Anomaly Detection With Multiple Hypotheses Predictions

Duc Tam Nguyen[1,2], Zhongyu Lou[2], Michael Klar[2], Thomas Brox[1]

[1]Computer Vision Lab - University of Freiburg, Germany

[2]Robert Bosch GmbH, Germany

Introduction

- Anomaly Detection is needed when a subset of classes is extremely rare or some classes are unknown at training time.
- Typically: Learn to approximate “normal” data distribution and measure deviation at test time.
- However:
  - Images inputs are high-dimensional → capturing the complete data density is **difficult** and **data-intensive**
  - In autoencoders, blurry reconstructions have the highest likelihood. But **blurry images** are also **anomalies**!

Introduction

- Anomaly Detection is needed when a subset of classes is extremely rare or some classes are unknown at training time.
- Typically: Learn to approximate “normal” data distribution and measure deviation at test time.
- However:
    - Images inputs are high-dimensional → capturing the complete data density is **difficult** and **data-intensive**
    - In autoencoders, blurry reconstructions have the highest likelihood. But **blurry images** are also **anomalies**!

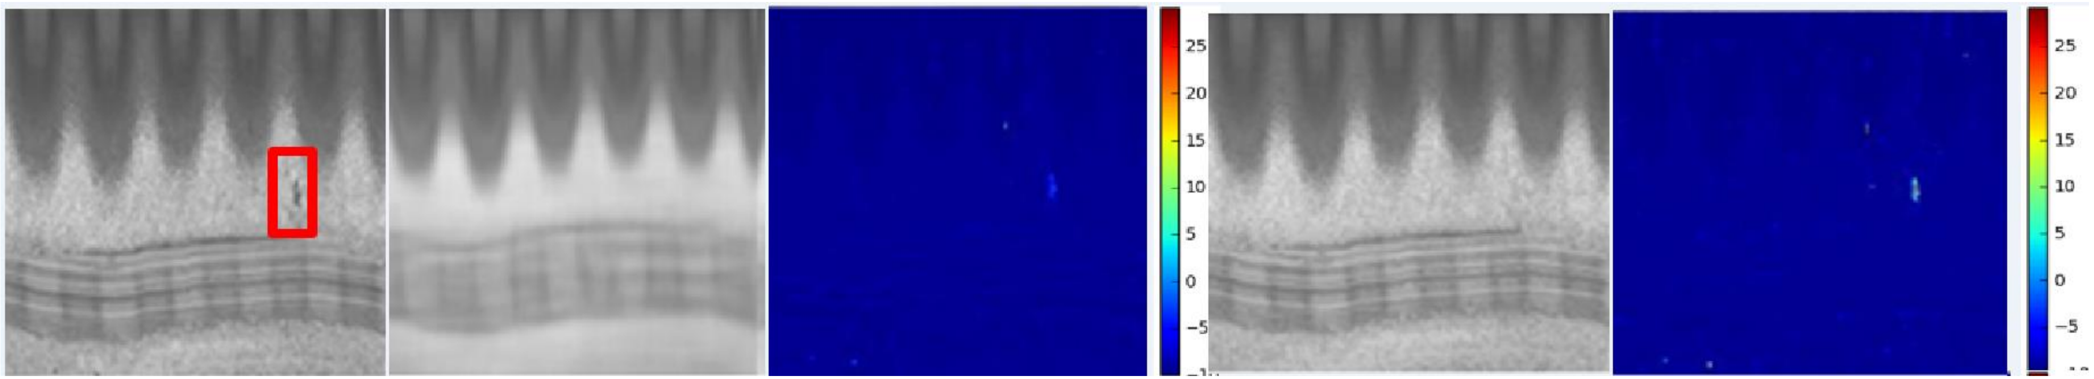

Capture multiple data modes with Multi-hypotheses-networks

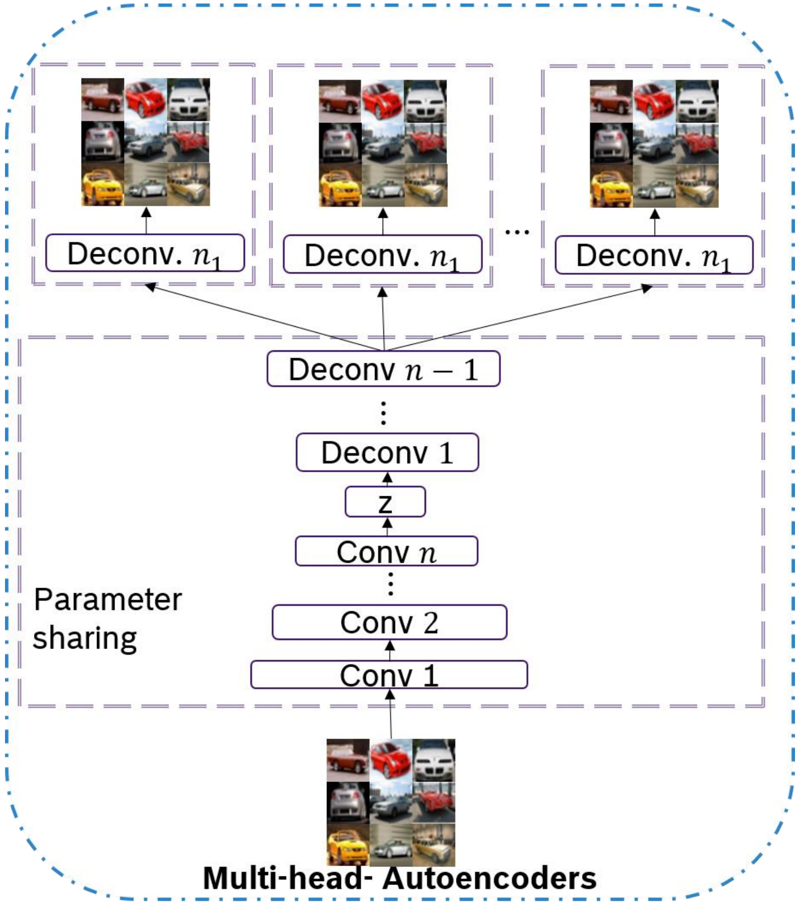

Capture multiple data modes with Multi-hypotheses-networks

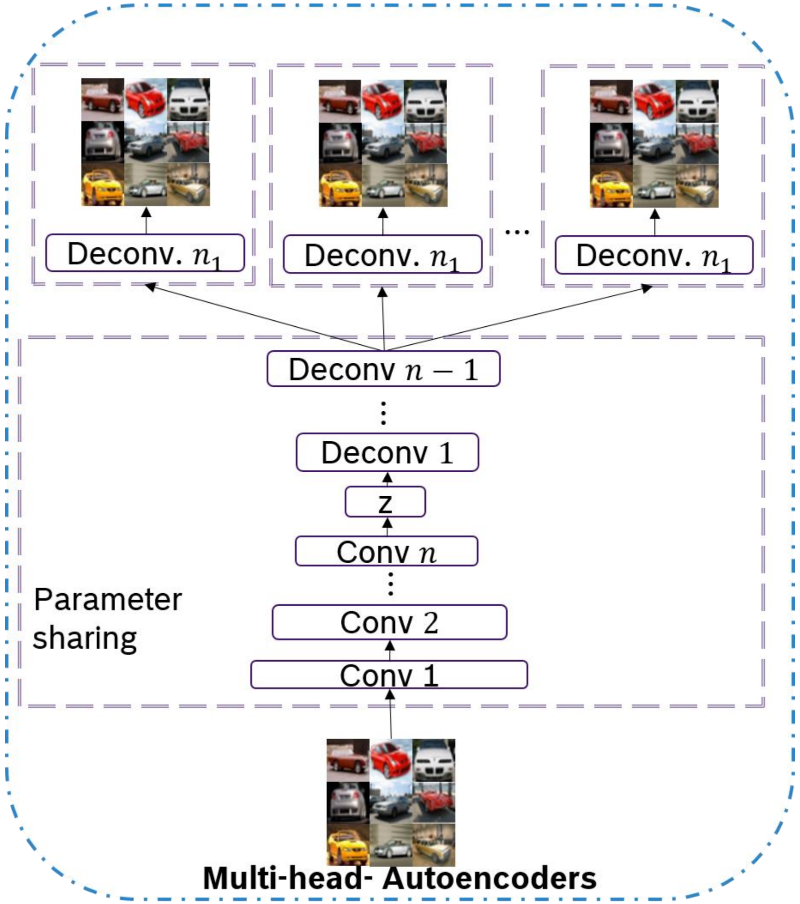

← Hypotheses could support non-existing data mode!

← Diversity among hypotheses?

→ Unsuitable for anomaly detection in this form !

