities to OOD objects, and OOD detectors assign higher uncertainty to FPs even though they are not explicitly meant to do so. These circumstances provide an opportunity to use softmax probability and one or more types of uncertainty scores combined in a manner that would boost both regular object detection and OOD detection performance. Finding an optimal way to combine multiple scores would require another set of experiments; this is beyond the scope of this work.



Figure 6: Faster R-CNN NAPTRON bounding box uncertainty score vs softmax confidence. Class 'bottle'.

6 Conclusion

In this work, we introduce the NAPTRON algorithm to estimate the uncertainty of predicted bounding boxes, based on the analysis of binary activation patterns. Our method enables the identification of the outlying test samples by computing the Hamming distance between



Figure 7: 1. OOD prediction - traffic sign \rightarrow train 2. FP prediction background (building) \rightarrow bottle.



(a) FP prediction background (quad) \rightarrow cow



(b) 1. OOD prediction - skateboard \rightarrow bicycle 2. FP prediction background (skatepark) \rightarrow boat.

Figure 8: Example of overconfident OOD and FP predictions with high NAP uncertainty. Each bounding box is attributed with softmax confidence and NAP uncertainty, respectively.

the sample's binary NAP and the most similar training NAP. Not only is our approach simple and intuitive but also outperforms all existing state-of-the-art OOD detectors. NAPTRON does not affect the underlying object detector's training or inference processes, which is another vital asset since we experimentally showed that altering default setups tends to lower object detection quality. Our experiments are conducted for three different architectures representing all common kinds of object detectors, which proves the universality of our approach.

Acknowledgements

This work was supported by the National Centre for Research and Development under the project LIDER/51/0221/L-11/19/NCBR/2020.

References

- Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom, 'nuscenes: A multimodal dataset for autonomous driving', in CVPR, (2020). 1
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko, 'End-to-end object detection with transformers', in *European conference on computer vision*, pp. 213–229. Springer, (2020). 1
- [3] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin, 'Mmdetection: Open mmlab detection toolbox and benchmark', *ArXiv*, abs/1906.07155, (2019). 4
- [4] Chih-Hong Cheng, Georg Nührenberg, and Hirotoshi Yasuoka, 'Runtime monitoring neuron activation patterns', in 2019 Design, Automation & Test in Europe Conference & Exhibition (DATE), pp. 300–303. IEEE, (2019). 3
- [5] Jiwoong Choi, Dayoung Chun, Hyun Kim, and Hyuk-Jae Lee, 'Gaussian yolov3: An accurate and fast object detector using localization uncertainty for autonomous driving', in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, (October 2019). 2, 5
- [6] Akshay Dhamija, Manuel Gunther, Jonathan Ventura, and Terrance Boult, 'The overlooked elephant of object detection: Open set', in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, (March 2020). 1
- [7] Xuefeng Du, Zhaoning Wang, Mu Cai, and Sharon Li, 'Towards unknown-aware learning with virtual outlier synthesis', in *International Conference on Learning Representations*, (2022). 1, 2, 5
- [8] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, 'The pascal visual object classes (voc) challenge', *International Journal of Computer Vision*, 88(2), 303–338, (June 2010). 4
- [9] Yarin Gal and Zoubin Ghahramani, 'Dropout as a bayesian approximation: Representing model uncertainty in deep learning', in *international conference on machine learning*, pp. 1050–1059. PMLR, (2016). 2
- [10] David Hall, Feras Dayoub, John Skinner, Haoyang Zhang, Dimity Miller, Peter Corke, Gustavo Carneiro, Anelia Angelova, and Niko Sünderhauf, 'Probabilistic object detection: Definition and evaluation', in *The IEEE Winter Conference on Applications of Computer Vision*, pp. 1031–1040, (2020). 6
- [11] Jiaming Han, Yuqiang Ren, Jian Ding, Xingjia Pan, Ke Yan, and Gui-Song Xia, 'Expanding low-density latent regions for open-set object detection', in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (2022). 2, 5
- [12] Boris Hanin and David Rolnick, 'Deep relu networks have surprisingly few activation patterns', in *Neural Information Processing Systems*, (2019). 1
- [13] D Hartmann, D Franzen, and S Brodehl, 'Studying the evolution of neural activation patterns during training of feed-forward relu networks', *Front. Artif. Intell*, 4, 642374, (2021). 1
- [14] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun, 'Deep residual learning for image recognition', 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778, (2015). 1
- [15] Dan Hendrycks and Kevin Gimpel, 'A baseline for detecting misclassified and out-of-distribution examples in neural networks', *Proceedings* of International Conference on Learning Representations, (2017). 1
- [16] Xiaowei Huang, Daniel Kroening, Wenjie Ruan, James Sharp, Youcheng Sun, Emese Thamo, Min Wu, and Xinping Yi, 'A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability', *Computer Science Review*, **37**, 100270, (2020). 1
- [17] K J Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Bal-

asubramanian, 'Towards open world object detection', in *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2021), (2021). 2, 5

- [18] Dahun Kim, Tsung-Yi Lin, Anelia Angelova, In So Kweon, and Weicheng Kuo, 'Learning open-world object proposals without learning to classify', *IEEE Robotics and Automation Letters (RA-L)*, (2022). 2, 5
- [19] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar, 'Focal loss for dense object detection', in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, (Oct 2017). 1, 4
- [20] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick, 'Microsoft coco: Common objects in context', in *European Conference on Computer Vision*, (2014). 4
- [21] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li, 'Energybased out-of-distribution detection', Advances in Neural Information Processing Systems, (2020). 2, 5
- [22] Dimity Miller, Feras Dayoub, Michael Milford, and Niko Sünderhauf, 'Evaluating merging strategies for sampling-based uncertainty techniques in object detection', 2019 International Conference on Robotics and Automation (ICRA), 2348–2354, (2019). 2
- [23] Dimity Miller, Niko Sünderhauf, Michael Milford, and Feras Dayoub, 'Uncertainty for identifying open-set errors in visual object detection', *IEEE Robotics and Automation Letters*, 1–1, (2021). 2, 5
- [24] Dimity Miller, Niko Sunderhauf, Haoyang Zhang, David Hall, and Feras Dayoub, 'Benchmarking sampling-based probabilistic object detectors', in *Proceedings of the IEEE/CVF Conference on Computer Vi*sion and Pattern Recognition (CVPR) Workshops, (June 2019). 2, 5
- [25] Bartłomiej Olber, Krystian Radlak, Adam Popowicz, Michal Szczepankiewicz, and Krystian Chachuła, 'Detection of out-ofdistribution samples using binary neuron activation patterns', in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3378–3387, (June 2023). 1, 3, 4
- [26] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi, 'You only look once: Unified, real-time object detection', in *Proceedings* of the IEEE conference on computer vision and pattern recognition (CVPR), pp. 779–788, (2016). 1
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun, 'Faster r-cnn: Towards real-time object detection with region proposal networks', in Advances in Neural Information Processing Systems, eds., C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, volume 28. Curran Associates, Inc., (2015). 4
- [28] Walter J. Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E. Boult, 'Toward open set recognition', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **35**(7), 1757–1772, (2013). 1
- [29] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov, 'Scalability in perception for autonomous driving: Waymo open dataset', in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (June 2020). 1
- [30] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He, 'Fcos: Fully convolutional one-stage object detection', in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, (October 2019). 4
- [31] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu, 'Generalized out-of-distribution detection: A survey', *arXiv preprint arXiv:2110.11334*, (2021). 1
- [32] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell, 'Bdd100k: A diverse driving dataset for heterogeneous multitask learning', in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (June 2020). 1, 4