Gradient-Based Quantification of Epistemic Uncertainty for Deep Object Detectors

Tobias Riedlinger*, Matthias Rottmann*, Marius Schubert† and Hanno Gottschalk*

University of Wuppertal, IZMD, School of Mathematics and Natural Sciences, Stochastics* and Applied Computer Science† Group

Abstract

We study *gradient-based* epistemic uncertainty metrics [1] for deep object detectors to obtain reliable confidence estimates. We show that they contain predictive information and that they *capture information orthogonal* to that of common, output-based uncertainty estimation methods. Moreover, we use meta classifiers as a *post-processing filter mechanism* to the object detection pipeline and compare object detection performance. Our results show that gradient-based uncertainty is itself on par with output-based methods across different datasets and that combined meta classifiers based on gradient and output-based metrics outperform the standalone models. This suggests that variational inference may be supplemented by gradient-based uncertainty to obtain improved confidence measures, contributing to down-stream applications of deep object detectors by *improving their probabilistic reliability.*

Gradient Uncertainty for Object Detection

Given an input image $\boldsymbol{x} \in \mathcal{X}$ and model weights $\boldsymbol{w}$, we regard the box prediction of $\hat{N}_{\boldsymbol{x}}$ instances

$$\hat{y}(\boldsymbol{x}, \boldsymbol{w}) = (\hat{y}^1, \ldots, \hat{y}^{N_{\boldsymbol{x}}}), \quad \hat{y}^j = (\hat{\xi}^j, \hat{s}^j, \hat{p}^j) \in \mathbb{R}^{4+1+C}, \qquad (1)$$

consisting of a bounding box $\hat{\xi}^j \in \mathbb{R}^4$, an objectness score $\hat{s}^j \in (0, 1)$ and a vector of class probabilities $\hat{p}^j \in (0, 1)^C$. We compute the candidate-restricted self-learning gradient

$$g^{\mathrm{cand}}(\boldsymbol{x}, \boldsymbol{w}, \hat{y}^j) := \partial_{\boldsymbol{w}} \mathscr{L}\left(\mathrm{cand}[\hat{y}^j](\boldsymbol{x}, \boldsymbol{w}), \overline{y}^j\right) \qquad (2)$$

of $\hat{y}^j$ and generate scalars by measuring the magnitude of (2). The resulting uncertainty metrics are used as metrics $\boldsymbol{\varphi}$ for meta classification and meta regression.

Theorem (Computational Complexity). *The number of FLOP required to compute the last layer ($t = T$) gradient is $\mathcal{O}(k_T hw + k_T k_{T-1}(2s_T + 1)^4)$. Similarly, for earlier layers t, we have $\mathcal{O}(k_{t+1}k_t + k_t k_{t-1})$, provided that we have previously computed the gradient for the consecutive layer $t + 1$. Performing variational inference only on the last layer requires $\mathcal{O}(k_T k_{T-1} hw)$ FLOP per sample.*

References

[1] Tobias Riedlinger, Matthias Rottmann, Marius Schubert and Hanno Gottschalk *Gradient-Based Quantification of Epistemic Uncertainty for Deep Object Detectors.* preprint arXiv:2107.04517 (2021)

