Anomaly Detection in Images

Giacomo Boracchi,
DEIB, Politecnico di Milano

Università Bicocca, February 26th 2025

giacomo.boracchi@polimi.it
https://boracchi.faculty.polimi.it/

POLITECNICO
MILANO 1863

Giacomo Boracchi

Mathematician (Università Statale degli Studi di Milano 2004),
PhD in Information Technology (DEIB, Politecnico di Milano 2008)
Associate Professor since 2019 at DEIB (Computer Science), Polimi

Research Interests are mathematical and statistical methods for:

- Image / Signal analysis and processing
- Unsupervised learning, change / anomaly detection

Major Courses:

- Mathematical Models and Methods for Image Processing (MSc, Polimi)
- Artificial Neural Networks and Deep Learning (MSc Polimi, Bocconi)
- Advanced Deep Learning Models And Methods (PhD, Polimi)
- Computer Vision (MSc Bocconi, USI 2020)

The Team

We are 3 faculties, 10 PhD students, 1 Research Assistant... and 20+ MSc students!

Giacomo Boracchi

Luca Magri

Federica Arrigoni

Riccardo Margheritti

Michele Craighero

Edoardo Peretti

Andrea Diecidue

Roberto Basla

Rakshith Madhavan

Andrea Ferraris

Diego Martin

Luca Alessandrini

Andrea Marelli

Carlo Sgaravatti

Computer Vision Research

Giacomo Boracchi, Luca Magri, Federica Arrigoni

  

3 faculties
10 PhD students
1 Research Assistants

Research themes and relevant projects:

- **Deep Learning for Visual Recognition:** Object Detection and Segmentation (CCD, SPAD, X-ray, 2D/3D Medical, Aerial Images), Anomaly Detection (images, signals), Learning under limited supervision.

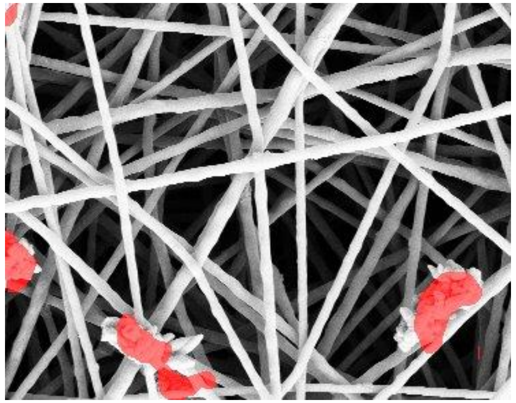

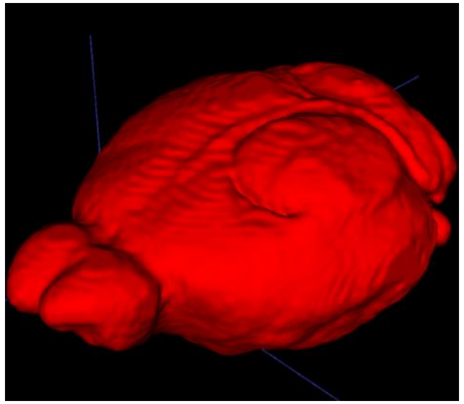

with Prof. Piero Fraternali

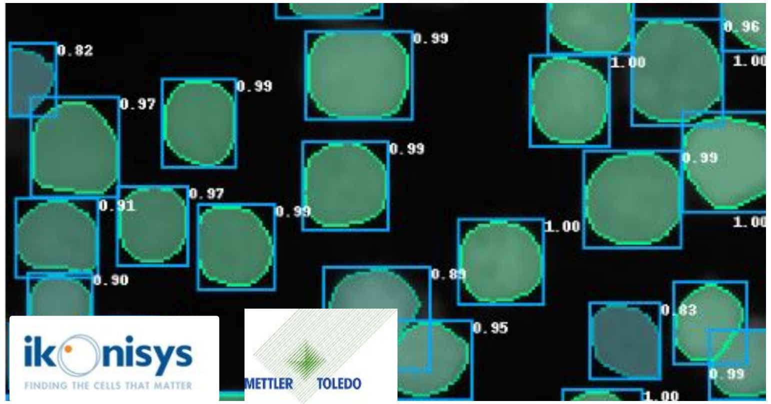

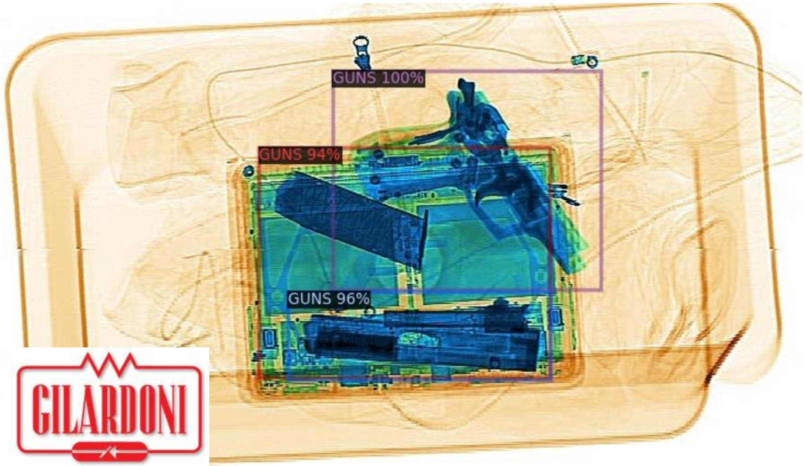

Computer Vision Research

Giacomo Boracchi, Luca Magri, Federica Arrigoni



3 faculties
10 PhD students
1 Research Assistants

Research themes and relevant projects:

- **Deep Learning for Visual Recognition:** Object Detection and Segmentation (CCD, SPAD, X-ray, 2D/3D Medical, Aerial Images), Anomaly Detection (images, signals), Learning under limited supervision.
- **Multi-view Geometry:** 3D Reconstruction, Calibration (Conventional / Event Cameras, X-ray, LiDar, hybrid systems), 6DOF Pose Estimation, Foundational Research.

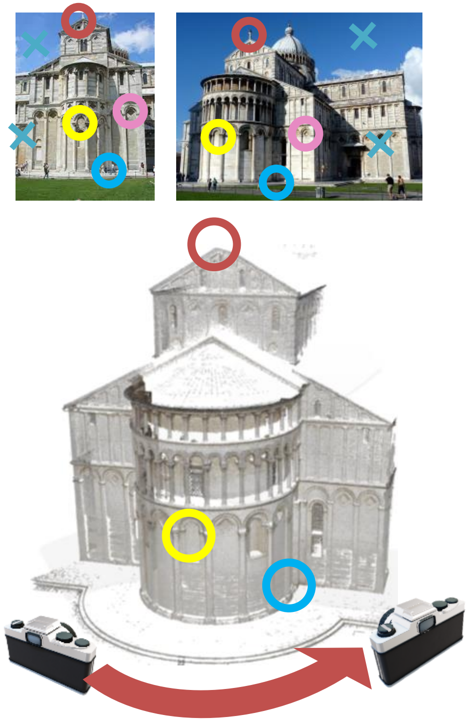

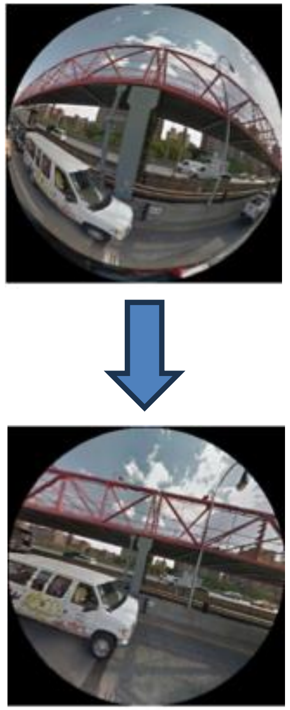

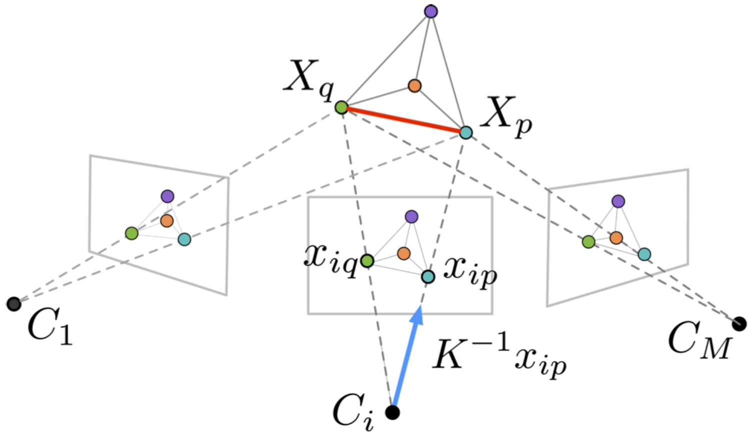

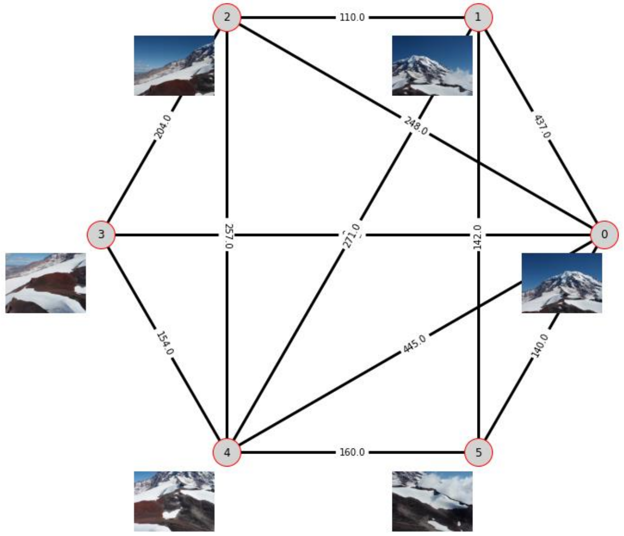

Computer Vision Research

Giacomo Boracchi, Luca Magri, Federica Arrigoni

3 faculties
10 PhD students
1 Research Assistants

Research themes and relevant projects:

- **Deep Learning for Visual Recognition:** Object Detection and Segmentation (CCD, SPAD, X-ray, 2D/3D Medical, Aerial Images), Anomaly Detection (images, signals), Learning under limited supervision.
- **Multi-view Geometry:** 3D Reconstruction, Calibration (Conventional / Event Cameras, X-ray, LiDar, hybrid systems), 6DOF Pose Estimation, Foundational Research.
- **Pattern Recognition:** Clustering, Robust Model Fitting (Scan2BIM, Template Detection), Quantum Computer Vision.

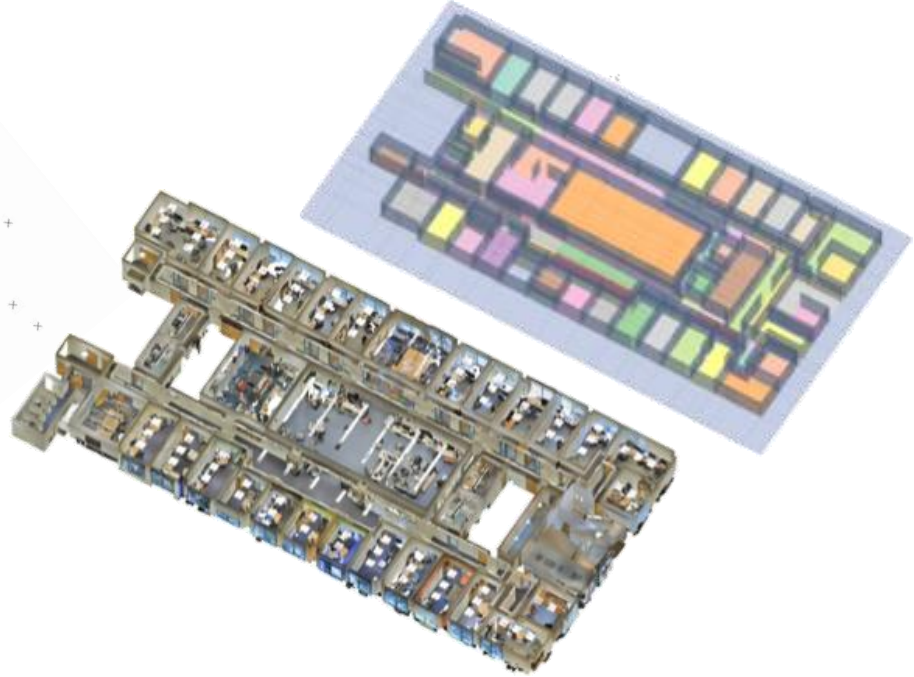

Computer Vision Research

Giacomo Boracchi, Luca Magri, Federica Arrigoni



3 faculties
10 PhD students
4 Research Assistants

Research themes and relevant projects:

- **Deep Learning for Visual Recognition:** Object Detection and Segmentation (CCD, SPAD, X-ray, 2D/3D Medical, Aerial Images), Anomaly Detection (images, signals), Learning under limited supervision.
- **Multi-view Geometry:** 3D Reconstruction, Calibration (Conventional / Event Cameras, X-ray, LiDar, hybrid systems), 6DOF Pose Estimation, Foundational Research.
- **Pattern Recognition:** Clustering, Robust Model Fitting (Scan2BIM, Template Detection), Quantum Computer Vision.
- **Image/Signal Processing and Analysis:** Image Enhancement and Restoration (X-ray, CCD, SPAD), Change and Outlier Detection (Optical fiber signals)

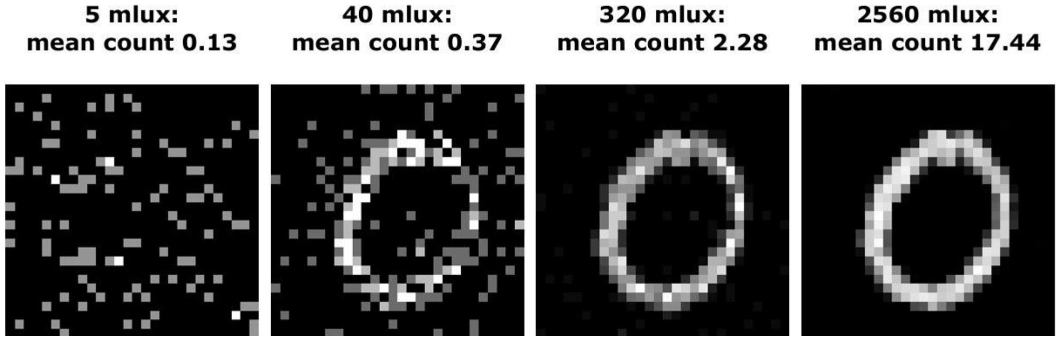

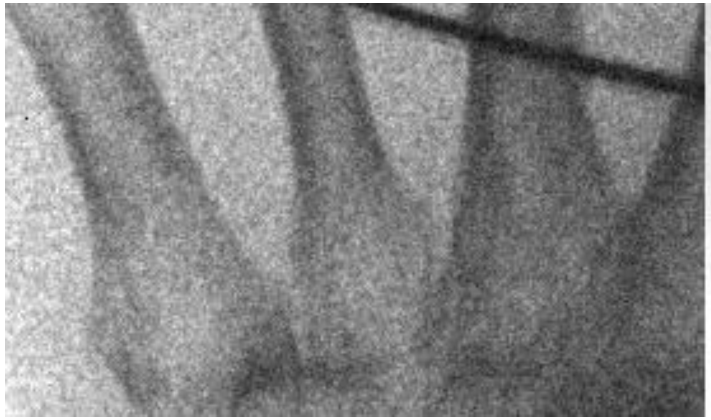
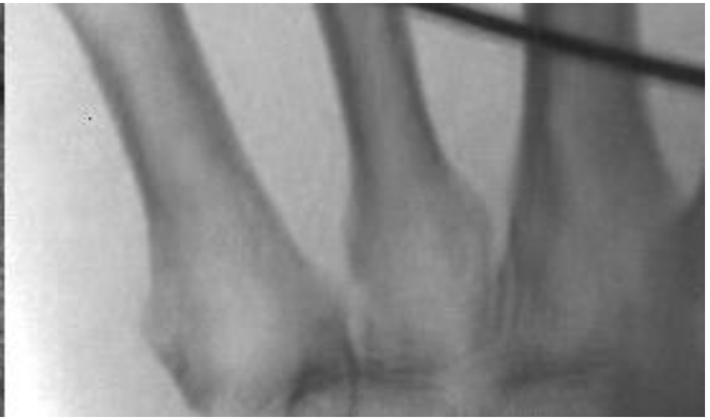

Computer Vision Research

Giacomo Boracchi, Luca Magri, Federica Arrigoni



3 faculties
10 PhD students
1 Research Assistants

Research themes and relevant projects:

- **Deep Learning for Visual Recognition:** Object Detection and Segmentation (CCD, SPAD, X-ray, 2D/3D Medical, Aerial Images), Anomaly Detection (images, signals), Learning under limited supervision.
- **Multi-view Geometry:** 3D Reconstruction, Calibration (Conventional / Event Cameras, X-ray, LiDar, hybrid systems), 6DOF Pose Estimation, Foundational Research.
- **Pattern Recognition:** Clustering, Robust Model Fitting (Scan2BIM, Template Detection), Quantum Computer Vision.
- **Image/Signal Processing and Analysis:** Image Enhancement and Restoration (X-ray, CCD, SPAD), Change and Outlier Detection (Optical fiber signals)
- **Change and Anomaly Detection:** Design of change-detection tests for high-dimensional data (QuantTree), Sequential Monitoring, Datastream Mining, Anomaly Detection (tree based methods).

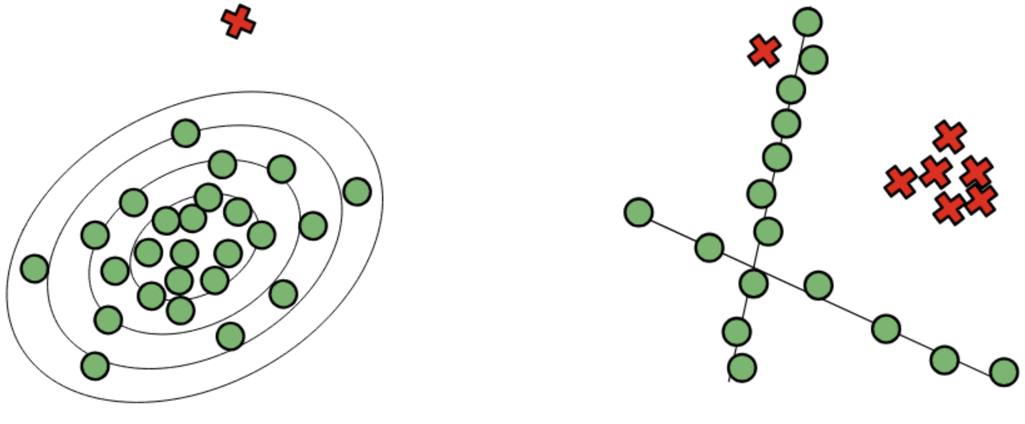

(a) *Density.* **(b)** *Structure.*

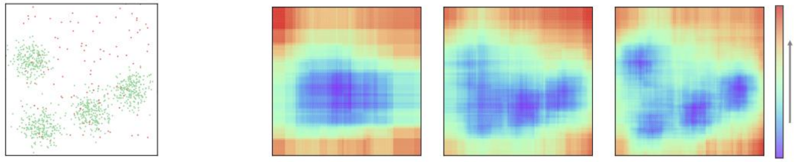

(a) *Data stream $x_1, \dots, x_t \in \mathbb{R}^d$.* **(b)** *Anomaly scores s at different time instants t, from left to right.*

