Optical surface inspection: A novelty detection approach based on CNN-encoded features

INSTITUTE FOR OPTICAL SYSTEMS

Hochschule Konstanz
University of Applied Sciences

Michael Grunwald, Matthias Hermann,
Fabian Freiberg, Pascal Laube,
Matthias O. Franz

HTWG Konstanz
Institute for Optical Systems

SPIE OPTICS + Photonics
10752 Applications of Digital Image Processing XLI
22 August 2018

INSTITUTE FOR OPTICAL SYSTEMS

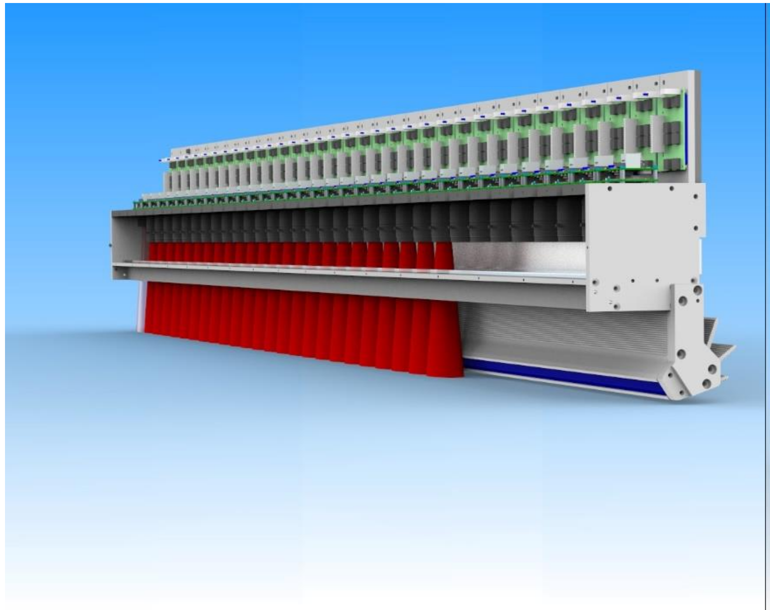
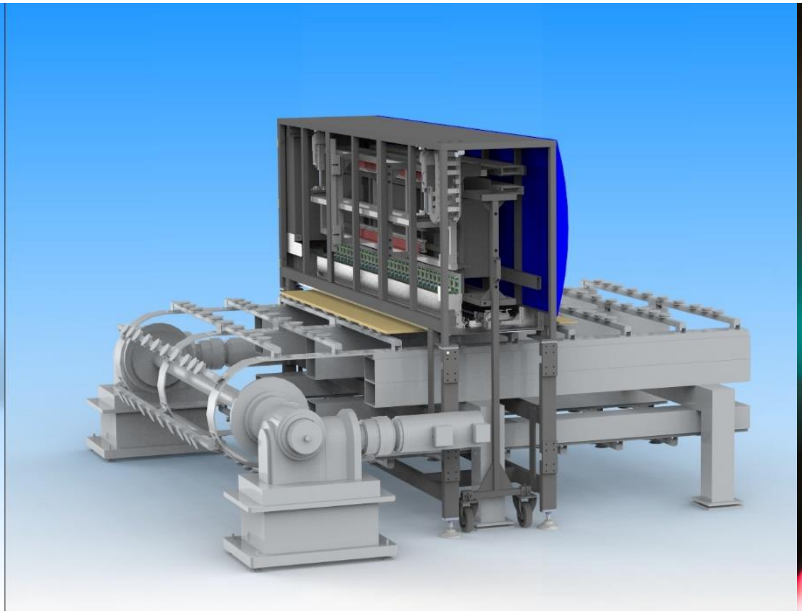
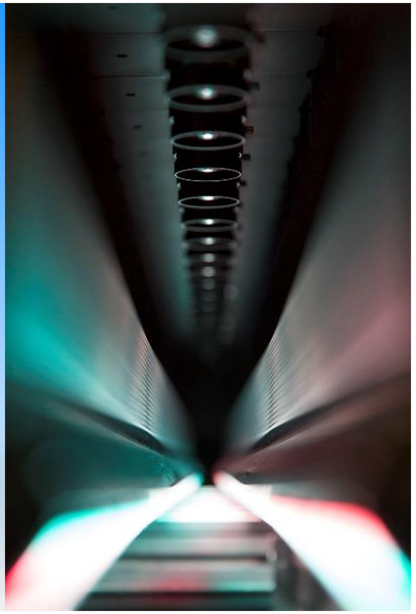

Introduction

Optical surface inspection of printed textured surfaces is challenging as a reference texture is typically only available in a digital format and there is only little information about potential anomalies (often texture-dependent).

Ensuring production quality and detecting unwanted visual anomalies gets even more attention through individualization, customization and personalization of surface textures like floors and decors.