Gradient-based Confidence Assignment

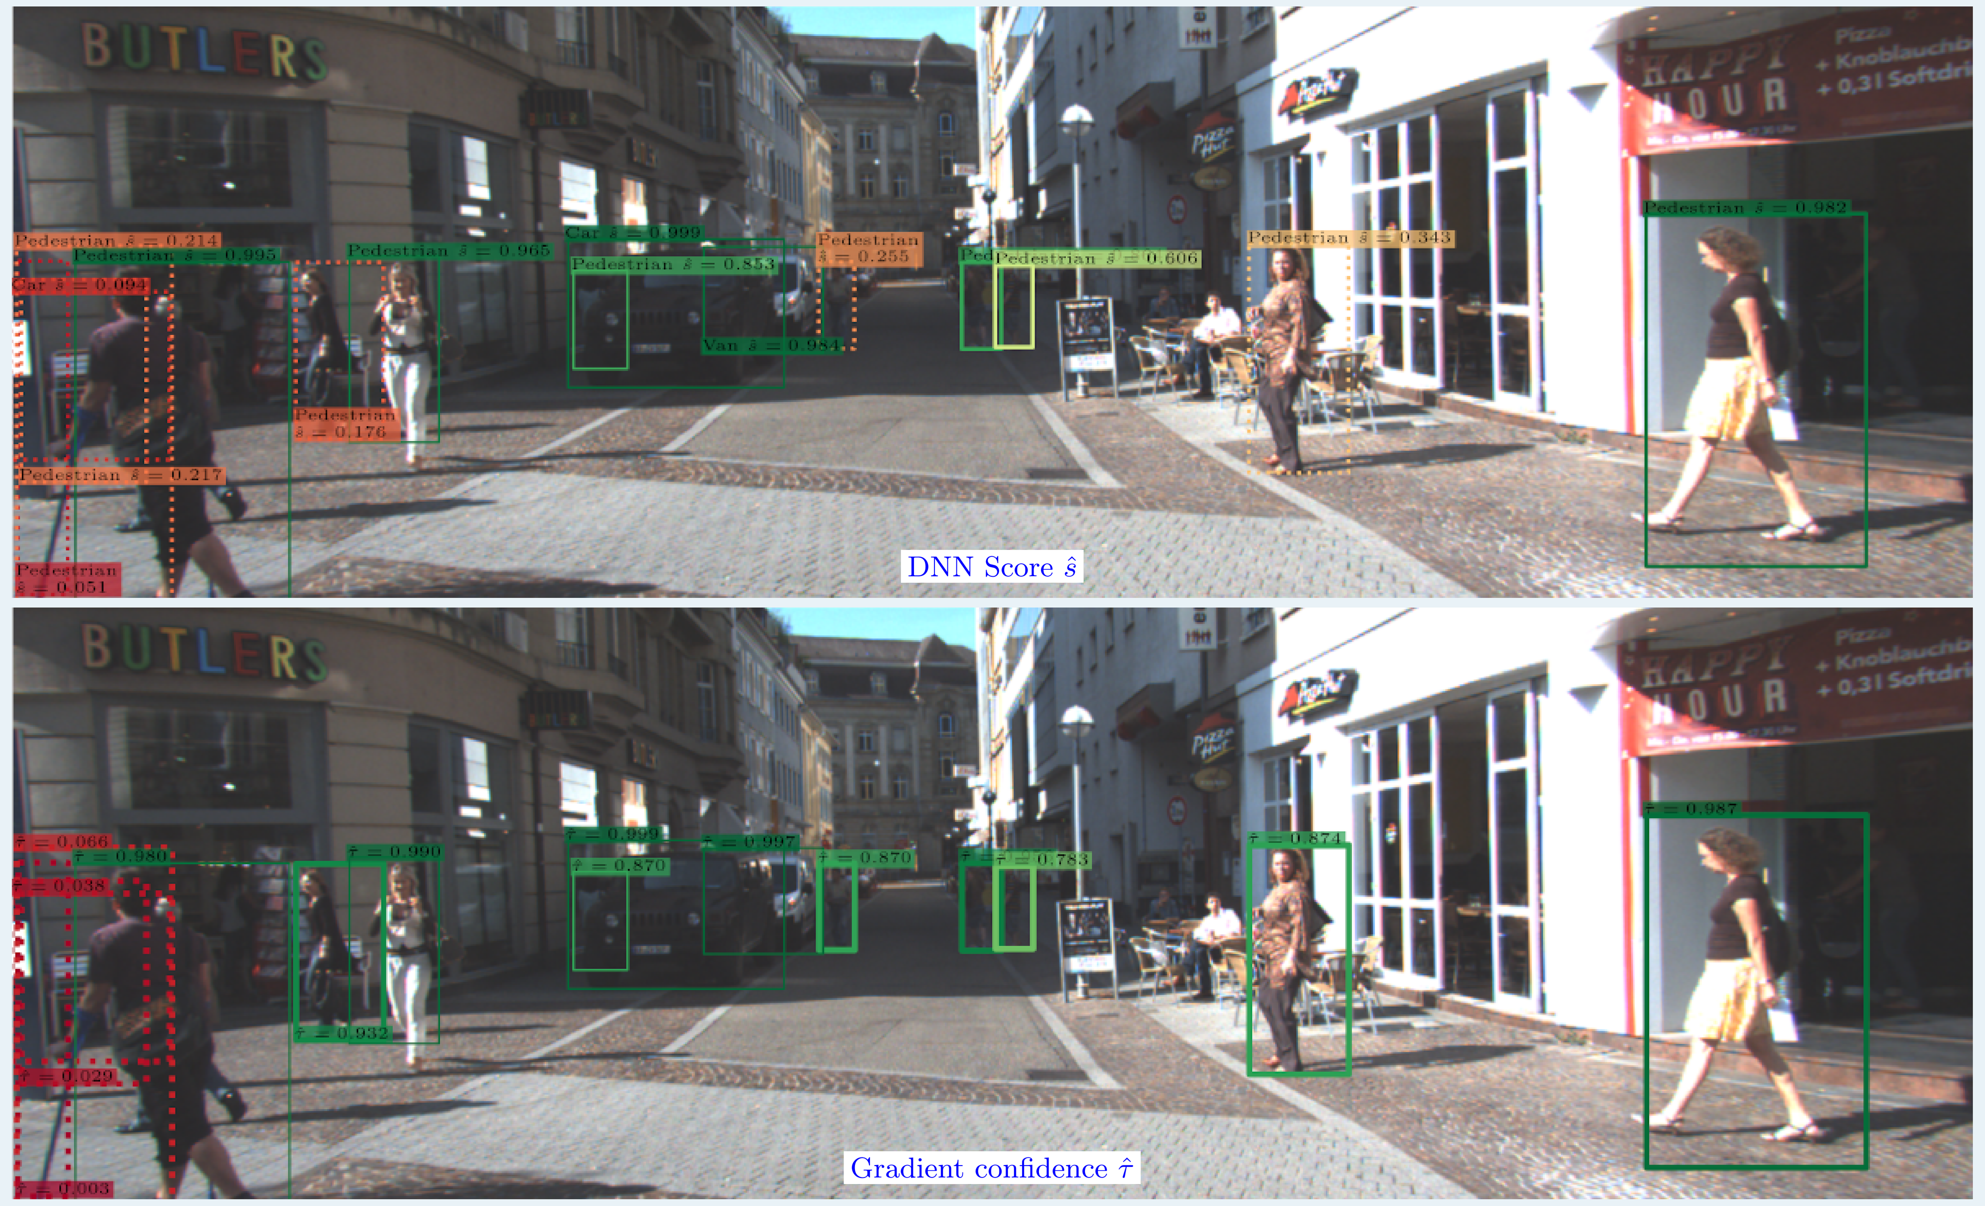

Meta Classification and Meta Regression

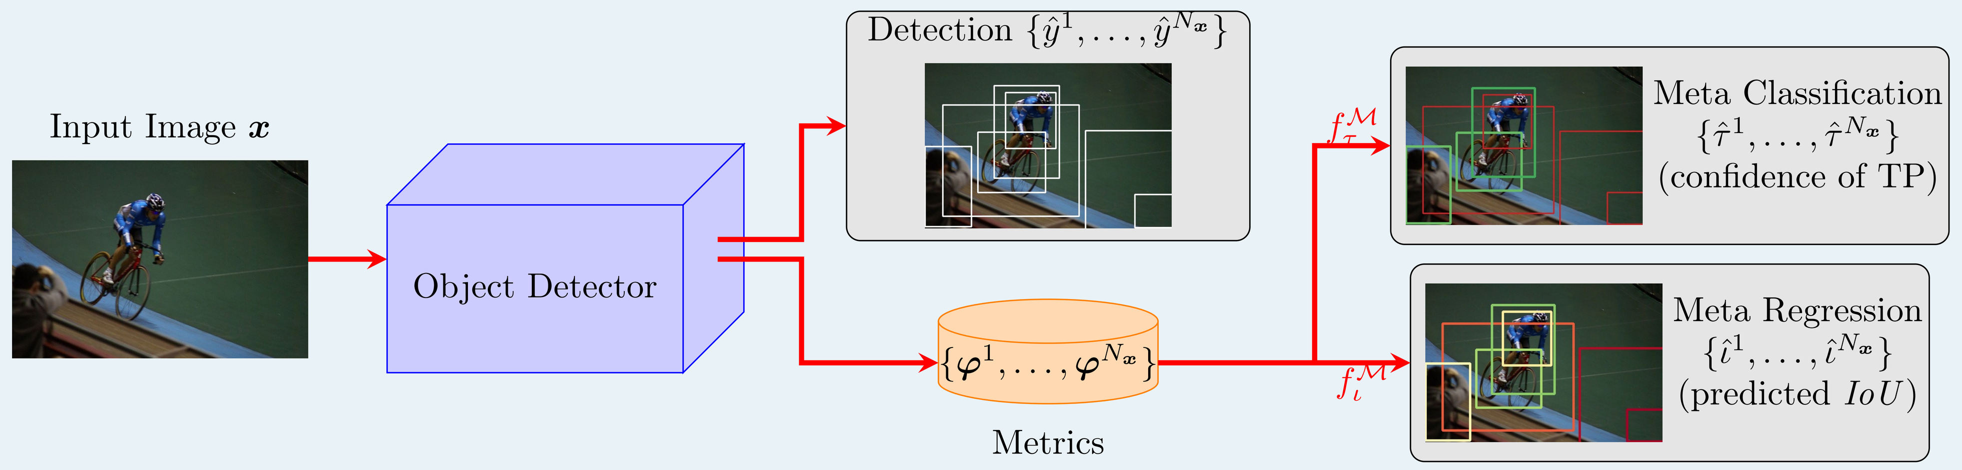

YOLOv3	Pascal VOC		COCO		KITTI	
Meta Classification	*AuROC*	*AP*	*AuROC*	*AP*	*AuROC*	*AP*
Score	90.68 ± 0.06	69.56 ± 0.12	82.97 ± 0.04	62.31 ± 0.05	96.53 ± 0.05	96.87 ± 0.03
Entropy	91.30 ± 0.02	61.94 ± 0.06	76.52 ± 0.02	42.52 ± 0.04	94.79 ± 0.06	94.83 ± 0.05
Energy Score	92.59 ± 0.02	64.65 ± 0.06	75.39 ± 0.02	39.72 ± 0.06	95.66 ± 0.02	95.33 ± 0.03
