

Automated Model Evaluation for Object Detection via Prediction Consistency and Reliability

Seungju Yoo¹ Hyuk Kwon¹ Joong-Won Hwang² Kibok Lee¹

¹Yonsei University ²ETRI

¹{seungju_yoo, kh12043, kibok}@yonsei.ac.kr ²jwhwang@etri.re.kr

Abstract

Recent advances in computer vision have made training object detectors more efficient and effective; however, assessing their performance in real-world applications still relies on costly manual annotation. To address this limitation, we develop an automated model evaluation (AutoEval) framework for object detection. We propose Prediction Consistency and Reliability (PCR), which leverages the multiple candidate bounding boxes that conventional detectors generate before non-maximum suppression (NMS). PCR estimates detection performance without ground-truth labels by jointly measuring 1) the spatial consistency between boxes before and after NMS, and 2) the reliability of the retained boxes via the confidence scores of overlapping boxes. For a more realistic and scalable evaluation, we construct a meta-dataset by applying image corruptions of varying severity. Experimental results demonstrate that PCR yields more accurate performance estimates than existing AutoEval methods, and the proposed meta-dataset covers a wider range of detection performance. The code is available at <https://github.com/YonseiML/autoeval-det>.

1. Introduction

Effective evaluation of machine learning models is essential before deployment, particularly when the target domain differs from the source domain because of environmental shifts or discrepancies in data distribution. However, assessing model performance in a new environment is challenging, as annotating test data is often costly and time-consuming. Automated model evaluation (AutoEval) [6] addresses this issue by estimating performance on unlabeled test datasets.

Research on AutoEval has primarily focused on image classification, which serves as a foundation for the quick assessment of new methods [6, 10, 15, 40]. In image classification, the limitations of evaluations on standard benchmarks often stem from domain discrepancies at the feature extraction stage, referred to as covariate shift, where fea-




Figure 1. **Visual example of PCR.** Green boxes represent the ground-truth bounding boxes, and red and blue boxes denote an incorrect detection with low confidence and a correct detection with high confidence, respectively. Orange boxes show the pre-NMS candidate boxes, where the overlaid numbers indicate their confidence scores. **(a) Consistency.** The red box overlaps many pre-NMS boxes, yielding high consistency. Our consistency score measures spatial consistency with a merged pre-NMS box, motivated by the observation that **boxes with low confidence and high consistency** correlate with lower mAP. **(b) Reliability.** The blue box overlaps many pre-NMS boxes with high confidence scores, yielding high reliability. Our reliability score measures the proportion of overlapping pre-NMS boxes with high confidence scores, motivated by the observation that **boxes with high confidence and high reliability** correlate with higher mAP.

tures learned during training become less effective when deployed in new environments. Recent studies have attempted to address this by measuring distances between the feature distributions of the source and target datasets [6, 15]. Although covariate shift is an important factor, we argue that it is insufficient for AutoEval in other computer vision tasks such as object detection, which introduce additional challenges requiring more advanced evaluation strategies.

Unlike image classification, object detection performance is sensitive to various factors such as variations in object scales, occlusion, background clutter, and object interactions; covariate shift cannot capture these intricate spatial relationships. Hence, we argue that estimat-

ing covariate shift alone cannot address the full spectrum of challenges in object detection. To address these challenges, we adopt a bottom-up strategy by 1) collecting several measures from model outputs, 2) analyzing their correlations with performance variations, and 3) developing a method based on these insights. Specifically, we focus on the observation that conventional object detection models first generate many candidate bounding boxes, and then retain high-confidence boxes while discarding overlapping low-confidence boxes through non-maximum suppression (NMS) [18, 31, 38, 39, 43]. We examine relationships between pre- and post-NMS boxes by measuring their geometric proximity and confidence scores, capturing both localization and classification aspects.

Building on these observations, we propose **Prediction Consistency and Reliability (PCR)** as an effective AutoEval method for object detection. Figure 1 illustrates the concepts of consistency and reliability in PCR. For low-confidence boxes, the high spatial consistency between pre- and post-NMS boxes indicates repeated mislocalization, suggesting that the model has misidentified objects. Because the associated pre-NMS boxes have even lower confidence than nearby post-NMS boxes, such consistent localization of low-confidence boxes is prone to incorrect detection results. In contrast, for high-confidence boxes, the reliability of post-NMS boxes is assessed based on the confidence scores of nearby pre-NMS boxes, capturing both localization and classification accuracy. Pre-NMS boxes with high confidence imply that the model is confident in both aspects, thereby promoting correct detection results.

Furthermore, we construct a meta-dataset to facilitate a more realistic and scalable evaluation of AutoEval methods for object detection. A recently proposed meta-dataset [45] relies on data augmentation techniques, such as sharpness adjustment, equalization, color temperature shifts, solarization, autocontrast, brightness modification, and rotation; however, we argue that such strong augmentations may not adequately reflect the environmental shifts or discrepancies in data distribution observable in real-world applications, and that such a meta-dataset may fail to capture the full range of detection performance in practice. Instead, we apply image corruptions of varying severity, adopting the transformations used to generate ImageNet-C [19], which are specifically designed to simulate real-world corruptions. The proposed meta-dataset is *realistic* in that it employs real-world corruptions rather than artificial augmentations, and *scalable* in that varying severity levels enable evaluation across a broader spectrum of detection performance.

We summarize our contributions as follows:

- **Prediction Consistency and Reliability (PCR).** We propose PCR, an effective AutoEval method for object detection that leverages the spatial alignment and confidence scores of the bounding boxes before and after NMS.

