Cost-Sensitive Uncertainty-Based Failure Recognition for Object Detection

Paper #176

Moussa Kassem Sbeyti[1,2], Michelle Karg[1], Christian Wirth[1], Nadja Klein[3], and Sahin Albayrak[2]

[1]Continental AG, Germany [2]DAI-Labor, Technische Universität Berlin, Germany
[3]Technische Universität Dortmund, Germany

Background and Motivation

Challenges and objectives for utilizing uncertainty in failure recognition:

I. **Uncertainty Overlap:** Overlapping uncertainty of correct detections (CDs) and false detections (FDs).
⇒ Address the cost implications of the overlap when thresholding.

II. **Manual Thresholding:** Subjectivity and lack of generalizability of manual thresholds, especially when simultaneously considering multiple uncertainties.
⇒ Develop a cost-sensitive, automated, adaptive thresholding method.

III. **Uncertainty Combination:** Difficulty in effectively combining multiple uncertainties with different ranges and contribution to failure recognition.
⇒ Define an optimal combination strategy for uncertainties with a range ∈ $[0, \infty)$.

Conclusion

- ✓ **Model-Agnostic Failure Recognition:** Addresses the trade-off safety vs. performance in object detection via a post-processing pipeline.
- ✓ **Application-Agnostic Budget-Based Thresholding:** Budget on removed FDs or maintained CDs.
- ✓ **Performance Enhancement Post-Thresholding.**
- ✓ **Efficient Uncertainty Estimation:** No added inference time; uses only loss attenuation.
- ✓ **Minimal Model Expansion:** Only 0.07% increase in parameters due to extending the localization head.
- ✓ **Transparent Evaluation:** Utilizes specifically defined requirements and metrics.