Full Softmax	93.81 ± 0.06	72.08 ± 0.15	82.91 ± 0.06	58.65 ± 0.10	97.07 ± 0.03	96.85 ± 0.03
MC Dropout (MC)	$\underline{96.72 \pm 0.02}$	78.15 ± 0.09	$\mathbf{89.04 \pm 0.02}$	64.94 ± 0.11	97.60 ± 0.07	97.17 ± 0.10
Ensemble (E)	$\mathbf{96.87 \pm 0.02}$	77.86 ± 0.11	$\underline{88.97 \pm 0.02}$	64.05 ± 0.12	97.63 ± 0.04	97.63 ± 0.05
MetaDetect (MD)	95.78 ± 0.05	$\mathbf{78.64 \pm 0.08}$	87.16 ± 0.04	$\underline{69.41 \pm 0.07}$	$\mathbf{98.23 \pm 0.02}$	$\mathbf{98.06 \pm 0.02}$
Grad. Score$_{\lVert\cdot\rVert_2}$ (GS$_{\lVert\cdot\rVert_2}$; ours)	94.76 ± 0.03	74.86 ± 0.10	86.05 ± 0.04	64.25 ± 0.06	97.31 ± 0.05	96.86 ± 0.10
Grad. Score$_{\mathrm{full}}$ (GS$_{\mathrm{full}}$; ours)	95.80 ± 0.04	$\underline{78.57 \pm 0.11}$	88.07 ± 0.03	$\mathbf{69.62 \pm 0.07}$	$\underline{98.04 \pm 0.03}$	$\underline{97.81 \pm 0.06}$
MC+E+MD (ours)	97.66 ± 0.02	85.13 ± 0.12	91.14 ± 0.02	73.82 ± 0.05	98.56 ± 0.03	98.45 ± 0.03
GS$_{\mathrm{full}}$+MC+E+MD (ours)	$\mathbf{97.95 \pm 0.02}$	$\mathbf{86.69 \pm 0.09}$	$\mathbf{91.65 \pm 0.03}$	$\mathbf{74.88 \pm 0.07}$	$\mathbf{98.74 \pm 0.02}$	$\mathbf{98.62 \pm 0.01}$