- **Corruption-based meta-dataset.** We construct a meta-dataset using image corruptions of varying severity levels, yielding a realistic and scalable benchmark that enables AutoEval for object detection.
- **Extensive empirical validation.** PCR consistently outperforms state-of-the-art AutoEval methods across various object detection models, including RetinaNet [31] and Faster R-CNN [39] paired with ResNet-50 [17] and Swin Transformer [32] backbones, across both augmentation- and corruption-based meta-datasets.

2. Related Work

Object Detection aims to localize and classify objects within an image by simultaneously predicting bounding boxes and class labels. Object detectors are typically categorized into 1) two-stage methods, such as R-CNN variants [12, 13, 39], which generate and refine region proposals, and 2) one-stage methods, such as YOLO [37] and RetinaNet [31], which perform detection in a single pass. While these models are commonly evaluated on standard benchmark datasets including COCO [29] and Pascal VOC [8], distribution shifts in real-world scenarios, such as varying lighting conditions, often lead to discrepancies between benchmark performance and actual deployment performance. Moreover, acquiring high-quality annotations in real-world scenarios is often costly and time-consuming. To address these challenges, we explore a framework for estimating detector performance on unlabeled test datasets without relying on manual annotations.

Automated Model Evaluation (AutoEval) aims to estimate model performance on unlabeled test datasets [6]. In the absence of ground-truth labels, AutoEval measures scores that correlate with performance, which are then used to predict performance via regression. Early studies utilized confidence-based measures, such as maximum softmax probabilities [10, 15, 20] and predictive entropy [40], demonstrating that simple scalar indicators derived from model outputs can effectively estimate classification performance. Subsequent studies investigated prediction disagreement, either between independently trained models [22] or between model outputs with and without dropout [1, 28], observing that higher disagreement often indicates poorer generalization. Other studies introduced self-supervised surrogate tasks, such as rotation prediction [7] and contrastive learning [35], providing label-free proxies for representation quality. Another line of research leveraged dataset-level feature statistics, such as Fréchet distances [6] and average feature energy [36], demonstrating correlations with model performance. While these methods have proven effective in AutoEval for image classification, the extension to other computer vision tasks, such as object detection, remains underexplored.

AutoEval for Object Detection extends the concept of AutoEval to object detection, where performance depends on both object localization and classification, requiring the method to capture intricate spatial relationships among objects. Recently, Yang et al. [45] proposed a method applicable to AutoEval for object detection, which compares bounding boxes generated by a detector with and without Monte Carlo dropout [9] perturbations applied to feature maps. The bounding box stability is measured by the difference between the original and perturbed predictions, which serves as a score for estimating detection performance. However, this method has several limitations. First, it is stochastic at test time, leading to inconsistent estimates across trials. Second, it requires an additional forward pass to compute perturbed predictions, effectively doubling the inference cost. Third, it does not incorporate model confidence, which reflects prediction uncertainty [16], leading to an overestimation of unreliable predictions. Furthermore, the meta-dataset constructed in their work relies on strong data augmentations with a fixed set of hyperparameters, which often yield unrealistic images and may fail to capture the full spectrum of detection performance in real-world scenarios. In response, we propose an efficient and effective AutoEval method for object detection that explicitly leverages confidence scores to estimate detection performance. In contrast to the prior work incorporating an additional forward pass with dropout [45], our proposed method utilizes pre-NMS boxes that are generated in a single forward pass, while containing valuable information about localization [2]. Moreover, we construct a meta-dataset using image corruptions of varying severity, enabling a more realistic and scalable evaluation.

Model Confidence in Object Detection is quantified as confidence scores per class associated with each bounding box predicted by object detection models. These scores are typically interpreted as classification probabilities, and prior work has primarily focused on calibrating them [27, 34], with limited exploration of alternative interpretations or uses. In contrast, Sun et al. [42] proposed a confidence-based bounding box localization strategy for small objects. However, their approach considered confidence scores a measure of uncertainty, and the explicit relationship between confidence scores and localization quality remains unexplored. Motivated by our observation that confidence scores reflect localization quality, we propose a confidence-driven approach to estimate detector performance.

3. Prediction Consistency and Reliability

In this section, we introduce PCR, an effective AutoEval method for object detection. PCR consists of two components: for each final prediction—or equivalently, each post-NMS box—1) the *consistency score* measures the spatial




Figure 2. The average IoU between ground-truth boxes and final predictions, grouped by confidence level using a threshold of 0.5 across datasets. Predictions with low confidence generally exhibit lower IoU than those with high confidence, indicating a correlation between confidence and localization quality.

consistency with its corresponding merged pre-NMS box, and 2) the *reliability score* measures the proportion of overlapping pre-NMS boxes with high confidence scores.

Conventional object detection models generate multiple candidate bounding boxes, followed by a filtering step such as Non-Maximum Suppression (NMS) [18, 31, 38, 39, 43]. Although NMS is essential for removing redundant boxes to produce final predictions, discarded pre-NMS boxes still contain valuable information about localization [2]. This motivates us to investigate pre-NMS boxes associated with final predictions to facilitate AutoEval for object detection.

3.1. Consistency

The consistency score is motivated by the observation in Figure 2 that the Intersection over Union (IoU) between a final prediction and its ground-truth box is correlated with the confidence score, *i.e.*, predictions with low confidence tend to exhibit low IoU with ground-truth boxes. This suggests that model confidence, which is often used as a proxy for classification accuracy [16], also serves as an indicator of localization performance. This becomes more evident when considering the pre-NMS boxes associated with a final prediction: these boxes always have even lower confidence, implying that the detector consistently focuses on a region that does not contain any ground-truth object, which in turn correlates with low mAP. However, measuring consistency for each individual pre-NMS box may introduce redundancy and lead to unstable assessments. Instead, we merge the pre-NMS boxes and compare the resulting merged box with the final prediction, as illustrated in Figure 3.