IOS
INSTITUTE FOR OPTICAL SYSTEMS

Overview

- Introduction
- Machine learning
  - Novelty detection
  - Loss functions for neural networks
- CNN-based texture descriptors
- Results
- Conclusion

IOS
INSTITUTE FOR OPTICAL SYSTEMS

Schematic Surface Inspection System

32 camera array

30k x 50k px

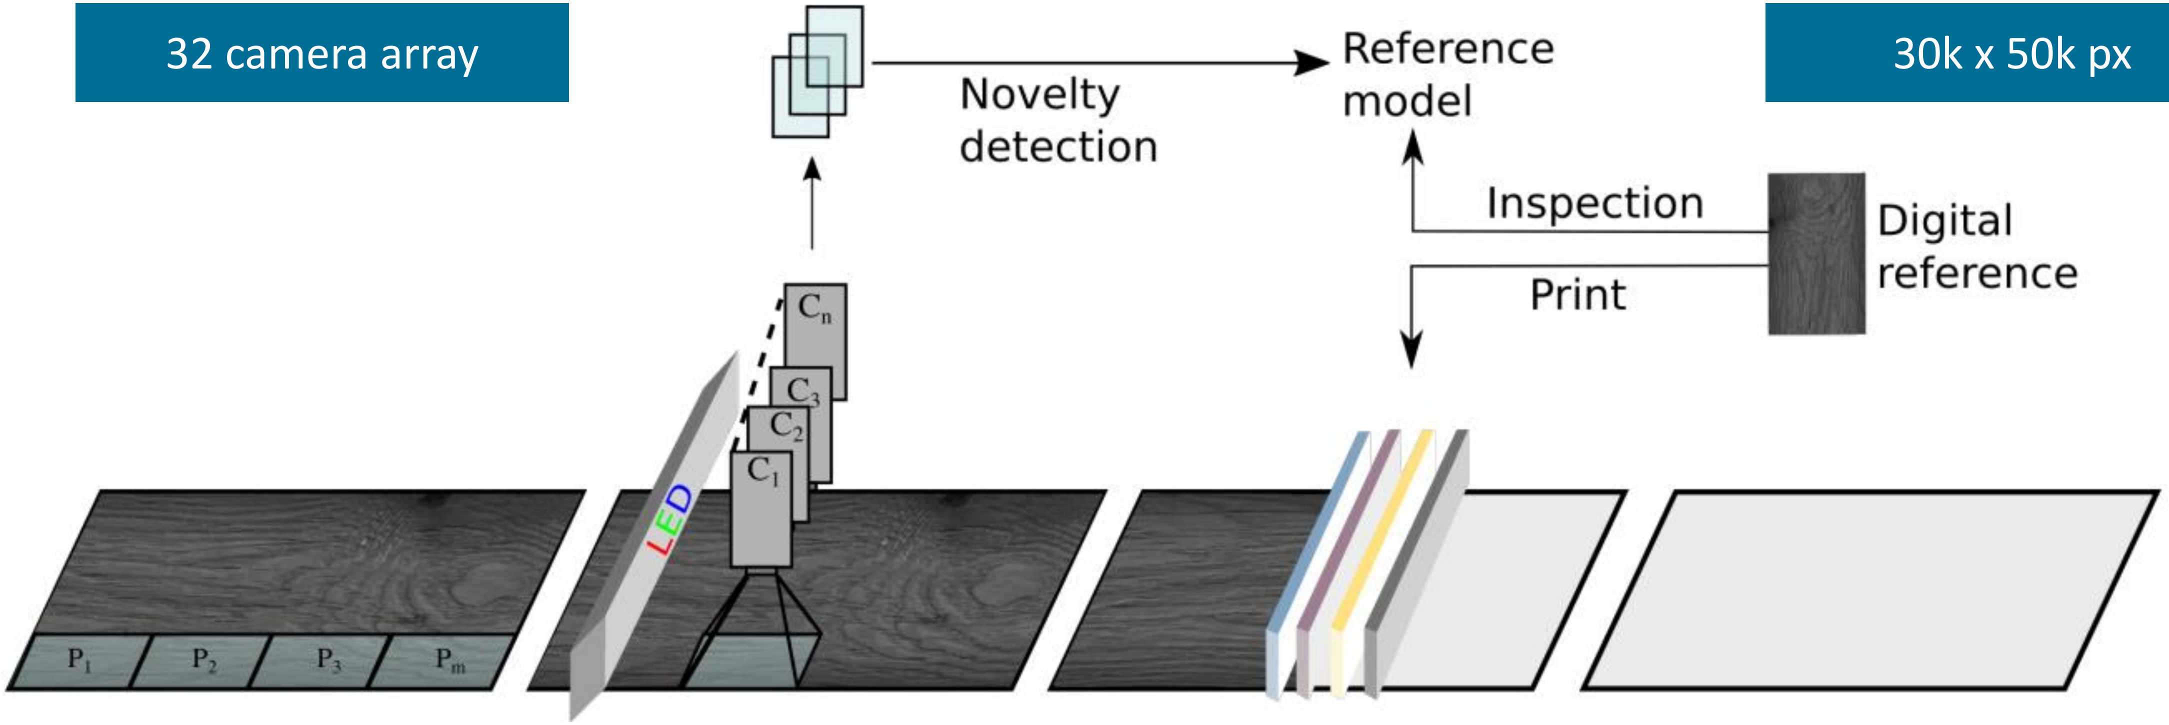

Top 5 errors in Digital Printing Systems

- Nozzle failure
- Contamination like water spots or dust
- Too much color ink
- Substrate failure
- Unwanted color ink

Example nozzle failures

INSTITUTE FOR OPTICAL SYSTEMS

Machine Learning

Machine learning is used to learn a model of the digital reference for a printed decor to enable detecting failures efficiently without registration.

