# Detecting the Unexpected via Image Resynthesis

Krzysztof Lis  Krishna Nakka  Pascal Fua  Mathieu Salzmann - CVLab EPFL

ICCV 2019 Seoul, Korea

Drive for All

This work was supported in part by the International Chair Drive for All: MINES ParisTech – Peugeot-Citroën – Safran – Valeo, and by the Swiss National Science Foundation.

## Motivation

Standard semantic segmentation fails to capture unusual objects which fall outside any of the known semantic classes. We attempt to detect such anomalies and warn about potential dangers in a self-driving scenario.

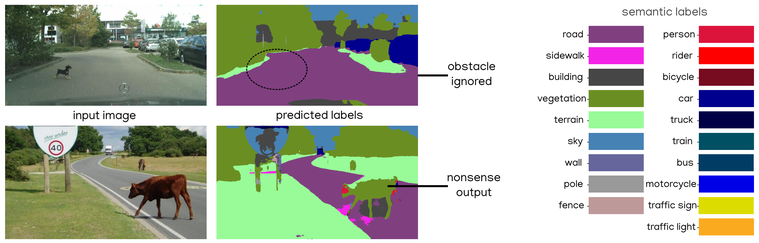


## Approach

We synthesize a plausible image solely from the predicted semantic labels using pix2pixHD[1]. Anomalies, not expressed by the known classes, will differ strongly between input and resynthesis, and will be detected as discrepancies.

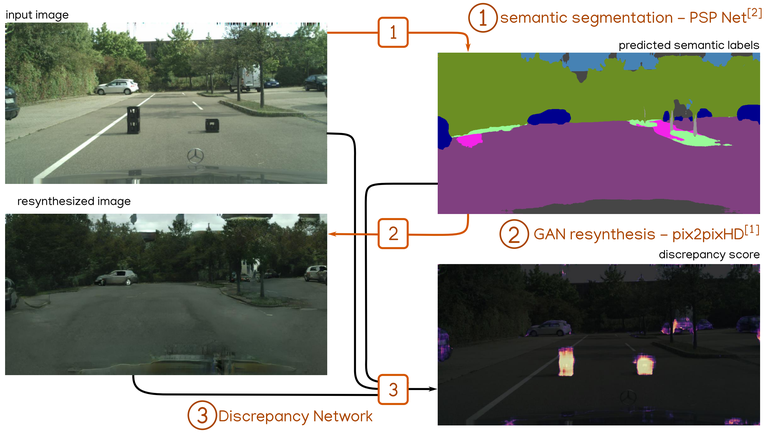


## Discrepancy Network

The discrepancy network learns to detect meaningful differences while ignoring synthesis artifacts and variation between objects of a given class.

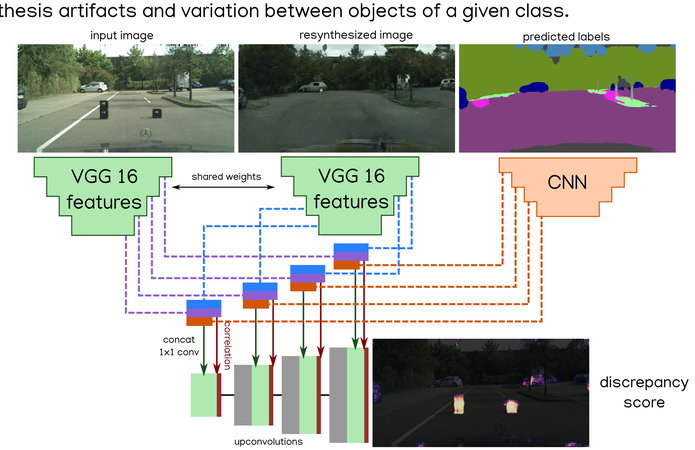


## Synthetic Discrepancy Training

We train the system without any prior knowledge of anomalies. We generate artificial discrepancies from Cityscapes[4] data by altering the semantics of selected object instances.

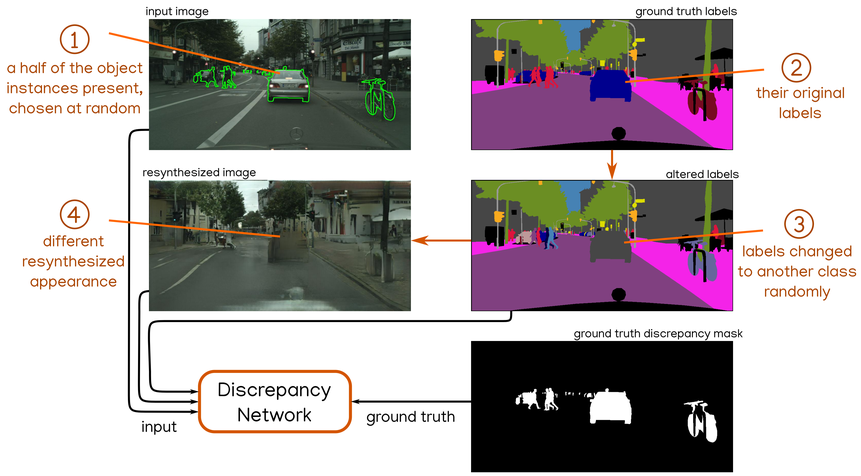


## New dataset: Road Anomaly

We contribute a new dataset of unusual objects in traffic scenes, with pixel mask annotations.

www.epfl.ch/labs/cvlab/data/road-anomaly

## Adversarial attack detection

Adversarial attacks[5][6] alter the image infinitesimally to fool the segmentation network. We detect the resulting nonsensical label map using a non-differentiable histogram of oriented gradients (HOG) feature distance between input and synthesized images.