Let $\mathcal{B}^{(i)} = \{B^{(ij)}\}_{j=1}^{K_i}$ denote the set of pre-NMS boxes overlapping with the i -th final prediction $B_{\text{final}}^{(i)}$, where K_i is the number of pre-NMS boxes. The merged box $B_{\text{merge}}^{(i)} = (x_{\text{merge}}^{(i)}, y_{\text{merge}}^{(i)}, w_{\text{merge}}^{(i)}, h_{\text{merge}}^{(i)})$ is defined as follows:

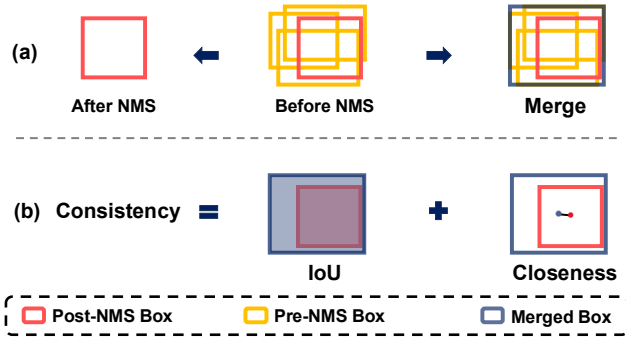


Figure 3. (a) A merged box tightly encloses all pre-NMS boxes associated with a post-NMS box. (b) Consistency is computed based on IoU and a closeness term measured by the normalized distance between the center points of a post-NMS box and its corresponding merged box.

$$\begin{aligned}
 x_{\text{merge}}^{(i)} &= \left(\min_j x_{\text{min}}^{(ij)} + \max_j x_{\text{max}}^{(ij)} \right) / 2, \\
 y_{\text{merge}}^{(i)} &= \left(\min_j y_{\text{min}}^{(ij)} + \max_j y_{\text{max}}^{(ij)} \right) / 2, \\
 w_{\text{merge}}^{(i)} &= \max_j x_{\text{max}}^{(ij)} - \min_j x_{\text{min}}^{(ij)}, \\
 h_{\text{merge}}^{(i)} &= \max_j y_{\text{max}}^{(ij)} - \min_j y_{\text{min}}^{(ij)},
 \end{aligned} \quad (1)$$

where $x_{\text{min}}^{(ij)}$, $x_{\text{max}}^{(ij)}$, $y_{\text{min}}^{(ij)}$, $y_{\text{max}}^{(ij)}$ denote the left, right, top, and bottom coordinates of the pre-NMS box $B^{(ij)}$, respectively. Then, the merged box tightly encloses all pre-NMS boxes, as shown in Figure 3(a).

Given a final prediction $B_{\text{final}}^{(i)}$ and its corresponding merged pre-NMS box $B_{\text{merge}}^{(i)}$, we measure their consistency using two metrics: IoU and the closeness between their centers, as illustrated in Figure 3(b). While IoU is a standard measure of spatial alignment, it can be misleading when the merged box is elongated, where small shifts along the shorter axis may cause significant drops in IoU, whereas large shifts along the longer axis may still yield high IoU despite localization errors. Moreover, because the merged box is generally large and tends to enclose the final prediction, IoU alone may not accurately capture their geometric relationship [48]. To address this, we introduce an additional measure that quantifies the Closeness between the Centers (CC) of the final prediction and the merged box:

$$\begin{aligned}
 \text{CC} \left(B_{\text{final}}^{(i)}, B_{\text{merge}}^{(i)} \right) = \\
 1 - \frac{\sqrt{\left(x_{\text{final}}^{(i)} - x_{\text{merge}}^{(i)} \right)^2 + \left(y_{\text{final}}^{(i)} - y_{\text{merge}}^{(i)} \right)^2}}{\sqrt{w_{\text{final}}^{(i)2} + h_{\text{final}}^{(i)2} / 2}}. \quad (2)
 \end{aligned}$$

Intuitively, CC is computed as one minus the normalized distance between the two centers, where the normalization

factor is the half-diagonal of the final prediction box, ensuring scale invariance. A higher CC value indicates that two centers are closer, with a value of 1 implying perfect alignment.¹ The consistency score of the i -th final prediction is defined as the average of IoU and CC:

$$S^{C(i)} = \frac{\text{IoU} \left(B_{\text{final}}^{(i)}, B_{\text{merge}}^{(i)} \right) + \text{CC} \left(B_{\text{final}}^{(i)}, B_{\text{merge}}^{(i)} \right)}{2}. \quad (3)$$

The consistency score of an image is computed as the weighted average of per-prediction consistency scores, where each weight is a scaled version of the confidence score of the corresponding final prediction:

$$S^C = \frac{1}{N} \sum_{i=1}^N S^{C(i)} \cdot \sigma_C \left(h \left(B_{\text{final}}^{(i)} \right) \right), \quad (4)$$

where $h(\cdot)$ denotes the confidence score of a prediction, $\sigma_C(x) = 1 / (1 + \exp(-k_C(x - c)))$ is a sigmoid function parameterized by a negative scale k_C and a confidence threshold c , and N is the number of final predictions. The function σ_C approaches one for low-confidence boxes and zero for high-confidence ones, thus emphasizing the final predictions with low confidence. A high consistency score indicates that the detector consistently focuses on a region without a ground-truth object, which is associated with low mAP. Figure 4(a) confirms that the consistency score S^C exhibits a strong negative correlation with mAP.