Deep neural networks

Deep neural networks consist of a hierarchy of layers, where each layer successively transforms the input data into more abstract representations (e.g. edge -> corners -> squares -> dice).

The output layer i.e. predictor uses these features to make predictions.

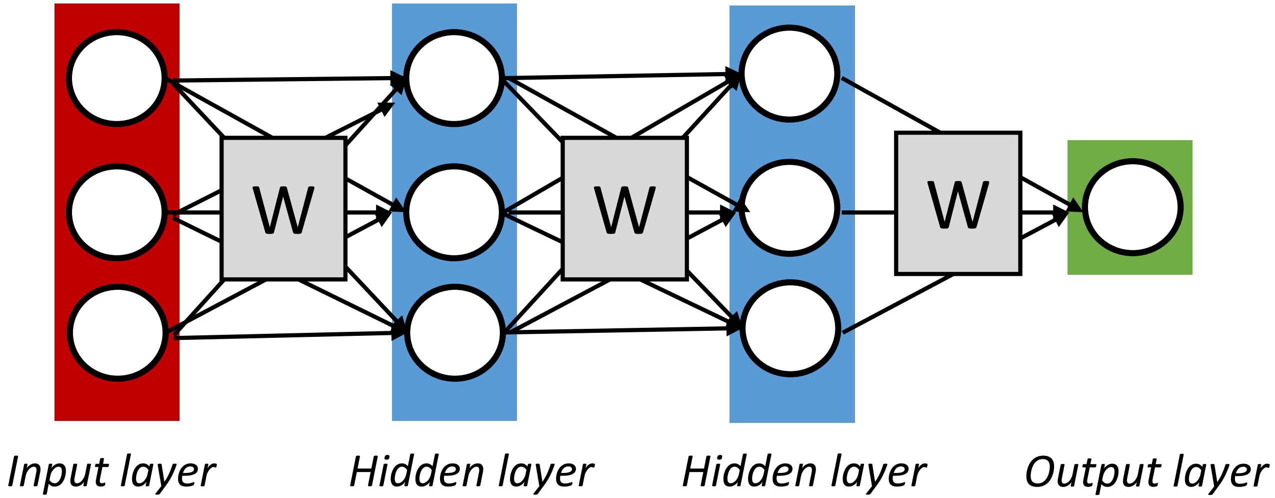

$$f(x) = f_{output}(f_{hidden}(f_{hidden}(X)))$$

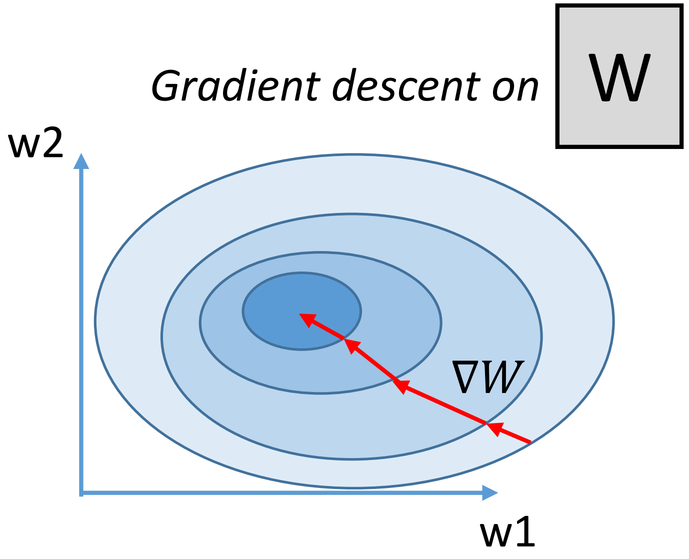

$$\arg\min_w \; Loss(X, W)$$

Novelty detection on textured surfaces

Novelty detection, Anomaly detection and Outlier detection are different names for the same technique.

- There is only one class of labeled data, the reference class.
- Typically we want to detect anomalous samples that do not belong to the reference class.
- Finding a good model of the reference class is key.

INSTITUTE FOR OPTICAL SYSTEMS

Novelty detection with neural networks

Hinge Loss (Chalapaty 2018)