Adversarial regularizer

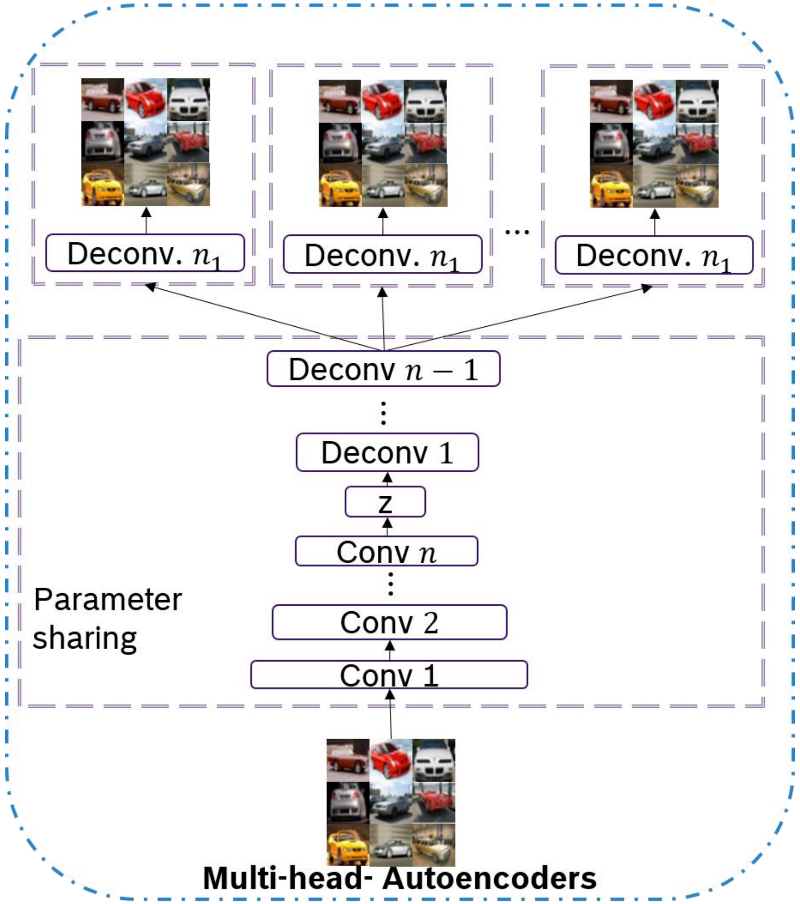

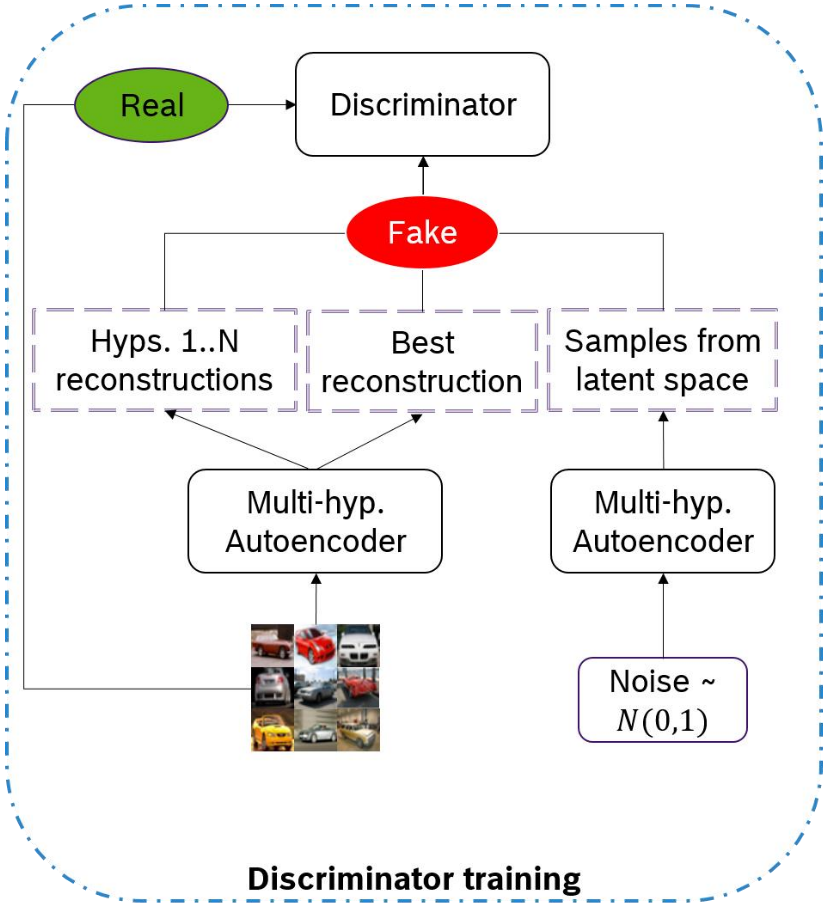

Anomaly detection with multiple hypotheses

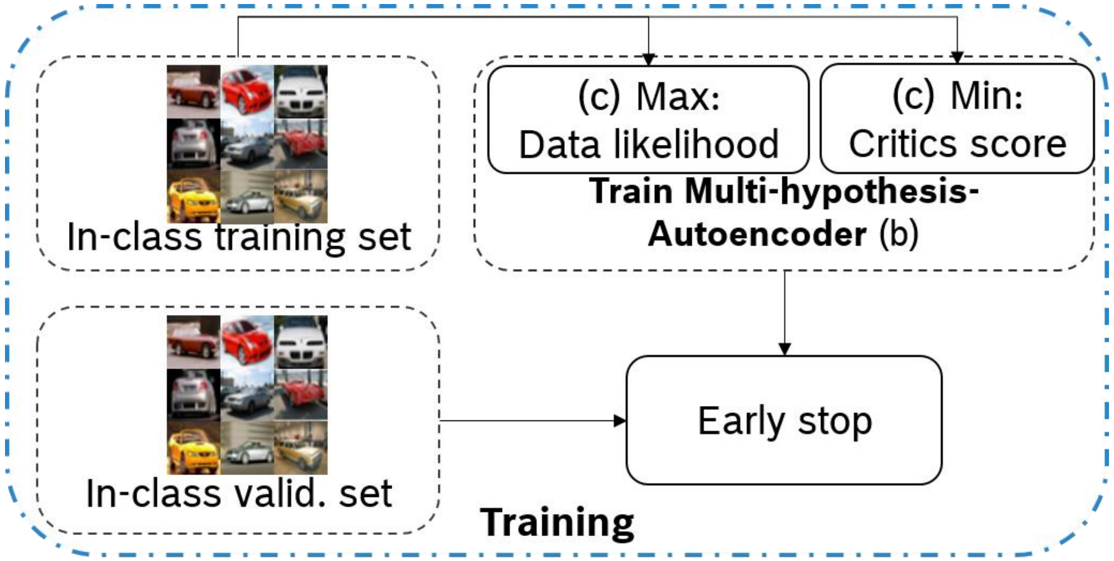

Anomaly detection with multiple hypotheses

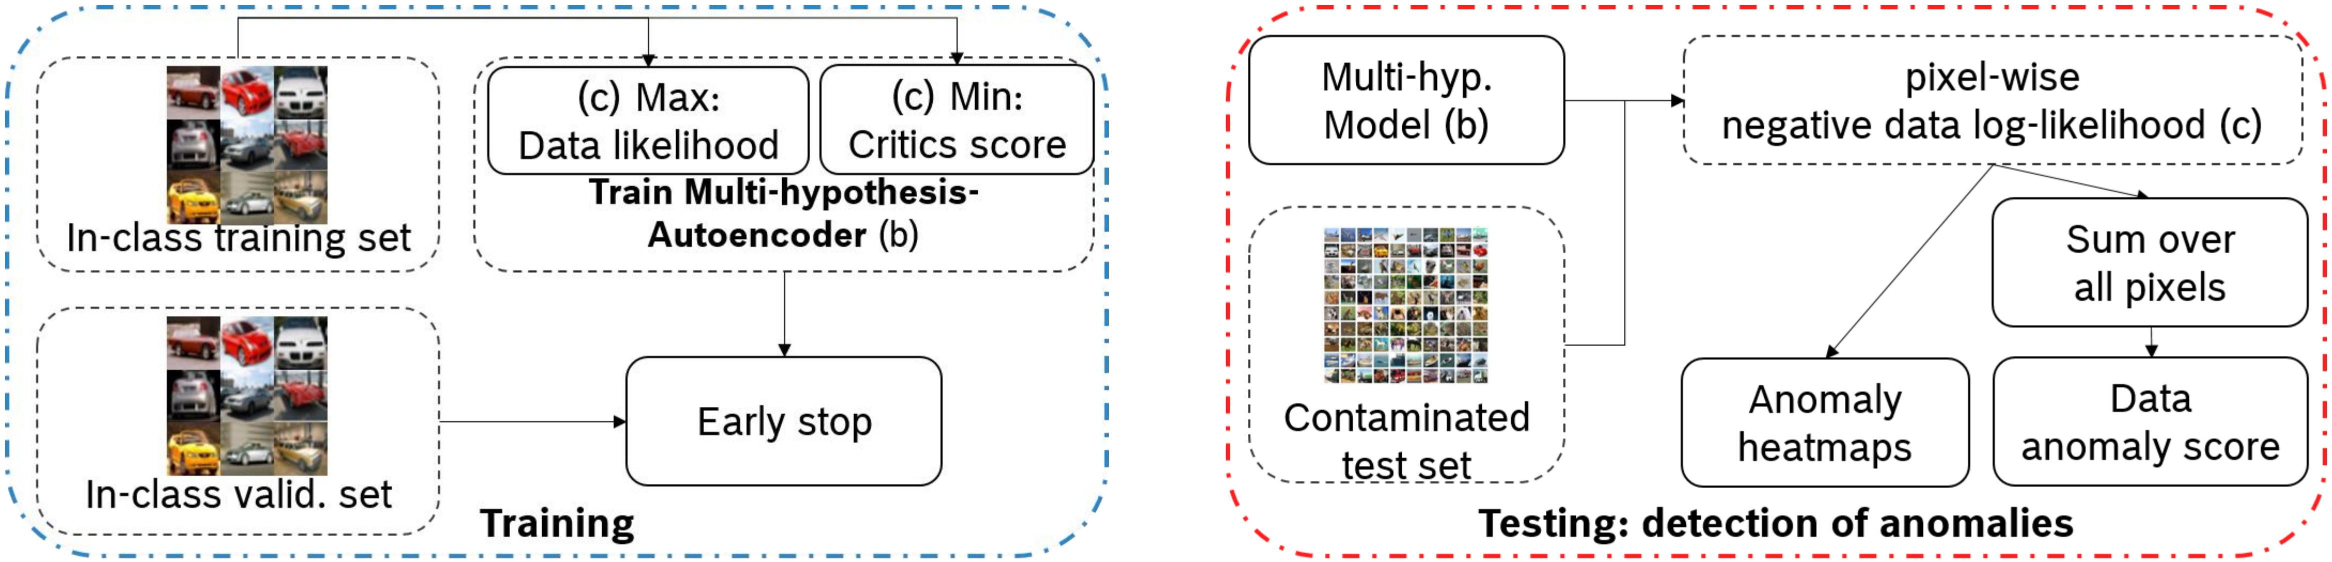

Experimental Results

Table 2. Anomaly detection on CIFAR-10, performance measured in AUROC. Each class is considered as the normal class once with all other classes being considered as anomalies, resulting in 10 one-vs-nine classification tasks. Performance is averaged for all ten tasks and over three runs each (see Appendix for detailed performance). Our approach significantly outperforms previous non-Deep Learning and Deep Learning methods.

TYPE	MODELS			