Note that our notion of consistency differs from prior works that assess the localization stability between model predictions and those obtained by additional forward passes with model transformations or image perturbations [21, 45]. Instead, our consistency score quantifies the spatial similarity among pre- and post-NMS boxes generated from a single forward pass. Indeed, we observe no correlation between our consistency score and the BoS [45] score, highlighting their conceptual distinction.

3.2. Reliability

The reliability score is also motivated by Figure 2, with a focus on final predictions with high confidence, which tend to exhibit high IoU with ground-truth boxes. Intuitively, NMS is employed in object detection to select the most probable bounding box among overlapping pre-NMS candidate boxes, based on confidence scores. A concentration of pre-NMS boxes with high confidence indicates that both the classification and regression heads repeatedly identify an object at the same location, thereby promoting the reliability of the prediction. However, this reliability diminishes if some overlapping pre-NMS boxes have low confidence, indicating uncertainty in the prediction around the

¹CC resembles the normalized distance in DIoU [48], differing in several aspects, *e.g.*, the normalization factor.

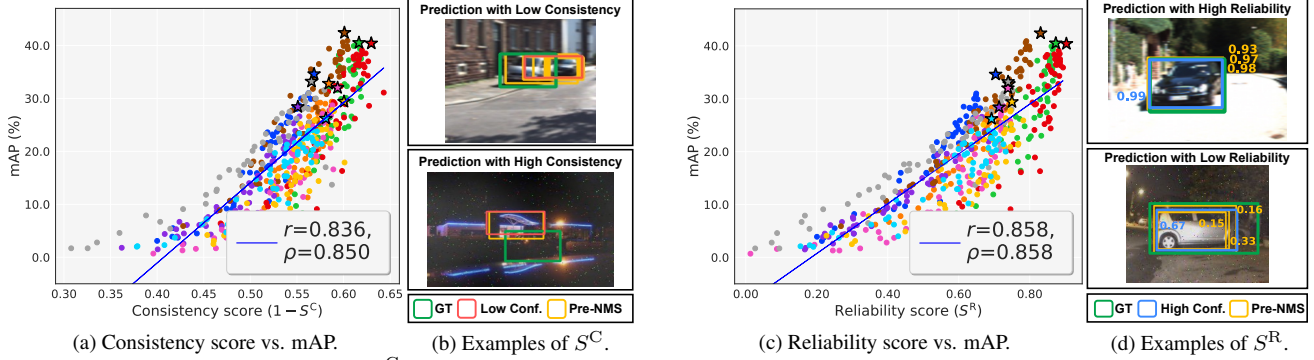


Figure 4. (a) The consistency score S^C shows a strong *negative* correlation with mAP. (b) Predictions with low confidence and high consistency suggest that the detector consistently localizes the same region without any object, indicating a detection failure. In contrast, predictions with low confidence and low consistency provide insufficient information to make a decision. (c) The reliability score S^R shows a strong *positive* correlation with mAP. (d) Predictions with high confidence and high reliability suggest that the detector repeatedly localizes and classifies the same object, indicating a successful detection. In contrast, predictions with high confidence and low reliability provide insufficient information to make a decision.

location. Hence, we define the reliability score as the proportion of pre-NMS boxes with confidence scores above the confidence threshold c .

Given the final predictions $\{B_{\text{final}}^{(i)}\}_{i=1}^N$ and the set of *all* pre-NMS boxes for an image $\mathcal{P} = \bigcup_i \mathcal{B}^{(i)}$ ², the reliability score of an image is defined as:

$$S^R = \frac{\sum_{i=1}^N \sum_{j=1}^{K_i} \mathbb{I}[h(B_{\text{final}}^{(i)}) > c] \sigma_R(h(B^{(ij)}))}{\sum_{B^{(ij)} \in \mathcal{P}} \sigma_R(h(B^{(ij)}))}, \quad (5)$$

where $\mathbb{I}[\cdot]$ is the indicator function and $\sigma_R(x) = \alpha + (1 - \alpha)/(1 + \exp(-k_R(x - c)))$ is a sigmoid function parameterized by a positive scale k_R , the confidence threshold c , and a floor value α . The function σ_R approaches one for high-confidence boxes and α for low-confidence ones, thus emphasizing the pre-NMS boxes with high confidence while retaining a minimal contribution from those with low confidence. A high reliability score indicates that the detector repeatedly assigns high confidence to a region, suggesting strong consensus in both classification and localization, which is associated with high mAP. Figure 4(b) confirms that the reliability score S^R exhibits a strong positive correlation with mAP.

3.3. PCR for AutoEval

We extend the notation of our scores to explicitly include an object detector f and an input image \mathcal{I} . Our proposed PCR consists of two scores: the consistency score $S^C(f, \mathcal{I})$ and the reliability score $S^R(f, \mathcal{I})$. These scores are evaluated over a meta-dataset $\mathcal{D} = \{\mathcal{D}_m\}_{m=1}^M$, where each dataset \mathcal{D}_m is a transformed version of the source dataset, *e.g.*, using augmentations or image corruptions. Given an object

²Some pre-NMS boxes may overlap multiple final predictions; to ensure that each pre-NMS box is counted only once, we take the union over all pre NMS-boxes.

detector f and the m -th dataset \mathcal{D}_m , the average scores are computed as:

$$\begin{aligned} \bar{S}^C(f, \mathcal{D}_m) &= \frac{1}{|\mathcal{D}_m|} \sum_{\mathcal{I} \in \mathcal{D}_m} S^C(f, \mathcal{I}), \\ \bar{S}^R(f, \mathcal{D}_m) &= \frac{1}{|\mathcal{D}_m|} \sum_{\mathcal{I} \in \mathcal{D}_m} S^R(f, \mathcal{I}). \end{aligned} \quad (6)$$

