HalluciDet: Hallucinating RGB Modality for Person Detection Through Privileged Information

Heitor R. Medeiros, Fidel G. Pena, Masih Aminbeidokhti, Eric Granger, Marco Pedersoli

Introduction & Motivation

Figure 1. IR and RGB images (LLVIP dataset). Images have complementary information; thus, one modality can help the other one.

Our work investigates image translation for object detection :

- Model adapts from pre-trained RGB to IR.
- Guide the image-to-image translation for the final detection task.

HalluciDet

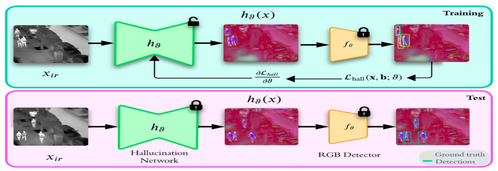

Figure 2. HalluciDet: During training, it is able to train the Hallucination Network with the knowledge from the RGB detector. During the test, it improves the detection of IR.

$$\mathcal{L}_{\text{hall}}(\mathbf{x}, \mathbf{b}, \vartheta) = \mathcal{L}_{\text{cls}}(f_\theta(h_\vartheta(\mathbf{x})), c) + \lambda \cdot \mathcal{L}_{\text{reg}}(f_\theta(h_\vartheta(\mathbf{x})), \mathbf{b})$$

Equation 1: Hallucination Loss.

Paper Code

Image-to-image Benchmark

Image-to-image translation	Learning strategy	AP@50↑ Test Set (Dataset: LLVIP) FCOS	RetinaNet	Faster R-CNN
Blur [10]	-	42.59 ± 4.17	47.06 ± 1.99	63.05 ± 1.96
Histogram Equalization [10]	-	33.10 ± 4.64	36.45 ± 2.02	51.47 ± 4.03
Histogram Stretching [10]	-	38.55 ± 4.25	41.97 ± 1.39	57.69 ± 2.78
Invert [10]	-	53.62 ± 2.07	55.43 ± 2.03	71.83 ± 3.04
Invert + Equalization [10]	-	50.03 ± 2.44	52.57 ± 1.50	68.69 ± 2.73
Invert + Equalization + Blur [10]	-	50.58 ± 2.41	52.62 ± 1.36	68.91 ± 2.74
Invert + Stretching [10]	-	51.48 ± 2.17	52.87 ± 1.80	69.34 ± 3.07
Invert + Stretching + Blur [10]	-	51.54 ± 1.92	52.96 ± 1.80	69.59 ± 2.90
Parallel Combination [10]	-	50.18 ± 2.25	52.52 ± 1.39	68.14 ± 2.98
U-Net [29]	Reconstruction	42.94 ± 4.14	47.35 ± 1.92	63.23 ± 2.03
CycleGAN [39]	Adversarial	22.76 ± 1.94	27.04 ± 4.23	38.92 ± 5.09
CUT [24]	Contrastive learning	19.16 ± 2.10	21.61 ± 2.09	35.17 ± 0.32
FastCUT [24]	Contrastive learning	46.87 ± 2.28	52.39 ± 2.31	67.73 ± 2.14
HalluciDet (ours)	Detection	**63.28 ± 3.49**	**56.48 ± 3.39**	**88.34 ± 1.50**

Table 1. Comparison of detection over different methods and HalluciDet.

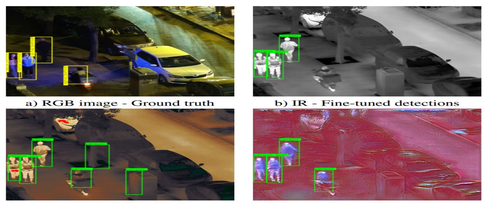

Figure 3. Different image-to-image on detection task.

Training Samples, Params. & Fine-Tuning

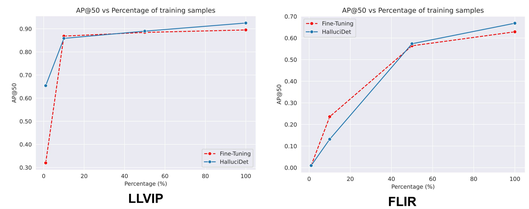

	No Adaptation	Fine-tuning	HalluciDet
FCOS	38.52 ± 0.79	42.22 ± 1.04	**49.18 ± 0.99**
RetinaNet	44.13 ± 2.01	47.87 ± 2.21	**49.01 ± 4.08**
Faster R-CNN	55.85 ± 1.19	61.48 ± 1.55	**70.90 ± 1.35**

Table 2. Comparison HalluciDet and Fine-tuning on FLIR dataset.

Method		Params.	AP@50↑
Faster R-CNN		41.3 M	84.83
HalluciDet	MobileNet_{v3s}	+ 3.1 M	85.20
	MobileNet_{v2}	+ 6.6 M	89.73
	ResNet_{18}	+ 14.3 M	90.42
	ResNet_{34}	+ 24.4 M	90.65

Table 3. Different Hallucination backbones were evaluated on the LLVIP dataset.

Qualitative Results

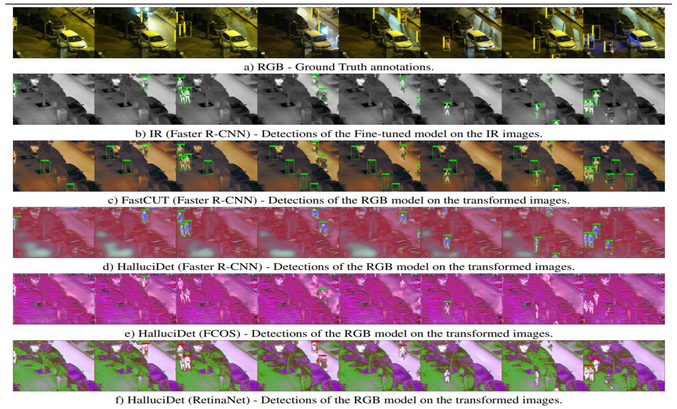

Figure 5. HalluciDet detections on IR LLVIP dataset over different detectors.

Conclusion

✓ We propose a **novel approach** that leverages privileged information from **pre-trained RGB detectors** and **adapts** it for **IR** detection without changing the detector performance on **RGB**.

✓ HalluciDet uses a straightforward yet **powerful image translation network** to **reduce the domain gap between IR-RGB** modalities, guided by the proposed **hallucination loss** function incorporating standard object detection terms.