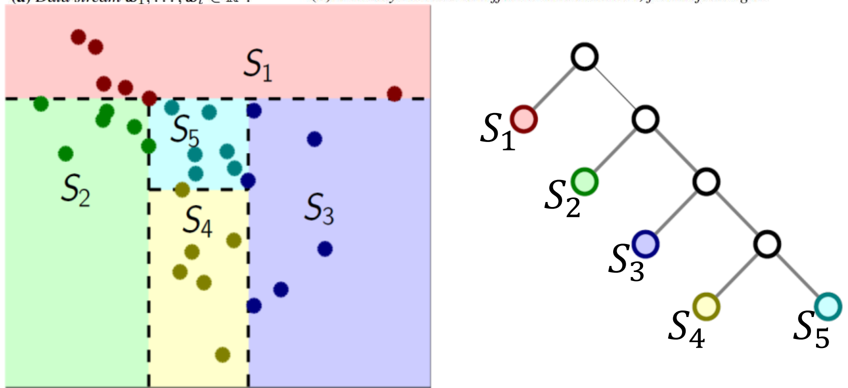

Anomalies

Anomaly detection in health

Mammograms

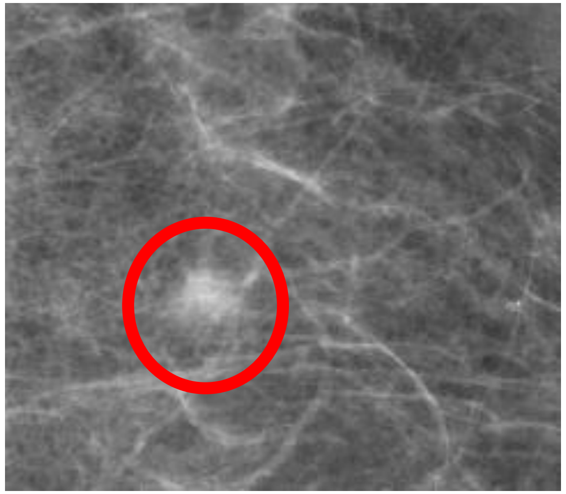

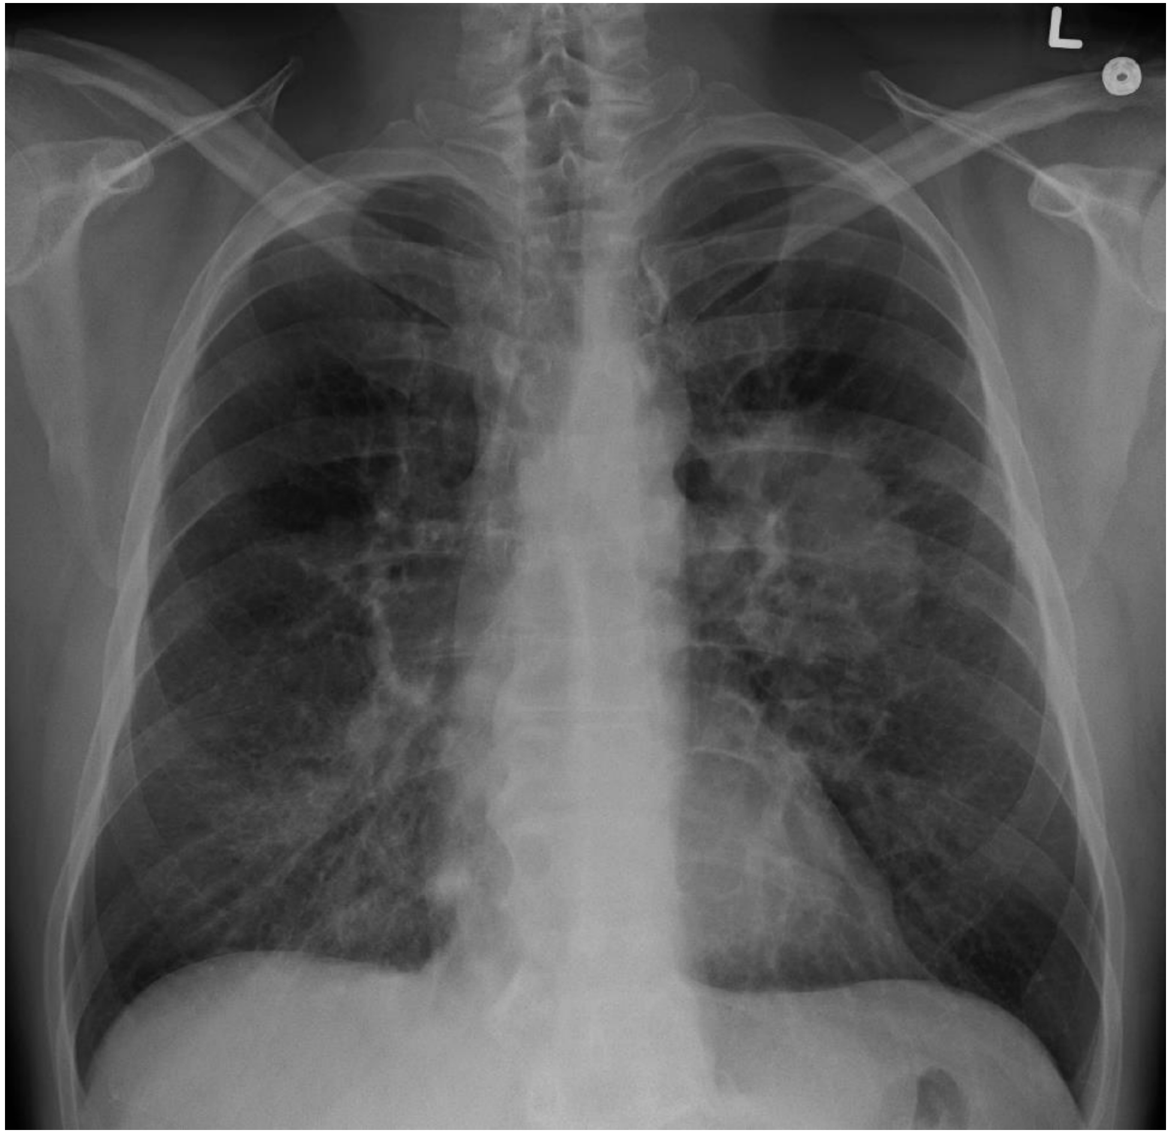

Our Running Example: Monitoring Nanofiber Production

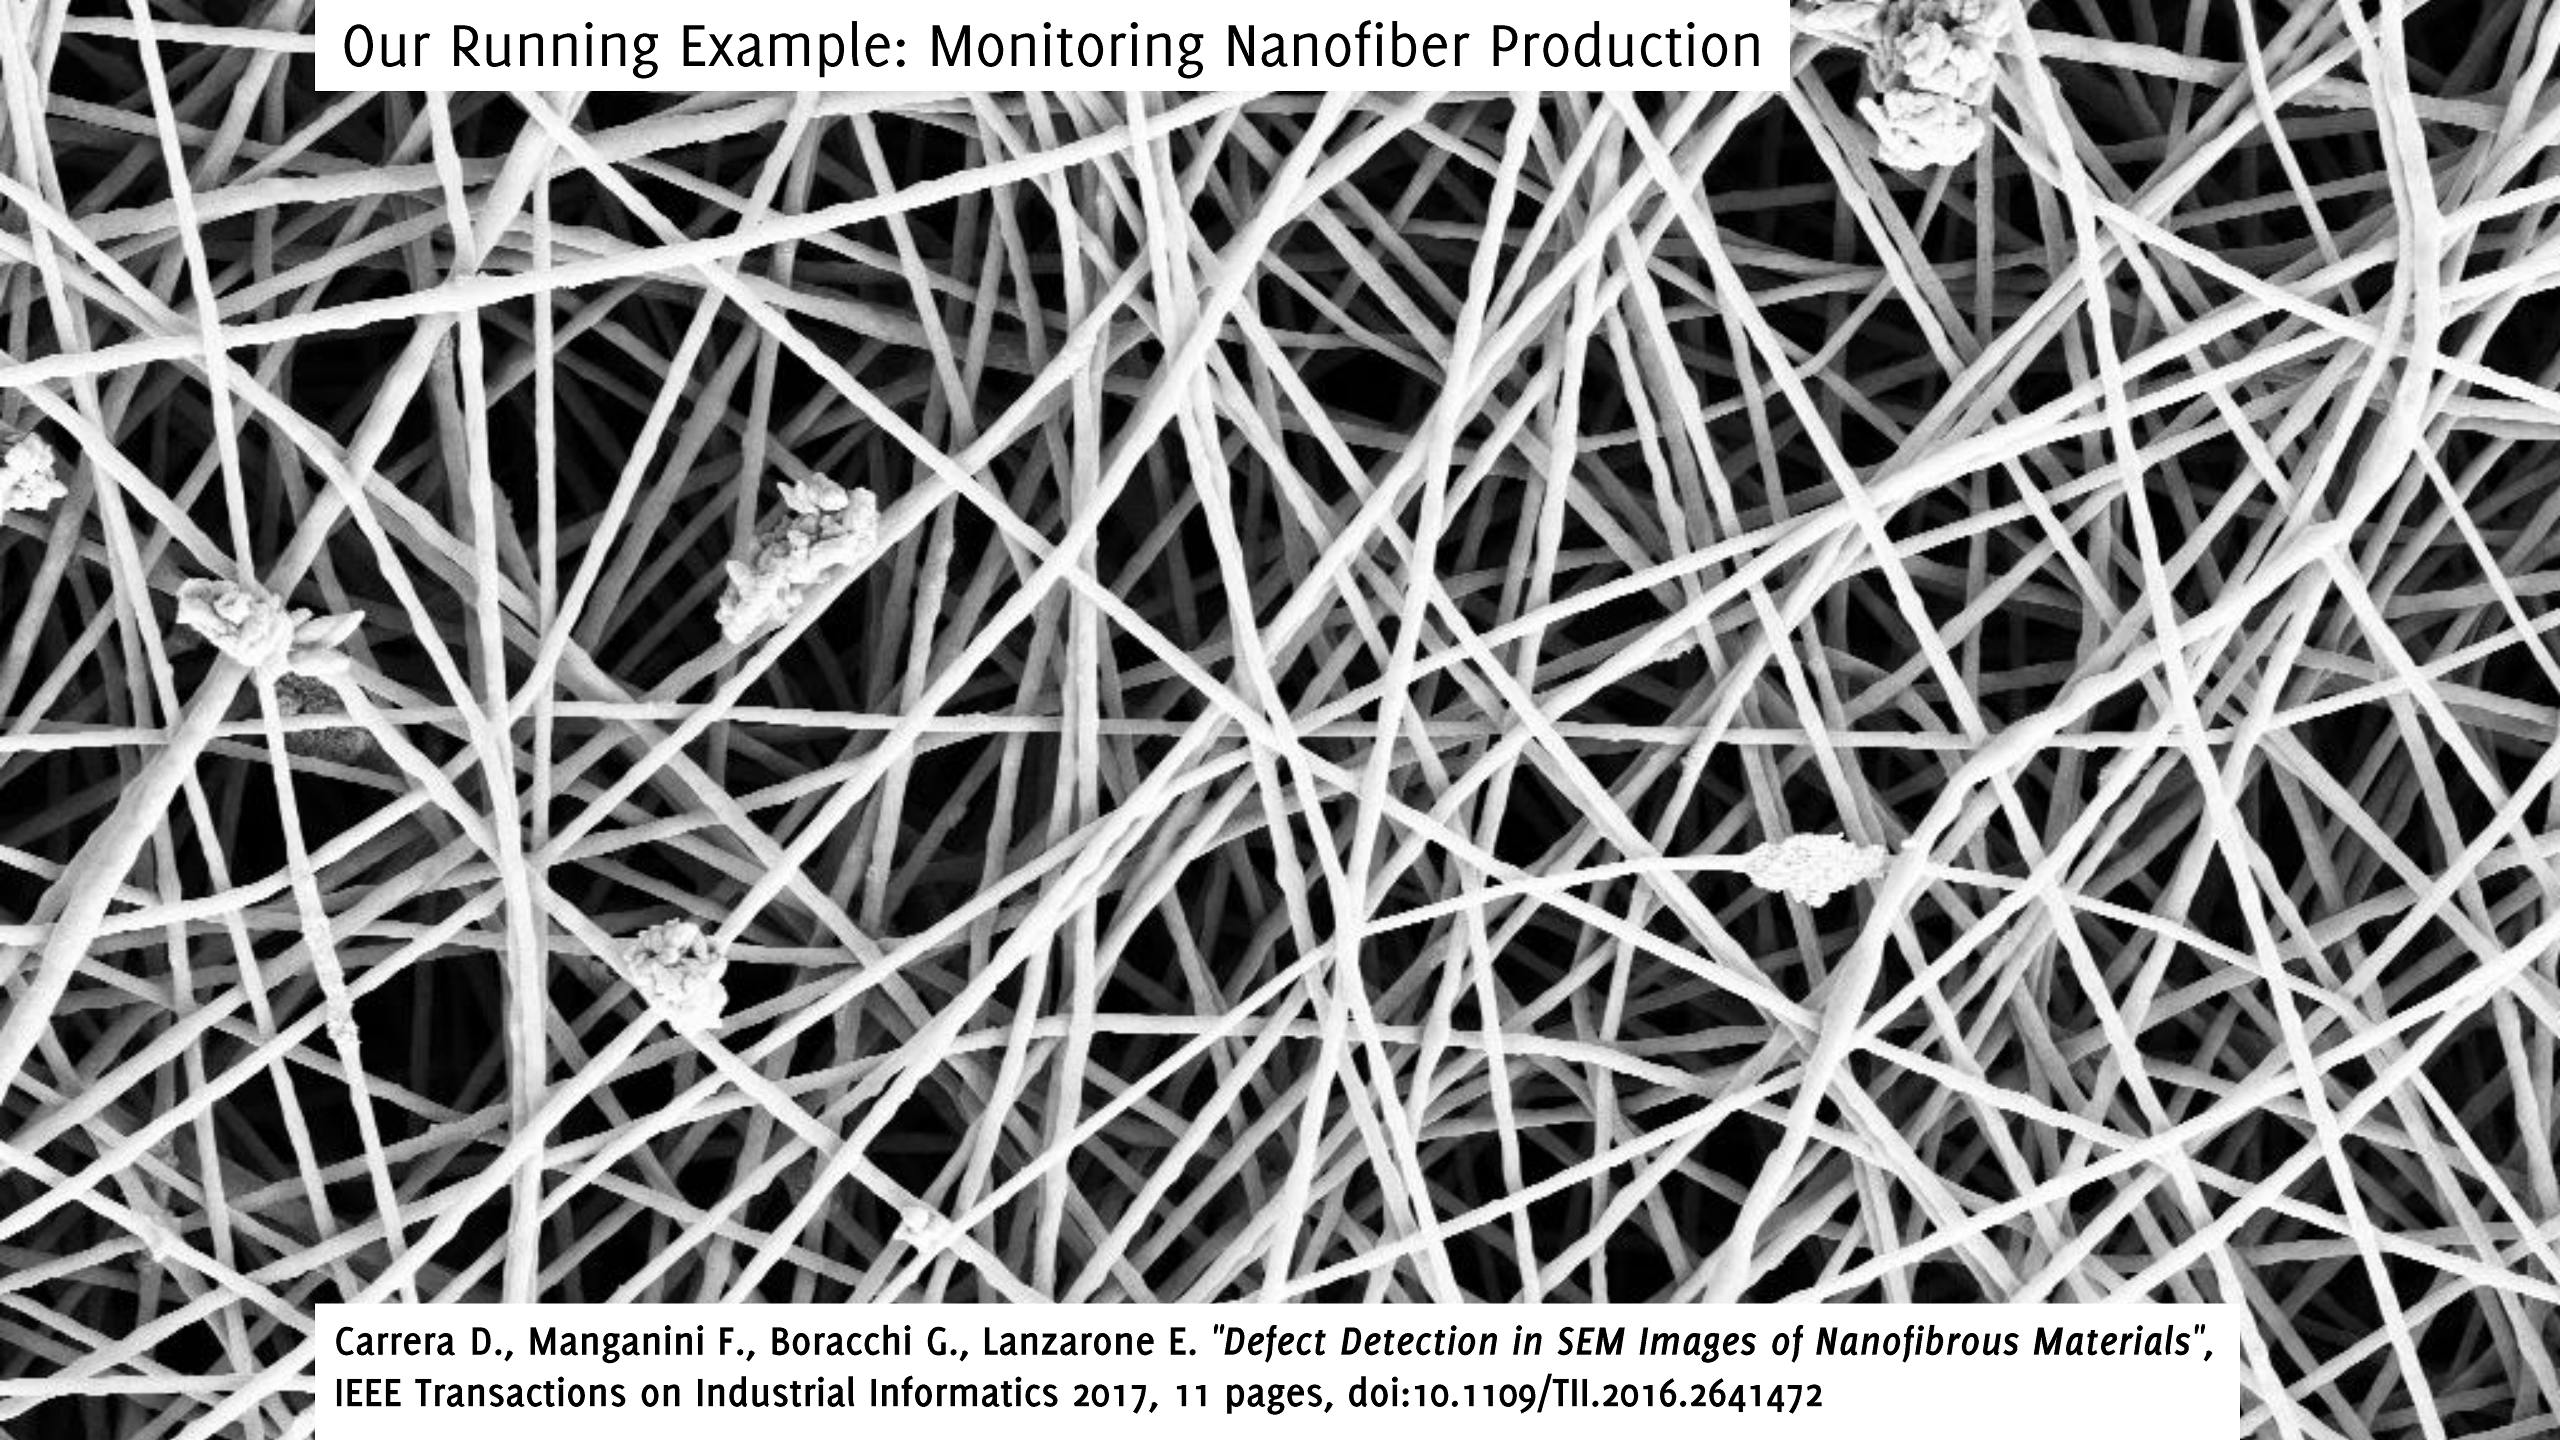

Carrera D., Manganini F., Boracchi G., Lanzarone E. *"Defect Detection in SEM Images of Nanofibrous Materials"*, IEEE Transactions on Industrial Informatics 2017, 11 pages, doi:10.1109/TII.2016.2641472

Our Running Example: Monitoring Nanofiber Production

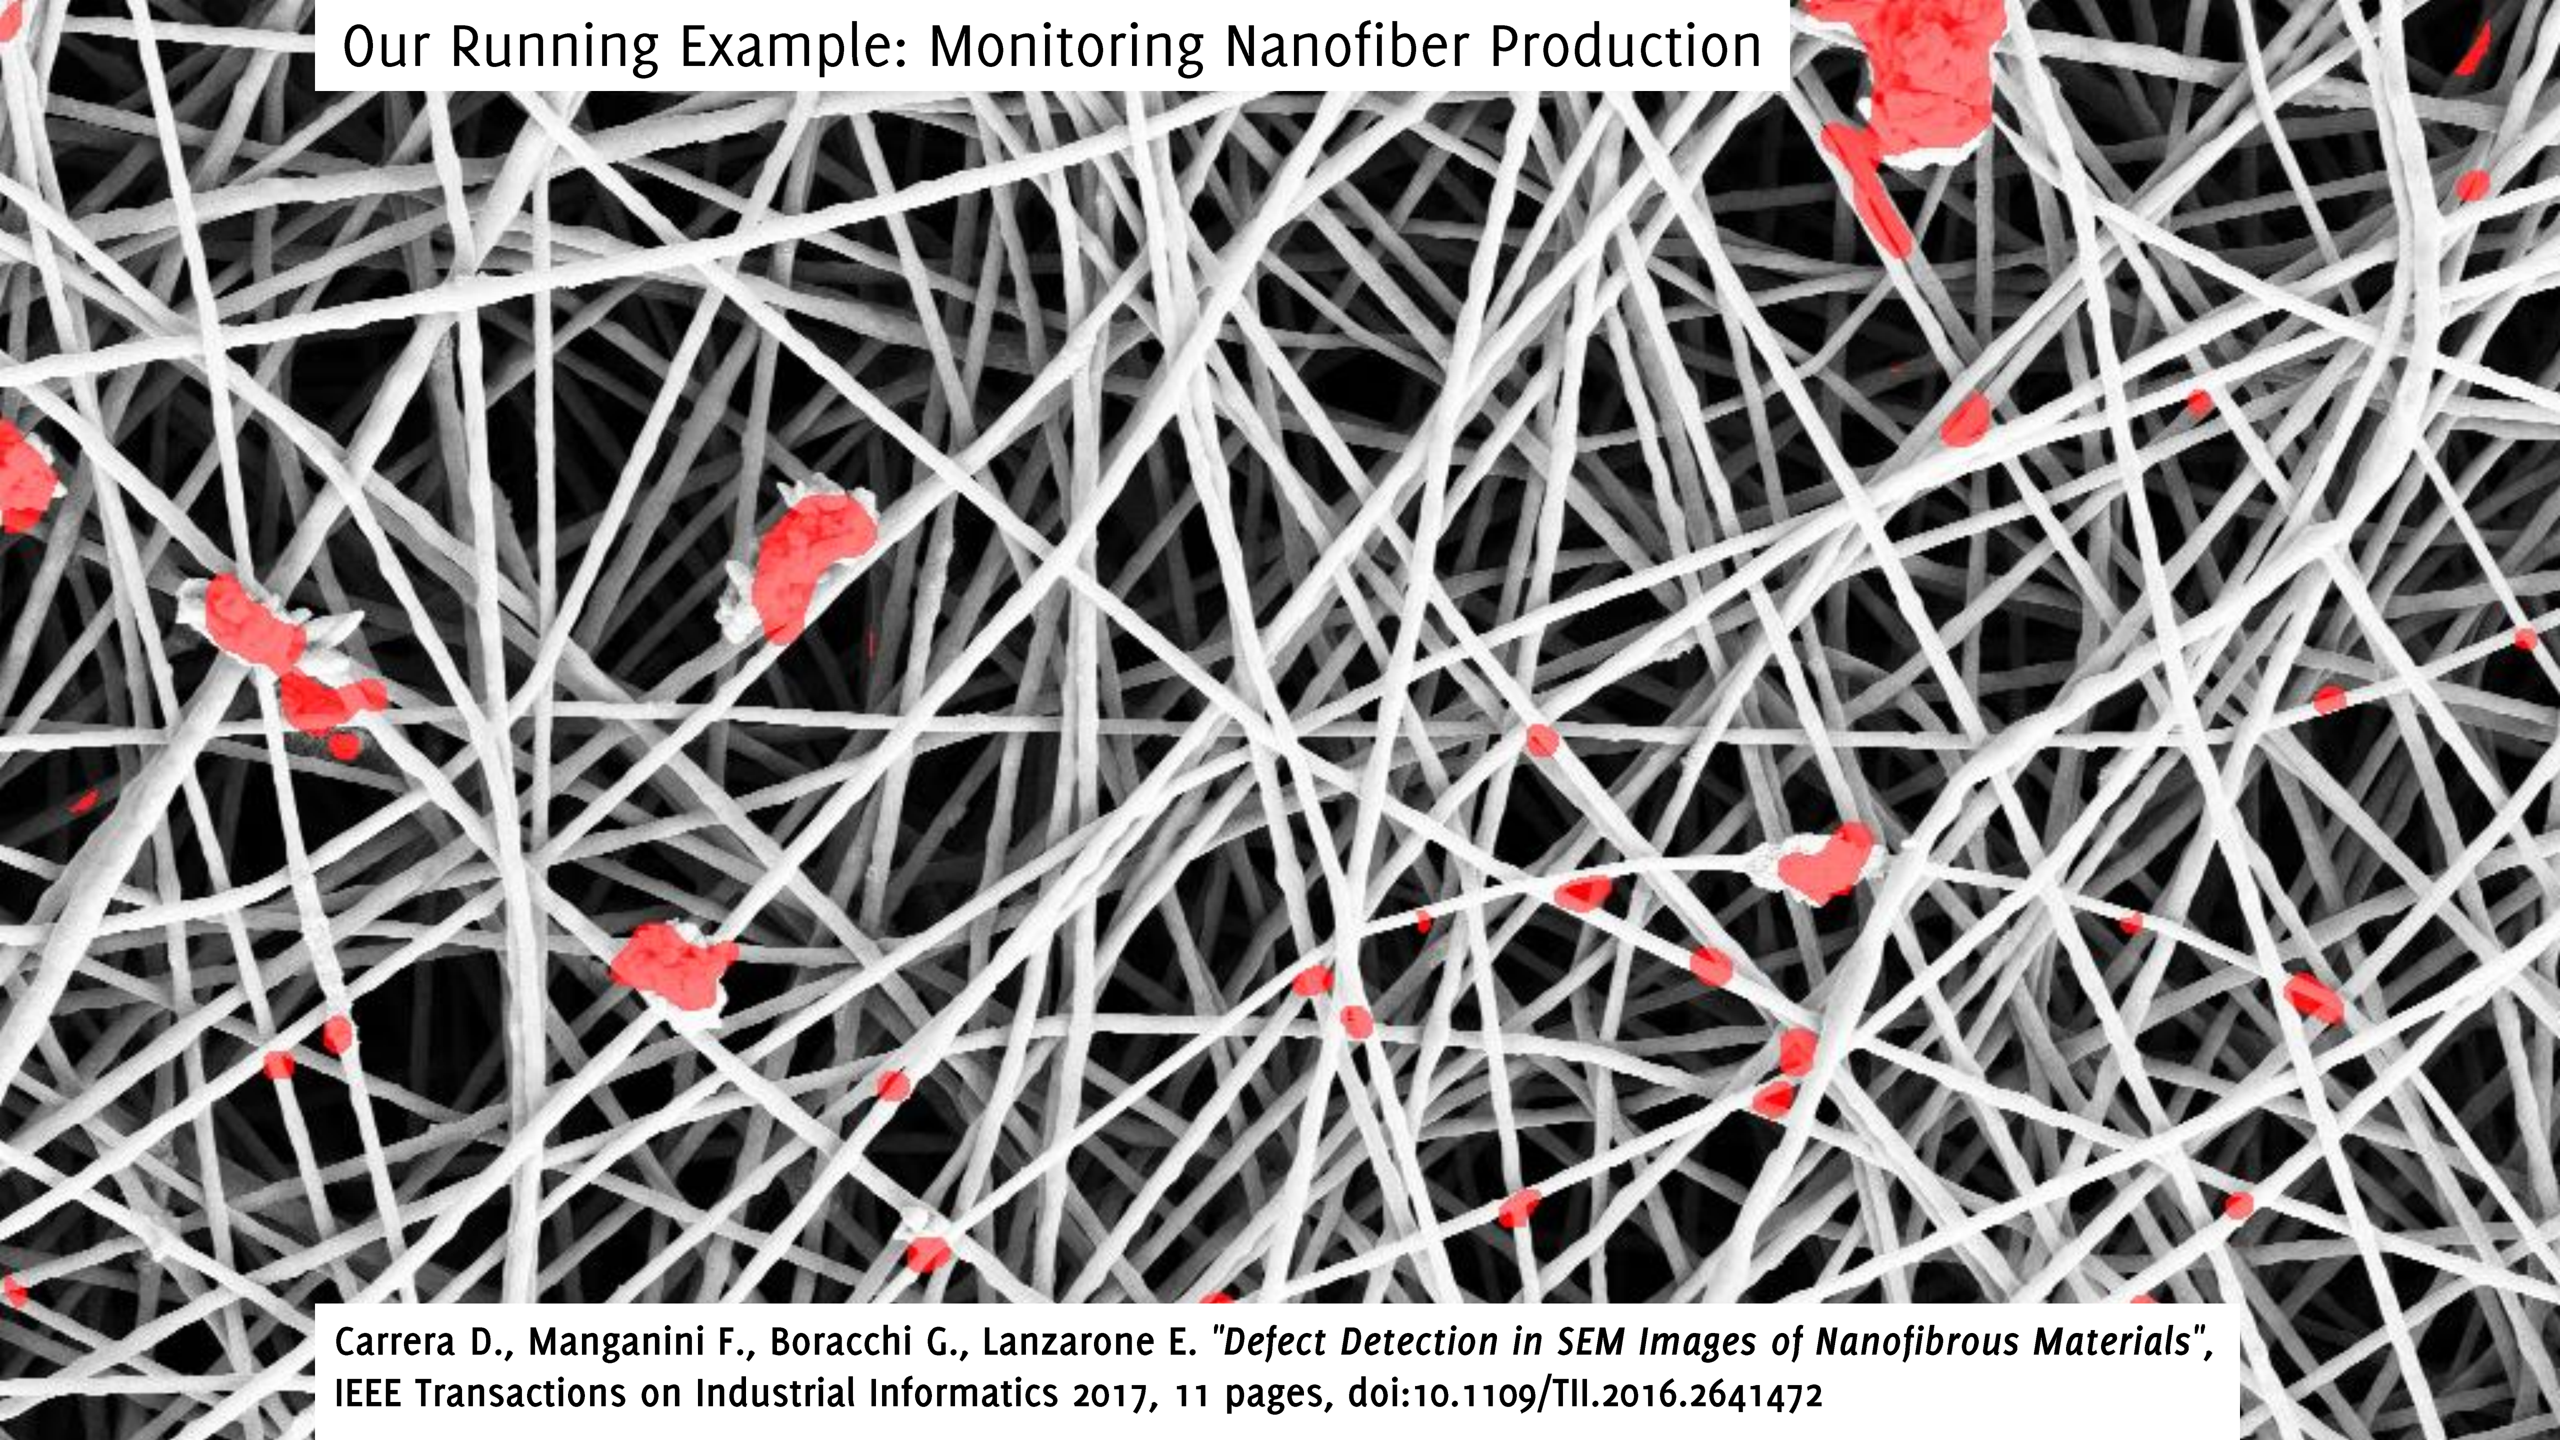

Carrera D., Manganini F., Boracchi G., Lanzarone E. *"Defect Detection in SEM Images of Nanofibrous Materials"*, IEEE Transactions on Industrial Informatics 2017, 11 pages, doi:10.1109/TII.2016.2641472

SYLICON WAFER MANUFACTURING

Defects detected as anomalies in microscope images

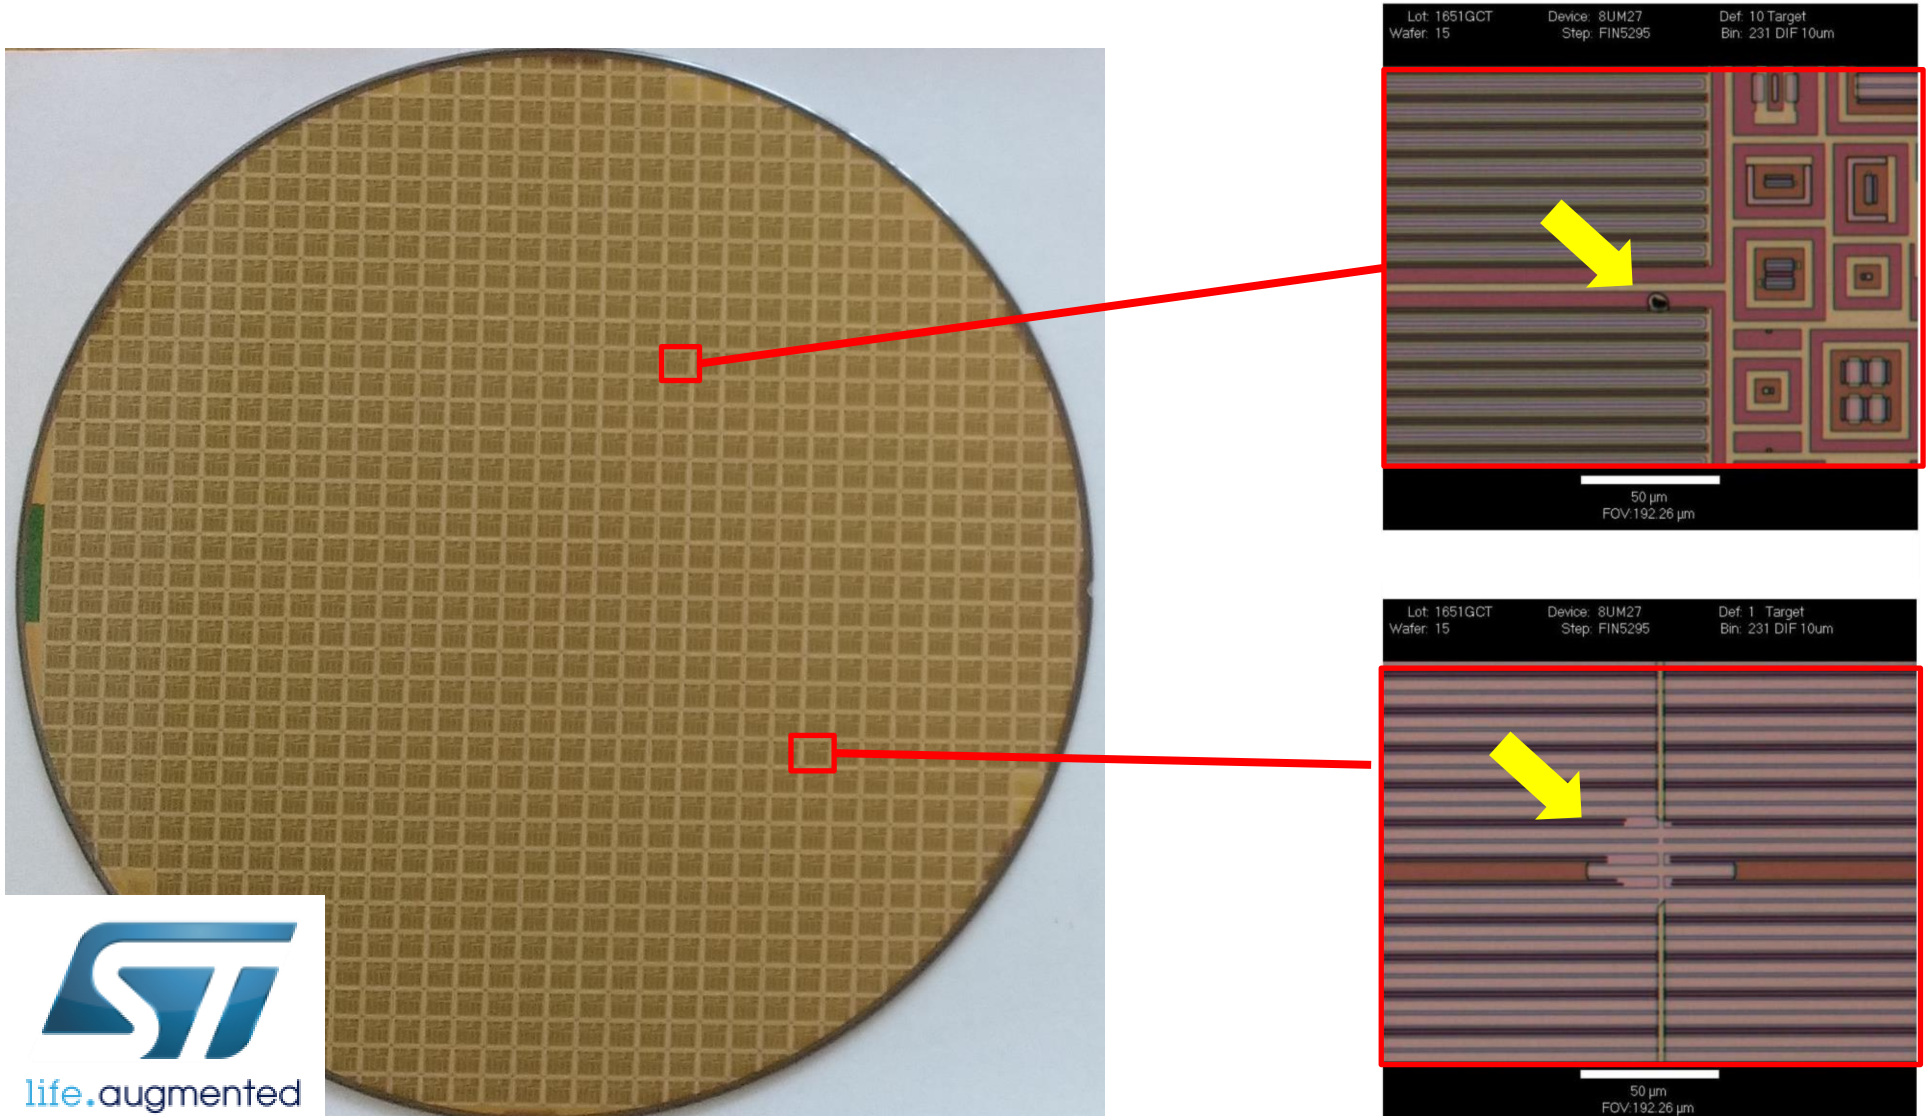

Automatic and long term Egc monitoring

Health monitoring / wearable devices:

Automatically analyze EGC tracings to detect arrhythmias or incorrect device positioning

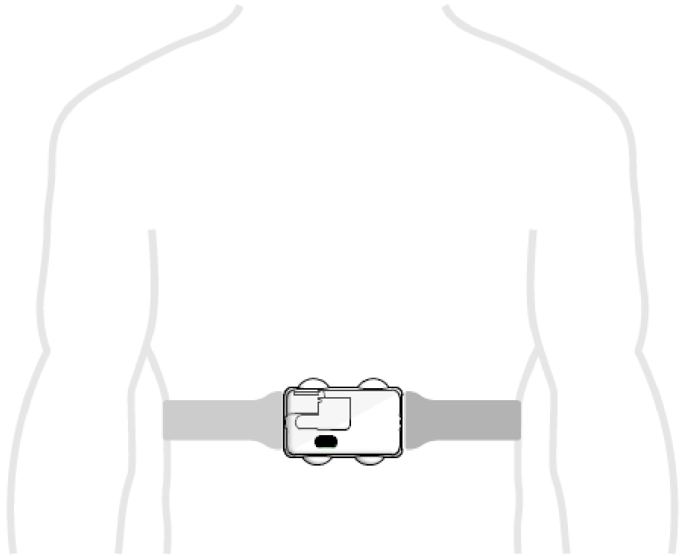

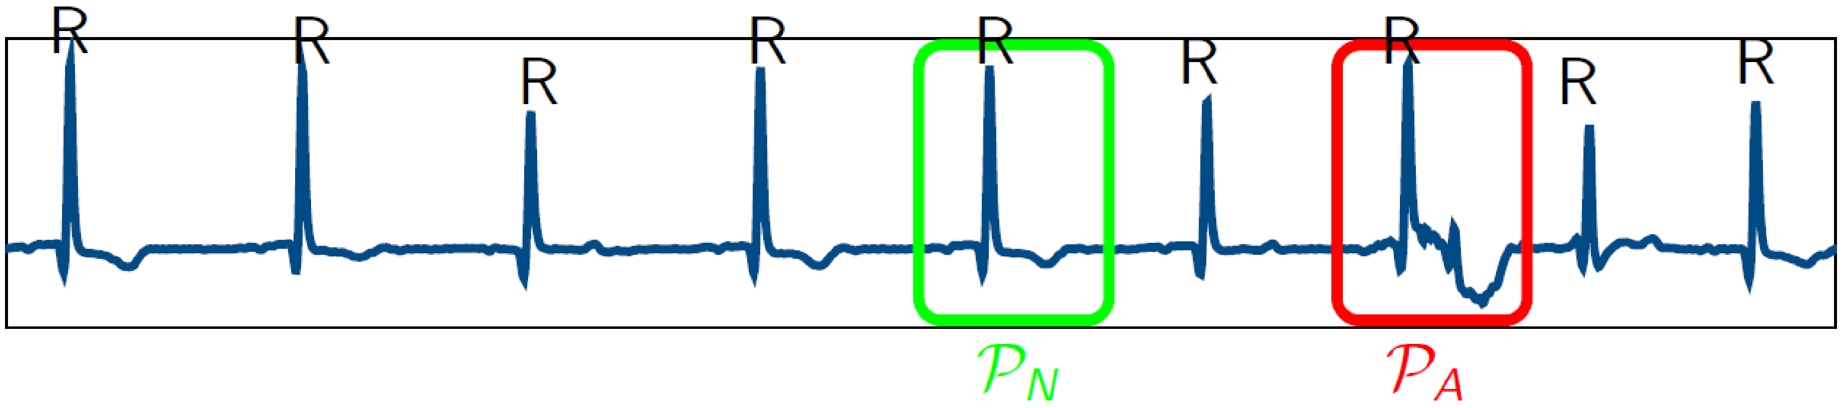

Carrera D., Rossi B., Fragneto P., Boracchi G. "*Online Anomaly Detection for Long-Term ECG Monitoring using Wearable Devices*", Pattern Recognition 2019