YOLOv3	Pascal VOC	COCO	KITTI
Meta Regression (R^2)			
Score	48.29 ± 0.04	32.60 ± 0.02	78.86 ± 0.05
Entropy	43.24 ± 0.03	21.10 ± 0.04	69.33 ± 0.04
Energy Score	47.18 ± 0.03	17.94 ± 0.02	71.53 ± 0.10
Full Softmax	53.86 ± 0.11	36.95 ± 0.13	78.92 ± 0.11
MC	$\underline{61.63 \pm 0.15}$	43.85 ± 0.09	82.10 ± 0.11
E	61.48 ± 0.07	43.53 ± 0.13	84.18 ± 0.12
MD	60.36 ± 0.14	$\underline{44.22 \pm 0.11}$	$\mathbf{85.88 \pm 0.10}$
GS$_{\lVert\cdot\rVert_2}$ (ours)	58.05 ± 0.13	38.77 ± 0.04	81.21 ± 0.05
GS$_{\mathrm{full}}$ (ours)	$\mathbf{62.50 \pm 0.11}$	$\mathbf{44.90 \pm 0.09}$	$\underline{85.40 \pm 0.11}$
MC+E+MD	69.38 ± 0.11	54.07 ± 0.08	87.78 ± 0.11
GS$_{\mathrm{full}}$+MC+E+MD	$\mathbf{72.26 \pm 0.08}$	$\mathbf{56.14 \pm 0.11}$	$\mathbf{88.80 \pm 0.07}$