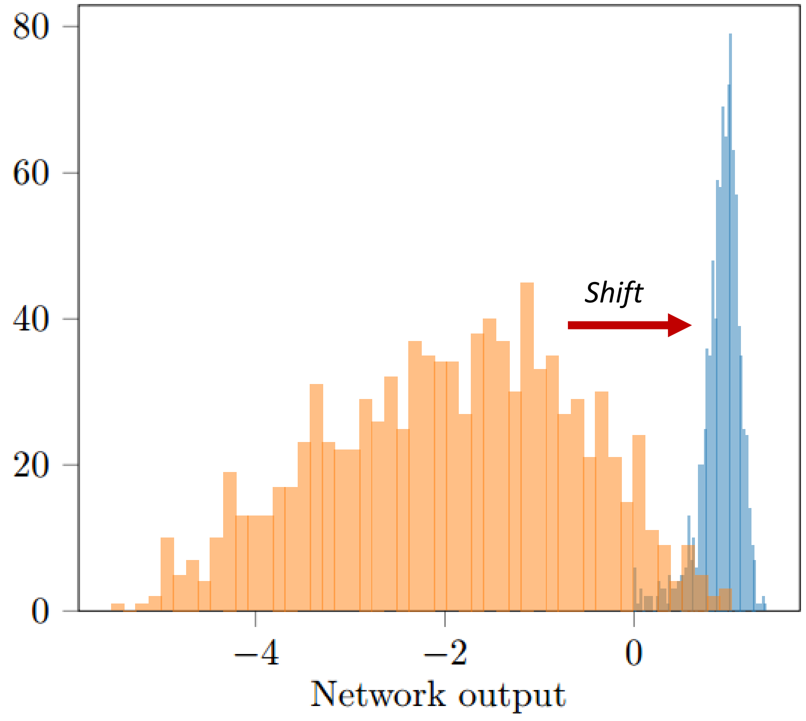

Distance Loss (*Ours)

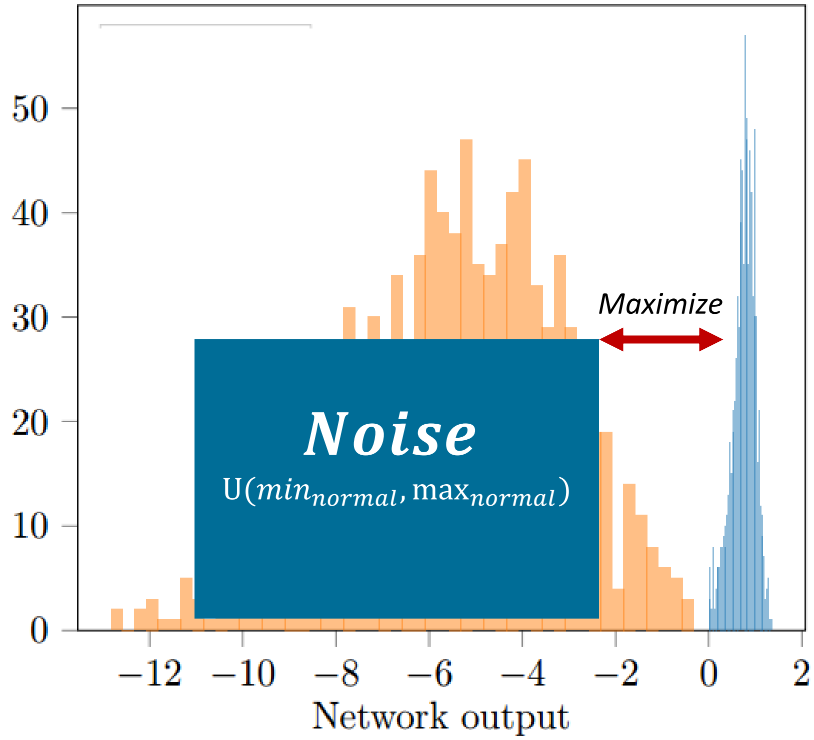

Experimental setup – Reference data

512x512 px patch
from Cut_T4
(600 dpi industrial example)

Full size: 1825x2335 px

512x512 px patch
from BleachedOakVeneer

Full size: 1194x1600 px

512x512 px patch
from Wood-0035

Full size: 512x512px

Experimental setup – Noise data

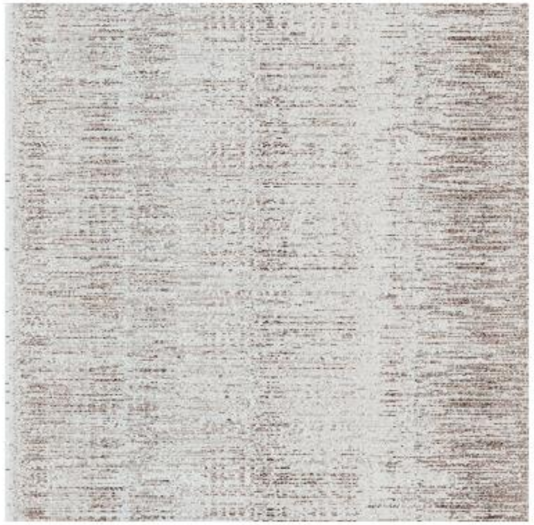

512x512 px patch
from Cut_T4
(600 dpi industrial example)