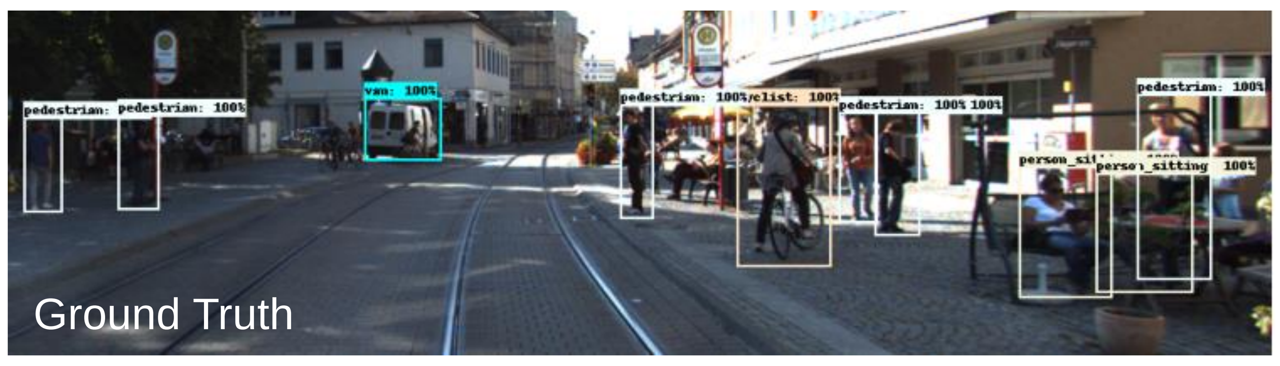

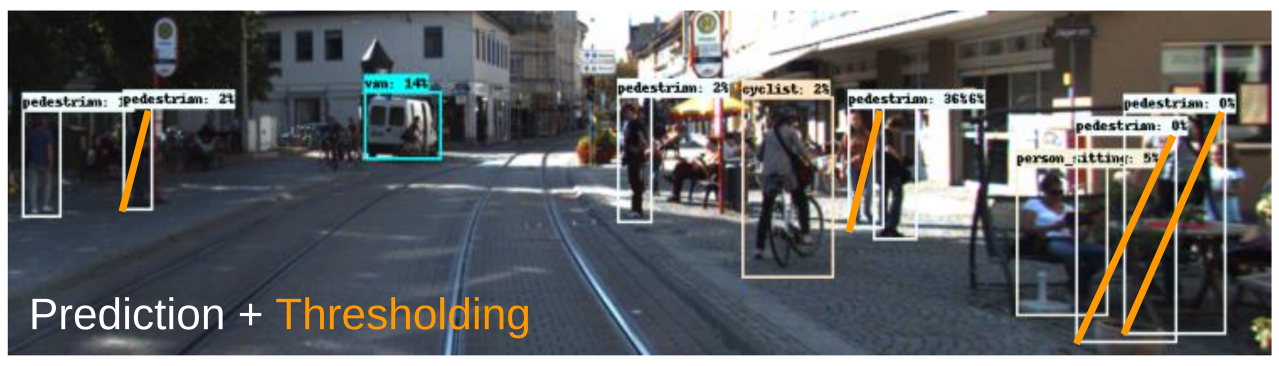

Pipeline

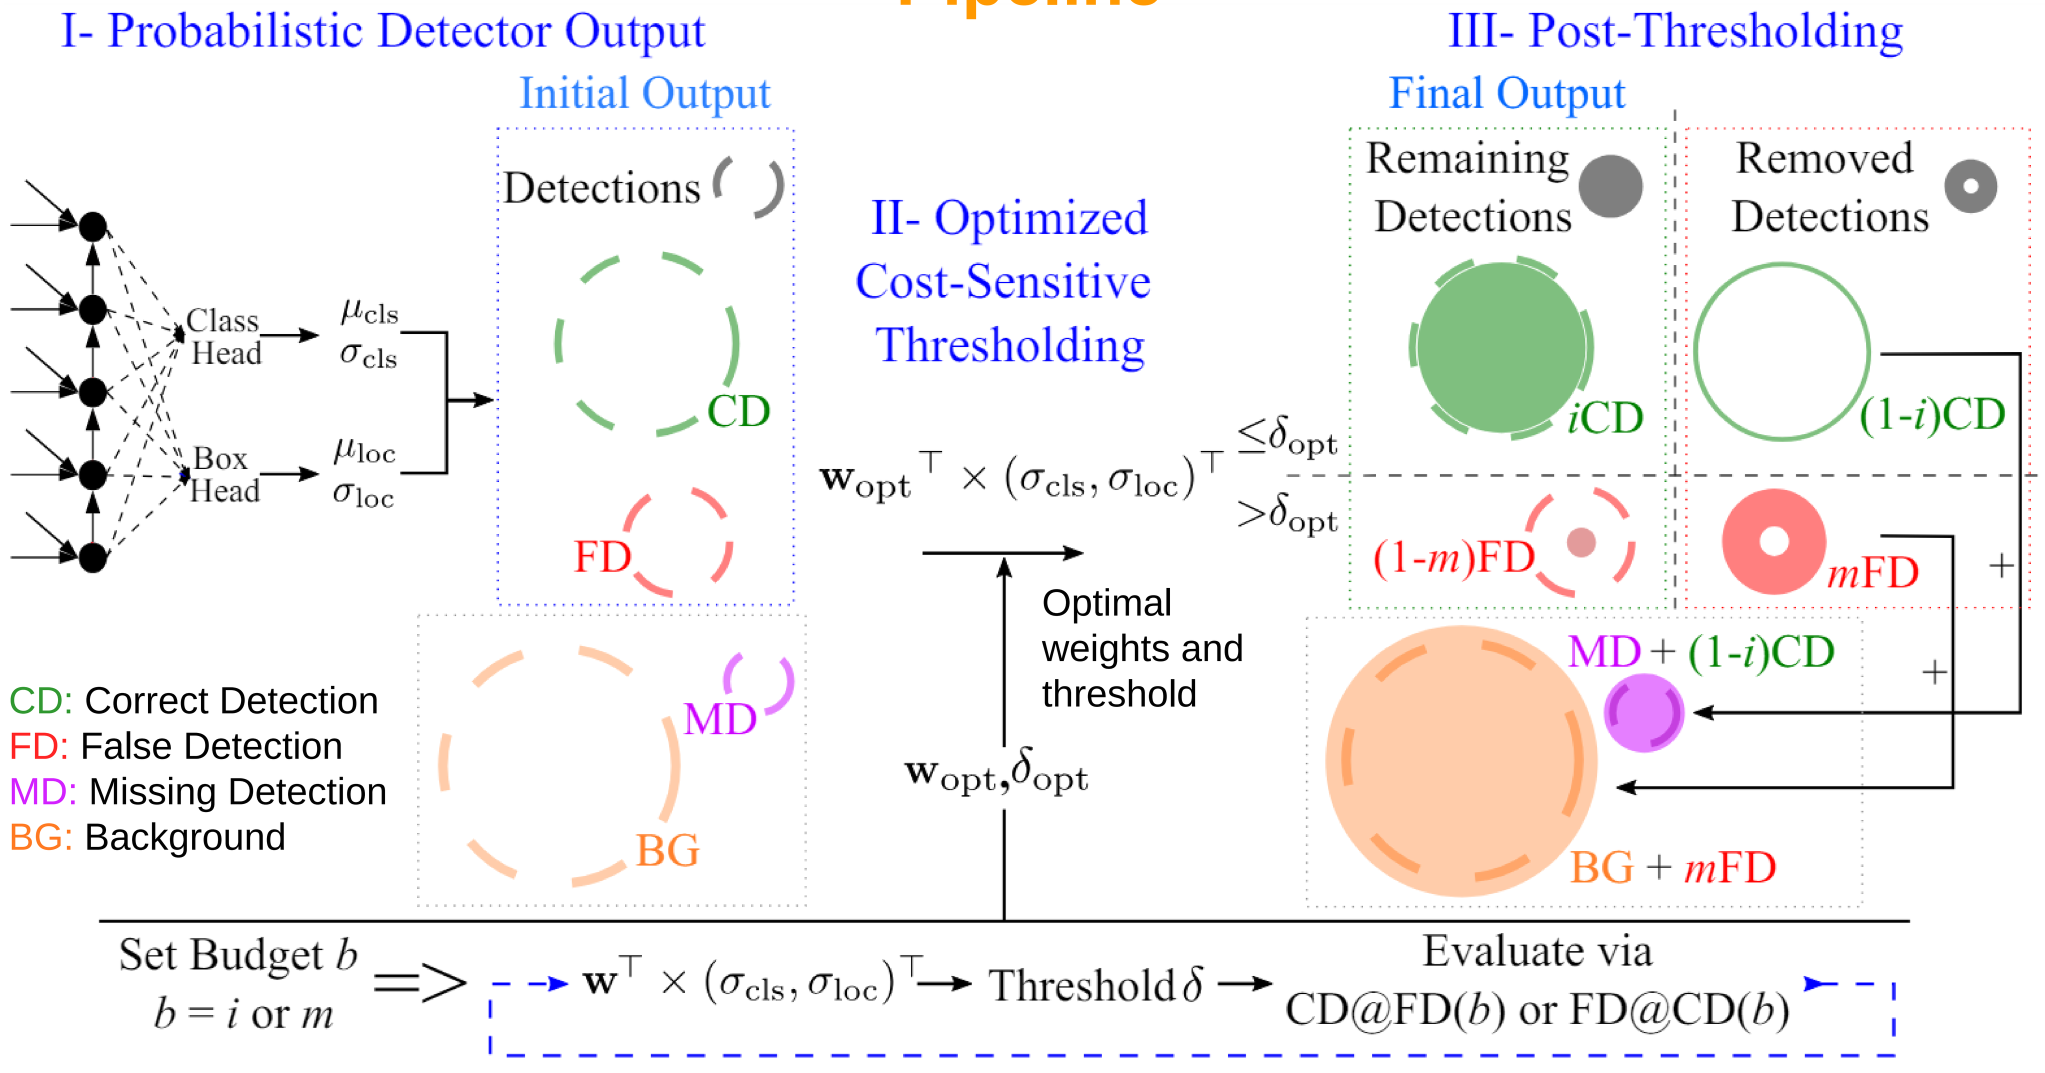

Our proposed metrics to measure thresholding performance

$$\mathrm{CD@FD}(b) = \sum_{\tau=0.5}^{0.75} \mathrm{TNR}(\mathrm{TPR}(\delta(b,\tau))) \quad \mathrm{FD@CD}(b) = \sum_{\tau=0.5}^{0.75} \mathrm{TPR}(\mathrm{FPR}(\delta(b,\tau)))$$