Meta Classifier Calibration

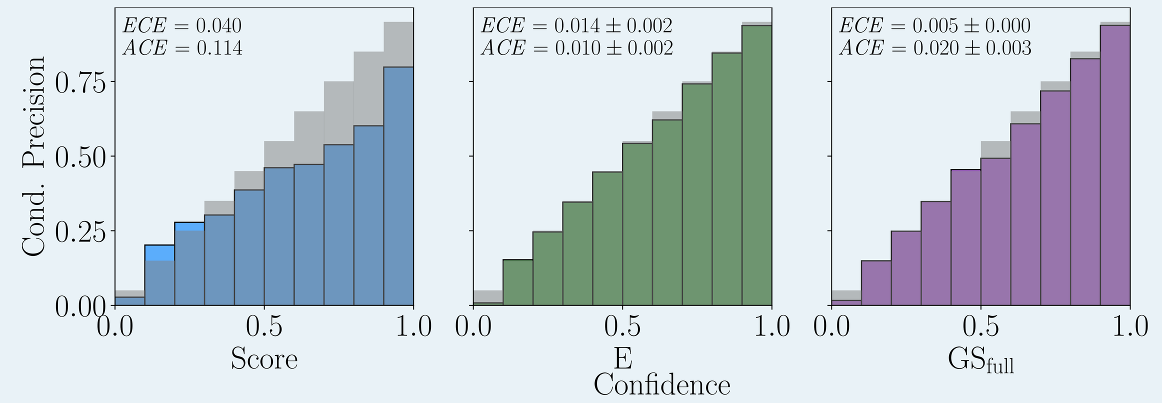

Object Detection Performance

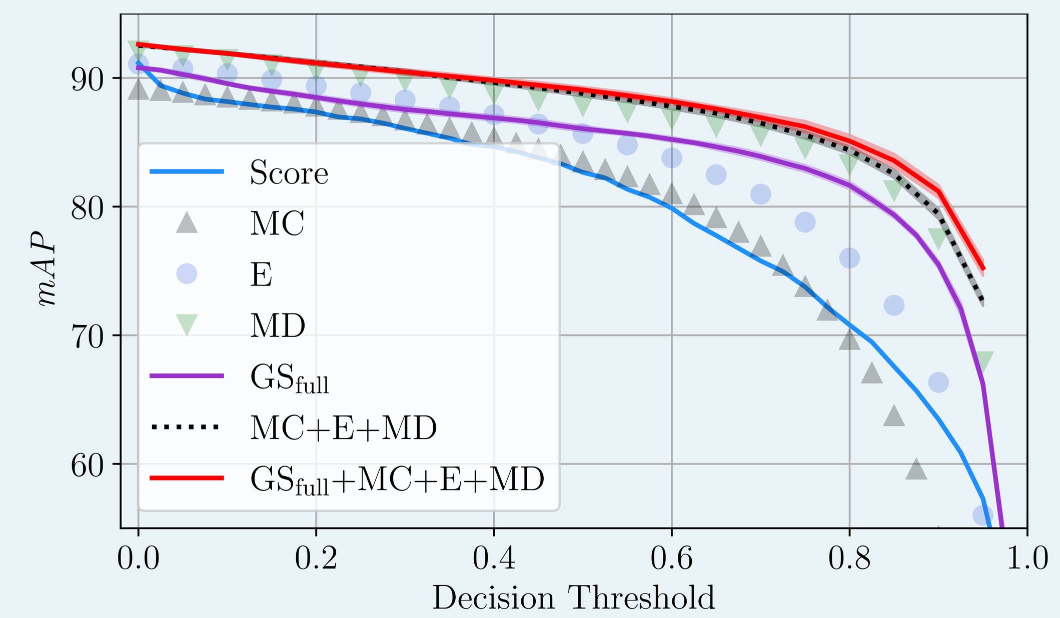

Decision Rule Tradeoff

Meta classifiers restricted to the "Pedestrian" class in KITTI:

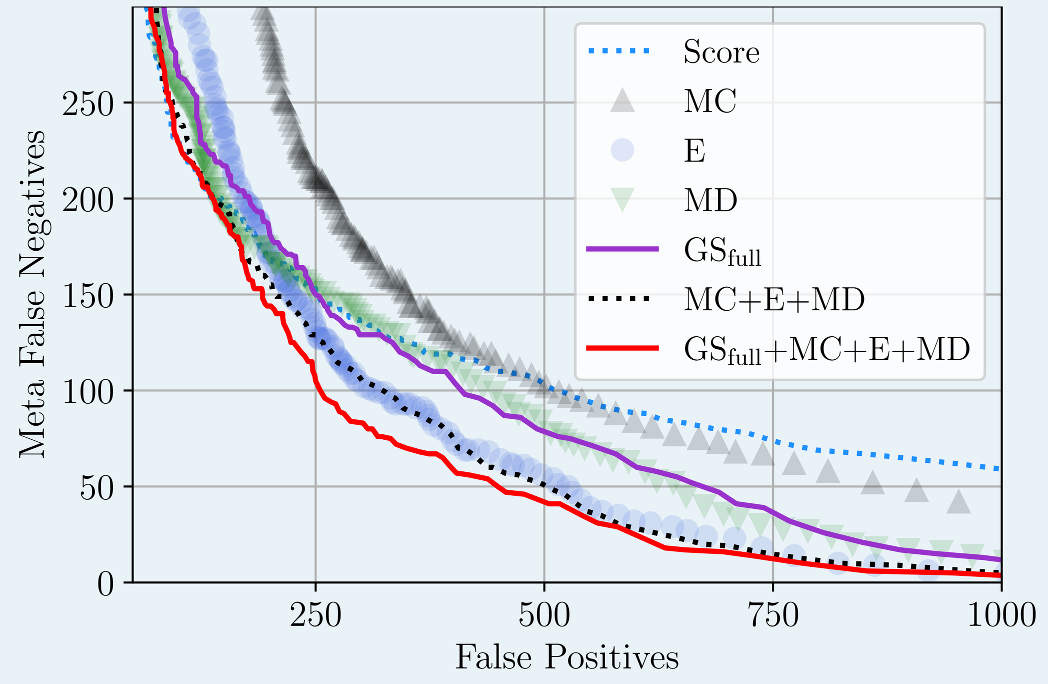

Compared with the Score baseline, we find a *reduction of upwards of 100 FPs at 100 FNs* when basing the decisions exclusively on gradient uncertainty. The combined model achieves a reduction of 250 FPs.

Acknowledgments. This work is funded by the German Federal Ministry for Economic Affairs and Energy.

Contact. `{riedlinger,rottmann,mschubert,hgottsch}@uni-wuppertal.de`

KI FAMILIE

KI ABSICHERUNG Safe AI for Automated Driving

BIT Bergische Innovationsplattform für Künstliche Intelligenz

IZMD Interdisziplinäres Zentrum Machine Learning and Data Analytics

BERGISCHE UNIVERSITÄT WUPPERTAL