Full size: 1825x2335 px

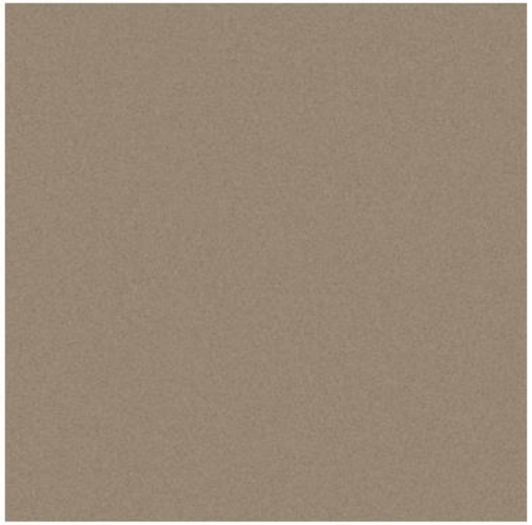

512x512 px patch
from BleachedOakVeneer

Full size: 1194x1600 px

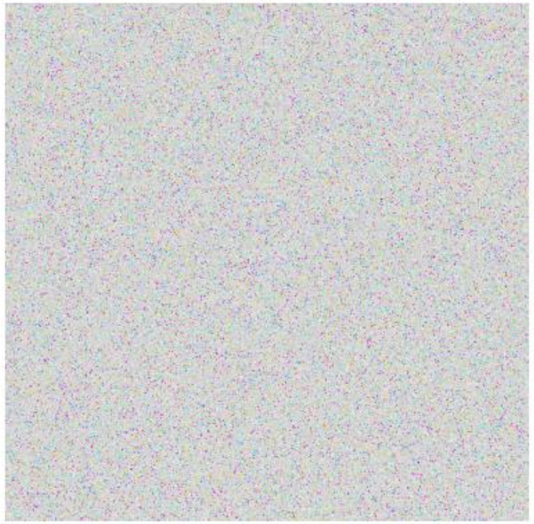

512x512 px patch
from Wood-0035

Full size: 512x512px

Performance Criterion - Area under the ROC / APR

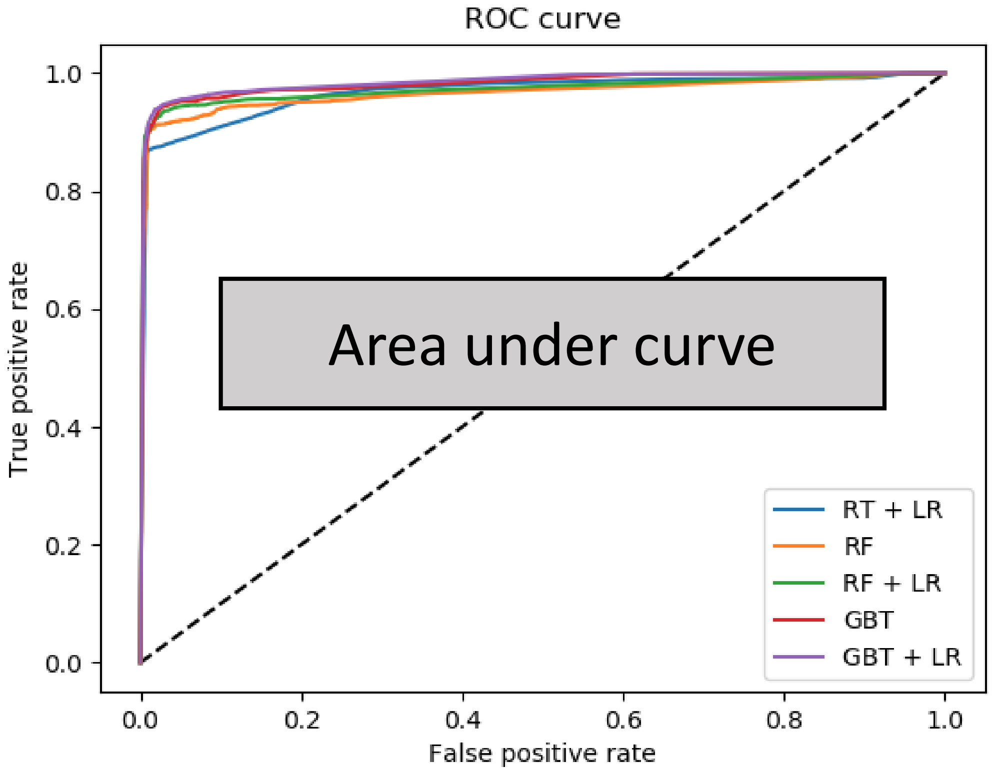

Decision boundary (Typically 0.1 % false positives)