To perform AutoEval, we fit a linear regression model using least squares over the meta-dataset. Given an object detector f and the m -th dataset \mathcal{D}_m , the mAP is estimated as:

$$\widehat{\text{mAP}}(f, \mathcal{D}_m) = w_0 + \sum_{t=1}^T w_t \cdot \bar{S}_t(f, \mathcal{D}_m), \quad (7)$$

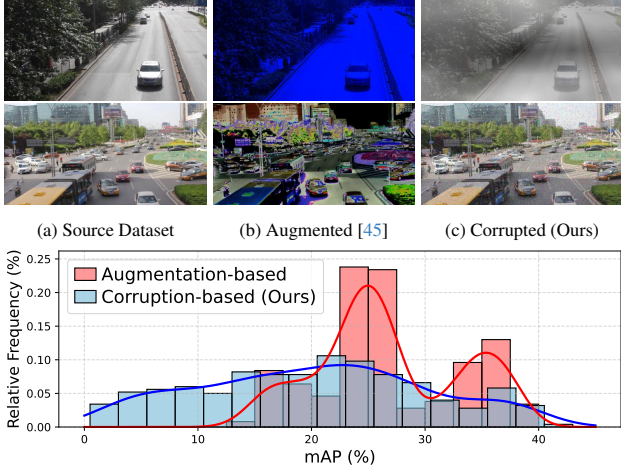
where $\{w_t\}_{t=0}^T$ are the regression coefficients and \bar{S}_t is the t -th AutoEval score. For PCR, the mAP estimate is:

$$\widehat{\text{mAP}}(f, \mathcal{D}_m) = w_0 + w_1 \cdot \bar{S}^C(f, \mathcal{D}_m) + w_2 \cdot \bar{S}^R(f, \mathcal{D}_m).$$

For training and evaluation, we adopt leave-one-out cross-validation, following Yang et al. [45]: a meta-dataset is constructed using all source datasets except one for training, while the remaining source dataset is held out for evaluation and used without any transformations.

4. Corruption-Based Meta-Dataset

To assess AutoEval methods for object detection, a meta-dataset is typically constructed by applying transformations to source datasets. While prior work [45] adopts strong data augmentation techniques for this purpose, we argue that such augmentations may not adequately capture the environmental shifts or discrepancies in data distribution observable in real-world scenarios. Consequently, such a



(d) **Distribution of mAP: Augmentation-based [45] vs. Corruption-based (Ours)**

Figure 5. (a–c) Sample images from: (a) the original DETRAC [44] dataset, (b) the augmentation-based meta-dataset, and (c) the corruption-based meta-dataset. (d) Histogram of mAP distributions for the augmentation-based meta-dataset [45] and our proposed corruption-based meta-dataset.

meta-dataset may fail to span the full spectrum of detection performance observable in practice.

To address this, we instead adopt the transformations applied to ImageNet-C [19], which simulate real-world corruptions. To capture scenarios of varying difficulty, we apply each corruption at different severity levels. Specifically, the meta-dataset consists of 50 datasets using ten corruptions {gaussian noise, shot noise, impulse noise, defocus blur, snow, frost, fog, contrast, pixelate, jpeg} with five severity levels {1, 2, 3, 4, 5}. We exclude corruptions that alter the coordinates of bounding boxes, *e.g.*, {zoom blur, elastic transformation}. This strategy yields a meta-dataset that is both realistic—as it reflects real-world corruptions rather than artificial augmentations—and scalable—as the varying severity levels enable evaluation across a broader spectrum of detection performance.

Figure 5 compares the augmentation-based meta-dataset from prior work [45] and our proposed corruption-based meta-dataset. As shown in Figure 5(b), images generated by strong augmentations often appear artificial, whereas those generated by corruptions are more realistic, as shown in Figure 5(c). In terms of performance coverage, Figure 5(d) shows that the augmentation-based meta-dataset exhibits a relatively narrow and skewed distribution of mAP, with limited coverage below 15% and a concentration around 25% and 35%. In contrast, our corruption-based meta-dataset covers a broader range of mAP, spanning from near 0% to 40% with a smoother distribution.

5. Experiments

5.1. Experimental Setup

Detectors. We compare our proposed PCR with baseline AutoEval methods using four object detectors: RetinaNet [31] and Faster R-CNN [39] paired with ResNet-50 [17] and Swin Transformer [32] backbones.

Datasets. Following the dataset sampling strategy in prior work [45], each dataset in the meta-dataset contains 250 images sampled from a source dataset. For **vehicle detection**, we use 10 source datasets: COCO [30], BDD [46], Cityscapes [5], DETRAC [44], ExDark [33], KITTI [11], Self-driving [23], Roboflow [26], Udacity [25], and Traffic [24]. For **pedestrian detection**, we use 9 source datasets: COCO [30], Caltech [14], CrowdHuman [41], Cityscapes [5], Self-driving [23], ExDark [33], EuroCity [3], KITTI [11], and CityPersons [47].

Baselines. We compare PCR with five AutoEval methods: 1) Prediction Score (PS) [20], 2) Entropy Score (ES) [40], 3) Average Confidence (AC) [15], 4) Average Thresholded Confidence (ATC) [10], and 5) Box Stability (BoS) [45]. The first four methods were originally proposed for image classification; we adapt them to object detection by computing their metrics using the confidence scores of the predicted bounding boxes. For BoS, we report the average over three runs to account for the stochasticity introduced by MC dropout. We use the same hyperparameters as reported in BoS [45] for all baselines and present results based on our own replications.