Effect on Detector

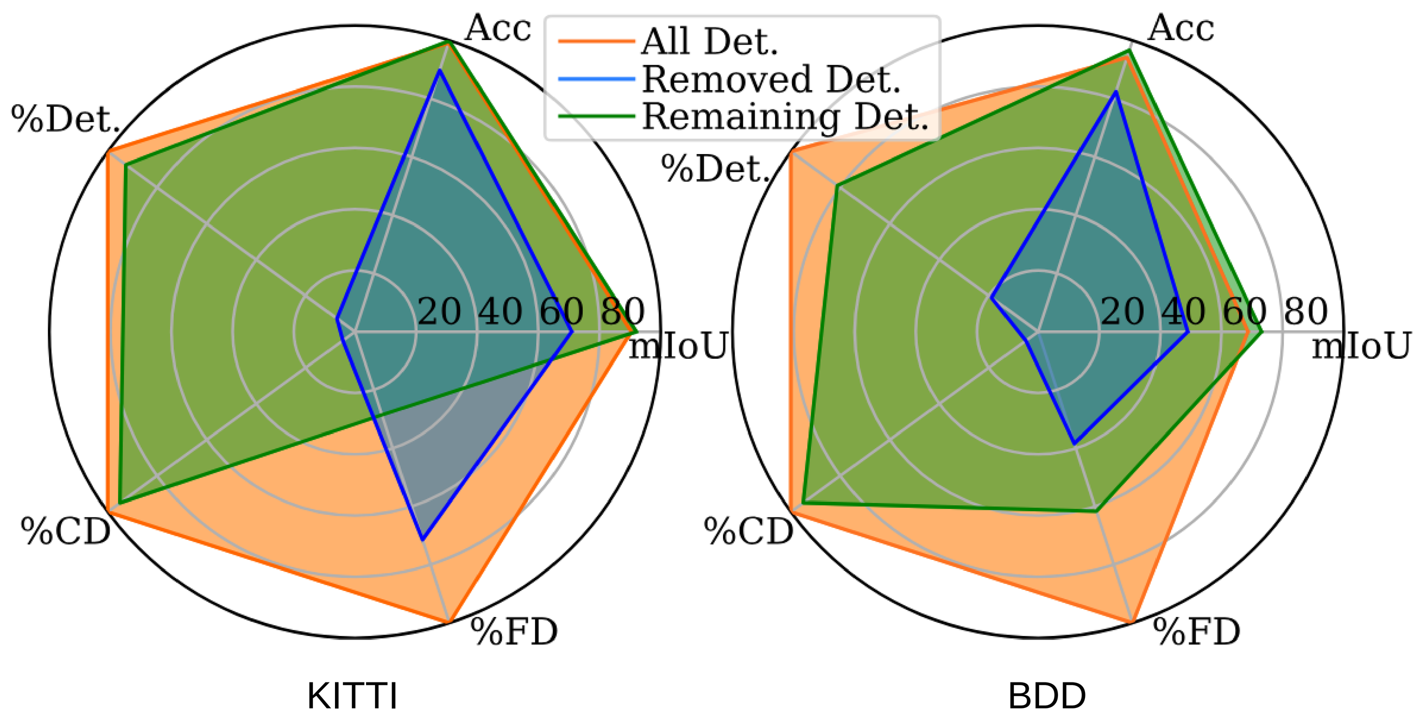

Impact of thresholding with a budget of maintaining ≥ 95% of CDs:

- **Considerable Decrease in FDs:** Targets detections with lower accuracy (Acc) and mean IoU (mIoU), enhancing performance.
- **Different Impact Across Datasets:** Notable reduction in the percentage of detections (% Det.) for BDD vs. KITTI due to the higher ratio of FDs to CDs, increasing the impact of thresholding.

Method and Application Insights

I. Uncertainty Overlap

Cost-Sensitive Approach

Defines cost-sensitivity for object detectors and reduces total cost indirectly via a budget on MDs or FDs.

Detector Cost Matrix

	CD	FD
$\mathrm{CD_T}$	~~$C_{CD} \cdot \lvert CD \rvert$~~ 0	$C_{FD} \cdot \lvert FD \rvert$
$\mathrm{FD_T}$	$C_{MD} \cdot \lvert MD \rvert$	~~$C_{BG} \cdot \lvert BG \rvert$~~ 0

$$C_{total} = C_{MD} \cdot |MD| + C_{FD} \cdot |FD|$$

Set a budget on either remaining correct detections ($|CD|_{PT}$) or removed false detections ($|FD|_{PT}$) post-thresholding:

Budget $b \in [0, 1]$
- ➢ $b \cdot |CD| \leq |CD|_{PT}$
- ➢ $b \cdot |FD| \leq |FD|_{PT}$

II. Manual Thresholding

Automation via ROC Curve

Automatically iterates through IoU thresholds from 0.5 to 0.75 in 0.05 increments to find the optimal threshold for a given budget.