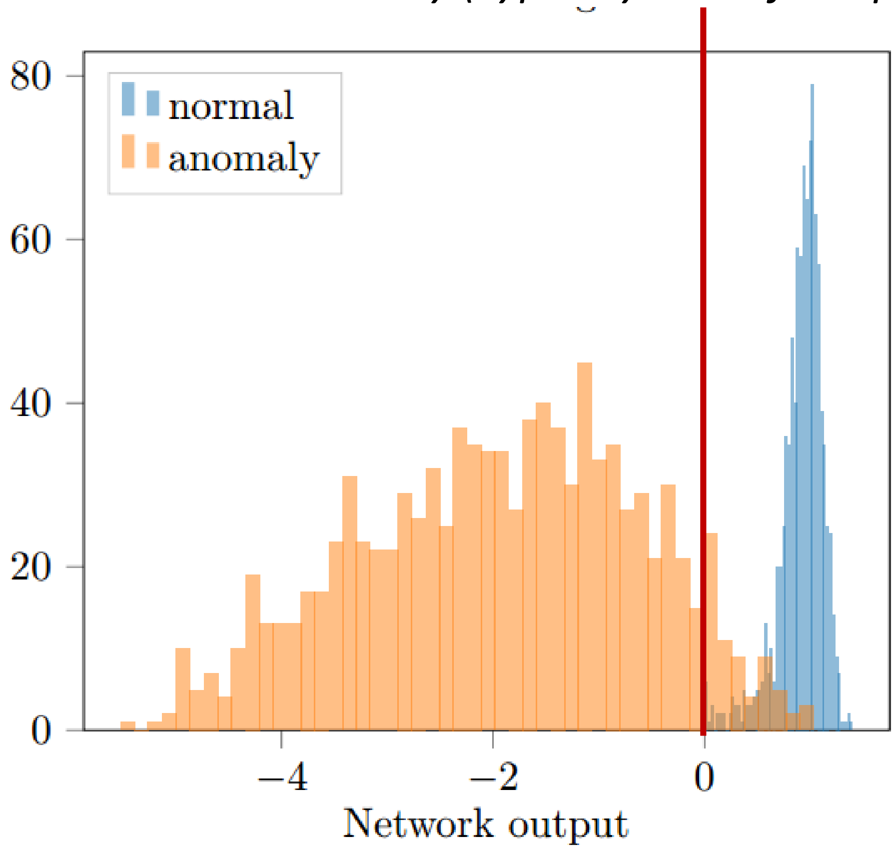

http://scikit-learn.org/stable/modules/generated/sklearn.metrics.roc_curve.html

OS
INSTITUTE FOR OPTICAL SYSTEMS

CNN-encoded features

Based on a pre-trained CNN model (VGG-19) we define several features on-top of intermediate layer activations. Those CNN-encoded features are used for novelty detection in an attached neural network.

Pre-trained Convolutional Deep Neural networks (CNN)

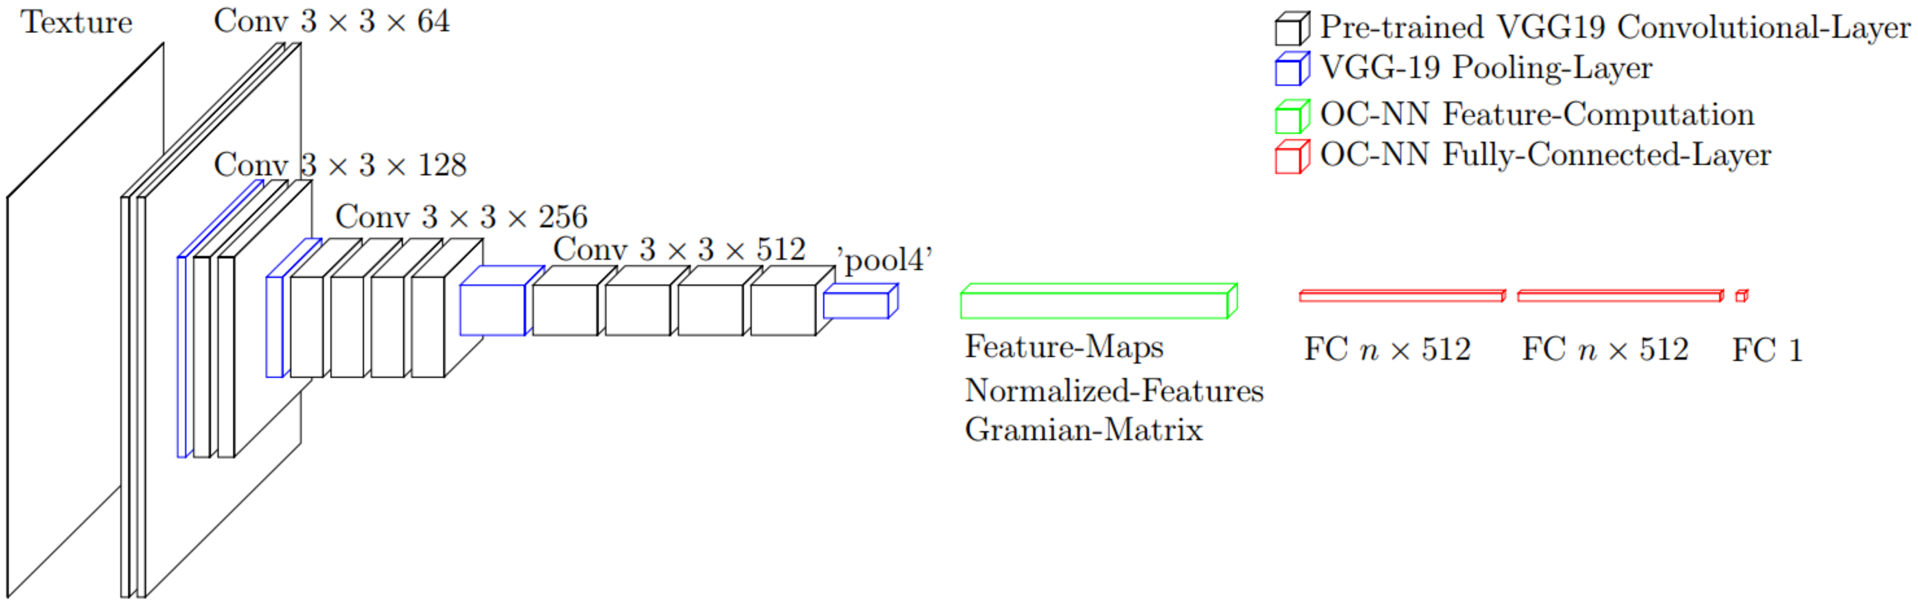

Karen Simonyan and Andrew Zisserman (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition.