Anomalous activities detection in videos

USCD Anomaly Dataset http://www.svcl.ucsd.edu/projects/anomaly/dataset.html

Anomaly Detection for automatic quality control

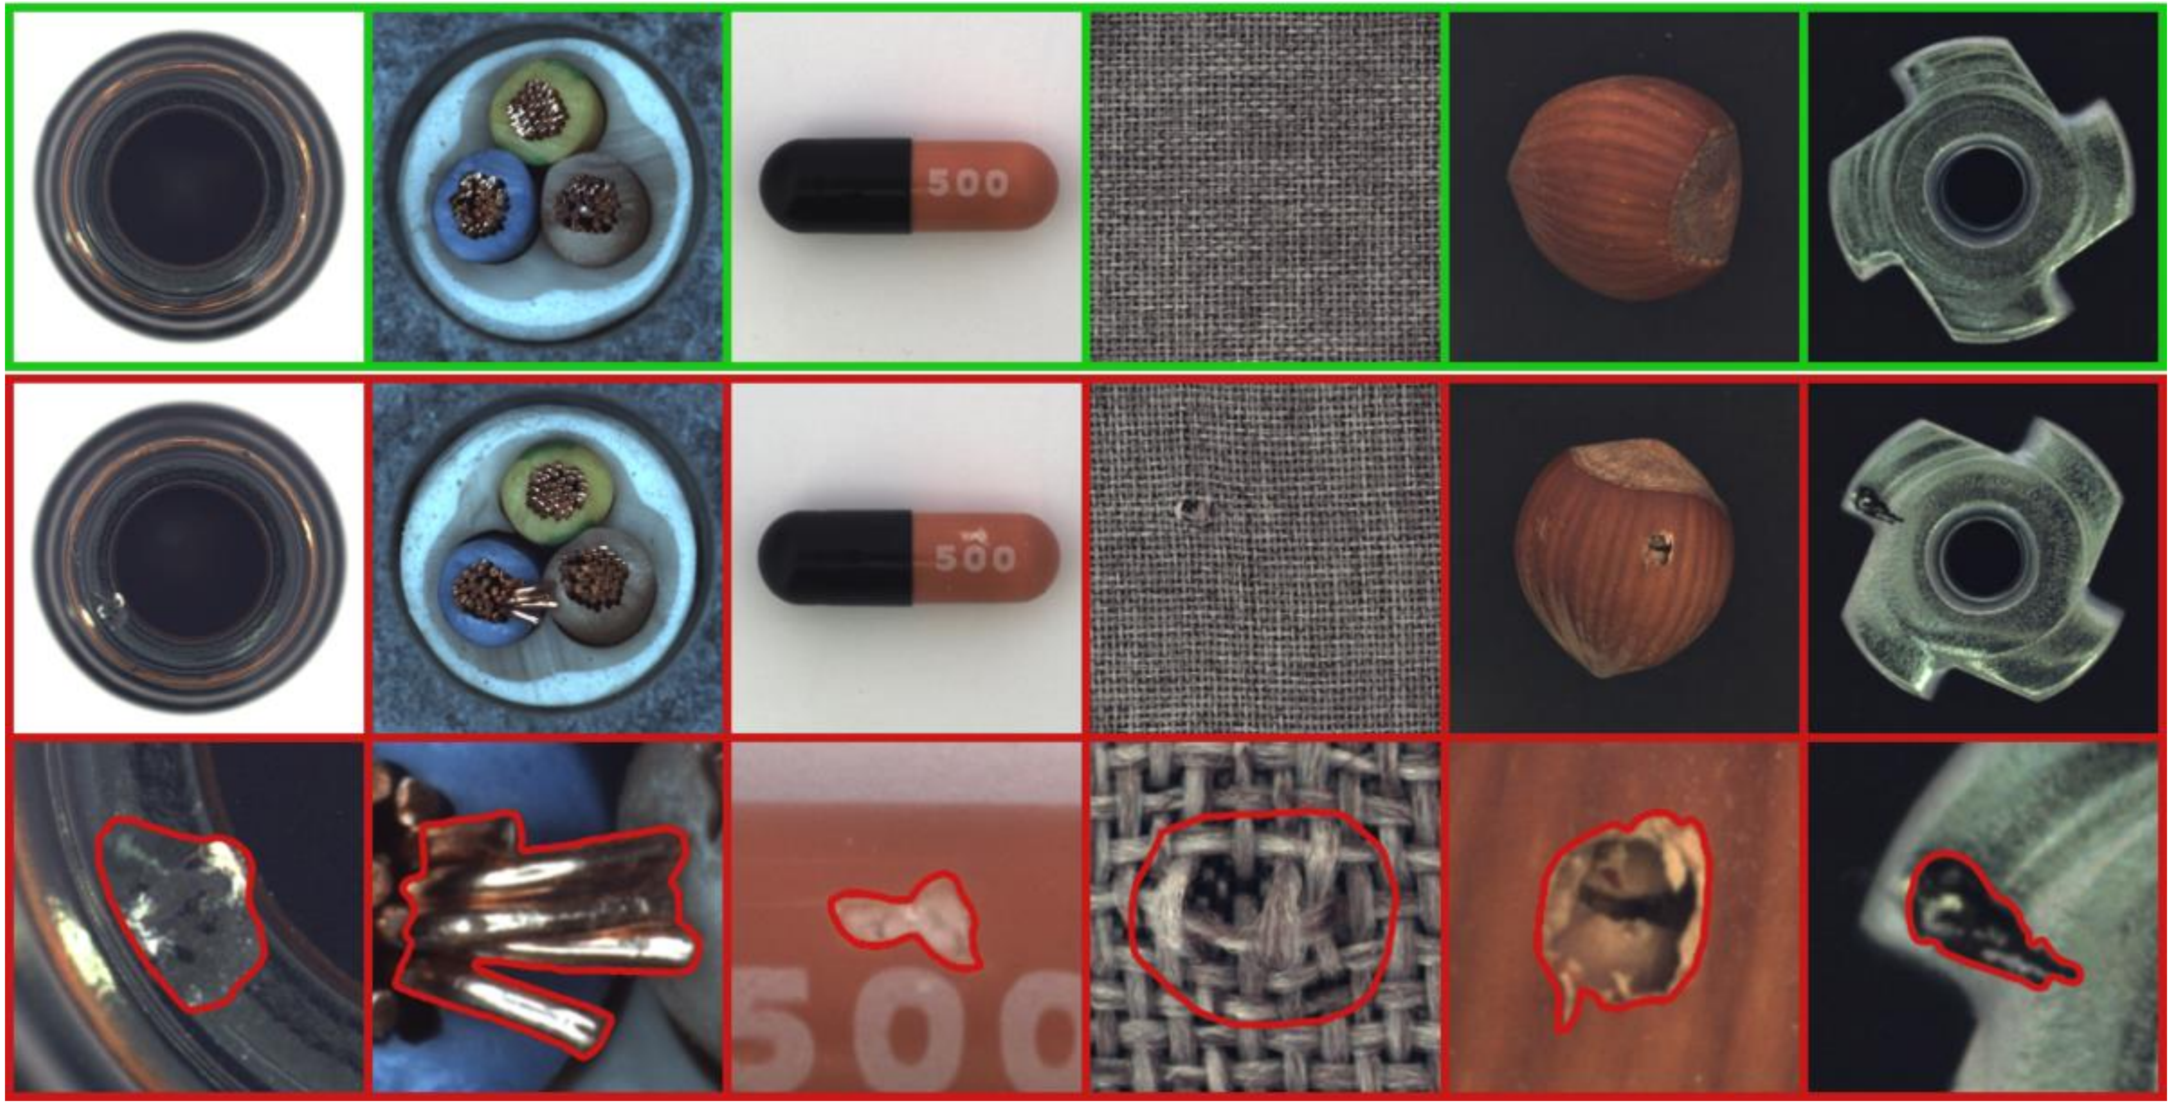

https://www.mvtec.com/company/research/datasets/mvtec-ad/

3D Anomaly Detection

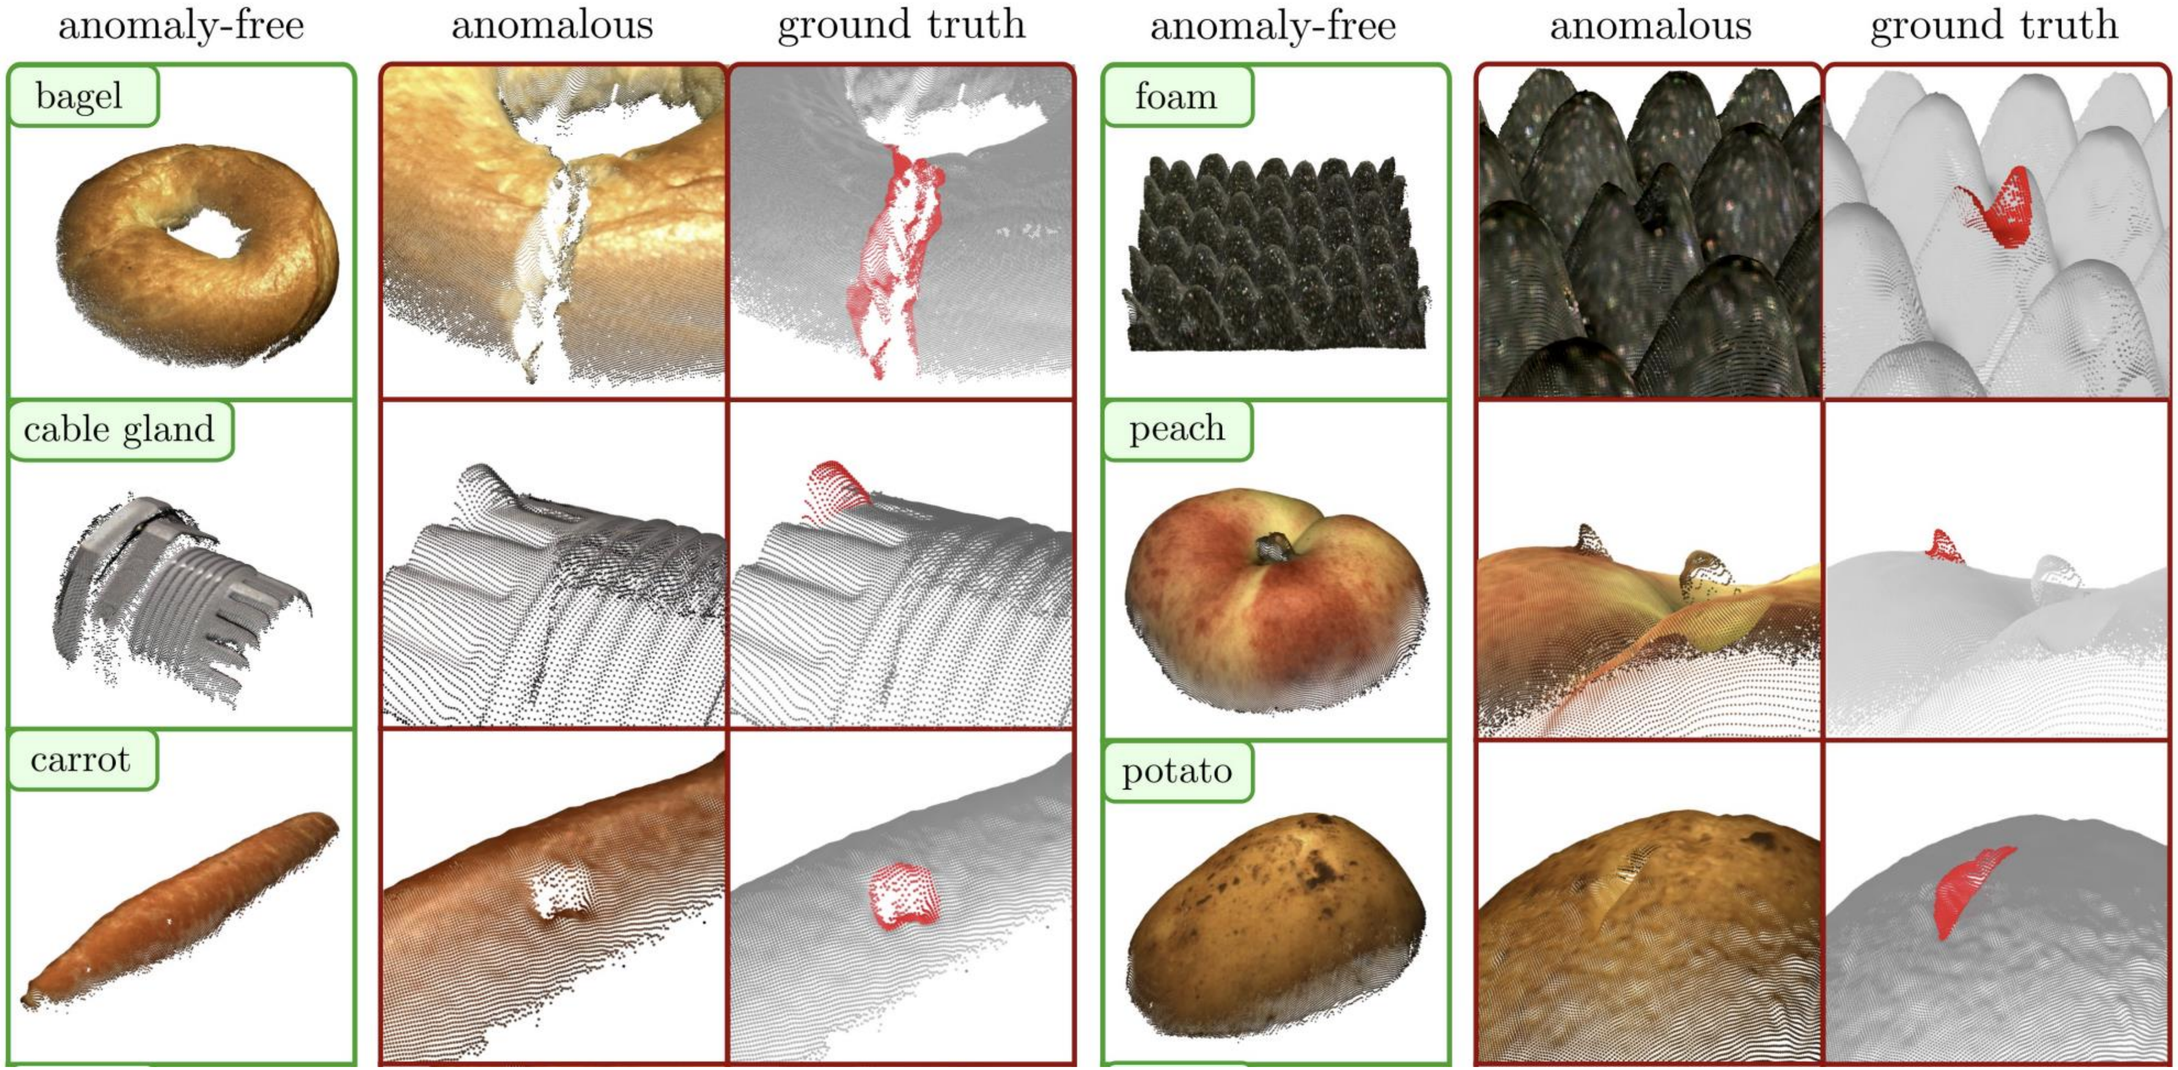

Bergmann, Paul, Xin Jin, David Sattlegger, and Carsten Steger. "The mvtec 3d-ad dataset for unsupervised 3d anomaly detection and localization." arXiv:2112.09045 (2021).

Logical constraints anomaly detection

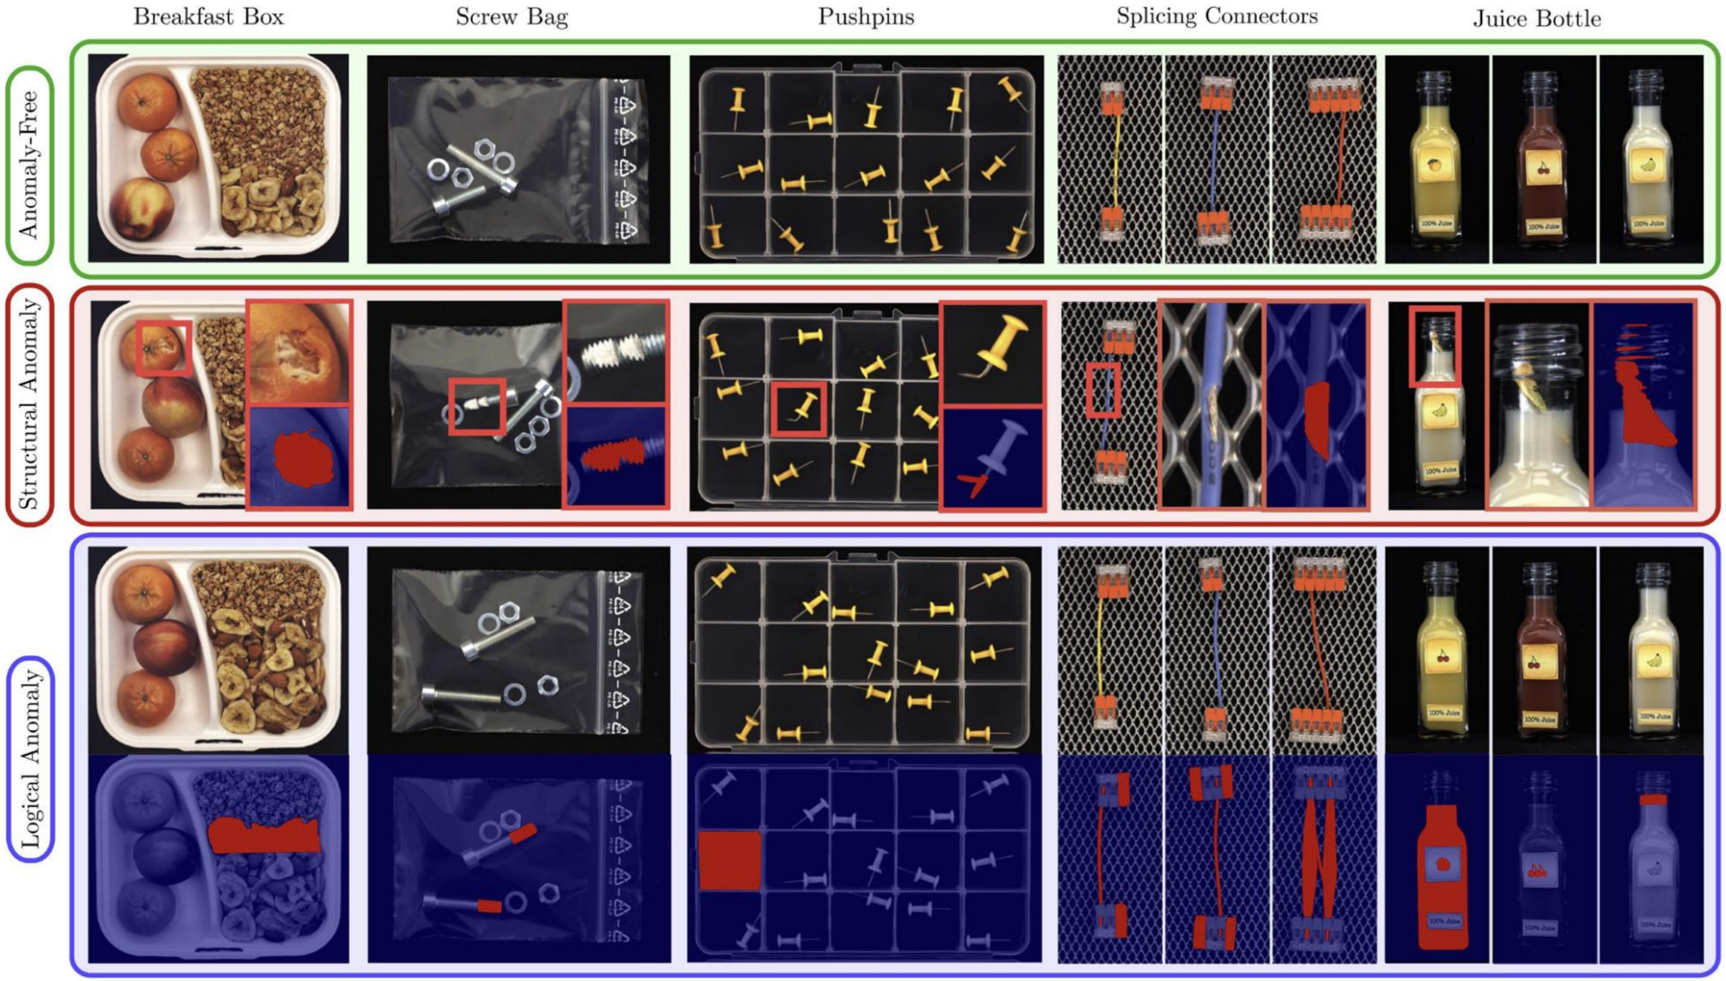

Fig. 3 Example images of the MVTec LOCO AD dataset for each of the five dataset categories. Each category contains anomaly-free train, validation, and test images. Additional test images contain various structural and logical anomalies. Pixel-precise ground truth annotations are provided for all anomalies

Bergmann, Paul, Kilian Batzner, Michael Fauser, David Sattlegger, and Carsten Steger. "Beyond dents and scratches: Logical constraints in unsupervised anomaly detection and localization." IJCV (2022): 947-969.

The Problem Formulation

Anomaly Detection in Images

Anomalies

"Anomalies are patterns in data that do not conform to a well defined notion of normal behavior"

Thus:

- **Normal data** are generated from a **stationary process** $\mathcal{P}_N$
- **Anomalies** are from a **different process** $\mathcal{P}_A \neq \mathcal{P}_N$

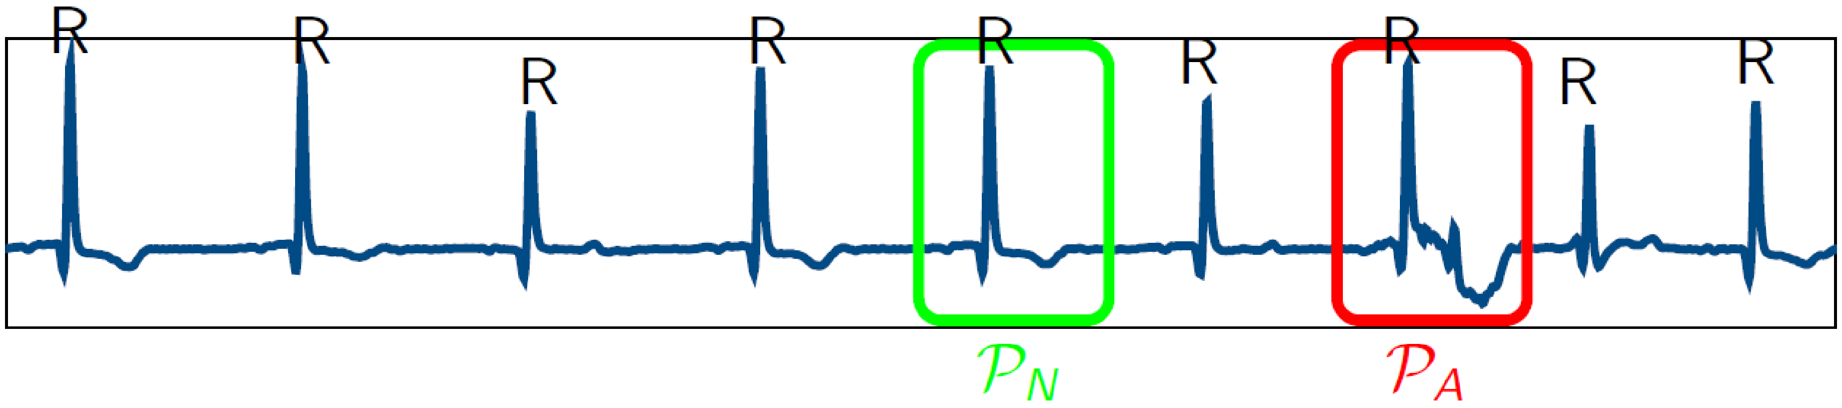

V. Chandola, A. Banerjee, V. Kumar. *"Anomaly detection: A survey"*. ACM Comput. Surv. 41, 3, Article 15 (July 2009), 58 pages.

Anomalies

"Anomalies are patterns in data that do not conform to a well defined notion of normal behavior"

Thus:

- **Normal data** are vectors drawn from a **stationary distribution** ϕ_0
- **Anomalies** are vectors drawn from a **different distribution** $\phi_1 \neq \phi_0$

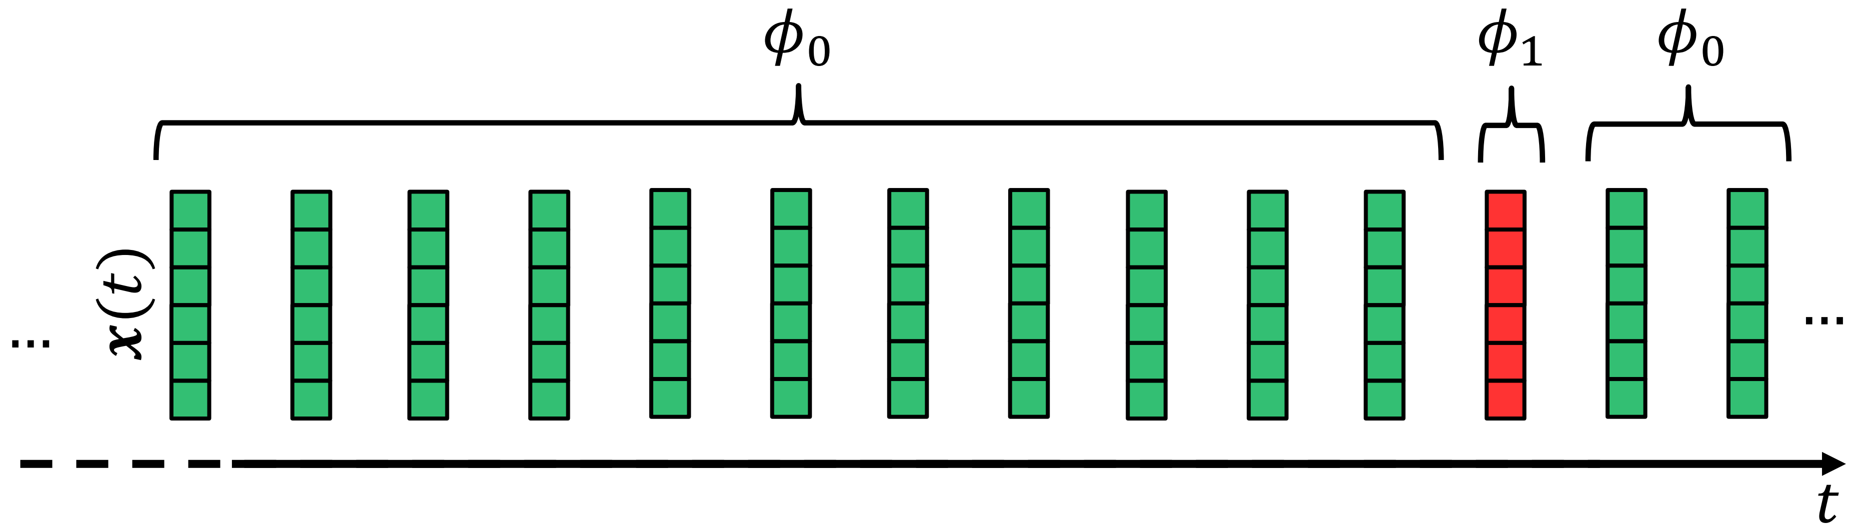

V. Chandola, A. Banerjee, V. Kumar. "*Anomaly detection: A survey*". ACM Comput. Surv. 41, 3, Article 15 (July 2009), 58 pages.

Anomalies

"Anomalies are patterns in data that do not conform to a well defined notion of normal behavior"

Thus:

- **Normal data** are generated from a **stationary process** $\mathcal{P}_N$
- **Anomalies** are from a **different process** $\mathcal{P}_A \neq \mathcal{P}_N$

Examples of Anomalies:

- **Frauds** in the stream of all the credit card transactions
- **Arrhythmias** in ECG tracings
- **Defective regions in an image**, which do not conform a reference pattern
- Anomalies might appear as **spurious** elements, and are typically the **most informative** samples in the stream

V. Chandola, A. Banerjee, V. Kumar. "*Anomaly detection: A survey*". ACM Comput. Surv. 41, 3, Article 15 (July 2009), 58 pages.

Problem formulation: Anomaly Detection in images

Let s be an image defined over the pixel domain $\mathcal{X} \subset \mathbb{Z}^2$,

let $c \in \mathcal{X}$ be a pixel and $s(c)$ the corresponding intensity.

Our goal is to **locate any anomalous region** in s, i.e. **estimating the unknown anomaly mask $\boldsymbol{\Omega}$** defined as

$$\Omega(c) = \begin{cases} 0 & \text{if } c \text{ falls in a normal region} \\ 1 & \text{if } c \text{ falls in an anomalous region} \end{cases}$$

s

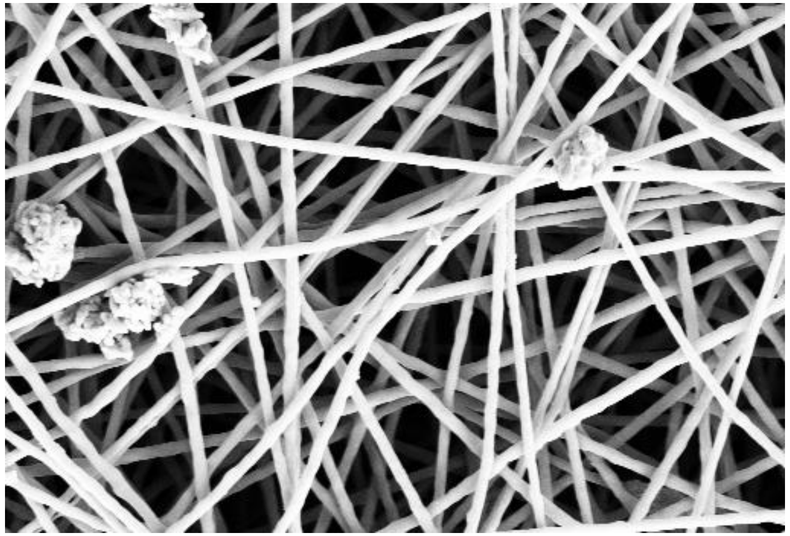

Problem formulation: Anomaly Detection in images

Let s be an image defined over the pixel domain $\mathcal{X} \subset \mathbb{Z}^2$,

let $c \in \mathcal{X}$ be a pixel and $s(c)$ the corresponding intensity.