NON-DL.	KDE-PCA	OC-SVM-PCA	IF	GMM
	59.0	61.0	55.8	58.5
DL	ANOGAN	OC-D-SVDD	ADGAN	CONAD
	61.2	63.2	62.0	**67.1**

Experimental Results

Table 2. Anomaly detection on CIFAR-10, performance measured in AUROC. Each class is considered as the normal class once with all other classes being considered as anomalies, resulting in 10 one-vs-nine classification tasks. Performance is averaged for all ten tasks and over three runs each (see Appendix for detailed performance). Our approach significantly outperforms previous non-Deep Learning and Deep Learning methods.

TYPE	MODELS			
NON-DL.	KDE-PCA	OC-SVM-PCA	IF	GMM
	59.0	61.0	55.8	58.5
DL	ANOGAN	OC-D-SVDD	ADGAN	CONAD
	61.2	63.2	62.0	**67.1**

Table 6. Anomaly detection performance and their standard variance on the Metal Anomaly dataset. To reduce noisy residuals due to the high-dimensional input domain, only 10% of maximally abnormal pixels with the highest residuals are summed to form the total anomaly score. AUROC is computed on an unseen test set, a combination of normal and anomaly data. For more detailed results see Appendix. The anomaly detection performance of plain MHP rapidly breaks down with an increasing number of hypotheses.

	HYPOTHESES			
MODEL	1	2	4	8
MHP	94.2 (1.4)	98.0 (0.5)	97.0 (1.0)	95.0 (0.2)
MHP+WTA		98.0 (0.9)	**98.0** (0.1)	94.6 (3.3)
MDN		90.0 (1.1)	91.0 (1.9)	91.6 (3.5)
MDN+GAN	93.6 (0.7)	94.2 (1.6)	91.3 (1.9)	94.3 (1.1)
CONAD		**98.5** (0.1)	97.7 (0.5)	**96.5** (0.2)

Visit our poster

- Our PosterID : #135 @ Pacific Ballroom