Implementation Details. For hyperparameters of PCR, we set the confidence threshold to 0.5, the scale in σ^C to $k_C = -60$, the scale in σ^R to $k_R = 10$, and the floor value in σ^R to $\alpha = 0.2$. Other implementation details and hyperparameter tuning results are provided in the supplementary material.

Evaluation Metric. We evaluate performance using the root mean squared error (RMSE) between the estimated mAP and the true mAP, where a lower RMSE indicates a more accurate estimation. Results in the correlation metric are reported in the supplementary material.

5.2. Vehicle Detection

Table 1 summarizes the mAP estimation results for vehicle detection. We report results on both the augmentation-based meta-dataset [45] and our proposed corruption-based meta-dataset. Each cell presents the average RMSE between the estimated mAP and the true mAP for each method with a pair of detector and meta-dataset, where each detector is trained on the “car” class from COCO [30]. We also report the average RMSE across all four detectors and both meta-datasets, along with the average performance rank. PCR achieves the lowest average RMSE of 5.03 and the best average rank of 1.13, consistently outperforming all baselines.

Meta-dataset	Augmentation-based [45]				Corruption-based (Ours)				Avg. RMSE	Avg. Rank
Detector	RetinaNet		Faster R-CNN		RetinaNet		Faster R-CNN			
	ResNet-50	Swin-T	ResNet-50	Swin-T	ResNet-50	Swin-T	ResNet-50	Swin-T		
PS [20]	4.85	7.39	4.95	6.28	<u>11.30</u>	10.44	10.41	<u>6.30</u>	7.74	3.13
ES [40]	6.51	<u>5.58</u>	5.74	6.46	14.87	7.20	12.53	10.02	8.62	4.75
AC [15]	9.50	8.17	5.68	6.97	14.23	10.65	11.77	7.23	9.27	5.13
ATC [10]	5.52	11.20	4.89	6.86	14.10	12.52	11.98	6.29	9.17	4.38
BoS [45]	<u>3.11</u>	7.69	3.33	<u>5.57</u>	13.50	<u>5.18</u>	<u>10.32</u>	6.85	<u>6.94</u>	<u>2.50</u>
PCR (Ours)	3.01	5.47	<u>4.00</u>	5.08	6.43	4.80	6.90	4.51	5.03	1.13

Table 1. Comparison of AutoEval methods for vehicle detection using four detectors on two meta-datasets. The best result for each combination of detector and meta-dataset is highlighted in **bold** and the second-best is underlined.

Meta-dataset	Augmentation-based [45]				Corruption-based (Ours)				Avg. RMSE	Avg. Rank
Detector	RetinaNet		Faster R-CNN		RetinaNet		Faster R-CNN			
	ResNet-50	Swin-T	ResNet-50	Swin-T	ResNet-50	Swin-T	ResNet-50	Swin-T		
PS [20]	7.47	7.35	6.20	4.64	10.37	10.71	<u>6.84</u>	5.93	7.44	4.50
ES [40]	5.13	3.55	8.16	6.89	8.46	7.40	<u>10.66</u>	10.00	7.53	5.00
AC [15]	5.03	4.39	6.80	4.07	<u>7.51</u>	<u>5.90</u>	9.15	<u>5.37</u>	<u>6.03</u>	<u>3.13</u>
ATC [10]	6.21	7.46	6.98	3.88	7.59	6.51	8.33	5.95	6.62	4.13
BoS [45]	<u>3.77</u>	<u>3.17</u>	<u>5.88</u>	<u>3.57</u>	10.23	6.08	10.09	6.98	6.22	3.25
PCR (Ours)	3.62	2.90	2.79	2.38	3.56	4.30	5.10	4.07	3.60	1.00

Table 2. Comparison of AutoEval methods for pedestrian detection using four detectors on two meta-datasets. The best result for each combination of detector and meta-dataset is highlighted in **bold** and the second-best is underlined.

5.3. Pedestrian Detection

Table 2 summarizes the mAP estimation results for pedestrian detection. For this experiment, detectors are trained on CrowdHuman [41]. Again, we report results on both the augmentation-based meta-dataset [45] and our proposed corruption-based meta-dataset. PCR achieves the lowest average RMSE of 3.60 and the best average performance rank of 1.00. These results suggest that PCR offers accurate mAP estimates, consistently outperforming all baselines under fair evaluation settings.

5.4. Analysis

We conduct experiments to analyze both the proposed method and the meta-dataset. Unless otherwise specified, all experiments use RetinaNet [31] with a ResNet-50 [17] backbone evaluated on the corruption-based meta-dataset.

Robustness of PCR Across Varying Difficulty. Recall that the meta-dataset includes five severity levels of corruption to reflect varying difficulties observable in real-world applications. To assess the robustness of our proposed PCR across a broader spectrum of detection performance, we perform linear regression separately on each 10 datasets corresponding to a single severity level.




Figure 6. Severity vs. RMSE. Table 3. Ablation on S^C and S^R .

S^C	S^R	Avg. RMSE
✓		6.67
	✓	6.52
✓	✓	6.43

Figure 6 compares the average RMSE of BoS [45] and PCR. While the average RMSE increases with the severity level for both methods, PCR consistently achieves a lower RMSE and demonstrates robustness across different severity levels. This suggests that PCR not only performs well within a specific range of mAP, but also generalizes across diverse conditions, enabling reliable AutoEval across a wide range of performance distributions.