Our goal is to **locate any anomalous region** in s, i.e. **estimating the unknown anomaly mask $\boldsymbol{\Omega}$** defined as

$$\Omega(c) = \begin{cases} 0 & \text{if } c \text{ falls in a normal region} \\ 1 & \text{if } c \text{ falls in an anomalous region} \end{cases}$$

s

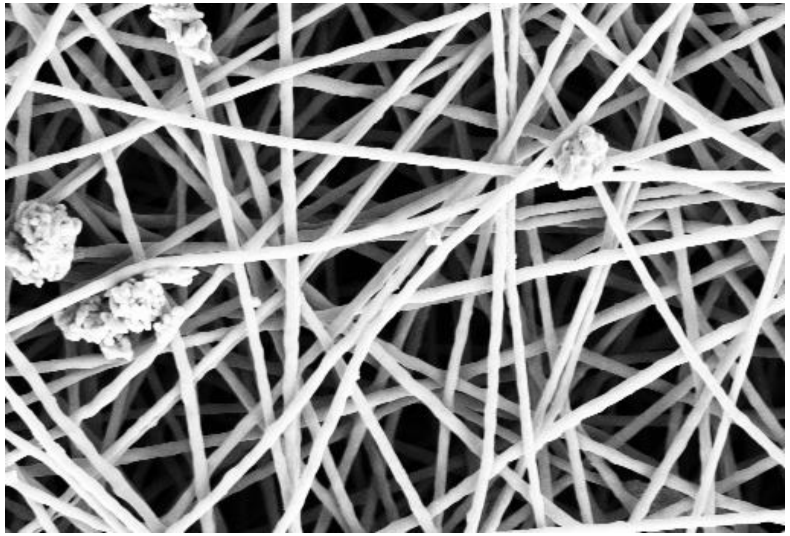

Ω

Patch-wise Anomaly detection

The goal not determining whether the whole image is normal or anomalous, but **locate/segment possible anomalies**

Therefore, it is convenient to

1. **Analyze the image patch-wise**
2. Isolate regions containing patches that are detected as as anomalies

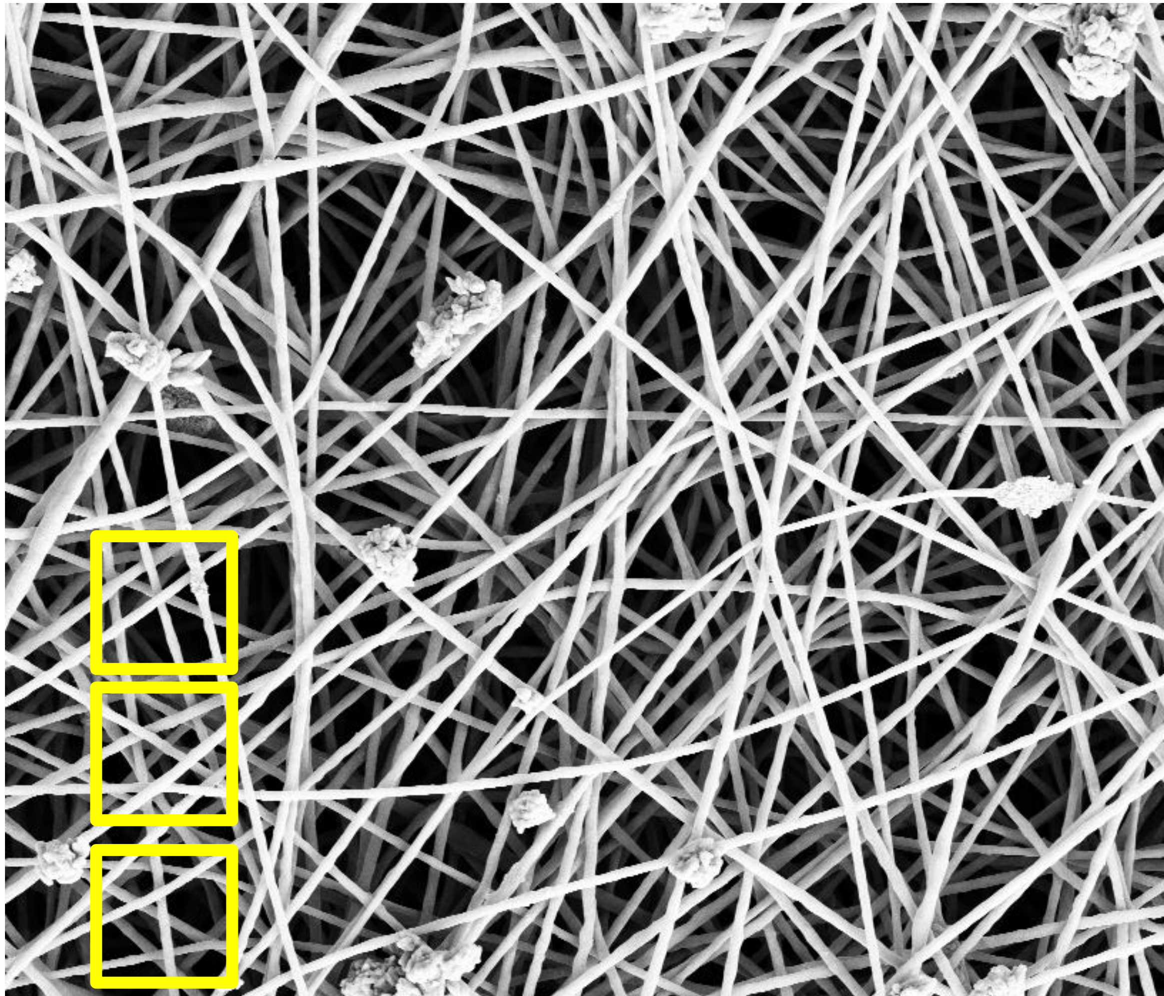

Patch-wise Anomaly detection

The goal not determining whether the whole image is normal or anomalous, but **locate/segment possible anomalies**

Therefore, it is convenient to

1. **Analyze the image patch-wise**
2. Isolate regions containing patches that are detected as as anomalies

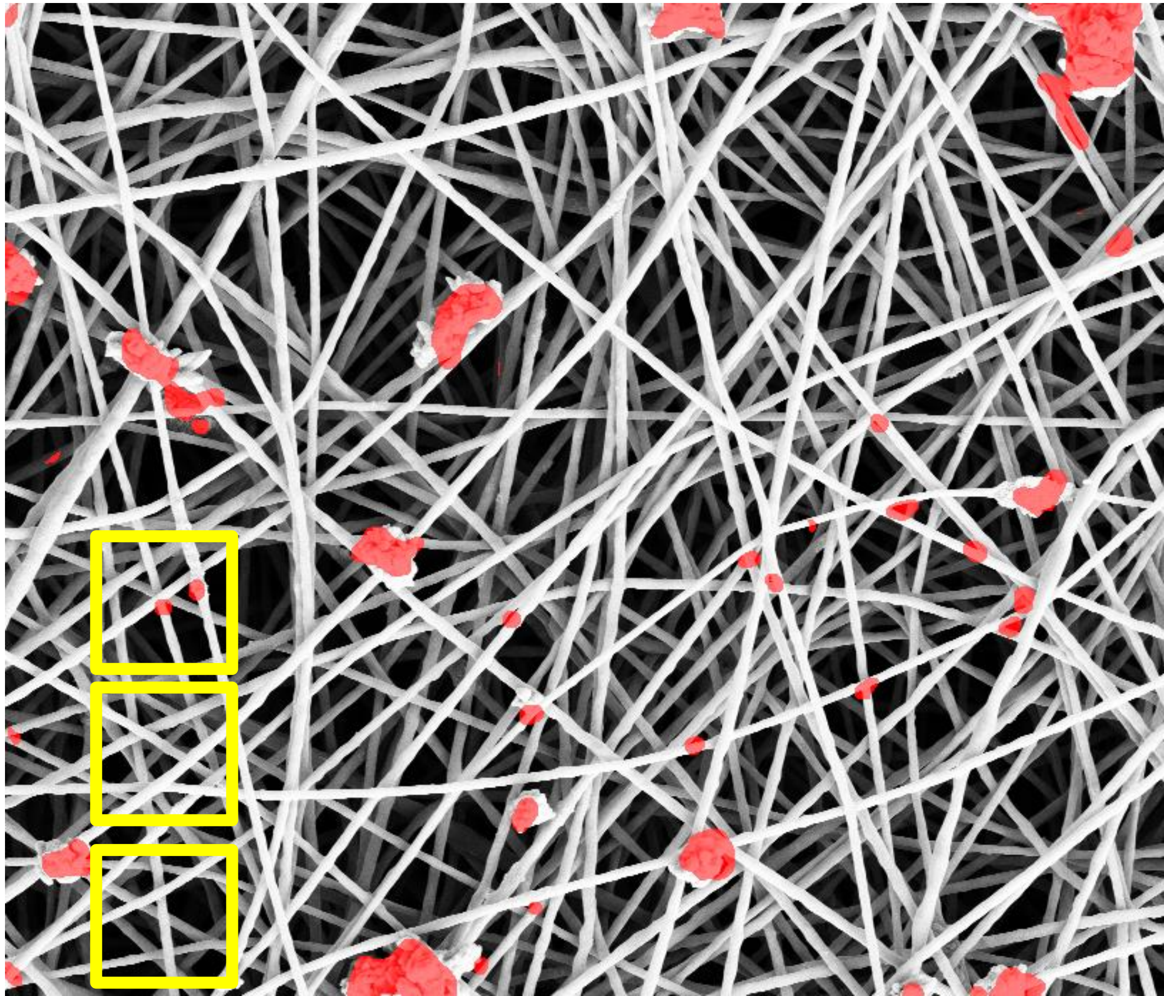

The Mainstream Solutions

Anomaly Detection in Images

The three major ingredients

Most detection algorithms have three major ingredients:

- The **background model** $\mathcal{M}$, learned from normal data
- The **statistic / anomaly score**: $\mathrm{err}(\boldsymbol{s}), \mathcal{L}(\boldsymbol{s}), \mathcal{A}(\boldsymbol{s}), \dots$
- **Decision rule** to detect, e.g. $\mathrm{err}(\boldsymbol{s}) \gtrless \gamma$ possibly controlling the FPR, as in other statistical detection methods

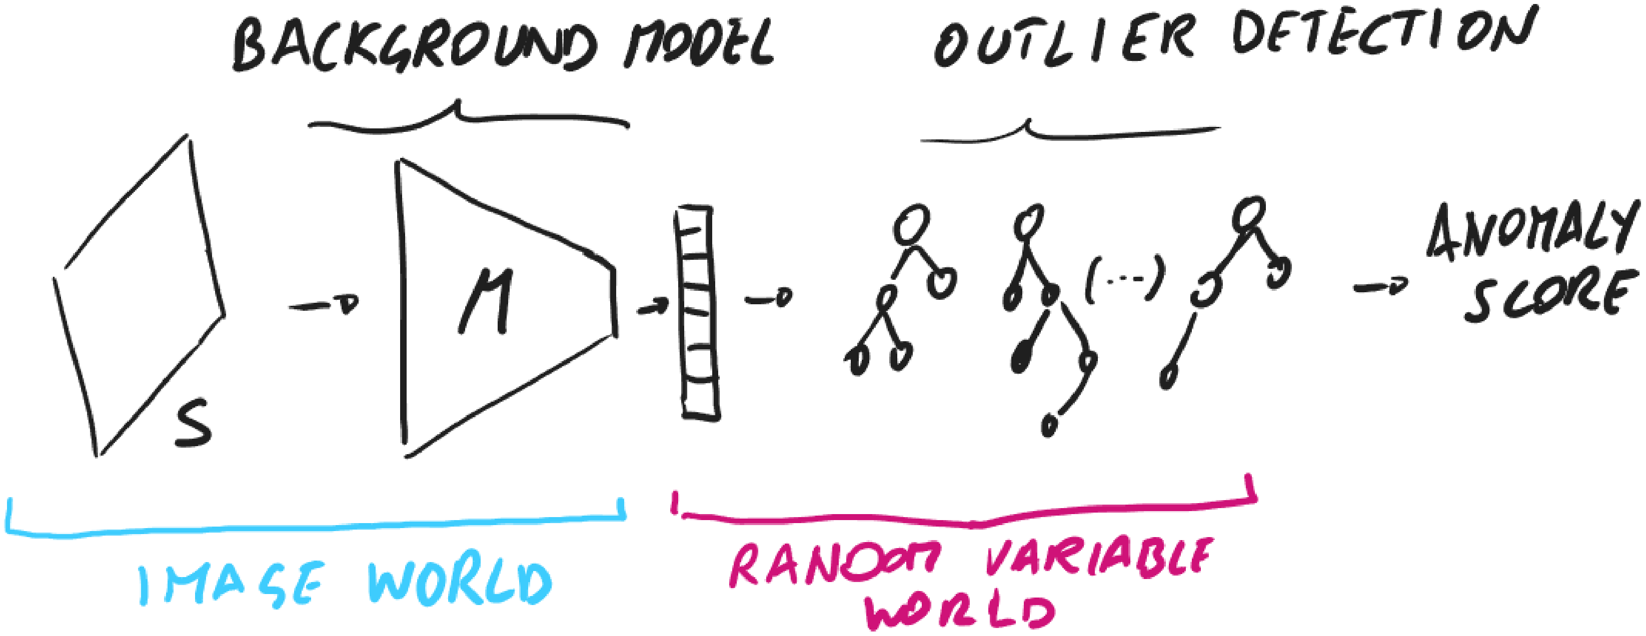

The three major ingredients

Most detection algorithms have three major ingredients:

- The **background model** $\mathcal{M}$, learned from normal data
- The **statistic / anomaly score**: $\mathrm{err}(\boldsymbol{s}), \mathcal{L}(\boldsymbol{s}), \mathcal{A}(\boldsymbol{s}), \dots$
- **Decision rule** to detect, e.g. $\mathrm{err}(\boldsymbol{s}) \gtrless \gamma$ possibly controlling the FPR, as in other statistical detection methods

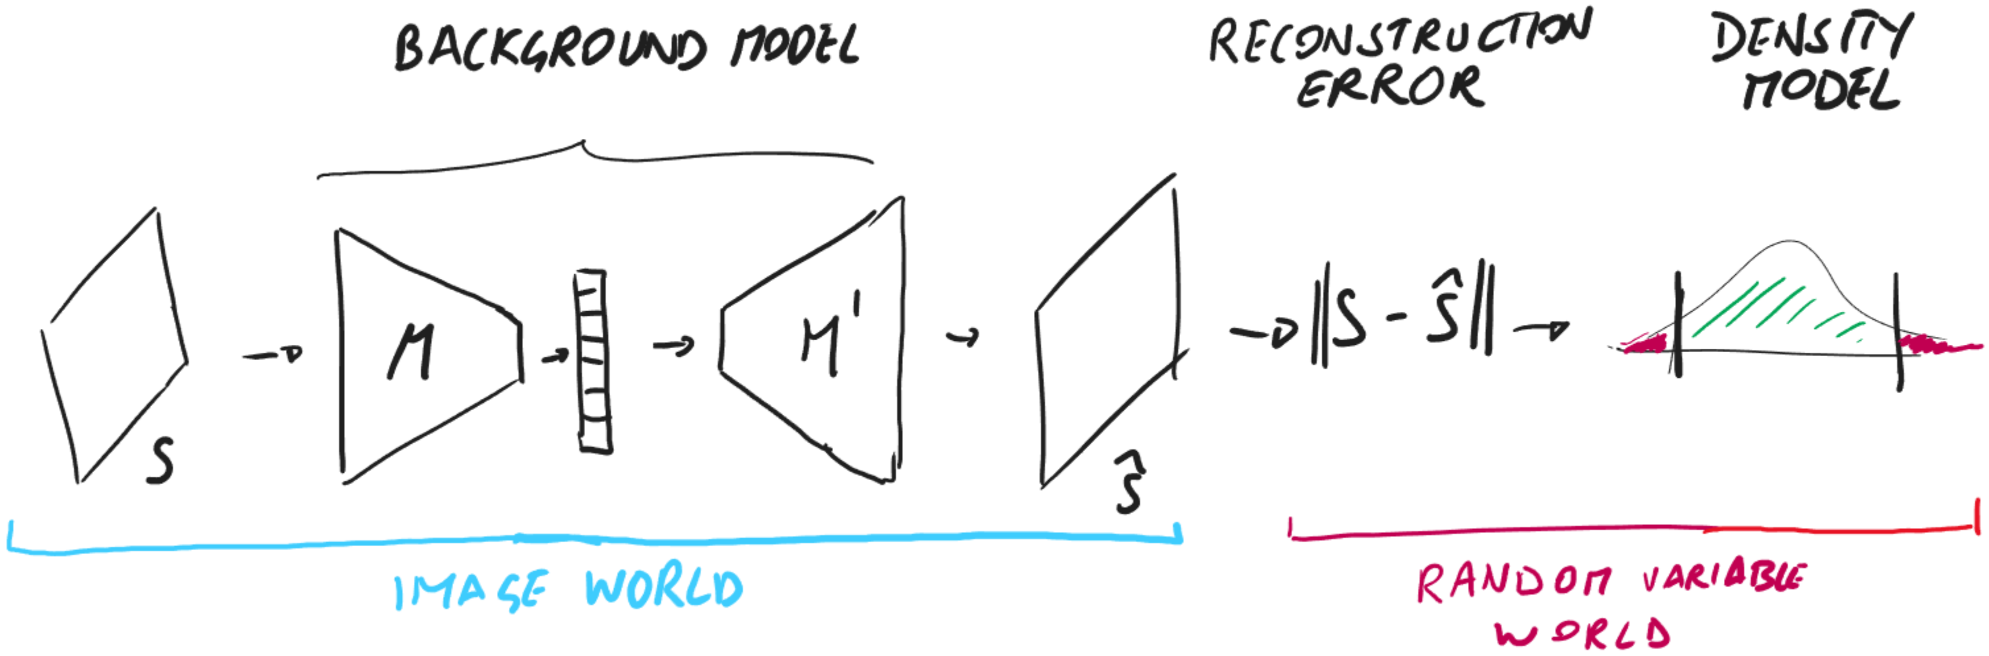

Diego Carrera

Anomaly Detection By Sparse Representations

Diego Carrera, Beatrice Rossi, Pasqualina Fragneto, Giacomo Boracchi *"Online Anomaly Detection for Long-Term ECG Monitoring using Wearable Devices"*, Pattern Recognition, 2019

Marco Longoni, Diego Carrera, Beatrice Rossi, Pasqualina Fragneto, Marco Pessione, Giacomo Boracchi, "*A Wearable Device for Online and Long-Term ECG Monitoring*" IJCAI 2018 - Demo Track

Carrera D., Rossi B., Fragneto P., and Boracchi G. *"Domain Adaptation for Online ECG Monitoring"* ICDM 2017,

Diego Carrera, Fabio Manganini, Giacomo Boracchi, Ettore Lanzarone, *"Defect Detection in SEM Images of Nanofibrous Materials"* IEEE TII, 2017

Diego Carrera, Giacomo Boracchi, Alessndro Foi and Brendt Wohlberg *"Scale-invariant Anomaly Detection With Multiscale Group-sparse Models"* IEEE ICIP 2016

Dictionaries Yielding Sparse Representations

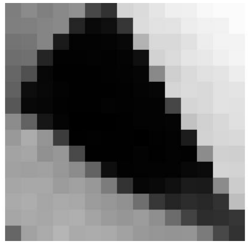

s

Dictionaries are just matrices! $D \in \mathbb{R}^{n \times m}$

Each column is called *an atom:*

- lives in the input space
- it is one of the learned building blocks to reconstruct the input signal

A good background model

- Unsupervised models
- Easy to plug in a **change/anomaly detection** framework
- Easy to **adapt**
- Simple and interpretable

$D =$

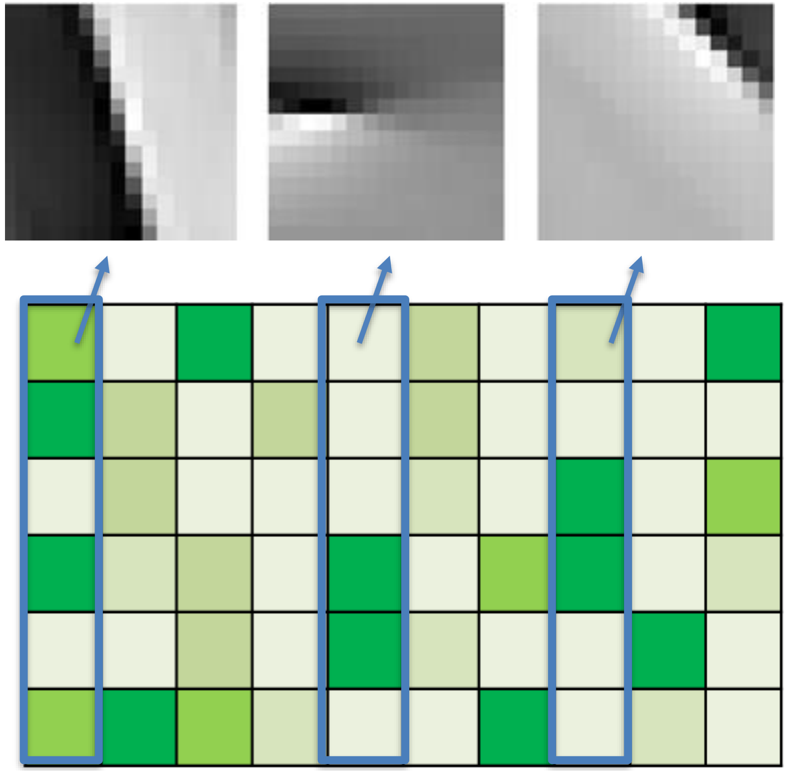

Sparse Representations

Let $\boldsymbol{s} \in \mathbb{R}^n$ be the input signal, a **sparse representation** is

$$\boldsymbol{s} = \sum_{i=1}^{M} x_i \, \boldsymbol{d_i},$$

A sparse representation is a **linear combination** of **few dictionary atoms** $\{\boldsymbol{d_i}\}$

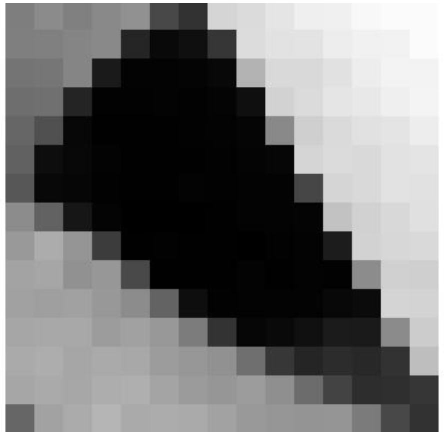

$= 0.7 \ *$

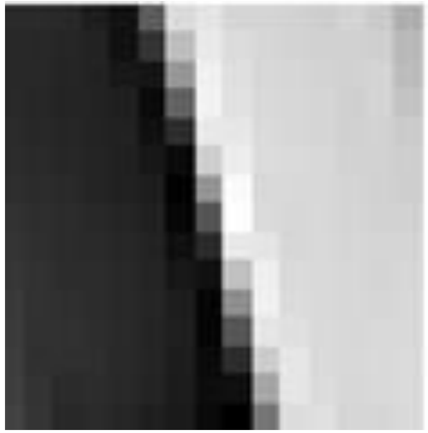

$+0.1 \ *$

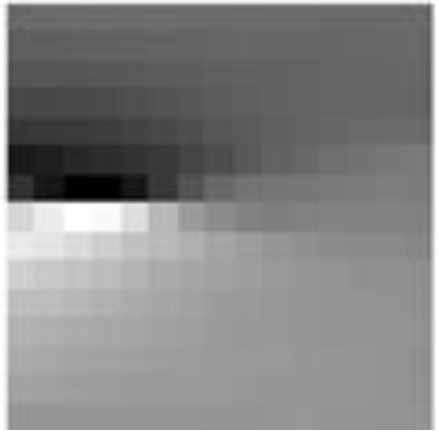

$-0.2 \ *$

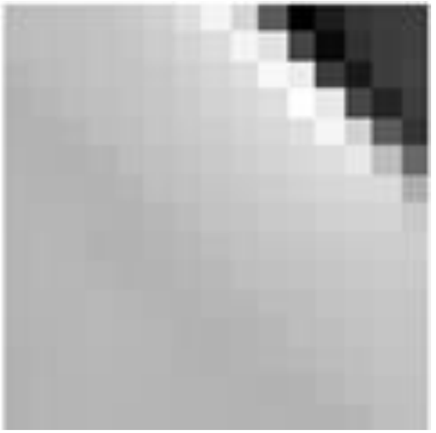

Sparse Representations

Let $\boldsymbol{s} \in \mathbb{R}^n$ be the input signal, a **sparse representation** is

$$\boldsymbol{s} = \sum_{i=1}^{M} x_i \, \boldsymbol{d_i} = D\boldsymbol{x}$$

A sparse representation is a **linear combination** of **few dictionary atoms** $\{\boldsymbol{d_i}\}$ and $\|\boldsymbol{x}\|_0 < L$, i.e. only a few coefficients are nonzero, i.e. $\boldsymbol{x}$ is **sparse**.

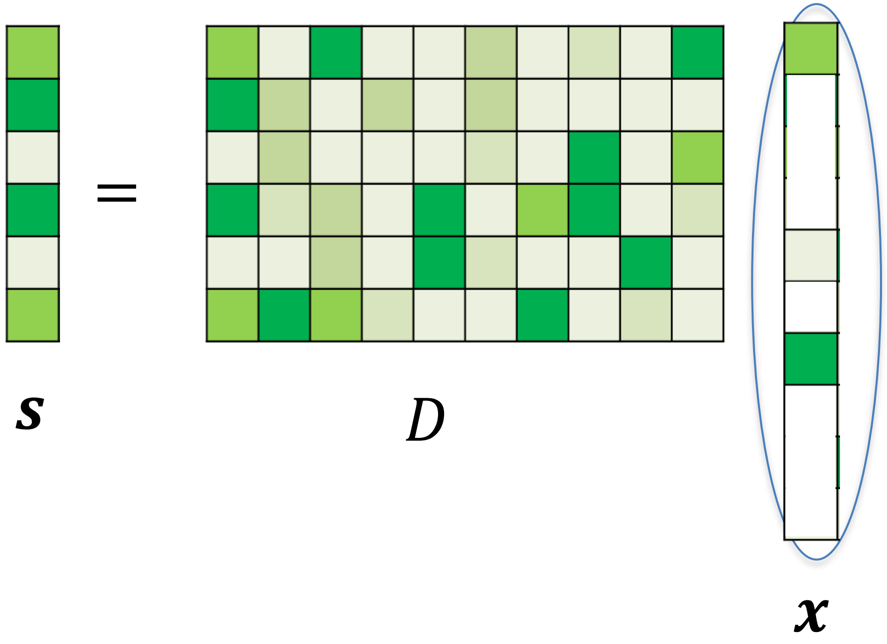

This vector $\boldsymbol{x} = [x_1, \quad \dots, \quad x_M]$ is **sparse**

Sparse Coding and Dictionary Learning

Sprase Coding: computing the sparse representation for an input signal $\boldsymbol{s}$ w.r.t. D

$$\boldsymbol{s} \in \mathbb{R}^n \quad \Rightarrow \quad \boldsymbol{x} \in \mathbb{R}^m$$

$\boldsymbol{s}$

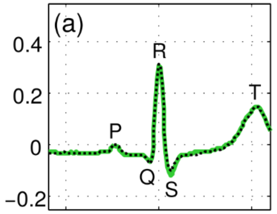

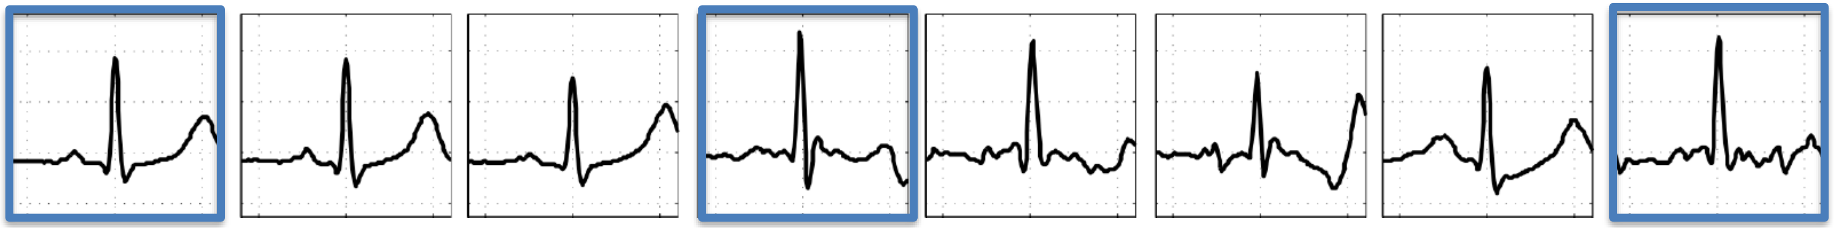

$$\boldsymbol{x} = \; 0.7, \quad 0, \quad 0, \quad 0.1, \quad 0, \quad 0, \quad 0, \quad -0.2$$

Dictionary Learning: estimate D from a training set of signals S

$$S = \{\boldsymbol{s_1}, \dots \boldsymbol{s_M}\} \quad \Rightarrow \quad D \in \mathbb{R}^{n \times m}$$

S

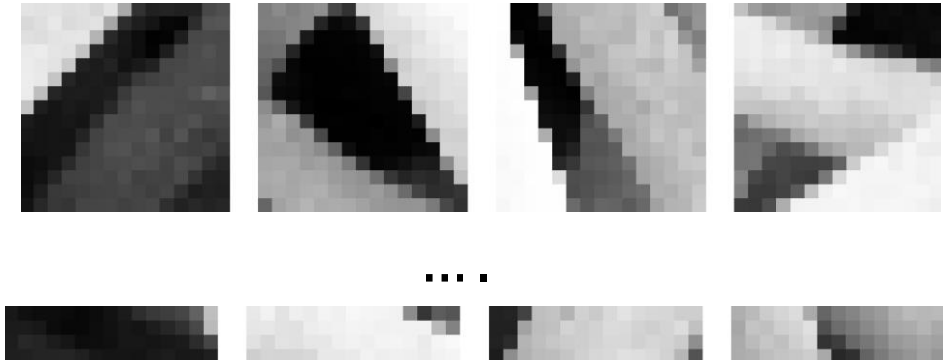

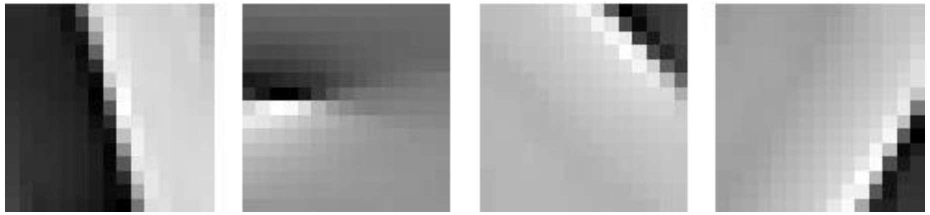

D

Online Monitoring through Sparse Representations

Training:

- **Learn a dictionary** D from a training set S containing **normal instances**

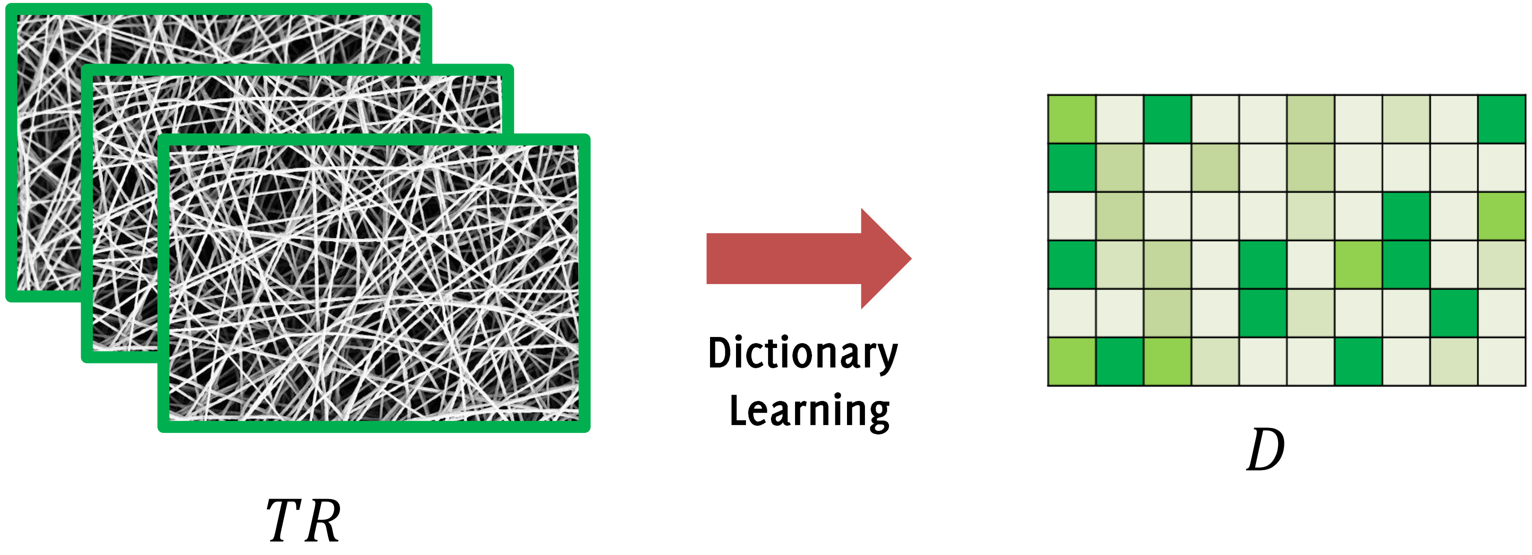

Online Monitoring through Sparse Representations

Training:

- **Learn a dictionary** D from a training set S containing **normal instances**
- **Learn how normal data are reconstructed** by D

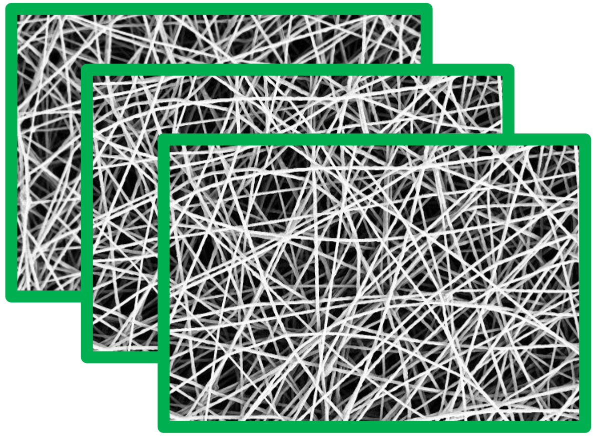

TR

Dictionary Learning

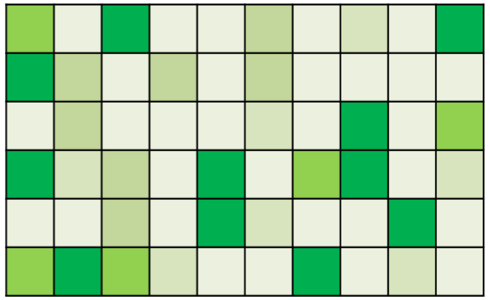

D

KDE

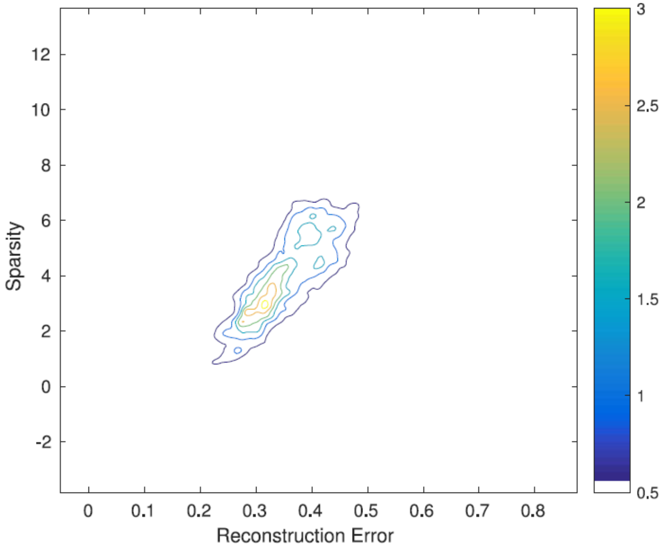

$\|\boldsymbol{x}\|_1$

$\|D\boldsymbol{x} - \mathbf{s}\|_2^2$

Online Monitoring through Sparse Representations

Training:

- **Learn a dictionary** D from a training set S containing **normal instances**
- **Learn how normal data are reconstructed** by D

Anomaly Detection:

Sparse Coding: encode each test signal $\boldsymbol{s}$ w.r.t. D, and assess its conformance with D.