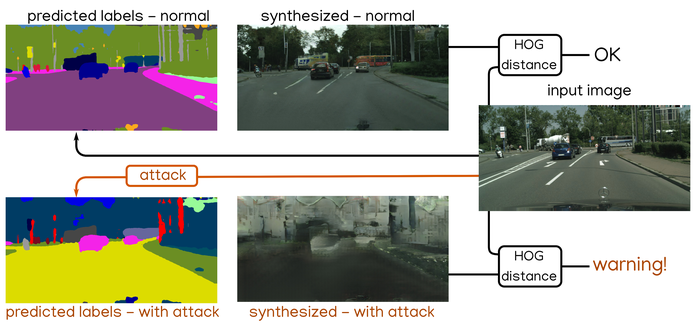


## Results

We evaluate anomaly detection performance on the Road Anomaly and Lost and Found[3] datasets. Baselines:

**Uncertainty (Ensemble)**
Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles.
B. Lakshminarayanan, A. Pritzel, and C. Blundell [NIPS 2017]

**Uncertainty (Dropout)**
Bayesian Segnet: Model Uncertainty in Deep Convolutional Encoder-Decoder Architectures for Scene Understanding
A. Kendall, V. Badrinarayanan, and R. Cipolla [ArXiv 2015]

**RBM**
Real-Time Small Obstacle Detection on Highways Using Compressive RBM Road Reconstruction.
C. Creuso and A. Munawar [Intelligent Vehicles Symposium 2015]

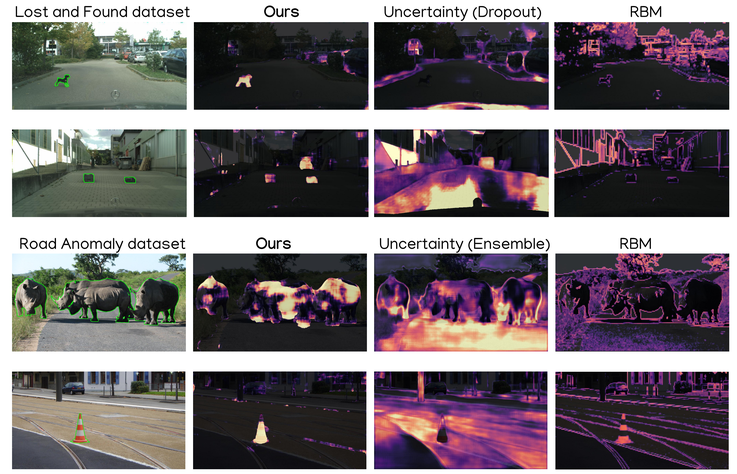


Quantitative anomaly detection results; the reported value is the AUROC score.

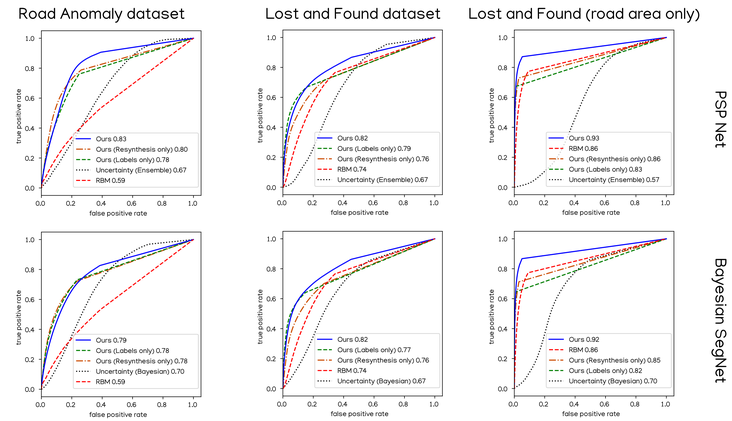


## References

[1] High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs.
T. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro [CVPR 2018]

[2] Pyramid Scene Parsing Network. H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia [CVPR 2017]

[3] Lost and Found: Detecting Small Road Hazards for Self-Driving Vehicles.
P. Pinggera, S. Ramos, S. Gehrig, U. Franke, C. Rother, and R. Mester [IROS 2016]

[4] The Cityscapes Dataset for Semantic Urban Scene Understanding.
M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele [CVPR 2016]

[5] Adversarial examples for semantic segmentation and object detection. C. Xie, J. Wang, Z. Zhang, Y. Zhou, L. Xie, and A. Yuille [ICCV 2017]

[6] Houdini: Fooling deep structured visual and speech recognition models with adversarial examples.
M. M. Cisse, Y. Adi, N. Neverova, and J. Keshet [NIPS 2017]