512x512 px

512x512 px

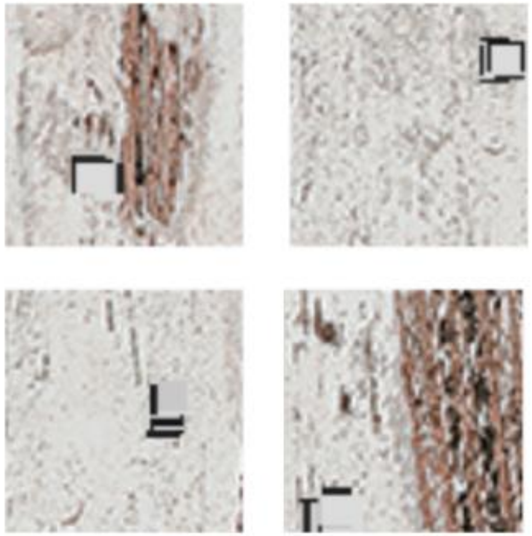

Anomaly ROI (50 x 50 px)

Results

Comparison of our results with our CNN-encoded features based on 32x32 resp. 512x512 px patches and synthetic anomalies (sub-millimeter) modeled by of 2x2 px resp. 16x16 black pixels and random generated continuous lines.

Comparison of different features and classifiers on BleachedOakVeneer

BleachedOakVeneerM
1000 × 32 × 32 normal patches; 1000 × 32 × 32 novelty patches with random 2 × 2 anomalies

Classifier	Feature	AUC	APR	AUC (DF)	APR(DF)
OC-SVM-Linear	Feature maps	0.4998	0.4925	0.4998	0.4925
	Gramian matrix	0.4998	0.4920	0.4998	0.4920
	Normalized features	0.4995	0.4980	0.4995	0.4980
OC-SVM-RBF	Feature maps	0.9805	0.9805	0.9805	0.9805
	Gramian matrix	0.6575	0.6575	0.6575	0.6575
	Normalized features	0.5310	0.5310	0.4997	0.4995
OC-NN-Hinge	Feature maps	0.5496	0.6478	0.6361	0.5788
	Gramian matrix	0.5500	0.4165	0.5772	0.5418
	Normalized features	0.9971	0.9957	0.9669	0.9380
OC-NN-Distance	Feature maps	0.7000	0.6210	0.5610	0.5325
	Gramian matrix	0.7170	0.6780	0.5385	0.5200
	Normalized features	1.0000	1.0000	**0.9915**	**0.9838**

** AUC is Area under ROC-Curve*
***APR is Average Precision Recall*
****DF is Decision function*

Normalized features work best

INSTITUTE FOR OPTICAL SYSTEMS

Comparison of different classifiers with Normalized features on Cut-T4 (600 dpi)

Anomaly		OC-SVM-Linear		OC-SVM-RBF		OC-NN-Hinge		OC-NN-Distance	
n_a	l_a	APR	AUC	APR	AUC	APR	AUC	APR	AUC
32	8	0.4992	0.4984	0.6658	**0.6658**	0.5367	0.5684	0.5623	0.6108
32	16	0.4992	0.4984	0.6658	0.6658	0.7001	0.7858	0.7653	**0.8465**
32	24	0.4992	0.4984	0.6658	0.6658	0.7559	0.8534	0.8386	**0.8998**
64	8	0.4992	0.4984	0.6658	0.6658	0.5498	0.5906	0.5861	**0.6470**
64	16	0.4992	0.4984	0.6658	0.6658	0.7304	0.8155	0.8200	**0.8900**
64	24	0.4992	0.4984	0.6658	0.6658	0.7924	0.8690	0.8948	**0.9409**
128	8	0.4992	0.4984	0.6658	0.6658	0.5548	0.5989	0.6048	**0.6732**
128	16	0.4992	0.4984	0.6658	0.6658	0.7842	0.8624	0.9095	**0.9499**
128	24	0.4992	0.4984	0.6658	0.6658	0.8797	0.9316	0.9705	**0.9844**

**AUC is Area under ROC-Curve*
***APR is Average Precision Recall*

**** n_a is number of randomly aligned continuous lines*
***** l_a is length of a single line*

***Distance Loss works best*!**

Conclusion

A novelty detection approach based on CNN-encoded features and neural networks.

- Apply decision function **before** computing performances in novelty detection scenarios
- Ensure that the loss function produces **spare gradient** with respect to the input

IOS
INSTITUTE FOR OPTICAL SYSTEMS

Research plan

This work

- Novelty detection in domain of optical surface inspection with neural networks using proposed Distance Loss
- Successfully introduced Normalized Features for novelty detection on complex non-ergodically textured surfaces based on a pre-trained neural network

Future work

- Incorporating production variance
- Comparing CNN-encoded features with standard models like Portilla and Simoncelli
- Synthesizing realistic errors on arbitrary textures for better offline evaluation

… why do things look as they do ?
K. Koffka, Principles of Gestalt Psychology, 1935

Thank you for your attention!