Check whether the representation is:

- Sparse $\|\boldsymbol{x}\|_1$
- Accurate $\|D\boldsymbol{x} - \mathbf{s}\|_2^2$

Anomalous

Normal

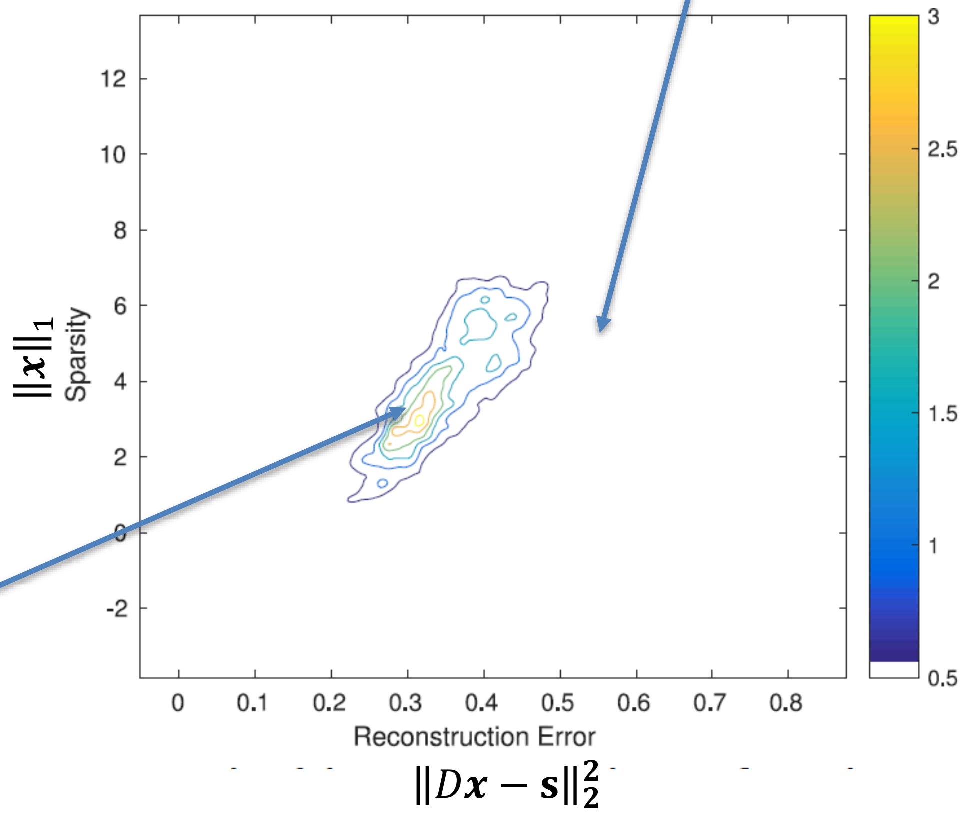

Anomaly Detection in Heartbeats

The solution is **general** and can be customized to different scenarios, including ECG monitoring.

Here we learn user-specific dictionaries to detect arrytthmias as heartbeats that do not conform to the user morphology

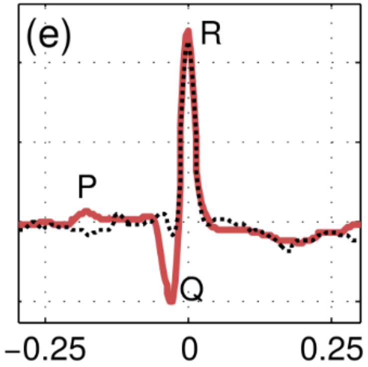

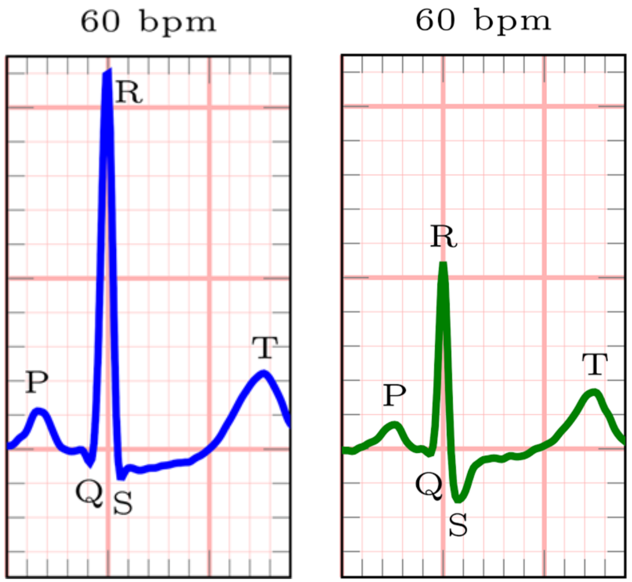

Different users feature different heartbeat morphology

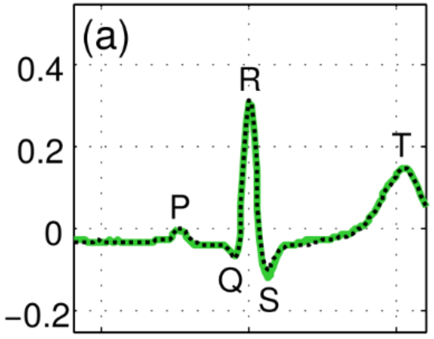

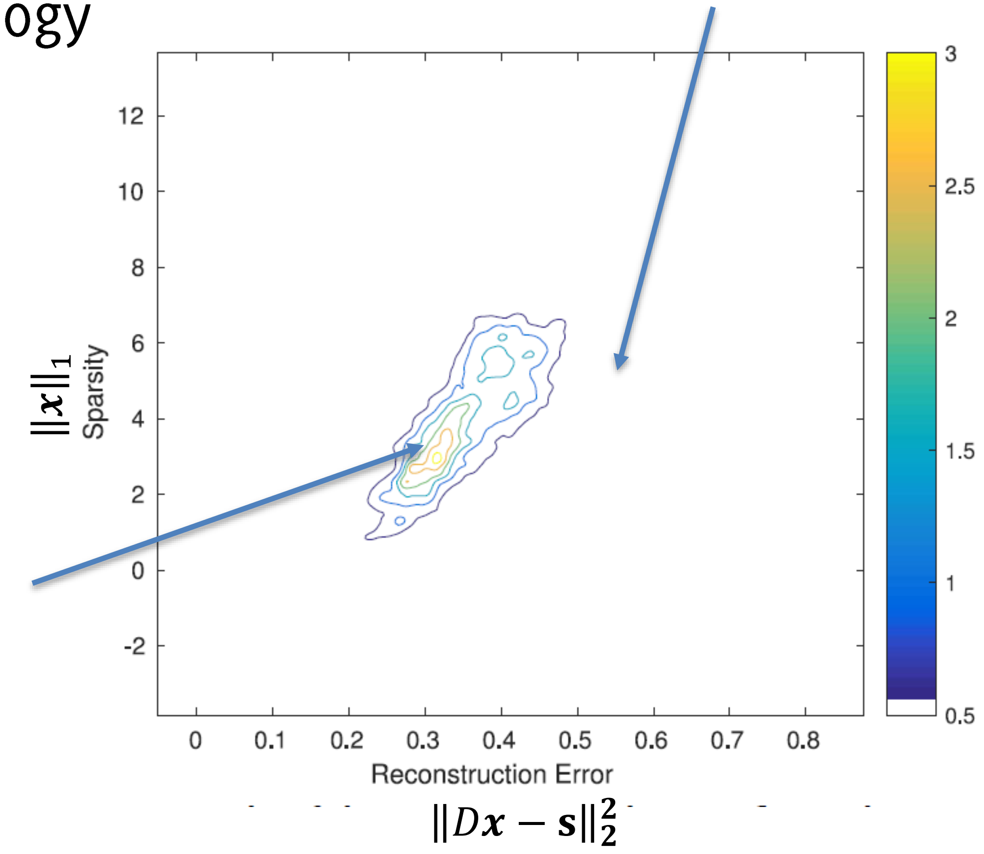

Diego Carrera, Beatrice Rossi, Pasqualina Fragneto, Giacomo Boracchi *"Online Anomaly Detection for Long-Term ECG Monitoring using Wearable Devices"*, Pattern Recognition, 2019

Domain Adaptation for Online ECG Monitoring

The issue:

- Dictionary has to be learned from each user
- ECG tracings for training can be only acquired in resting conditions
- During daily activities heart-rate changes and do not match the learned dictionary

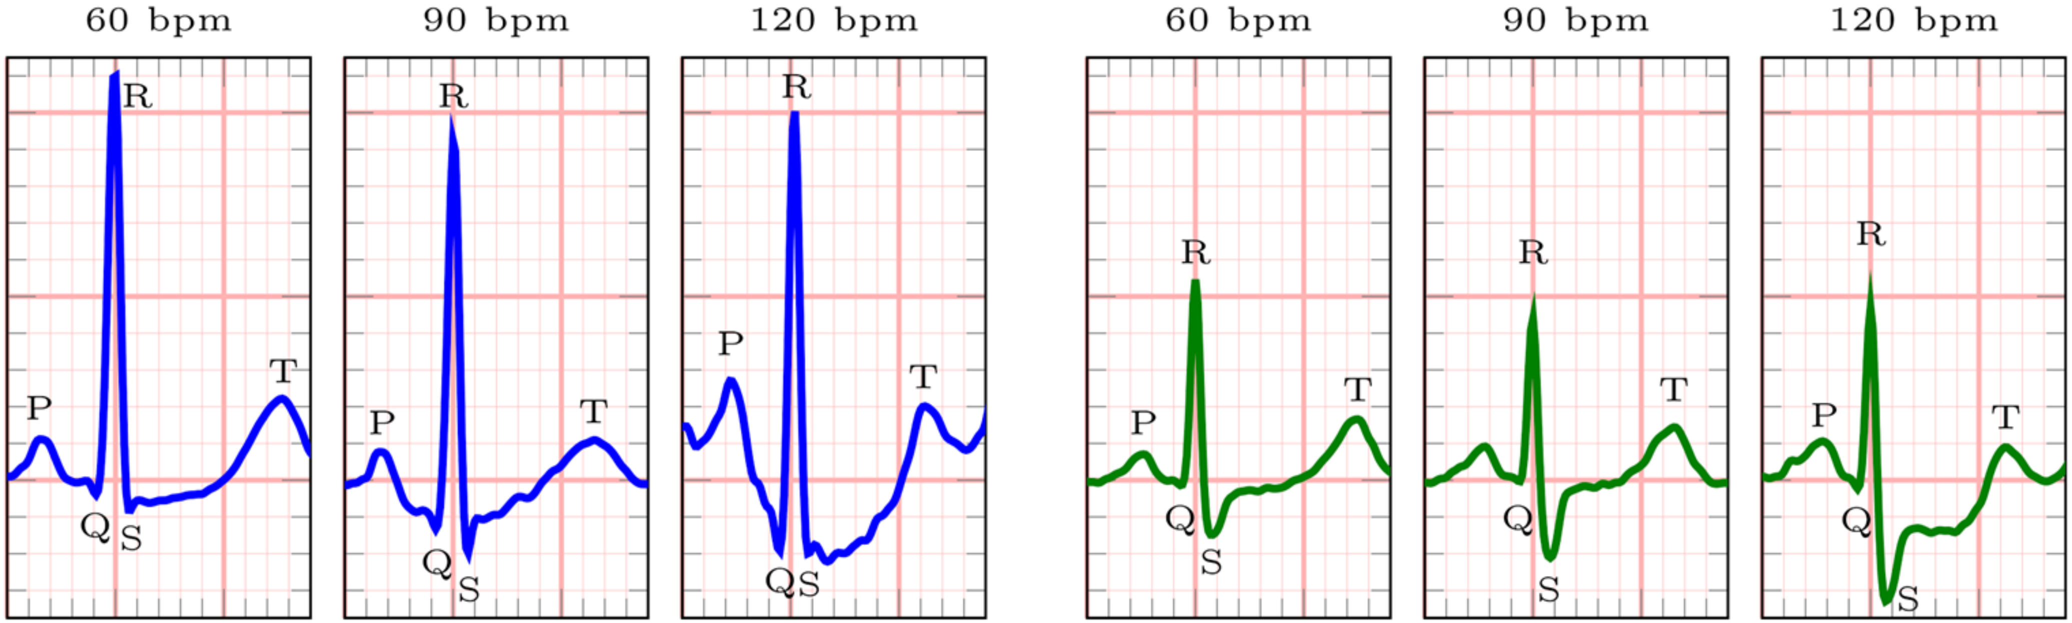

The **heartbeats** get **transformed** when the **heart rate changes**:
learned models have to be adapted according to the heart rate.

Domain Adaptation for Online ECG Monitoring

We propose to design linear transformations F_{r_1,r_0} to adapt user-specific dictionaries

$$D_{u,r_1} = F_{r_1,r_0} \cdot D_{u,r_0}, \qquad F_{r_0,r_1} \in \mathbb{R}^{m \times m}$$

Surprisingly these **transformations** can be **learned from a publicly available dataset** containing ECG recordings at different heart rates from several users

User-independent transformations enable accurate mapping of user-specific dictionaries

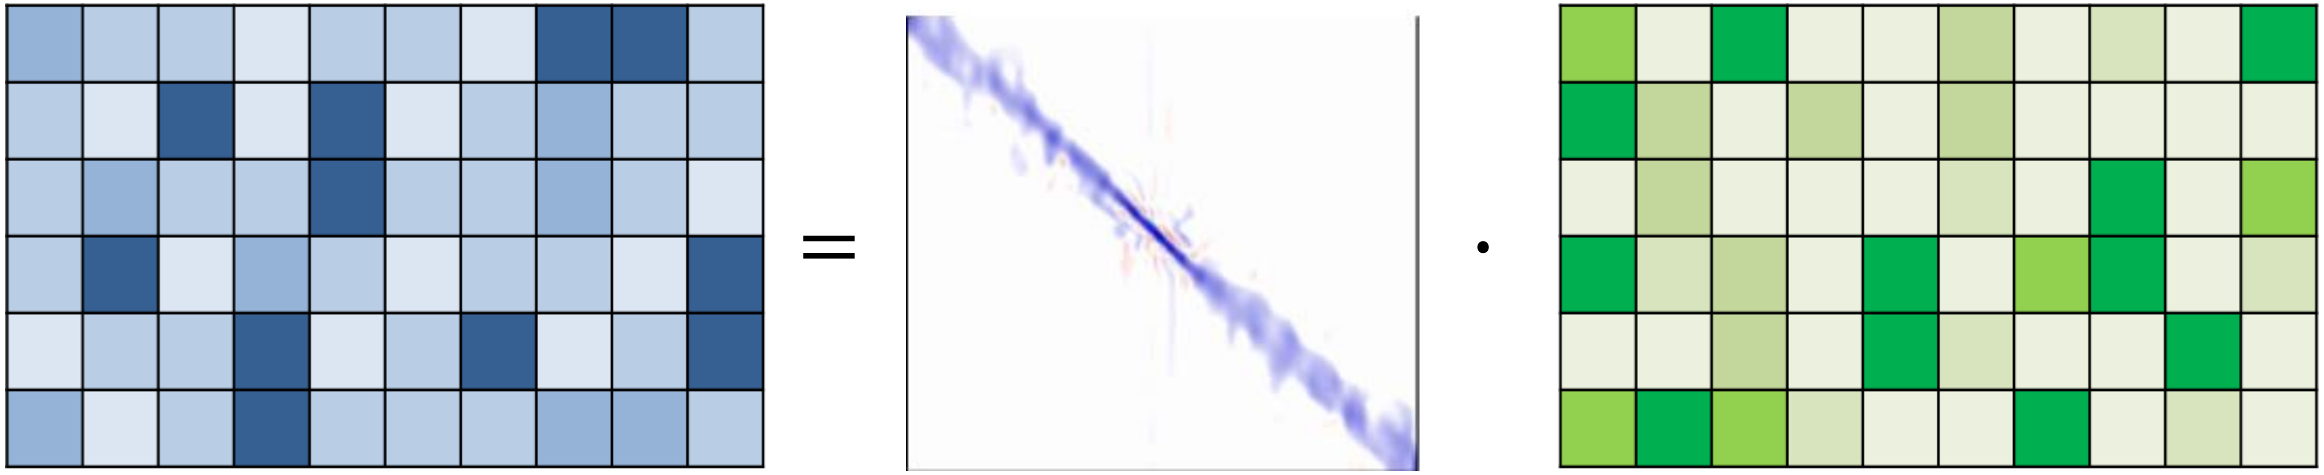

$\widehat{D}_{u,r_1}$ $\quad F_{r_1,r_0}$ $\quad D_{u,r_0}$

Carrera D., Rossi B., Fragneto P., and Boracchi G. "Domain Adaptation for Online ECG Monitoring" ICDM 2017,

Domain Adaptation for Online ECG Monitoring

We propose to design linear transformations F_{r_1,r_0} to adapt user-specific dictionaries

$$D_{u,r_1} = F_{r_1,r_0} \cdot D_{u,r_0}, \qquad F_{r_0,r_1} \in \mathbb{R}^{m \times m}$$

Surprisingly these **transformations** can be **learned from a publicly available dataset** containing ECG recordings at different heart rates from several users

User-independent transformations enable accurate mapping of user-specific dictionaries

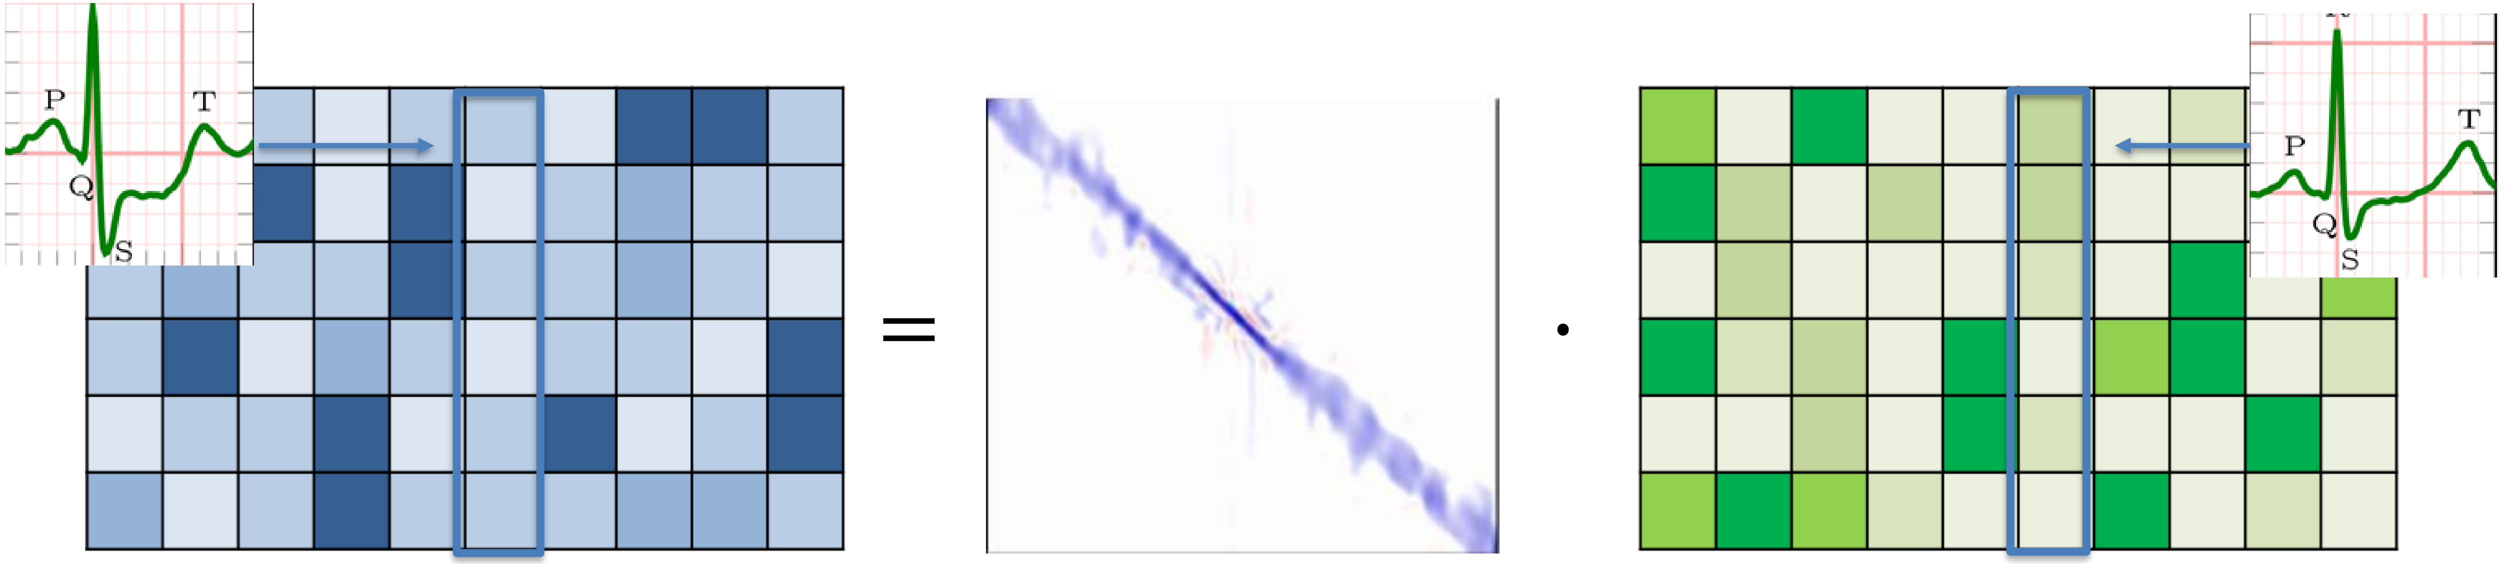

$\widehat{D}_{u,r_1}$ $\quad = \quad$ F_{r_1,r_0} $\quad \cdot \quad$ D_{u,r_0}

Carrera D., Rossi B., Fragneto P., and Boracchi G. "Domain Adaptation for Online ECG Monitoring" ICDM 2017,

ML: Long Term ECG monitoring in wearables

Efficient variant of the sparse coding has been implemented in an MCU, to enable **online and long term ECG monitoring**. Automatic **detection of anomalies** (e.g., arrhythmias) and **domain adaptation** to track heart rate variations. The model is based on **a learned and user-specific model** of heartbeats.

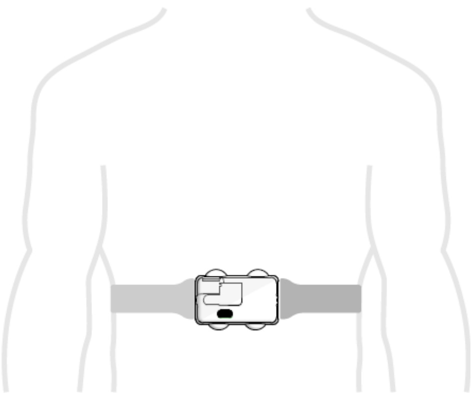

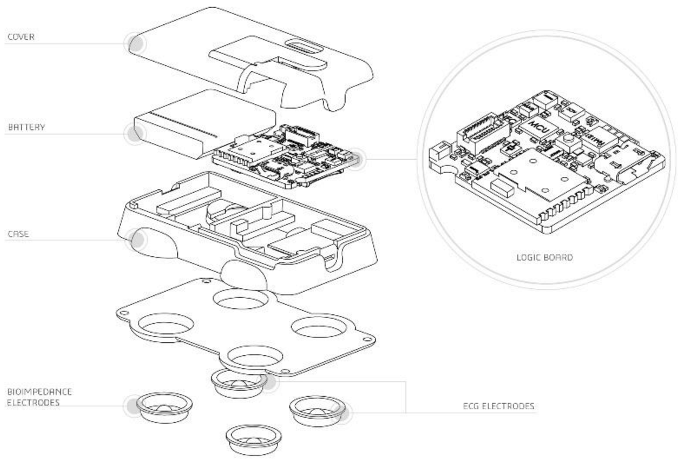

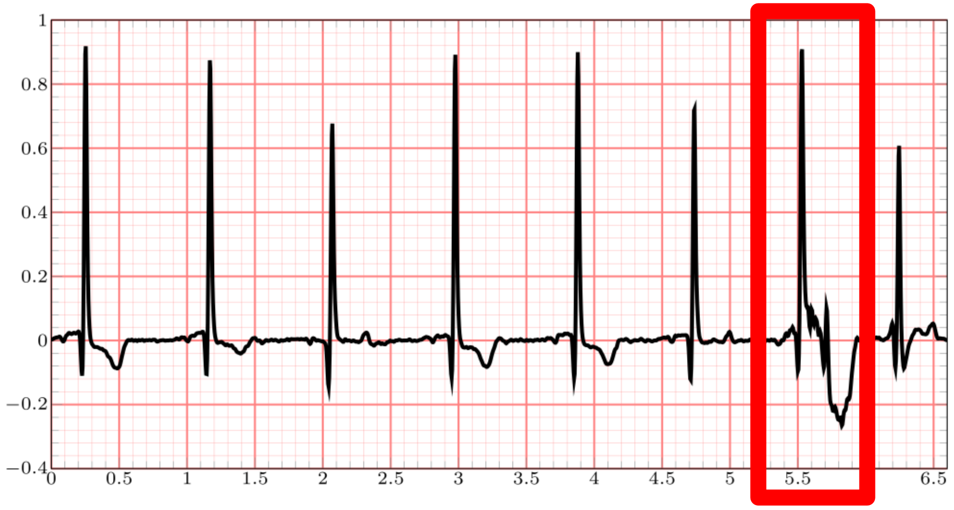

Marco Longoni, Diego Carrera, Beatrice Rossi, Pasqualina Fragneto, Marco Pessione, Giacomo Boracchi, "A Wearable Device for Online and Long-Term ECG Monitoring" IJCAI 2018 - Demo Track https://boracchi.faculty.polimi.it/Projects/ECG_Monitoring/IJCAI_2018_Demo.html

Domain Adaptation on Quality Inspection

The Issue:

- SEM images can be acquired at different zooming levels

Solution:

- **Synthetically generate** training images at **different zooming levels**
- Learn a dictionary for each scale
- Combine all the learned dictionaries in a **multiscale dictionary** D
- Perform **sparse-coding** including a penalized, **group sparsity term**

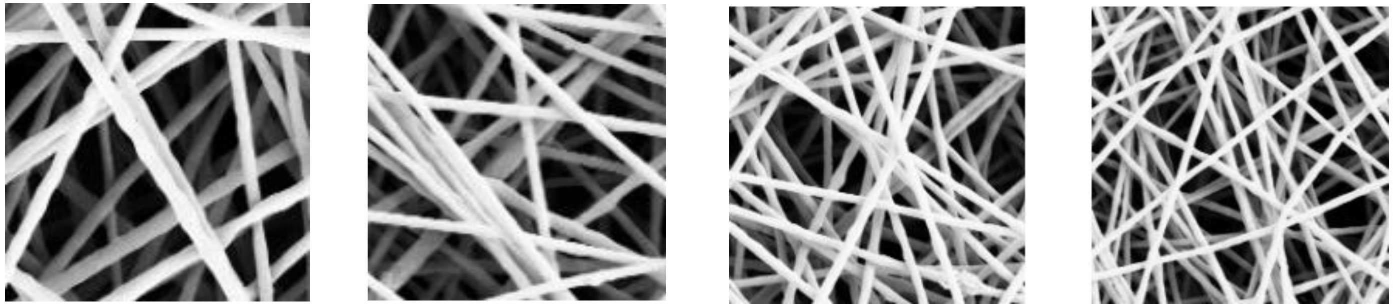

$$D = [\quad D_1 \qquad D_2 \qquad D_3 \qquad D_4 \quad]$$

Carrera D., Boracchi G., Foi A. and B. Wohlberg *"Scale-invariant Anomaly Detection With multiscale Group-sparse Models"* IEEE ICIP 2016

The Elephant in the Room….

