Extended Abstract of:

Gradient-Based Quantification of Epistemic Uncertainty for Deep Object Detectors

[Submission to DataNinja Spring School 2022 Poster Session]

Tobias Riedlinger   Matthias Rottmann   Marius Schubert   Hanno Gottschalk

School of Mathematics and Natural Sciences
University of Wuppertal
{riedlinger, rottmann, mschubert, hgottsch}@uni-wuppertal.de

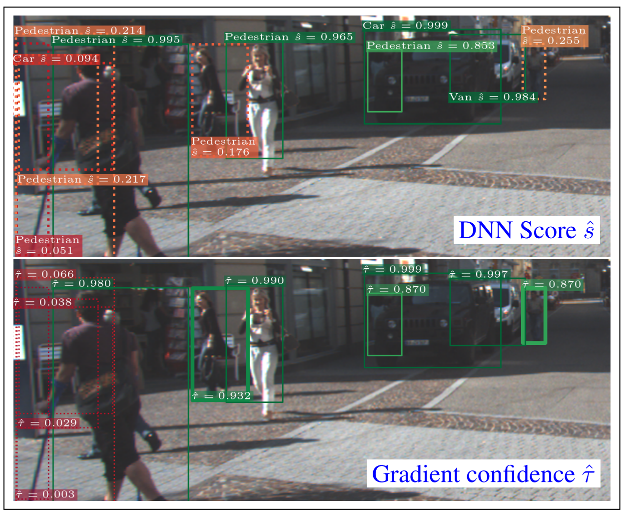

Figure 1. Object detection in a street scene. Top coloration: Score $\hat{s}$; bottom coloration: instance-wise gradient-based confidence $\hat{\tau}$ obtained by our method. Dashed boxes here indicate the discarding at any confidence threshold in $[0.3, 0.85]$. The top image contains FNs which are not separable from correctly discarded boxes based on the score (lower threshold would lead to FPs). In the bottom image, those $\hat{s}$-FNs are assigned higher confidences and there is a large range of thresholds with no FPs.

The vast majority of uncertainty quantification methods for deep object detectors such as variational inference are based on the network output. As such, they constitute straight-forward adaptations of methods originally developed for the classification task, were the prediction consists of only one probability distribution per input. Other methods which do not exclusively depend on the predicted probability distribution cannot be adapted to the object detection setting directly due to architectural restrictions and the need for instance-wise uncertainty quantification. Here [1], we study gradient-based epistemic uncertainty metrics for deep object detectors to obtain reliable confidence estimates. Gradient uncertainty is based on the loss gradient obtained by the network's own prediction as ground truth label. Deep object detectors generate a prediction

$$\boldsymbol{y} = (y^1, \dots, y^N) \tag{1}$$

consisting of a variable number of instances per input from one set of parameters which is shared among instances. The network's self-learning gradient, therefore, gives at best an uncertainty estimate for the entire prediction $\boldsymbol{y}$, but not for one individual instance y^j which is desired in many applications of uncertainty quantification. In this work, we introduce an instance-based generalization of the self-learning gradient and study its properties in comparison with various other uncertainty estimation methods. We show that they contain predictive information and that they capture information orthogonal to that of common, output-based epistemic uncertainty estimation / approximation methods like Monte-Carlo dropout and deep ensembles. To this end, we use the tools (cf. Fig. 2) of meta classification (estimation of true positive confidence) and meta regression (estimation of localization accuracy) to produce advanced confidence estimates using gradient metrics and other quantities obtained by uncertainty quantification methods which are in principle applicable to any object detection architecture. Specifically, we employ false positive detection and prediction of localization quality to investigate and determine the uncertainty content of our metrics. We investigate the calibration properties of our obtained meta classifiers and show that they generally have significantly smaller calibration errors than the objectness score of the associated object detector. Moreover, we use meta classifiers as a well-calibrated post-processing filter mechanism to the ob-

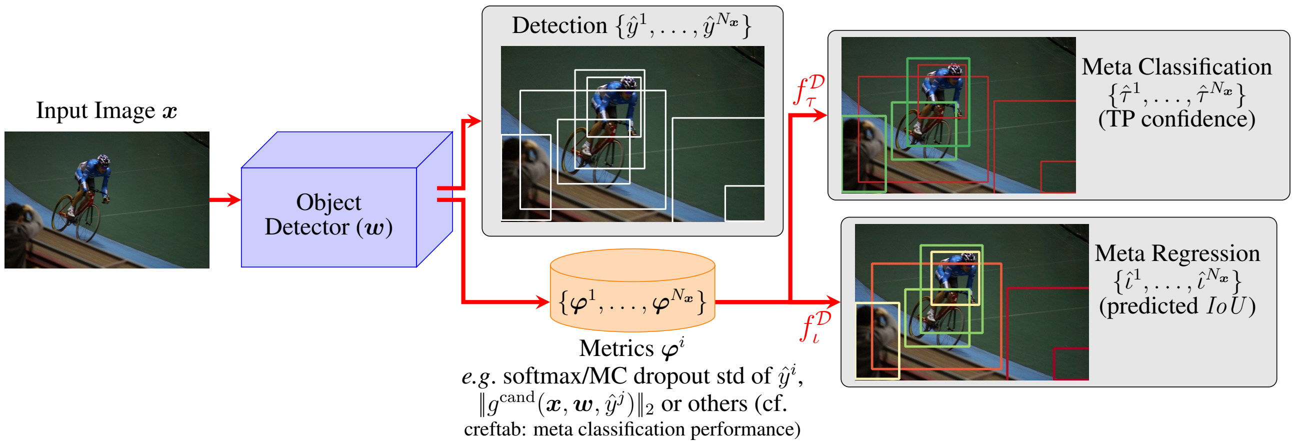

Figure 2. Meta classification and meta regression pipeline: An uncertainty feature vector $\boldsymbol{\varphi}^j$ is assigned to each detected box $\hat{y}^j$. During training, we fit $f_\tau^{\mathcal{D}}$ and $f_\iota^{\mathcal{D}}$ to map $\boldsymbol{\varphi}^j$ to τ^j (TP/FP) and max. IoU ι^j of $\hat{y}^j$, resp. At inference, $f_\tau^{\mathcal{D}}$ and $f_\iota^{\mathcal{D}}$ yield confidence and IoU estimates $\hat{\tau}^j$ and $\hat{\iota}^j$ for $\hat{y}^j$ based on $\boldsymbol{\varphi}^j$.

ject detection pipeline and compare object detection performance. Our results show that gradient-based uncertainty is itself on par with output-based methods across different detectors and datasets in terms of meta classification and meta regression performance. More significantly, combined meta classifiers based on gradient and output-based metrics outperform the standalone models. Based on this result, we conclude that gradient uncertainty adds orthogonal information to output-based epistemic uncertainty approximation methods. This suggests that variational inference such as MC dropout or deep ensemble sampling may be supplemented by gradient-based uncertainty to obtain improved confidence measures, contributing to down-stream applications of deep object detectors and improving probabilistic reliability in instance-based detection tasks.

References

[1] Tobias Riedlinger, Matthias Rottmann, Marius Schubert, and Hanno Gottschalk. Gradient-based quantification of epistemic uncertainty for deep object detectors. *arXiv preprint arXiv:2107.04517*, 2021. 1