Ablation on Components of PCR. To evaluate the contribution of each component in PCR, we conduct an ablation study by selectively removing either the consistency or the reliability score. Table 3 shows that using either score alone already outperforms the baseline methods, while combining both scores achieves the best performance. This result suggests that these two scores capture complementary aspects of detection performance: while the consistency score

Consistency Score	Avg. RMSE	Consistency Scaling	Avg. RMSE
IoU only	6.76	All (S_{all}^C)	8.24
+ Closeness	6.67	Low (S^C)	6.67

Table 4. Ablation on components of consistency.

Table 5. Ablation on confidence-based scaling.

Method	RetinaNet + ResNet-50		
	mAP	mAP ₅₀	mAP ₇₅
PS [20]	11.30	18.64	12.16
ES [40]	14.87	25.12	15.63
AC [15]	14.23	23.58	15.24
ATC [10]	14.10	23.33	15.10
BoS [45]	13.50	22.75	14.43
PCR (Ours)	6.43	9.89	7.78

Table 6. RMSE Comparison of AutoEval methods in estimating mAP, mAP₅₀ and mAP₇₅.

focuses on low-confidence predictions, the reliability score focuses on high-confidence predictions. Their combination allows PCR to effectively capture the relationships among pre- and post-NMS boxes based on the confidence scores.

Ablation on Components of Consistency. While IoU is a widely used metric to quantify the similarity between two boxes, our consistency score additionally incorporates the closeness between the centers of the final prediction and the corresponding merged pre-NMS box. To confirm the effectiveness of this additional component, we conduct an ablation study with and without it. As shown in Table 4, incorporating the closeness of centers improves the consistency measure, leading to improved performance.

Effect of Confidence-Based Scaling in Consistency. In PCR, the consistency score emphasizes low-confidence predictions by scaling with confidence scores, following the intuition discussed in Section 3.1. To validate this design choice, we conduct an ablation study comparing the proposed consistency score in Eq. (4) and a variant that does not apply confidence-based scaling, which is formulated as:

$$S_{\text{all}}^C = \frac{1}{N} \sum_{i=1}^N S^{C(i)}, \quad (8)$$

Table 5 presents the RMSE with and without confidence-based scaling. We observe that scaling the consistency scores to weight low-confidence predictions more results in better performance. This supports our design choice to focus on low-confidence predictions.

Performance on mAP₅₀ and mAP₇₅. So far, our experiments have focused on estimating mAP, which averages performance across 10 IoU thresholds ranging from 0.50 to 0.95. To assess whether our proposed PCR is also eligible

Meta-dataset	Corruption-based (Ours)				Avg. RMSE
	RetinaNet		Faster R-CNN		
Detector	ResNet-50	Swin-T	ResNet-50	Swin-T	
BoS [45]	13.50	5.18	10.32	6.85	8.96
PCR (Ours)	6.43	4.80	6.90	4.51	5.66
PCR + BoS	6.29	4.25	5.50	4.12	5.04

Table 7. Performance of the combined BoS [45] and PCR in vehicle detection with corruption-based meta-dataset (Ours). A detector trained on the COCO training set is used, and the average RMSE (%) across 10 datasets is reported.

to estimate other metrics, we conduct experiments on estimating mAP₅₀ and mAP₇₅, which correspond to fixed IoU thresholds of 0.50 and 0.75, respectively, reflecting specific preferences for localization precision. As shown in Table 6, PCR achieves the best estimate for both mAP₅₀ and mAP₇₅, consistently outperforming all baseline methods.

Combination of PCR and BoS. As discussed in Section 3.1, BoS [45] and our proposed PCR capture different aspects of consistency. As combining complementary components often leads to improved performance, evidenced by PCR incorporating consistency and reliability, we evaluate whether combining BoS and PCR yields further improvements. Table 7 compares BoS, PCR, and their combination, where we perform linear regression using scores from both methods for the combination. We observe that

the combination outperforms individual methods, suggesting that BoS and PCR capture distinct yet complementary aspects of consistency. This is further supported Figure 7, which show that the consistency scores defined in BoS and our work exhibit low correlation.




Figure 7. Consistency scores of PCR and BoS [45].

6. Conclusion

In this paper, we develop an AutoEval framework for object detection. We propose Prediction Consistency and Reliability (PCR) as an AutoEval method for object detection, which leverages the relationships between pre- and post-NMS boxes to estimate detection performance without ground-truth labels. PCR captures both localization and classification aspects through consistency and reliability measures conditioned on confidence scores. To enable realistic and scalable evaluation, we construct a meta-dataset using image corruptions of varying severity. We hope that our proposed method and meta-dataset provide a solid foundation for future research on AutoEval for object detection.

Acknowledgements

This work was partially supported by the National Research Foundation of Korea (NRF) grant funded by the Ministry of Science and ICT (MSIT) of the Korean government (RS2024-00341749), and Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by MSIT (RS-2022-II220124, RS-2023-00259934, RS-2025-02283048).