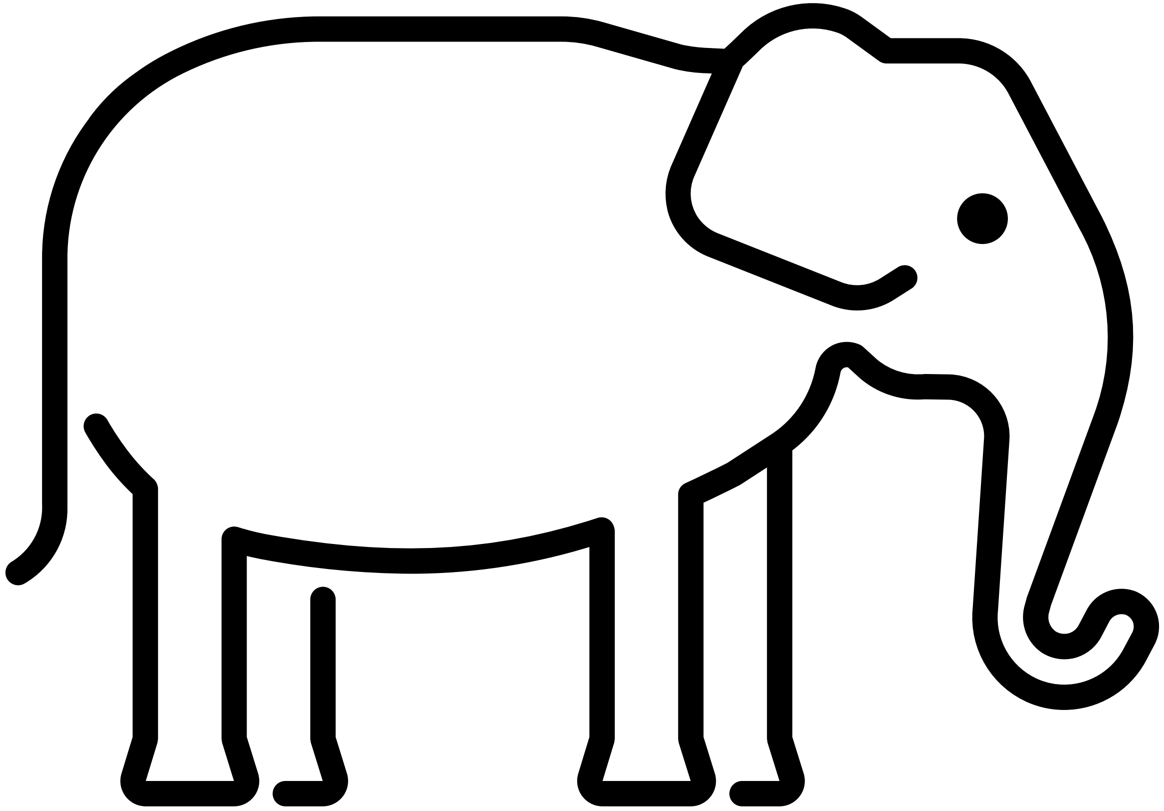

Anomaly Detection by Deep Learning Models

In 2018 the «deep learning» tsunami meets Anomaly Detection

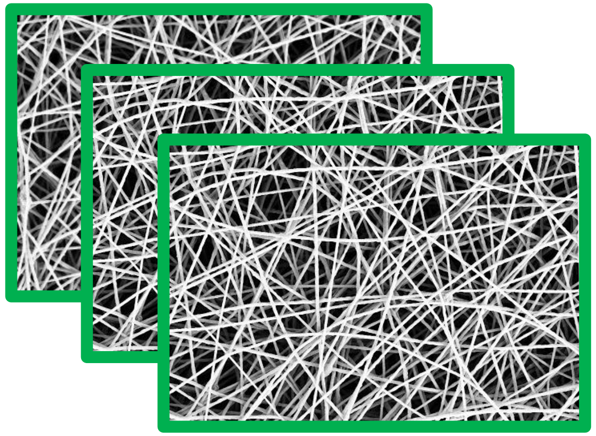

TR

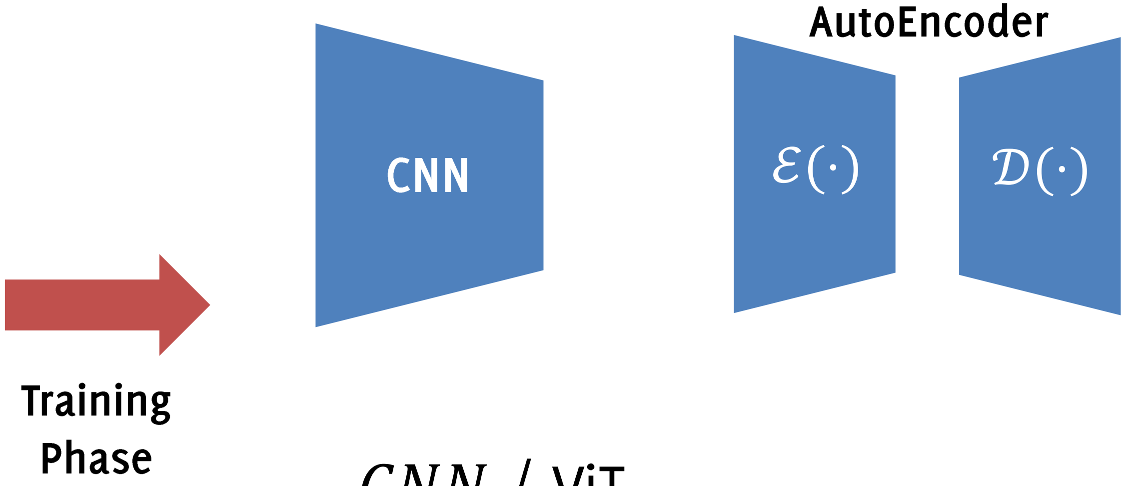

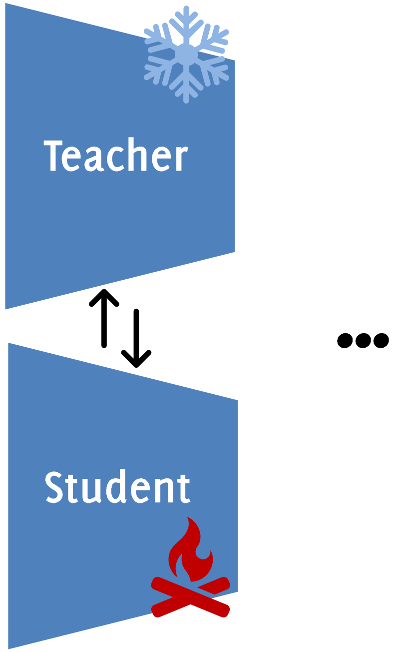

CNN / ViT

- Trained on normal data
- possibly pretrained
- Student teacher architectures

Napoletano P., Piccoli F., Schettini R., *"Anomaly Detection in Nanofibrous Materials by CNN-Based Self-Similarity"*, Sensors 2018

Zong et al, *“Deep Autoencoding Gaussian Mixture Model for Unsupervised Anomaly Detection”, ICLR 2018*

Sabokrou, Mohammad, et al. "Deep-anomaly: Fully convolutional neural network for fast anomaly detection in crowded scenes." *CVIU* 2018

Burlina, Philippe, Neil Joshi, and I. Wang. "Where's Wally now? Deep generative and discriminative embeddings for novelty detection." *CVPR 2019*

Defard, Thomas, et al. “PADIM: a patch distribution modeling framework for anomaly detection and localization.“ ICPR 2020

Bergmann, Paul, et al. "*Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings.*" CVPR 2020

Yunkang Cao

Zero-Shot Anomaly Detection

Yunkang Cao, Jiangning Zhang , Luca Frittoli, Yuqi Cheng, Weiming Shen, and Giacomo Boracchi
“AdaCLIP: Adapting CLIP with Hybrid Learnable Prompts for Zero-Shot Anomaly Detection” ECCV, 2024

Zero Shot Deep Learning

Learning perception *from the* ***supervision*** *contained in* ***natural language*** *paired with images.*

CLIP: a Text encoder and an Image encoder projecting latent representations in the same space.

The latent representations of a text and of an image are nearby when the two portraits similar content

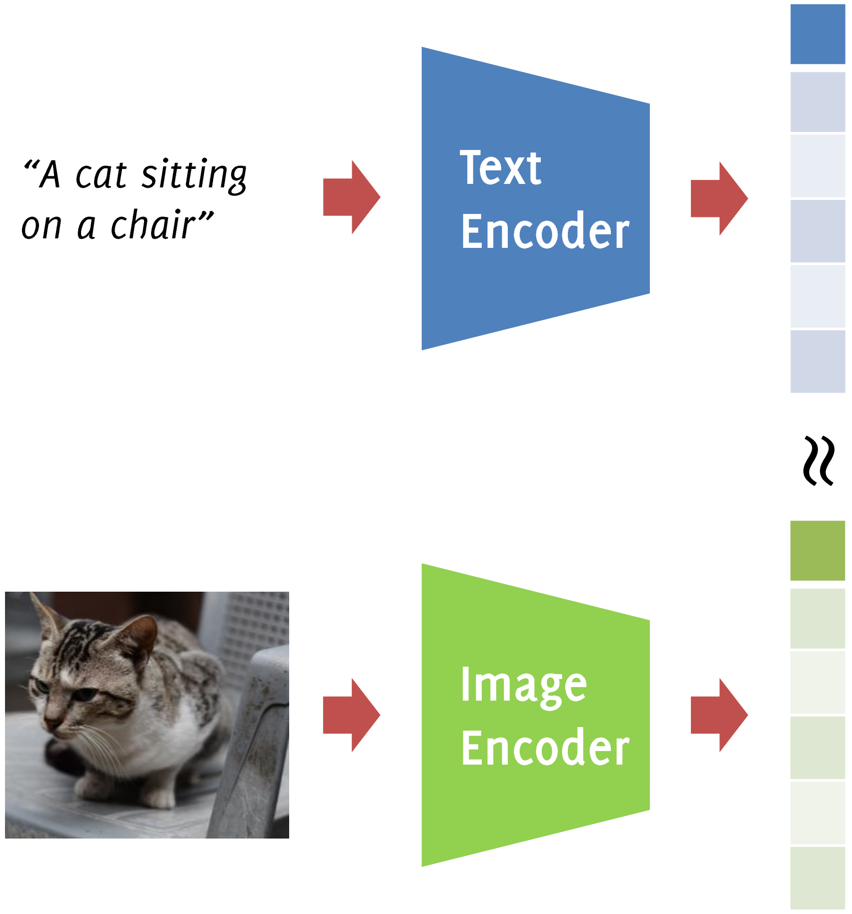

Zero-Shot CLIP

- Take pre-trained CLIP
- Consider the class names as text $\{t_i\}$
- Classify an image I by associating it to the text $\{t_i\}$ that is closest to I according to CLIP.

Radford, Alec, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry et al. *"Learning transferable visual models from natural language supervision."* ICML, 2021.

The Zero-Shot Paradigm in Anomaly Detection

Most AD methods require normal samples to train a (typically unsupervised) DL model

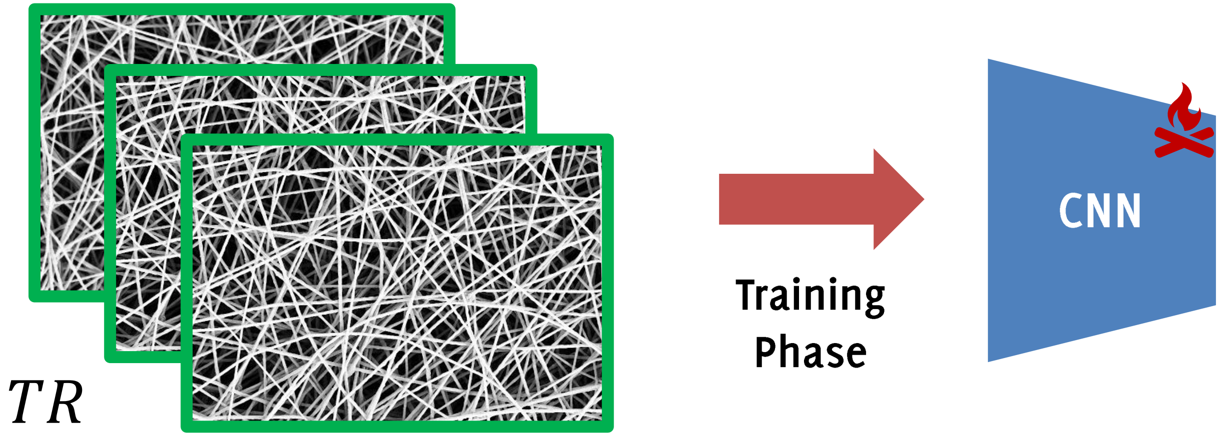

"Cold Start" limitation: **need to gather data before starting monitoring**

However, pre-trained **Vision-Language Models** *"know"* the visual pattern of defects.

Testing images may exhibit universal patterns, either normal or anomalous, that VLMs can identify.

Jeong, J., Zou, Y., Kim, T., Zhang, D., Ravichandran, A., Dabeer, O.: "Winclip: Zero-/few-shot anomaly classification and segmentation". CVPR 2023

The Zero-Shot Paradigm in Anomaly Detection

Zero-shot AD: Enable detecting anomalies even without normal images for training.

Rationale behind Zero-shot AD:

- Adopt a pretrained CLIP Image and Text Encoders
- Query for “defective”, “broken” or “damaged” prompts
- Detect as anomalous images that are close to these prompts in the latent space.
- Possibly adopt “normal” text prompts as well

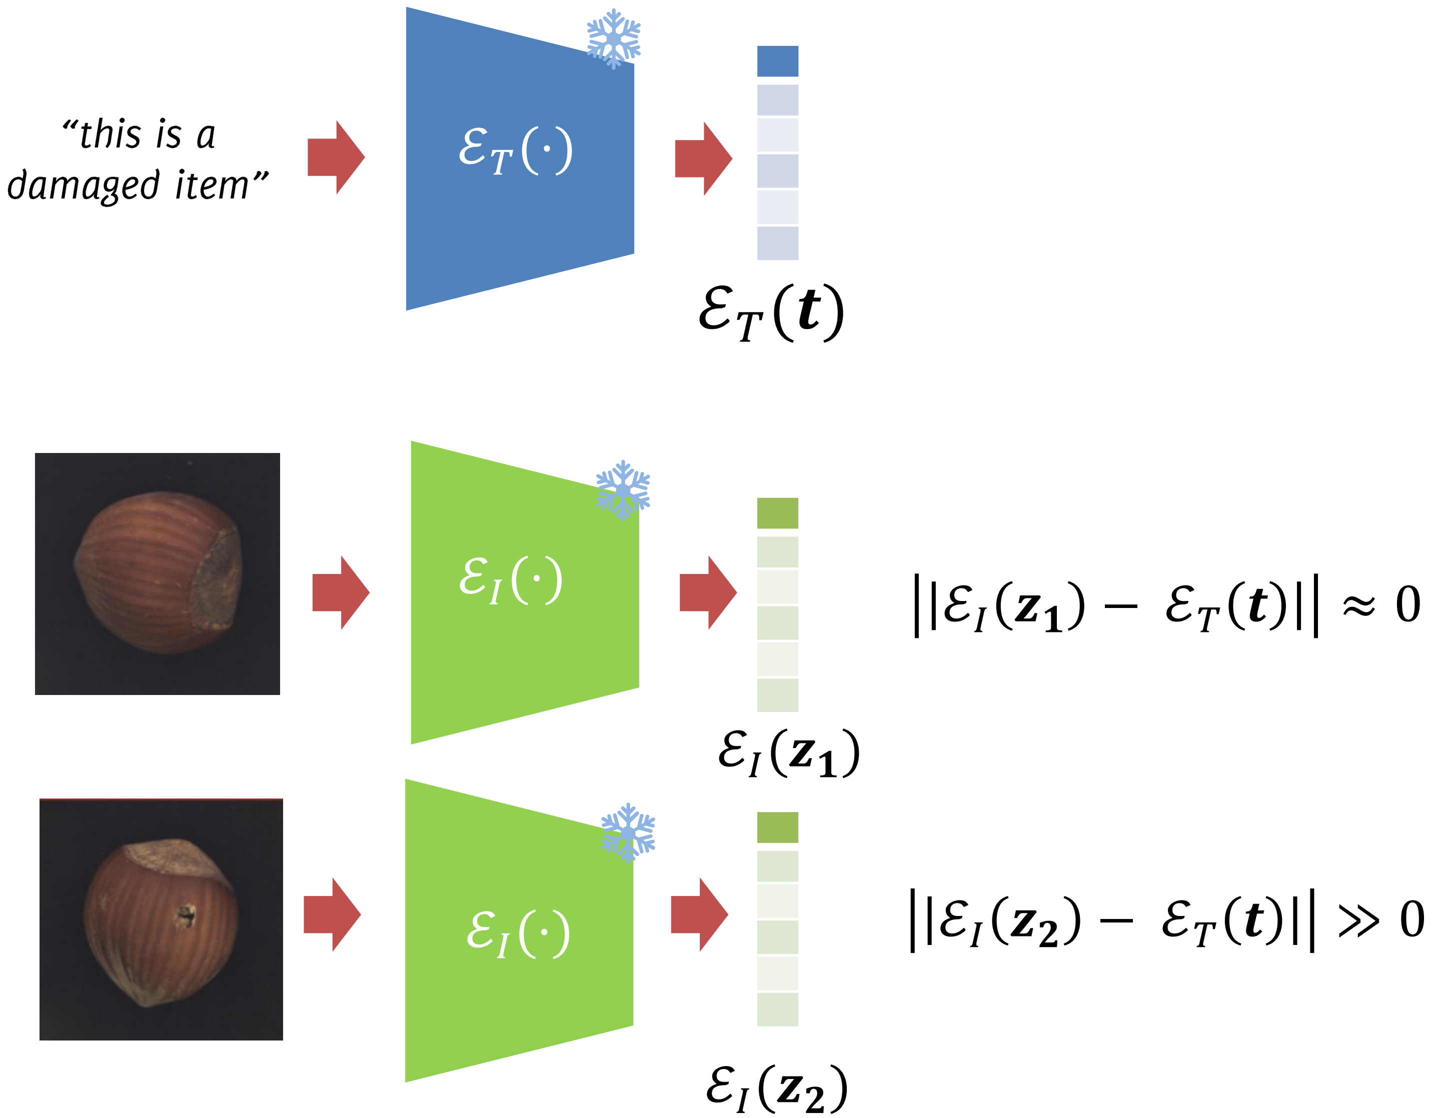

Jeong, J., Zou, Y., Kim, T., Zhang, D., Ravichandran, A., Dabeer, O.: “Winclip: Zero-/few-shot anomaly classification and segmentation”. CVPR 2023

AdaClip

Remarks:

- **CLIP** is trained on **natural image-text**, not specialized for anomaly detection.
- There is a large **availability** of **defect** detection annotated **datasets.**

Idea behind AdaCLIP:

- Leverage defect-detection datasets as auxiliary annotated data to **fine tune learnable prompts** into a zero-shot framework from CLIP.
- Adaptation adheres to the zero-shot learning paradigm, as long as **testing images do not belong to categories** presented in the **auxiliary AD dataset.**

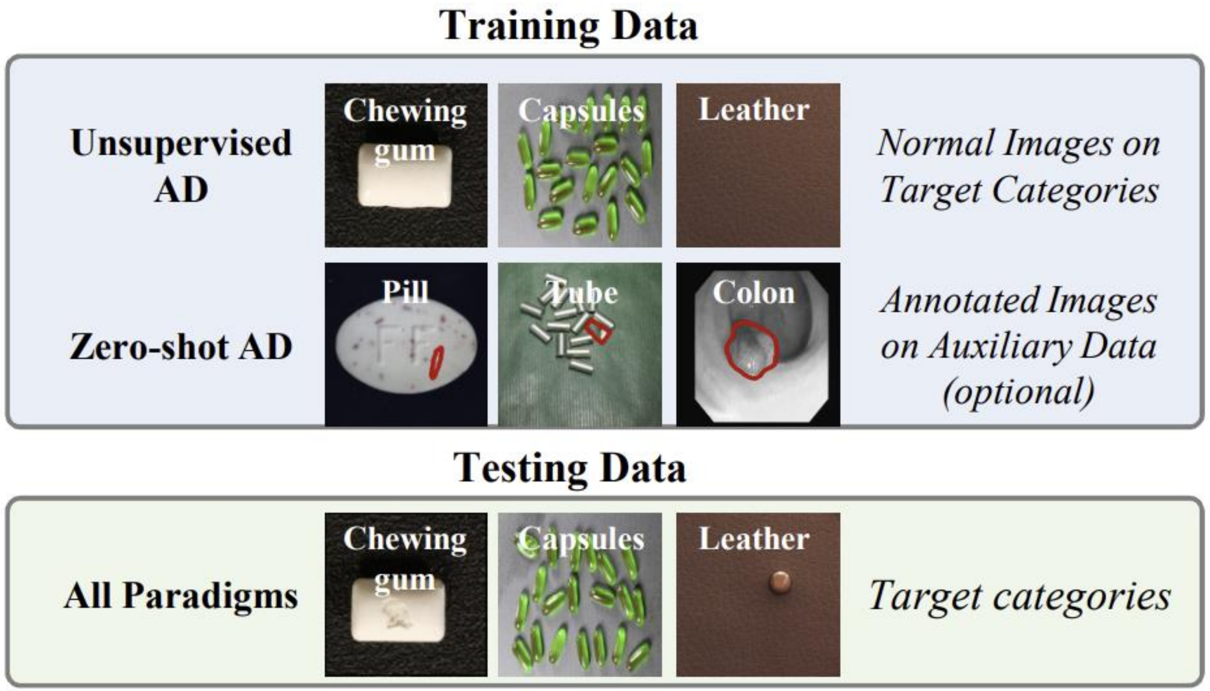

Yunkang Cao, Jiangning Zhang , Luca Frittoli, Yuqi Cheng, Weiming Shen, and Giacomo Boracchi *"AdaCLIP: Adapting CLIP with Hybrid Learnable Prompts for Zero-Shot Anomaly Detection"* ECCV, 2024

AdaClip

AdaCLIP fine tunes on auxiliary datasets **projection and prompting** layers. **Static prompts** are shared across all images and **dynamic prompts** are generated based on the testing image

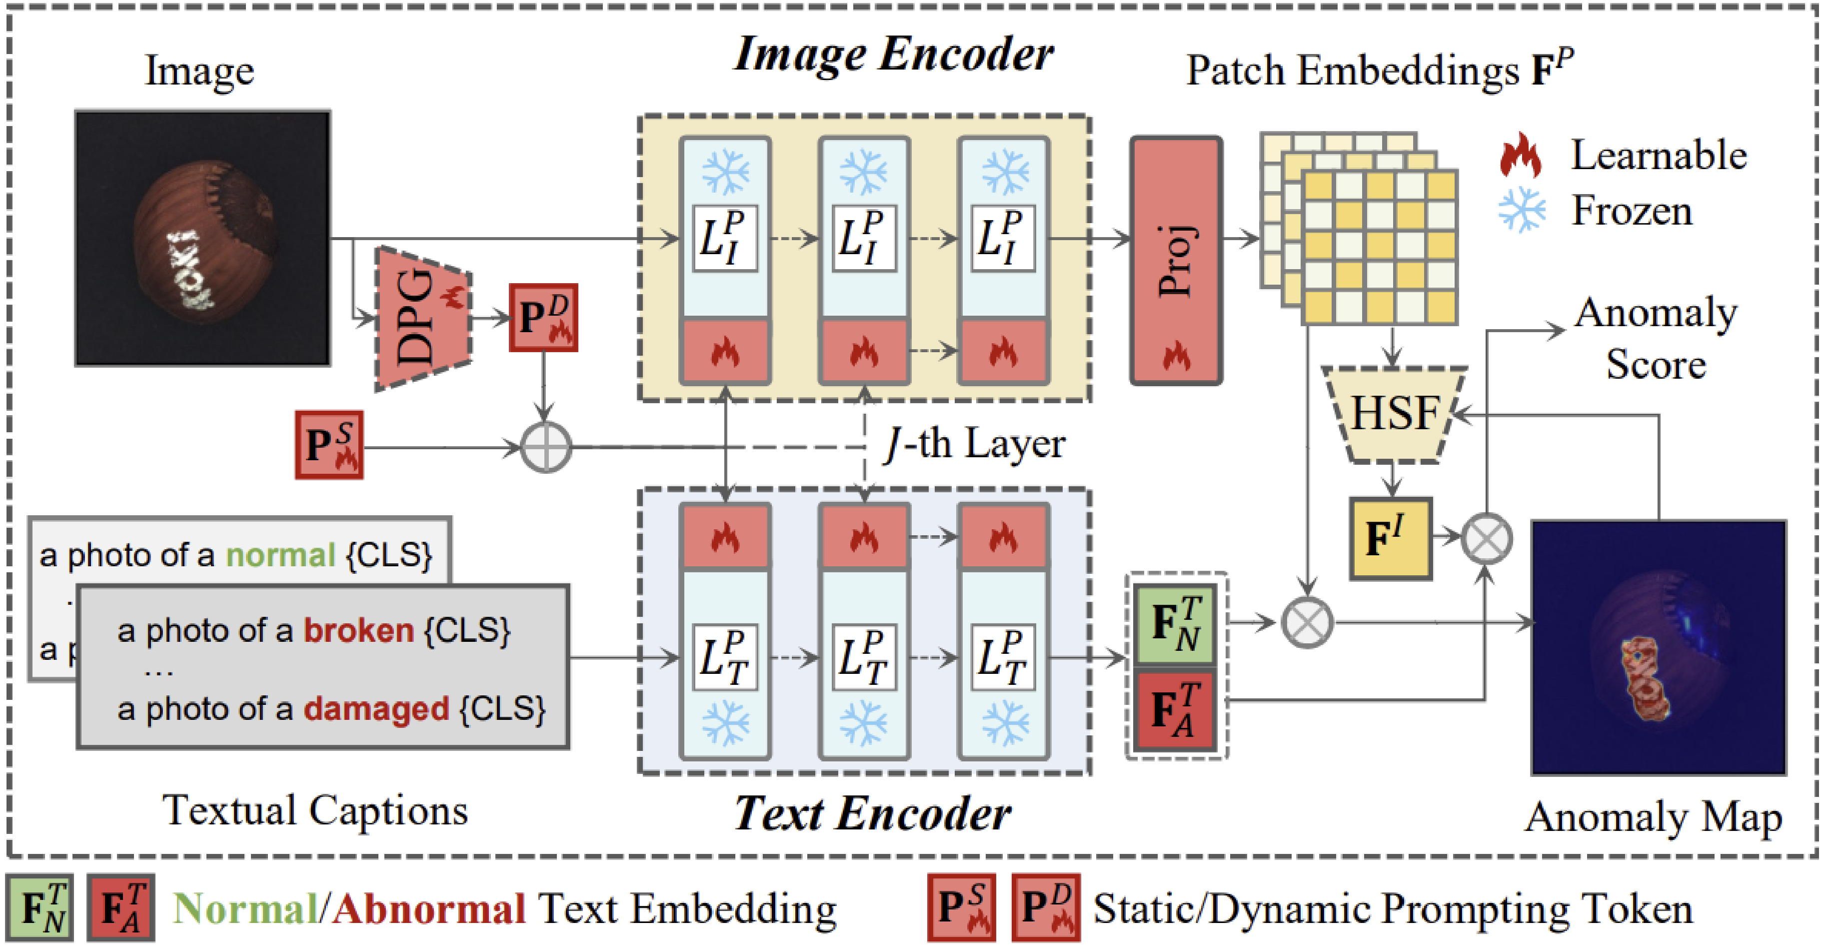

Yunkang Cao, Jiangning Zhang, Luca Frittoli, Yuqi Cheng, Weiming Shen, and Giacomo Boracchi *"AdaCLIP: Adapting CLIP with Hybrid Learnable Prompts for Zero-Shot Anomaly Detection"* ECCV, 2024

AdaClip

AdaCLIP fine tunes on auxiliary datasets **projection and prompting** layers. **Static prompts** are shared across all images and **dynamic prompts** are generated based on the testing image

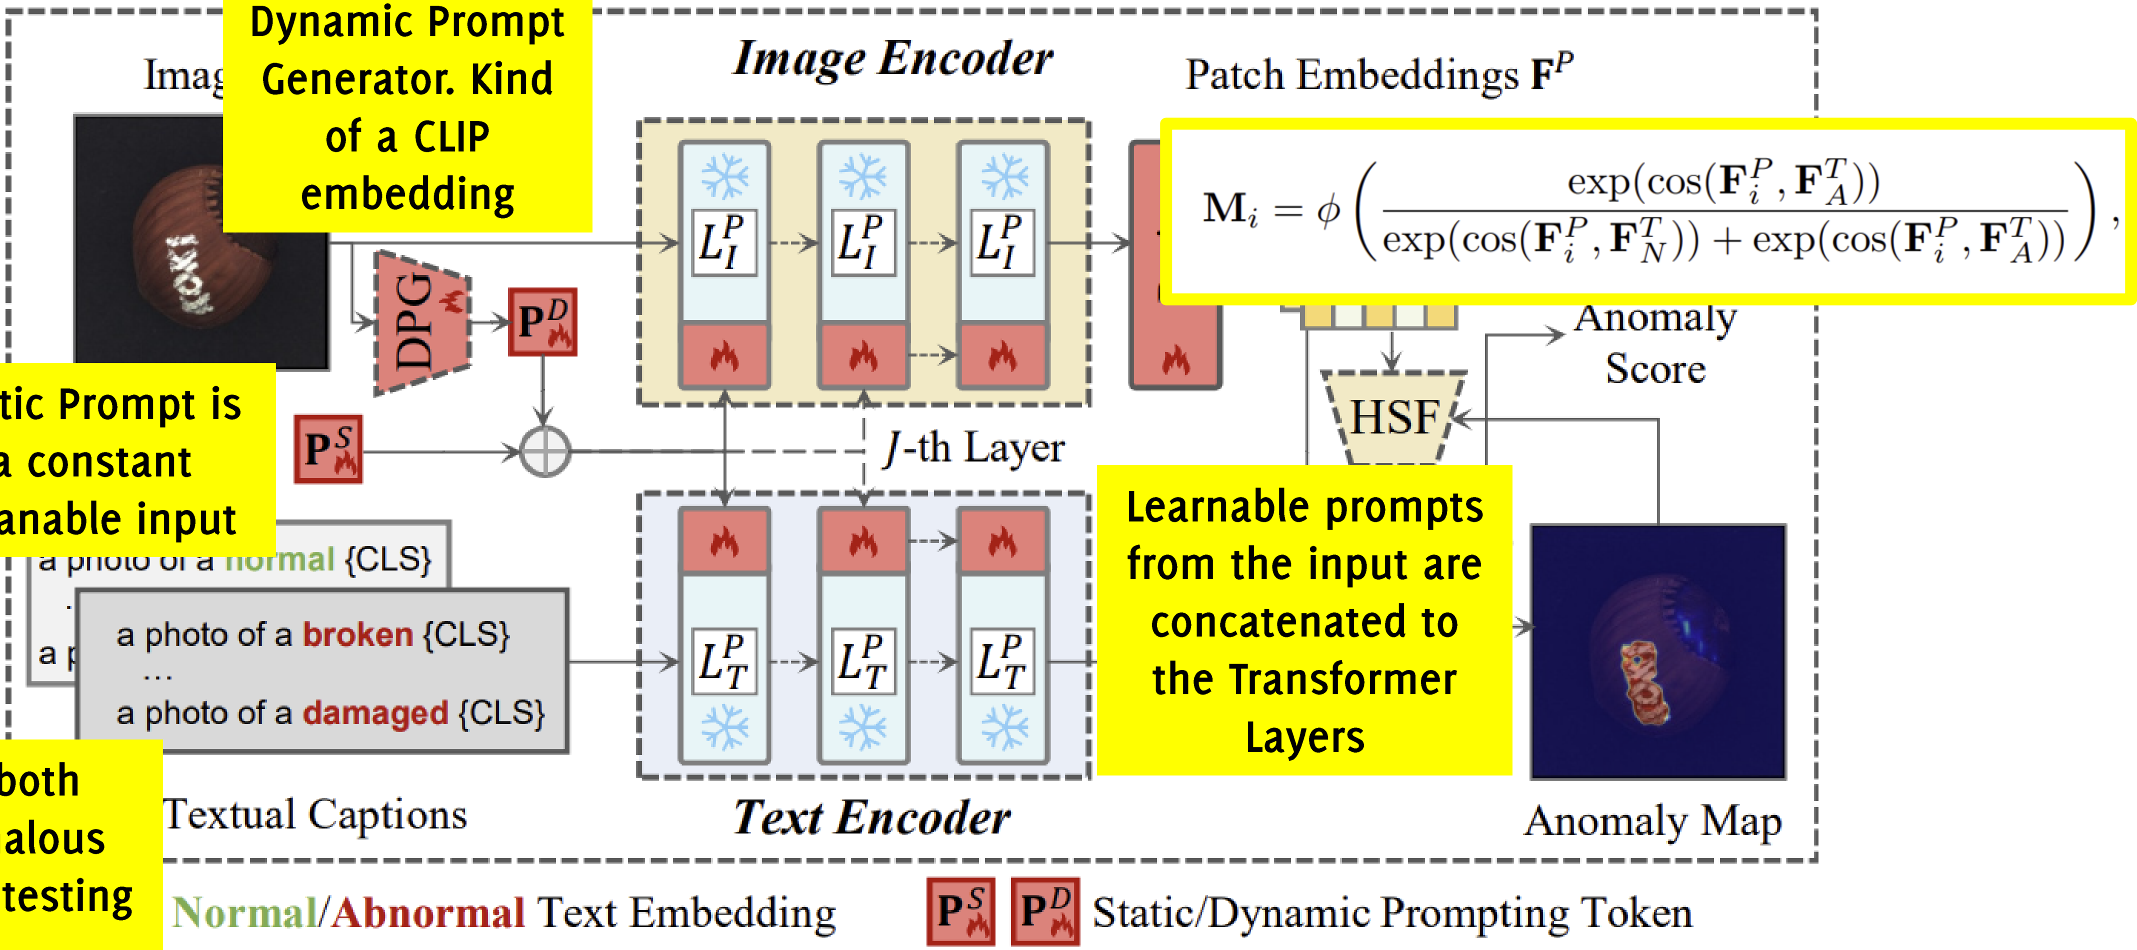

Dynamic Prompt Generator. Kind of a CLIP embedding

$$\mathbf{M}_i = \phi\left(\frac{\exp(\cos(\mathbf{F}_i^P, \mathbf{F}_A^T))}{\exp(\cos(\mathbf{F}_i^P, \mathbf{F}_N^T)) + \exp(\cos(\mathbf{F}_i^P, \mathbf{F}_A^T))}\right),$$

Static Prompt is a constant leranable input

Learnable prompts from the input are concatenated to the Transformer Layers

Prompts for both normal/ anomalous text. CLS is the testing category.