Matthias Hermann

Institute for Optical Systems
HTWG Konstanz

matthias.hermann@htwg-konstanz.de

www.ios.htwg-konstanz.de

INSTITUTE FOR OPTICAL SYSTEMS

References

- *Chalapathy, A., Halmetschlager-Funek, G., Prankl, J. and Vincze, M. (2018). Anomaly Detection using One-Class Neural Networks. CoRR.*
- *Gatys, L., Alexander S. Ecker, Matthias Bethge (2015). Texture Synthesis using Convolutional Neural Networks. NIPS.*
- *Portilla, Javier and Eero P. Simoncelli (2000). A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients. IJCV 40(1).*
- *Simonyan, K. and Andrew Zisserman (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition. CoRR.*

Features and Network Architecture

CNN-encoded features

- Gramian matrix (VGG-19-pool4) with **<u>512x512</u>** features
- Raw feature maps (VGG-19-pool4) with **<u>4x512</u>** features
- Normalized features (VGG-19-pool4) with **<u>512</u>** features (*Ours)
  → *Diagonal of Gramian matrix*

Training Parameters

- Input layer with specific feature size
- 3-layer architecture
- Hidden layer with 1000 neurons
- Output layer with 1 neuron
- Training with SGD ($\eta = 10^{-3}$ with 1000 epochs)
- Xavier-Initialization

Novelty Detection with Neural Networks

Hinge Loss (Chalapaty 2018)

$$L_{hinge} = \frac{1}{N}\sum \max(0, r - Output) - r + reg$$

with

$$r = \nu^{th} quantile\ of\ Output,$$
$$\nu \triangleq \%\ false\ positives,$$

and

$$reg \triangleq regularization\ terms$$

Distance Loss (Ours)

$$L_{dist} = \max(0, 1 - \tanh(\hat{R}_{ref} - \hat{R}_{noise}) + reg$$

with

$$\hat{R}_{ref} = \phi\ of\ samples\ in\ [\min(O_{ref}), Q_\rho(O_{ref})],$$
$$\hat{R}_{noise} = \phi\ of\ noise\ in\ [\min(O_{noise}), Q_\rho(O_{noise})]$$

and

$$O \triangleq Output\ and\ Q_\rho \triangleq \rho^{th} quantile,$$
$$reg \triangleq regularization\ terms$$

INSTITUTE FOR OPTICAL SYSTEMS

Prerequisites for inspecting digitally printed decors and wallpapers

Camera calibration

- Camera Errors
  - Varying pixel sensitivities to incoming photons.
  - Non-linearity in camera characteristic.
  - Vignetting is observed shading in image.
- Noise
  - Read noise caused by sensor electronics.

Printing system calibration

- Color Calibration
  - CMYK / RGB transformation
- Noise
  - Varying intensities caused by printing substrate
  - Varying print results because of variances in printing ink

Machine Learning Basics

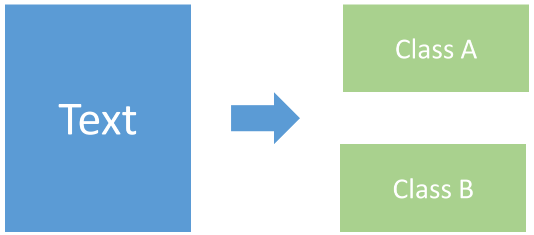

Classification
(supervised – predictive)

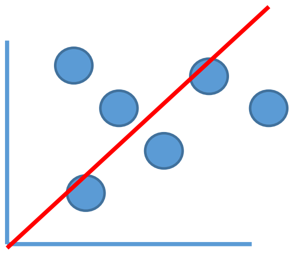

Regression
(supervised – predictive)

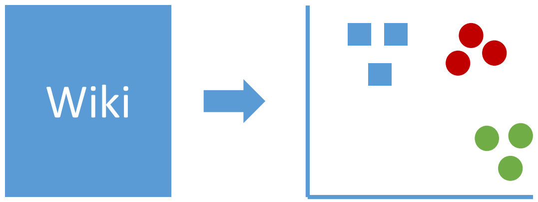

Clustering
(unsupervised – descriptive)

Anomaly Detection
(unsupervised – descriptive)

Detecting novel sub-millimeter anomalies on high-resolution textures (600 dpi)

512x512 px

512x512 px

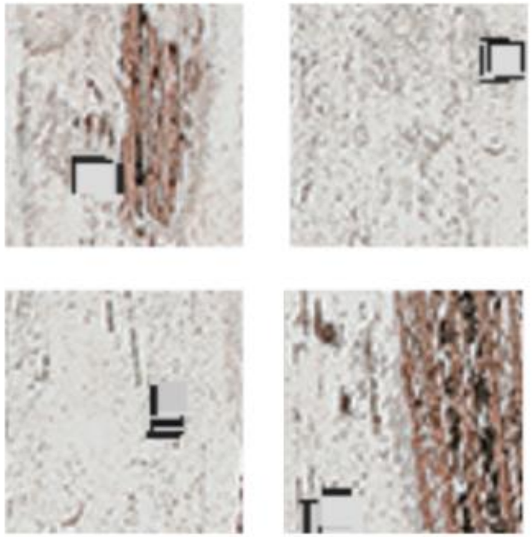

Anomaly ROI (50 x 50 px)

IOS
INSTITUTE FOR OPTICAL SYSTEMS

Detecting novel sub-millimeter anomalies on high-resolution textures (600 dpi)