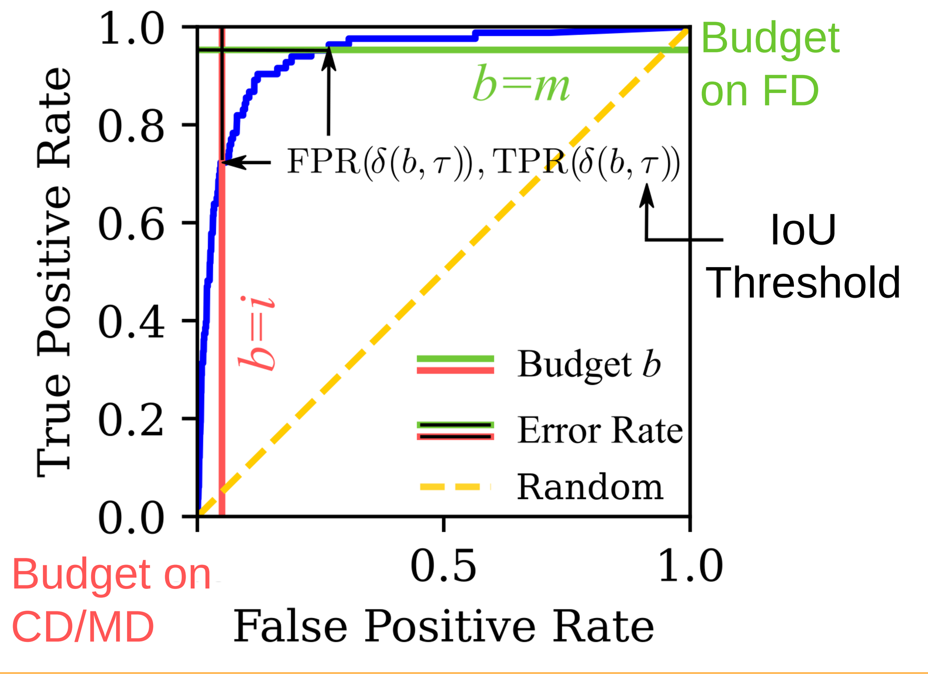

III. Uncertainty Combination

Black-Box Optimization

With entropy and aleatoric uncertainty only:
- **+2–11%** via optimized sum.
- **+36–60%** over conventional methods.

		FD@CD95↑	Binary Accuracy↑	
KITTI	$\sum \sigma_{mc+la}$	68.02±1.97	0.81±0.01	Standard
	$\sum * \sigma_{mc+la}$	**72.36±2.72**	**0.83±0.01**	Optimized
	$\sum \sigma_{la}$	65.86±3.43	0.80±0.02	
	$\sum * \sigma_{la}$	**70.93±1.47**	**0.83±0.01**	
BDD	$\sum \sigma_{mc+la}$	32.03±0.24	0.63±0.00	
	$\sum * \sigma_{mc+la}$	**37.98±0.90**	**0.67±0.00**	
	$\sum \sigma_{la}$	30.65±0.23	0.63±0.00	
	$\sum * \sigma_{la}$	**38.11±0.21**	**0.67±0.00**	
CODA	$\sum \sigma_{mc+la}$	40.60±0.21	0.68±0.00	
	$\sum * \sigma_{mc+la}$	**45.68±0.53**	**0.70±0.00**	
	$\sum \sigma_{la}$	38.49±0.96	0.67±0.00	
	$\sum * \sigma_{la}$	**43.95±0.43**	**0.69±0.00**	

Our baseline is EfficientDet-D0:
- Pre-trained on COCO.
- Fine-tuned on two autonomous driving datasets separately: KITTI and BDD100K.
- Evaluated on an additional corner cases dataset CODA.

We investigate the classification, localization, epistemic and aleatoric uncertainties, entropy, their calibrated and normalized versions.

Inference time with and without uncertainty estimation:
Baseline: ~35ms
Loss Attenuation (la): ~30ms Monte Carlo Dropout (mc): ~185ms