Yunkang Cao, Jiangning Zhang, Luca Frittoli, Yuqi Cheng, Weiming Shen, and Giacomo Boracchi *"AdaCLIP: Adapting CLIP with Hybrid Learnable Prompts for Zero-Shot Anomaly Detection"* ECCV, 2024

Filippo Leveni *Luca Magri*

Structure-based Anomaly Detection

Filippo Leveni, Luca Magri, Cesare Alippi, Giacomo Boracchi *“Preference Isolation Forest for Structure-based Anomaly Detection”* Under Review

Filippo Leveni, Luca Magri, Cesare Alippi and Giacomo Boracchi *“Hashing for Structure-based Anomaly Detection“* International Conference on Image Analisys and Processing (ICIAP) 2023 Best Paper Award

Filippo Leveni, Luca Magri, Giacomo Boracchi, Cesare Alippi, *“Anomaly detection via preference embedding”* International Conference on Pattern Recognition (ICPR) 2020,

Structure-based Anomaly Detection

Density-based Anomaly Detection

Anomalies are identified as samples that lie in *low-density regions*.

The **background model $\mathcal{M}$ is trained** to extract i.i.d. indicators or anomaly scores, to move the AD problem in the random-variable settings.

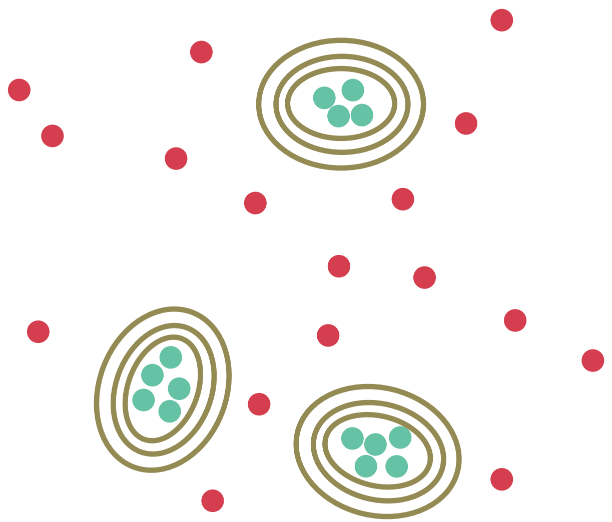

Structure-based Anomaly Detection

Density-based Anomaly Detection

Anomalies are identified as samples that lie in *low-density regions*.

The **background model $\mathcal{M}$ is trained** to extract i.i.d. indicators or anomaly scores, to move the AD problem in the random-variable settings.

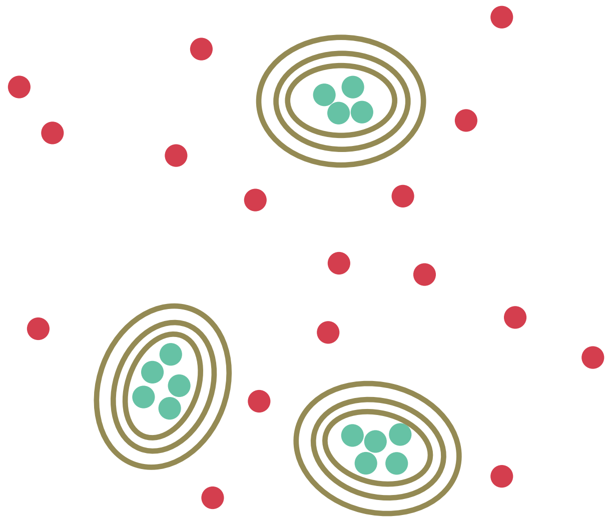

Structure-based Anomaly Detection

In some cases, anomalies are better identified as samples that do not conform to *low-dimensional manifolds*.

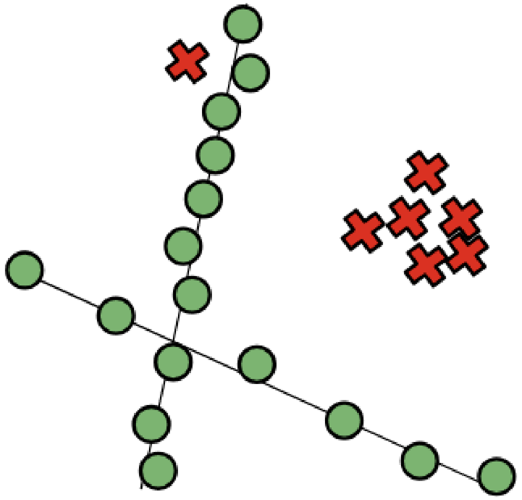

Structure-based Anomaly Detection

Structure-based Anomaly Detection

In some cases, anomalies are better identified as samples that do not conform to *low-dimensional manifolds*.

Such as *periodic texture* in 2D images

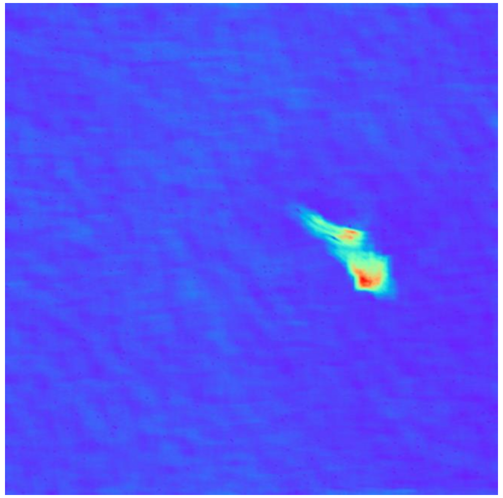

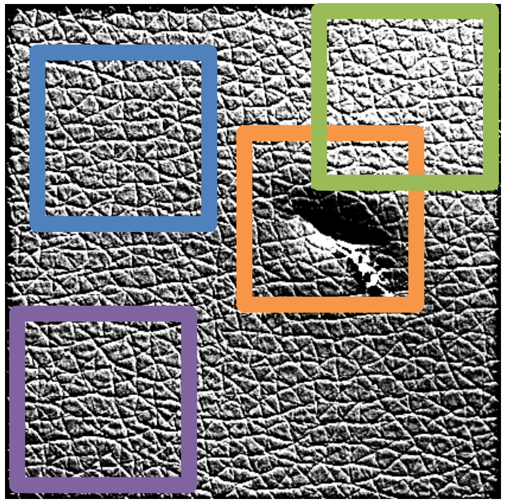

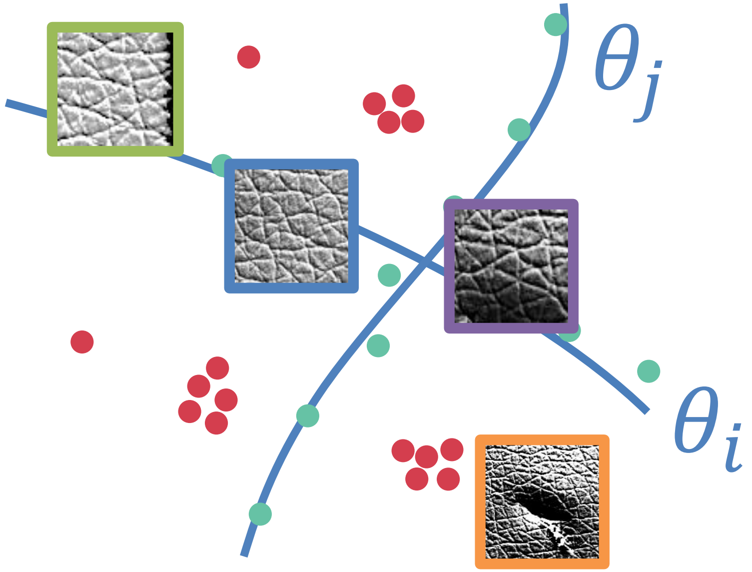

Structure-based Anomaly Detection

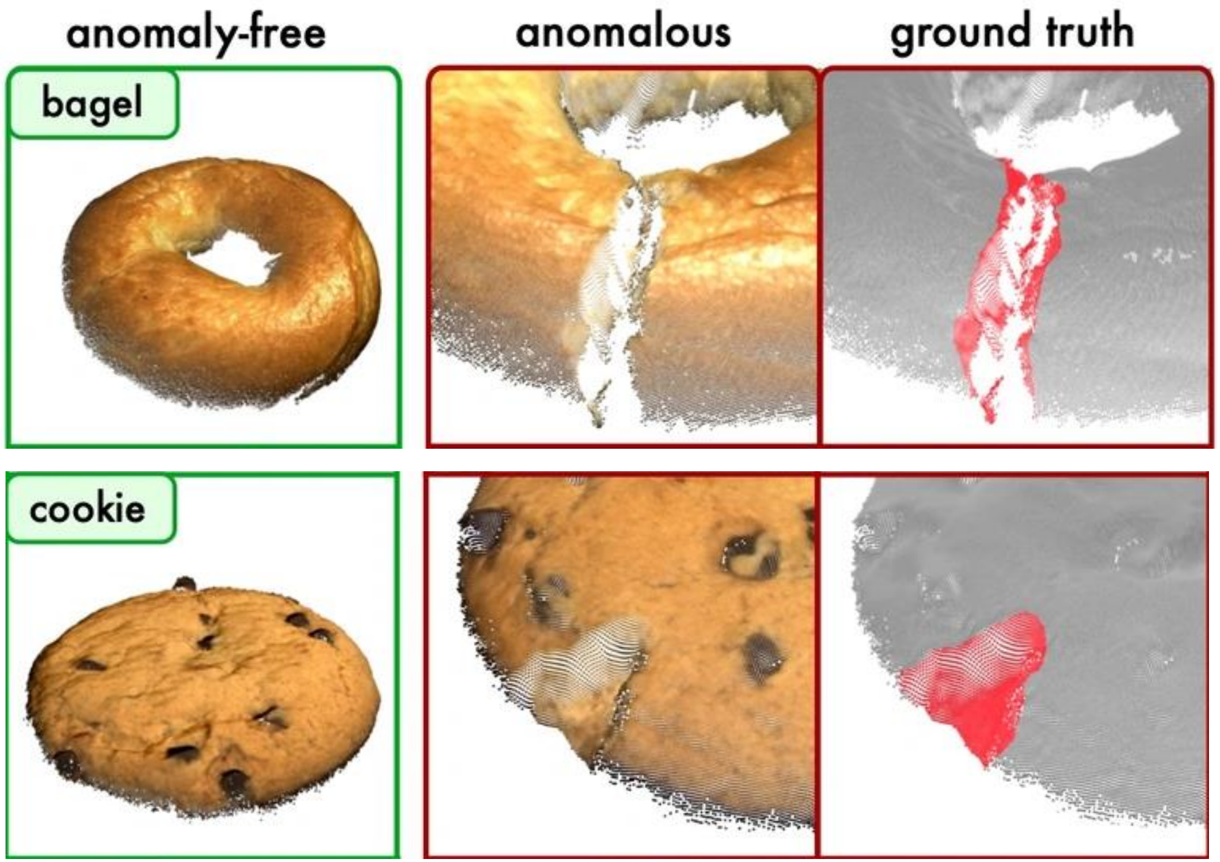

"The MVTec 3D-AD Dataset for Unsupervised 3D Anomaly Detection and Localization".

Normal data are locally approximated by parabolic shapes or even planes. Anomalies (cracks, spikes, holes) depart from these assumptions.

Structure-based Anomaly Detection

The only assumption: the *analytical expression* of the model family $\mathcal{F}$ describing normal data.

- Assume a *local structure* e.g., "locally approximated by planes".

Truly **"unsupervised settings"** no training data needed, neither pre-trained models!

PIF: Preference Isolation Forest

PIF has two major components:

i) **Preference Embedding**, that maps input data the (high-dimensional) Preference Space, by fitting models from the family $\mathcal{F}$.

ii) **Preference Isolation**, to detect anomalies as isolated points in the Preference Space.

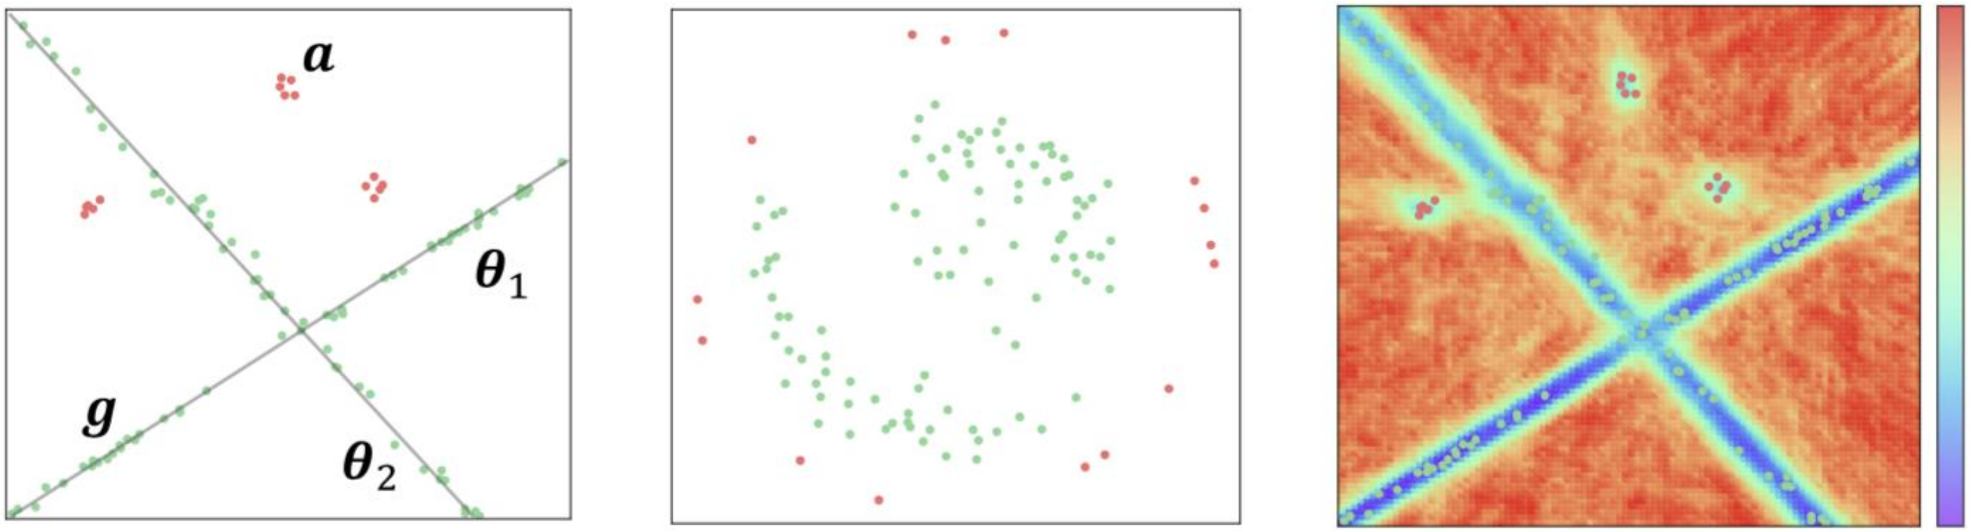

(a) Input data $X = G \cup A$. (b) Preference space $\mathcal{P}$. (c) Anomaly scores α.

MDS visualizaztion

Filippo Leveni, Luca Magri, Giacomo Boracchi, Cesare Alippi, *“Anomaly detection via preference embedding”* International Conference on Pattern Recognition (ICPR) 2020,

i) Preference Embedding

Map points from the ambient space to Preference Space.

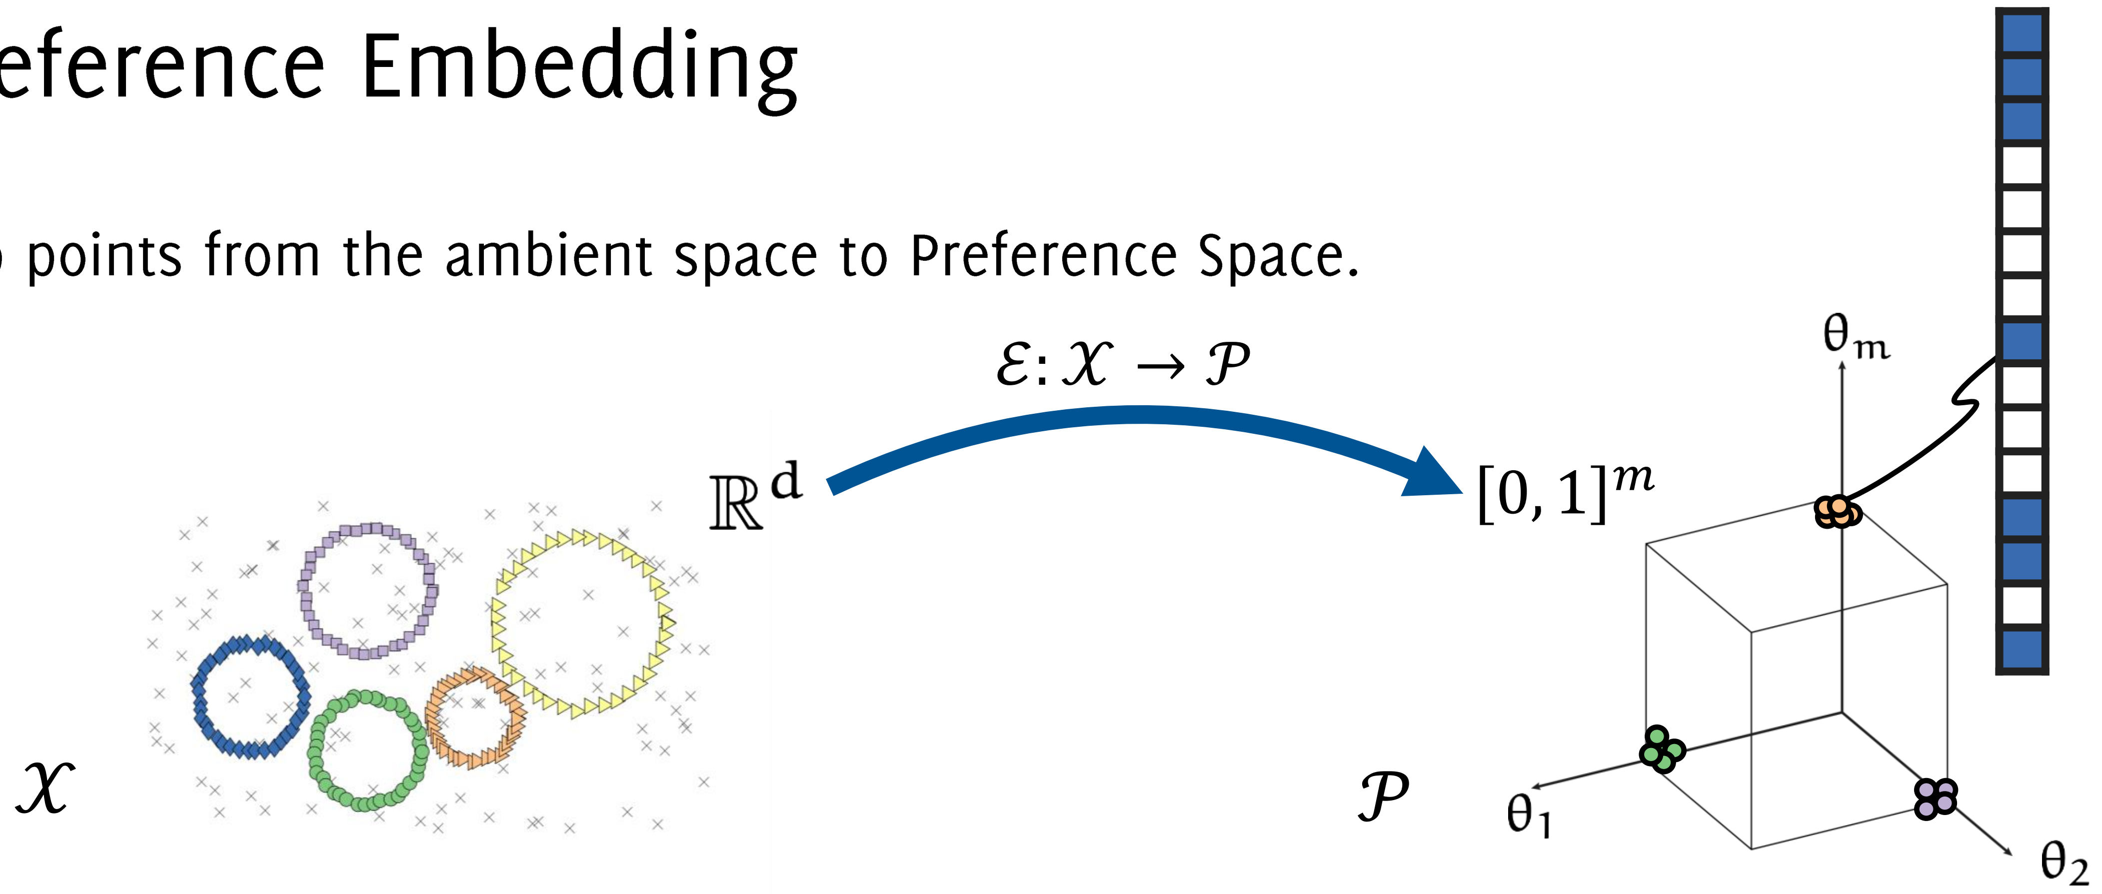

L. Magri, A. Fusiello, *"T-linkage: A continuous relaxation of j-linkage for multi-model fitting"*, IEEE CVPR, 2014

i) Preference Embedding

Map points from the ambient space to Preference Space.

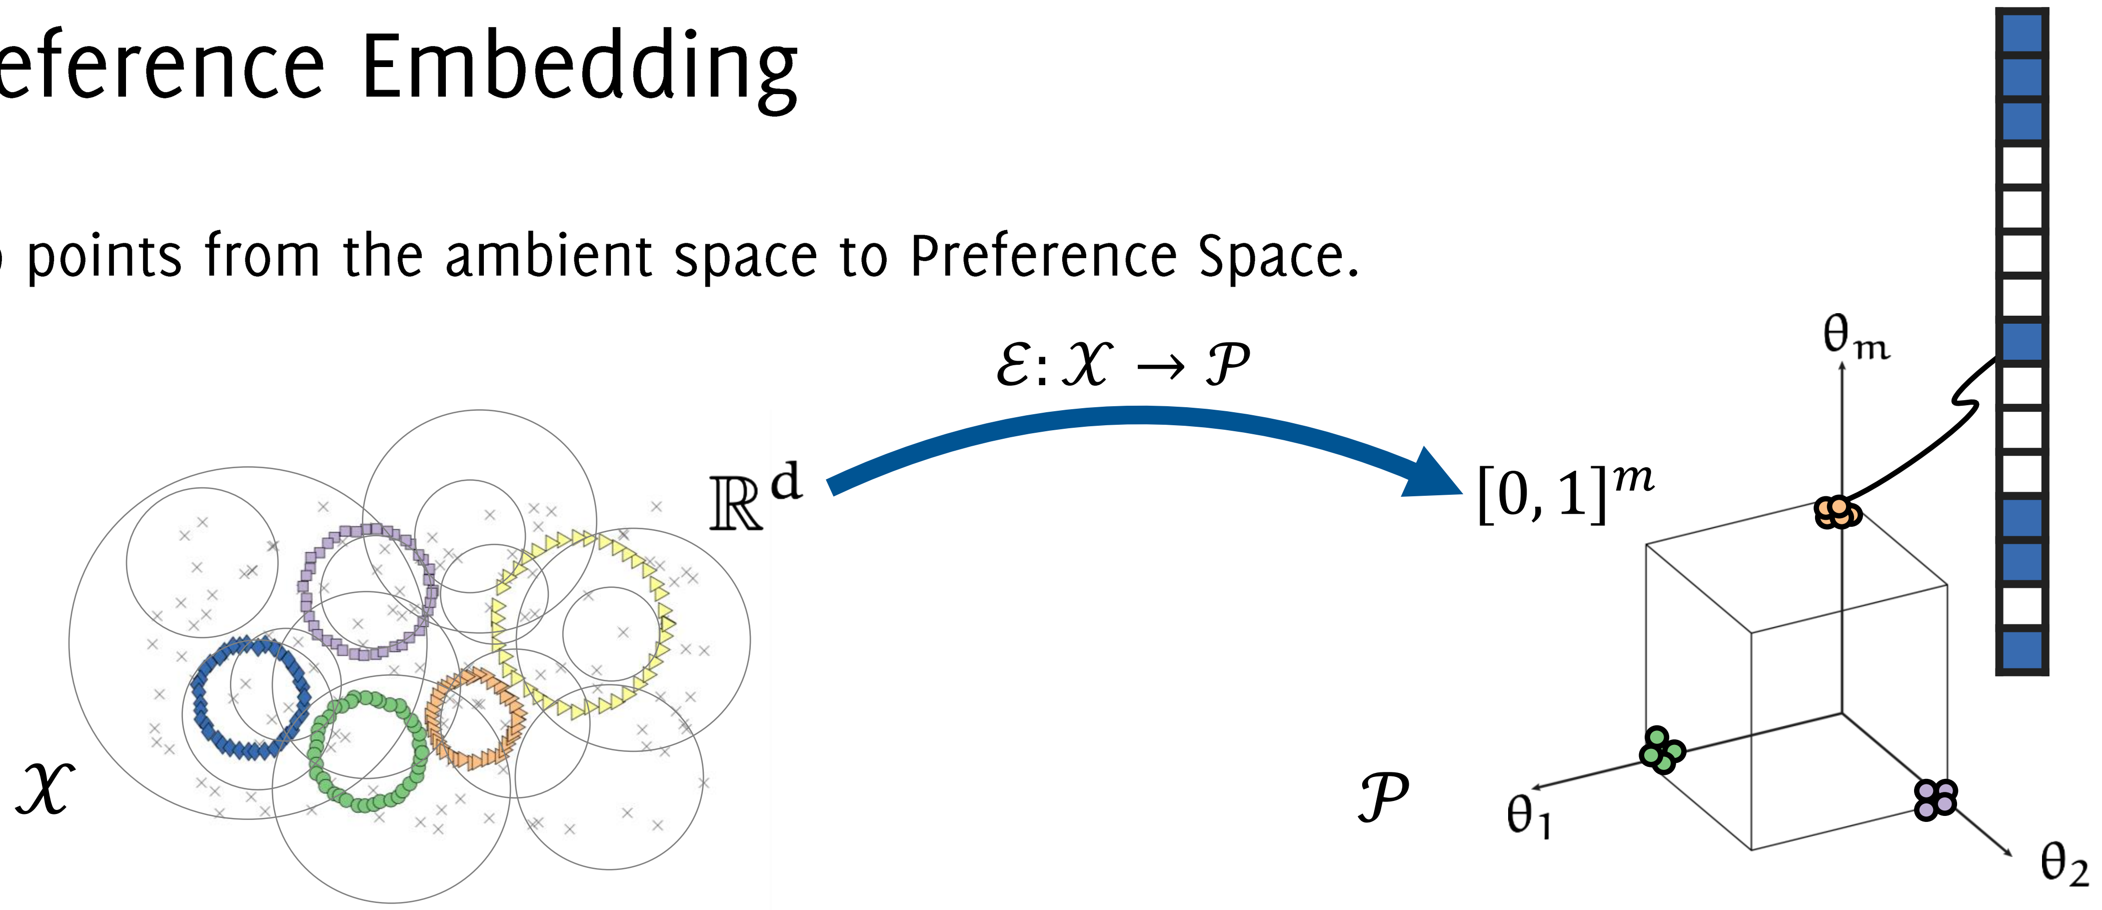

A set $\{\boldsymbol{\theta_1}, \dots, \boldsymbol{\theta_m}\}$ of models from family $\mathcal{F}$ are fit on $\mathcal{X}$ by randomly sampling minimum sets (with $m \gg 0$).