References

- [1] Christina Baek, Yiding Jiang, Aditi Raghunathan, and J. Zico Kolter. Agreement-on-the-line: Predicting the performance of neural networks under distribution shift. In *NeurIPS*, 2022. 2
- [2] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S Davis. Soft-nms—improving object detection with one line of code. In *Proceedings of the IEEE international conference on computer vision*, pages 5561–5569, 2017. 3
- [3] Markus Braun, Sebastian Krebs, Fabian Flohr, and Dariu M Gavrila. Eurocity persons: A novel benchmark for person detection in traffic scenes. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1844–1861, 2019. 6
- [4] 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 preprint arXiv:2303.01803*, 2019. 2
- [5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Scharwächter, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset. In *CVPR Workshop on the Future of Datasets in Vision*, page 1, 2015. 6
- [6] Weijian Deng and Liang Zheng. Are labels always necessary for classifier accuracy evaluation? In *CVPR*, 2021. 1, 2
- [7] Weijian Deng, Stephen Gould, and Liang Zheng. What does rotation prediction tell us about classifier accuracy under varying testing environments? In *ICML*, 2021. 2
- [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, 2010. 2
- [9] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016. 3
- [10] Saurabh Garg, Sivaraman Balakrishnan, Zachary C. Lipton, Behnam Neyshabur, and Hanie Sedghi. Leveraging unlabeled data to predict out-of-distribution performance. In *ICLR*, 2022. 1, 2, 6, 7, 8, 4, 5, 9
- [11] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 6
- [12] Ross Girshick. Fast r-cnn. In *ICCV*, 2015. 2
- [13] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *ICCV*, 2014. 2
- [14] Gregory Griffin, Alex Holub, Pietro Perona, et al. Caltech-256 object category dataset. Technical report, Technical Report 7694, California Institute of Technology Pasadena, 2007. 6
- [15] Devin Guillory, Vaishaal Shankar, Sayna Ebrahimi, Trevor Darrell, and Ludwig Schmidt. Predicting with confidence on unseen distributions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1134–1144, 2021. 1, 2, 6, 7, 8, 4, 5, 9
- [16] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *ICML*, 2017. 3
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 6, 7
- [18] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2, 3
- [19] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019. 2, 6, 3, 4
- [20] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017. 2, 6, 7, 8, 4, 5, 9
- [21] Jisoo Jeong, Seungeui Lee, Jeesoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. *Advances in neural information processing systems*, 32, 2019. 4
- [22] Yiding Jiang, Vaishnavh Nagarajan, Christina Baek, and J. Zico Kolter. Assessing generalization of sgd via disagreement. In *ICLR*, 2022. 2
- [23] Kaggle. Self-driving cars, 2020. <https://www.kaggle.com/datasets/alincijov/self-driving-cars>, 2020. 5, 13, 14, 17. 6
- [24] Kaggle. Traffic vehicles object detection, 2020. <https://www.kaggle.com/datasets/saumyapatel/traffic-vehicles-object-detection>, 2020. 5, 14. 6
- [25] Kaggle. Udacity self-driving car dataset, 2021. <https://www.kaggle.com/datasets/sshikamaru/udacity-self-driving-car-dataset>, 2021. 5, 14. 6
- [26] Kaggle. Car-person custom object detection v2 roboflow, 2022. <https://www.kaggle.com/datasets/owaiskhan9654/car-person-v2-roboflow>, 2022. 5, 14. 6
- [27] Fabian Kuppens, Jan Kronenberger, Amirhossein Shantia, and Anselm Haselhoff. Multivariate confidence calibration for object detection. In *CVPRW*, 2020. 3
- [28] Taekyung Lee, Sorn Chottanaturak, Taesik Gong, and Sung-Ju Lee. Aetta: Label-free accuracy estimation for test-time adaptation. In *CVPR*, 2024. 2

- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context. In *ECCV*, 2014. [2](#)
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. [6](#)
- [31] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. [2](#), [3](#), [6](#), [7](#)
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. [2](#), [6](#)
- [33] Yuen Peng Loh and Chee Seng Chan. Getting to know low-light images with the exclusively dark dataset. *Computer Vision and Image Understanding*, 178:30–42, 2019. [6](#)
- [34] Bimsara Pathiraja, Malitha Gunawardhana, and Muhammad Haris Khan. Multiclass confidence and localization calibration for object detection. In *CVPR*, 2023. [3](#)
- [35] Ru Peng, Qiuyang Duan, Haobo Wang, Jiachen Ma, Yanbo Jiang, Yongjun Tu, Xiu Jiang, and Junbo Zhao. Came: Contrastive automated model evaluation. In *ICCV*, 2023. [2](#)
- [36] Ru Peng, Heming Zou, Haobo Wang, Yawen Zeng, Zenan Huang, and Junbo Zhao. Energy-based automated model evaluation. In *ICLR*, 2024. [2](#)
- [37] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *ICCV*, 2016. [2](#)
- [38] 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*, pages 779–788, 2016. [2](#), [3](#)
- [39] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. [2](#), [3](#), [6](#)
- [40] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8050–8058, 2019. [1](#), [2](#), [6](#), [7](#), [8](#), [4](#), [5](#), [9](#)
- [41] Shuai Shao, Zijian Zhao, Boxun Li, Tete Xiao, Gang Yu, Xiangyu Zhang, and Jian Sun. Crowdhuman: A benchmark for detecting human in a crowd. *arXiv preprint arXiv:1805.00123*, 2018. [6](#), [7](#)
- [42] Huixin Sun, Baochang Zhang, Yanjing Li, and Xianbin Cao. Confidence-driven bounding box localization for small object detection. *arXiv preprint arXiv:2303.01803*, 2023. [3](#)
- [43] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, 2019. [2](#), [3](#)
- [44] Longyin Wen, Dawei Du, Zhaowei Cai, Zhen Lei, Ming-Ching Chang, Honggang Qi, Jongwoo Lim, Ming-Hsuan Yang, and Siwei Lyu. Ua-detrac: A new benchmark and protocol for multi-object detection and tracking. *Computer Vision and Image Understanding*, 193:102907, 2020. [6](#)
- [45] Yang Yang, Wenhai Wang, Zhe Chen, Jifeng Dai, and Liang Zheng. Bounding box stability against feature dropout reflects detector generalization across environments. In *ICLR*, 2024. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#)
- [46] 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*, pages 2636–2645, 2020. [6](#)
- [47] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3221, 2017. [6](#)
- [48] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In *Proceedings of the AAAI conference on artificial intelligence*, pages 12993–13000, 2020. [4](#)