L. Magri, A. Fusiello, *"T-linkage: A continuous relaxation of j-linkage for multi-model fitting"*, IEEE CVPR, 2014

i) Preference Embedding

To each model-point pair we associate a preference value in [0,1]:

- 0 means $\boldsymbol{x}$ does not belong the the model $\boldsymbol{\theta_i}$, namely $|\mathcal{F}(\boldsymbol{x}, \boldsymbol{\theta_i})| > \boldsymbol{\epsilon}$
- ≈ 1 means $\boldsymbol{x}$ has a strong preference for $\boldsymbol{\theta_i}$, namely $|\mathcal{F}(\boldsymbol{x}, \boldsymbol{\theta_i})| \approx 0$

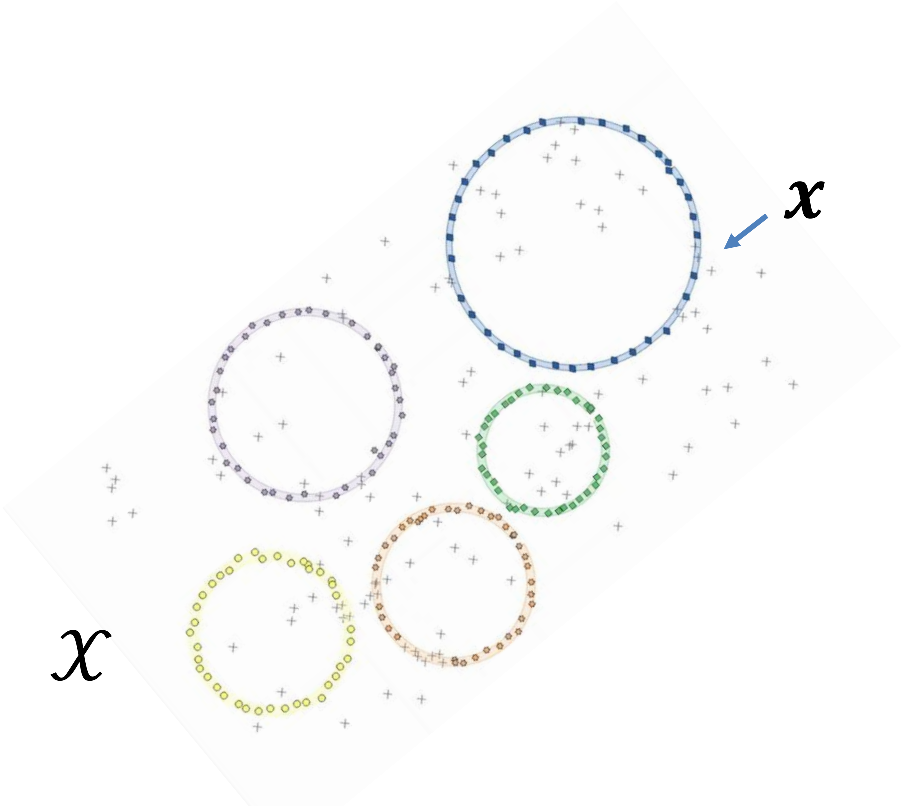

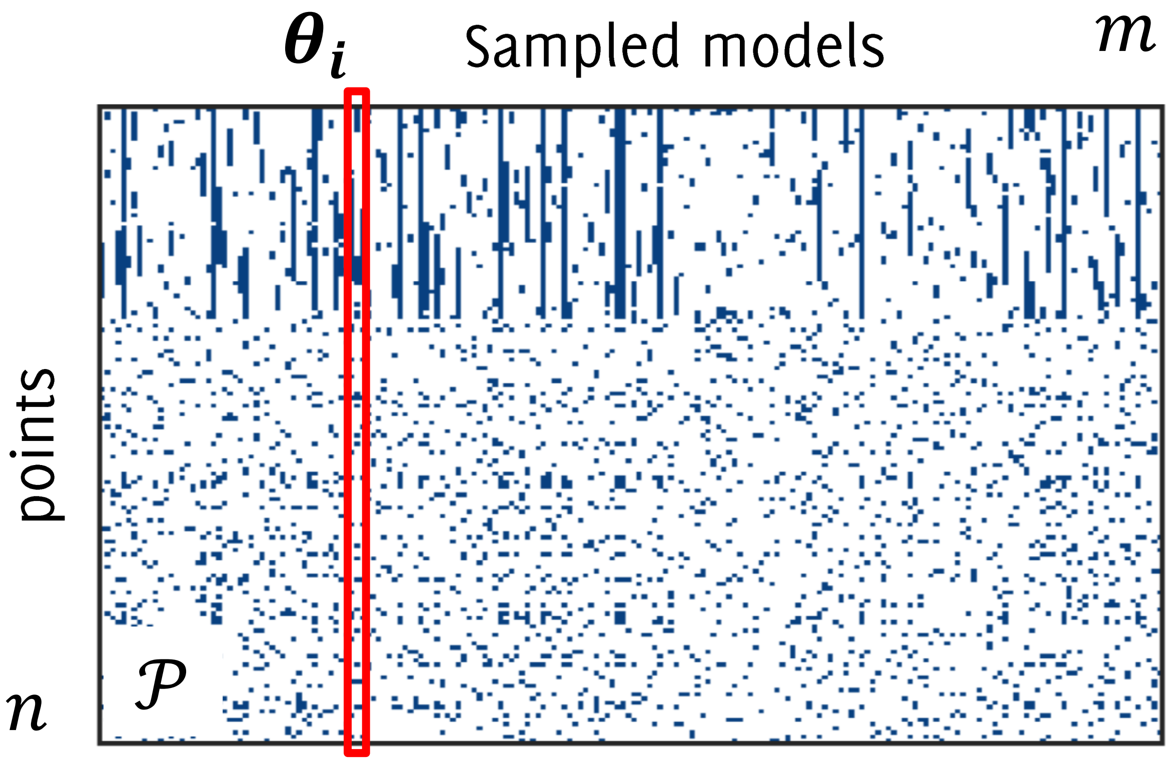

L. Magri, A. Fusiello, *"T-linkage: A continuous relaxation of j-linkage for multi-model fitting"*, IEEE CVPR, 2014

i) Preference Embedding

The prefence embdding consists in associating each point $\boldsymbol{x}$ to all its preferences w.r.t. the m sampled models.

$$\mathcal{E}: \boldsymbol{x} \mapsto \boldsymbol{p_x} \quad \text{such that } [\boldsymbol{p_x}]_{\boldsymbol{i}} = \begin{cases} \phi(\mathcal{F}(\boldsymbol{x}, \boldsymbol{\theta_i})), & \text{if } |\mathcal{F}(\boldsymbol{x}, \boldsymbol{\theta_i})| < \boldsymbol{\epsilon} \\ 0\,, & \textit{otherwise} \end{cases}$$

This is mapping $\mathcal{E}: \mathcal{X} \to [0,1]^m$ corresponds to the **preference embedding**

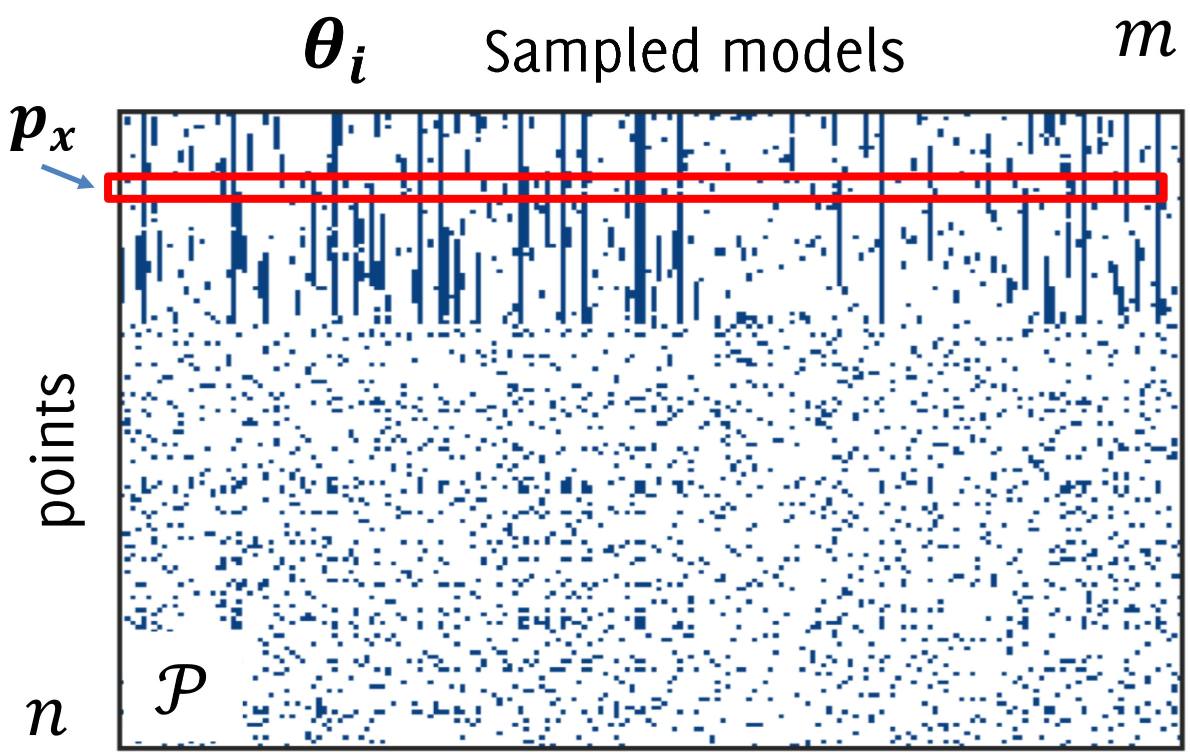

L. Magri, A. Fusiello, *"T-linkage: A continuous relaxation of j-linkage for multi-model fitting"*, IEEE CVPR, 2014

ii) Preference Isolation

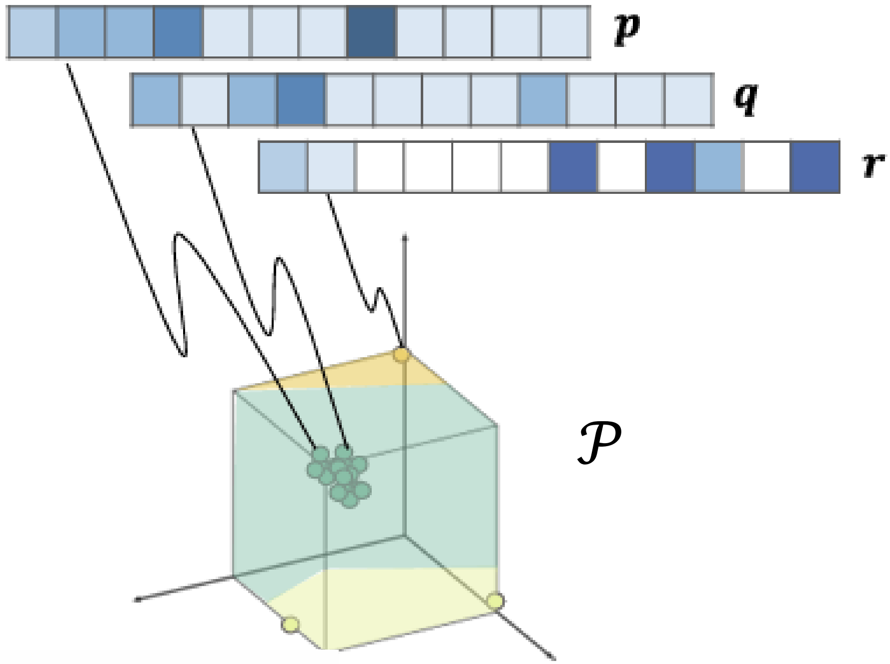

In the preference space:

- genuine points cluster together
- anomalies appear far apart

Idea: use isolation-based AD schemes to detect anomalies

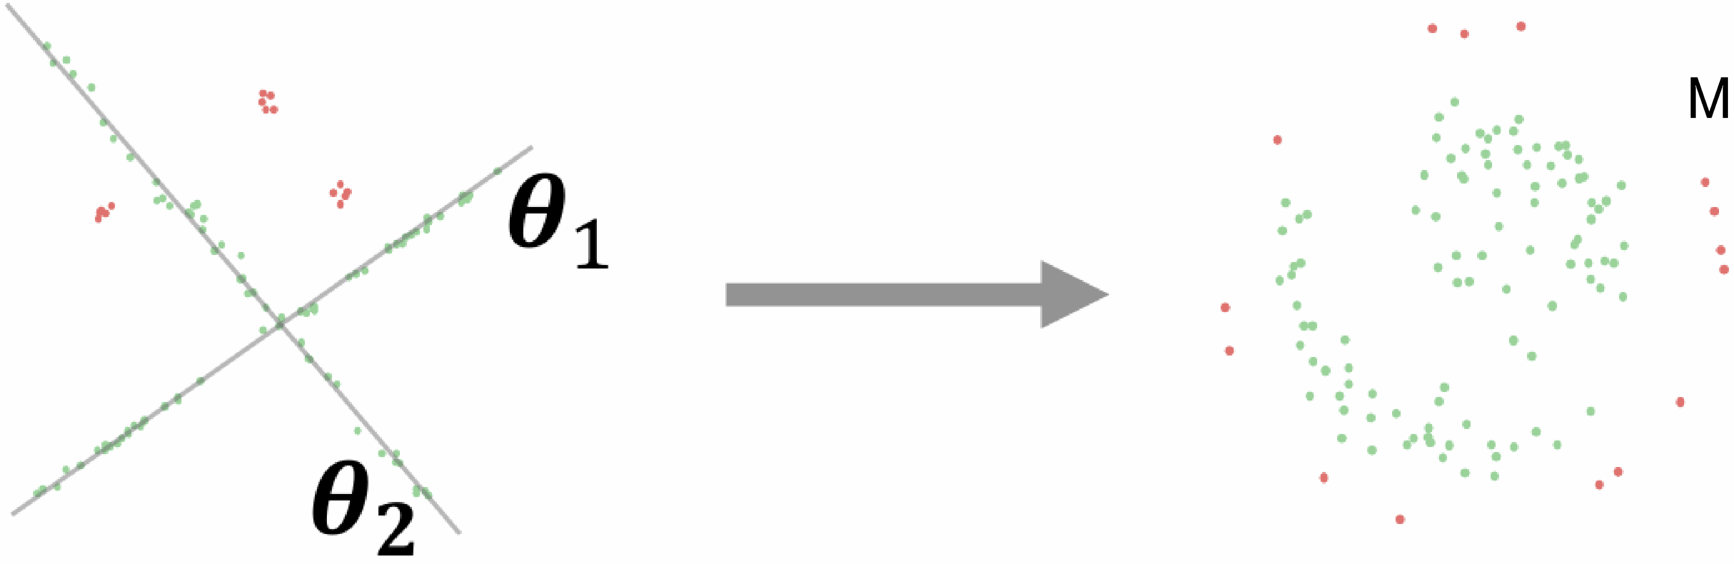

MDS($\mathcal{P}$)

Visualized via MDS, distances in $\mathcal{P}$ measured by Tanimoto

Filippo Leveni, Luca Magri, Giacomo Boracchi, Cesare Alippi, "*Anomaly detection via preference embedding*" International Conference on Pattern Recognition (ICPR) 2020,

Preference Isolation

Any **isolation-based AD** can be used in $\mathcal{P}$, returning an anomaly score.

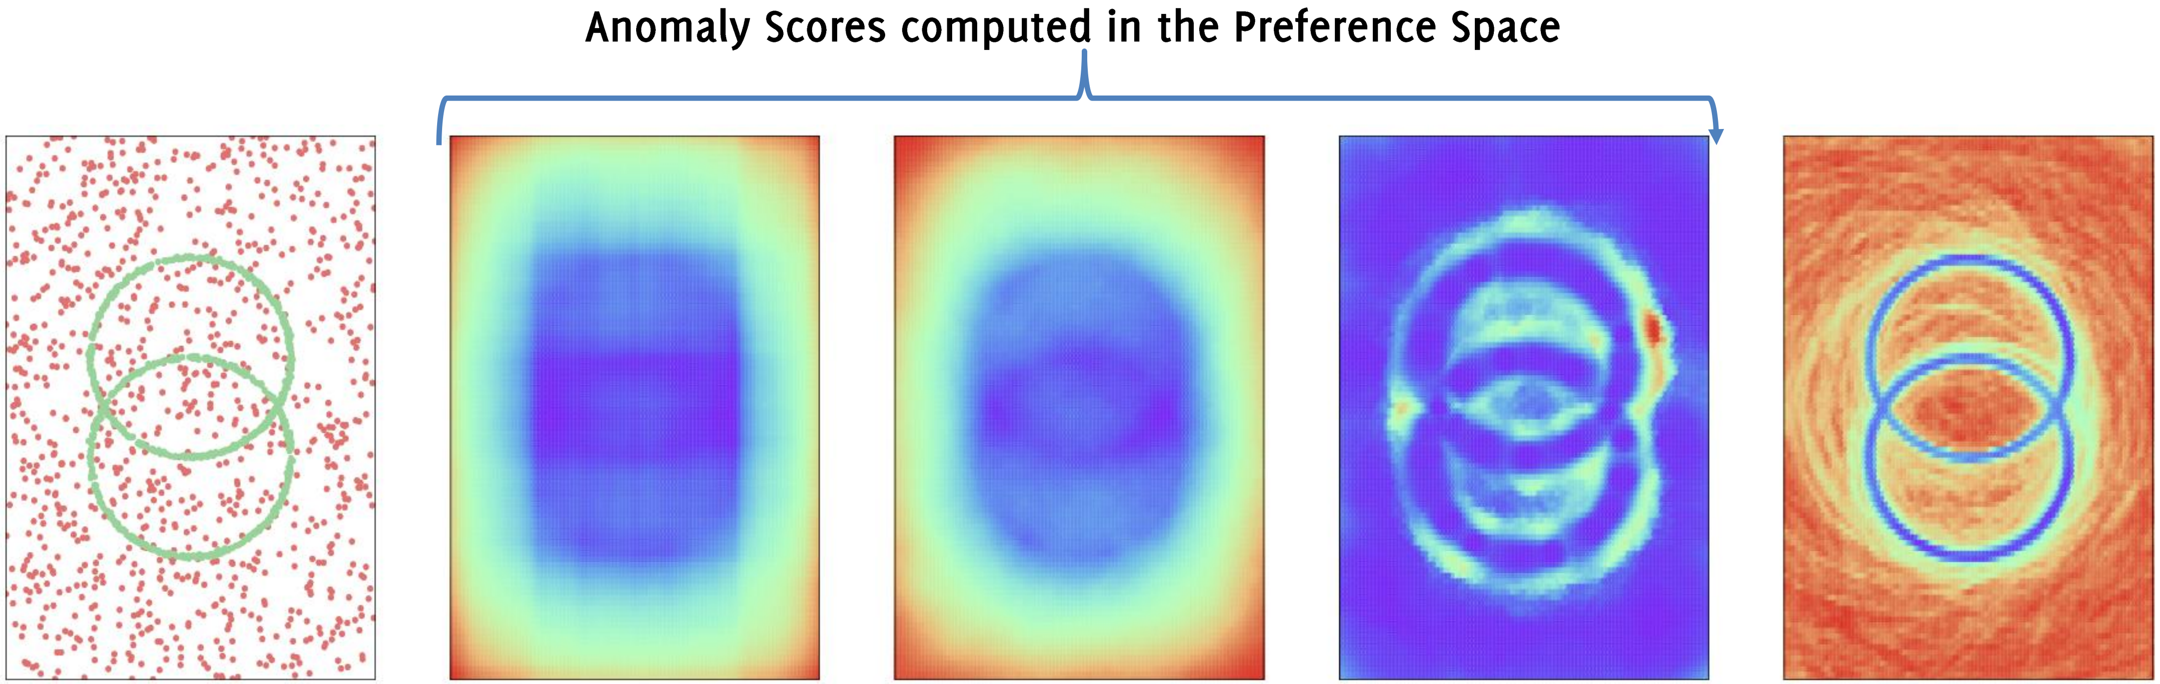

(a) Input data. (b) iFor [14] (c) EiFor [15] (d) LOF [13] (e) PIF

However, distances in the preference space are not the same as in $\mathcal{X}$, and isolation-based methods based on Euclidean distance fall short in AD.

Filippo Leveni, Luca Magri, Cesare Alippi, Giacomo Boracchi *"Preference Isolation Forest for Structure-based Anomaly Detection"* Under Review

Preference Isolation

Distances in the Preference Space are better measured by *Jaccad, Tanimoto, Ruzicka* distances

To effectively detect anomalies in $\mathcal{P}$ we construct a forest of Voronoi tesselations (Voronoi i-Tree), based on any of the above distances

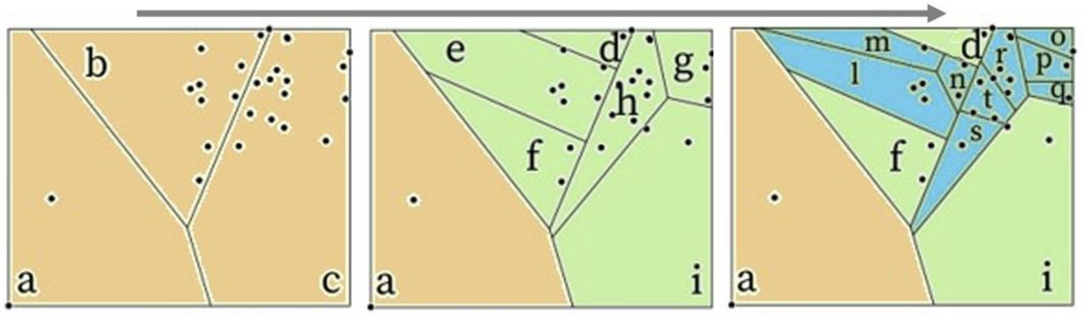

An efficient variant, where trees are constucted based on hashing of the Ruzicka distance has been also investigated.

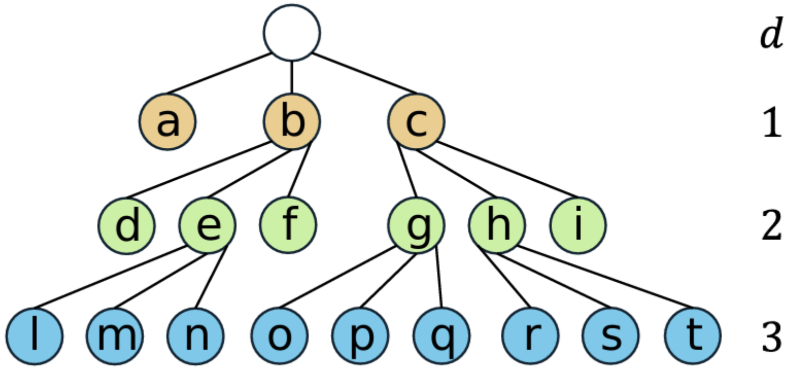

Filippo Leveni, Luca Magri, Cesare Alippi and Giacomo Boracchi ***"Hashing for Structure-based Anomaly Detection"*** **International Conference on Image Analisys and Processing (ICIAP) 2023 Best Paper Award**

Anomaly Score

The anomaly score is computed from the forest ψ as in IFOR,

$$\alpha_\psi(\boldsymbol{x}) = 2^{-\frac{E[D(\boldsymbol{x})]}{c(\psi)}}$$

where

- $E[D(\boldsymbol{x})]$ is the average depth reached by a point
- $c(\psi)$ is a normalizing factor depending on the forest ψ

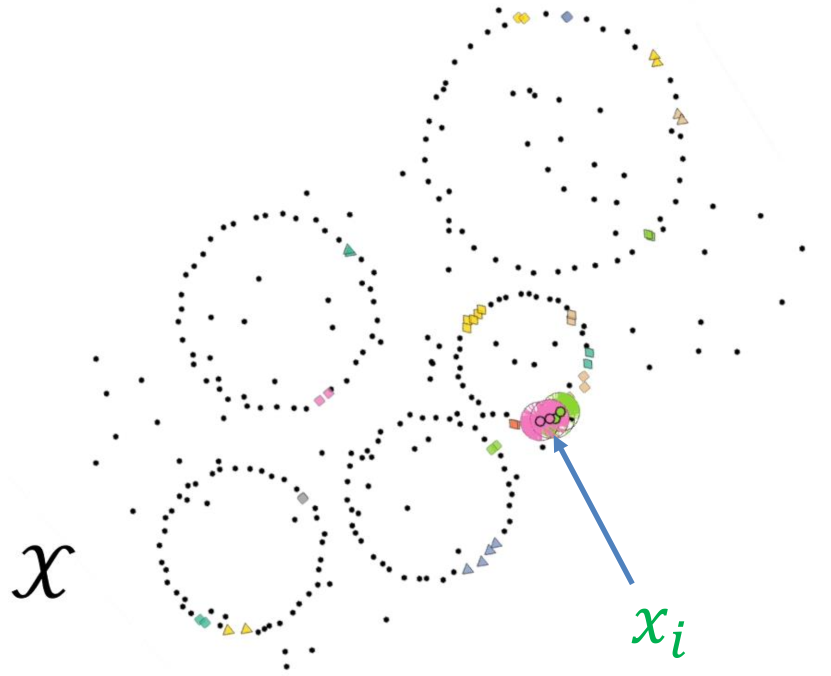

This score underpins the intuition that normal samples reaches leaves at larger depths than anomalous ones, which can be easily isolated by a few splits

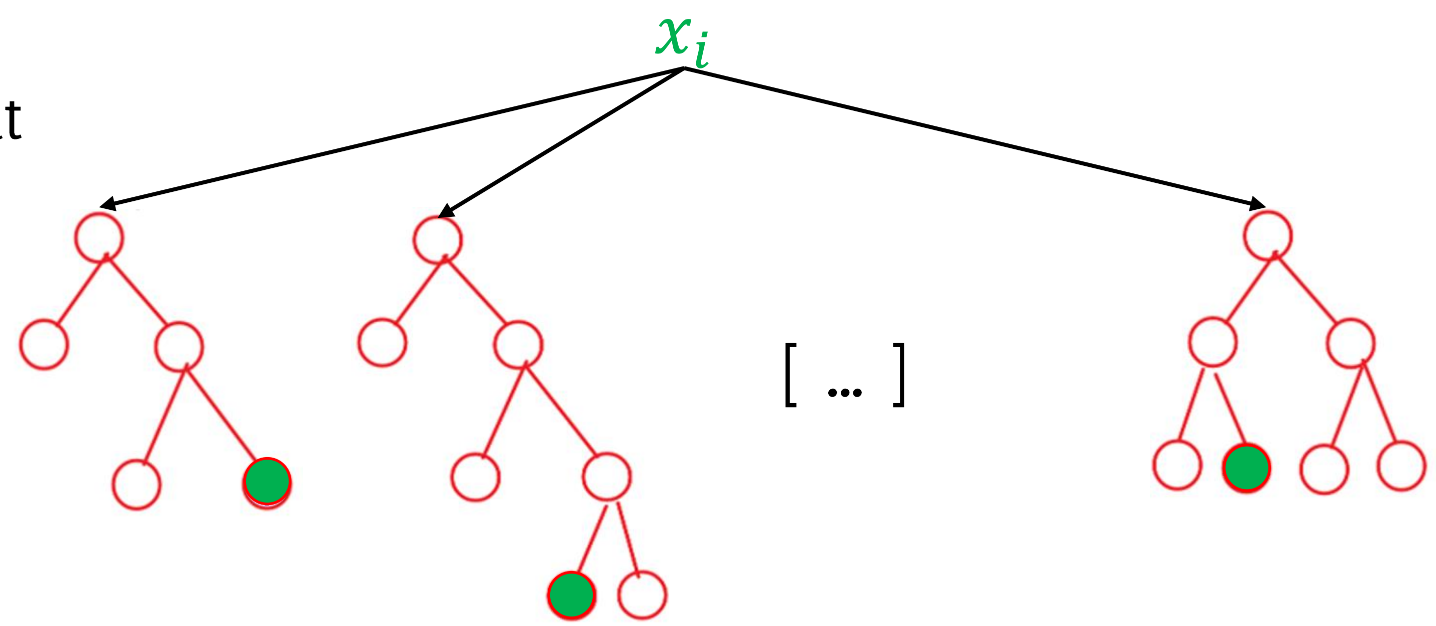

Anomaly Score

The anomaly score is computed from the forest ψ as in IFOR,

$$\alpha_{\psi}(\boldsymbol{x}) = 2^{-\frac{E[D(\boldsymbol{x})]}{c(\psi)}}$$

where

- $E[D(\boldsymbol{x})]$ is the average depth reached by a point
- $c(\psi)$ is a normalizing factor depending on the forest ψ

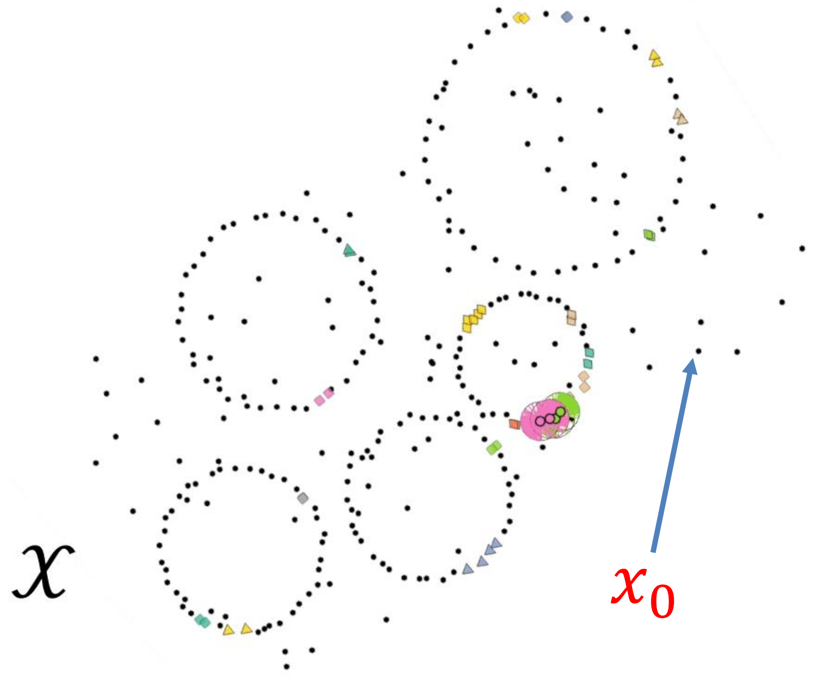

This score underpins the intuition that normal samples reaches leaves at larger depths than anomalous ones, which can be easily isolated by a few splits

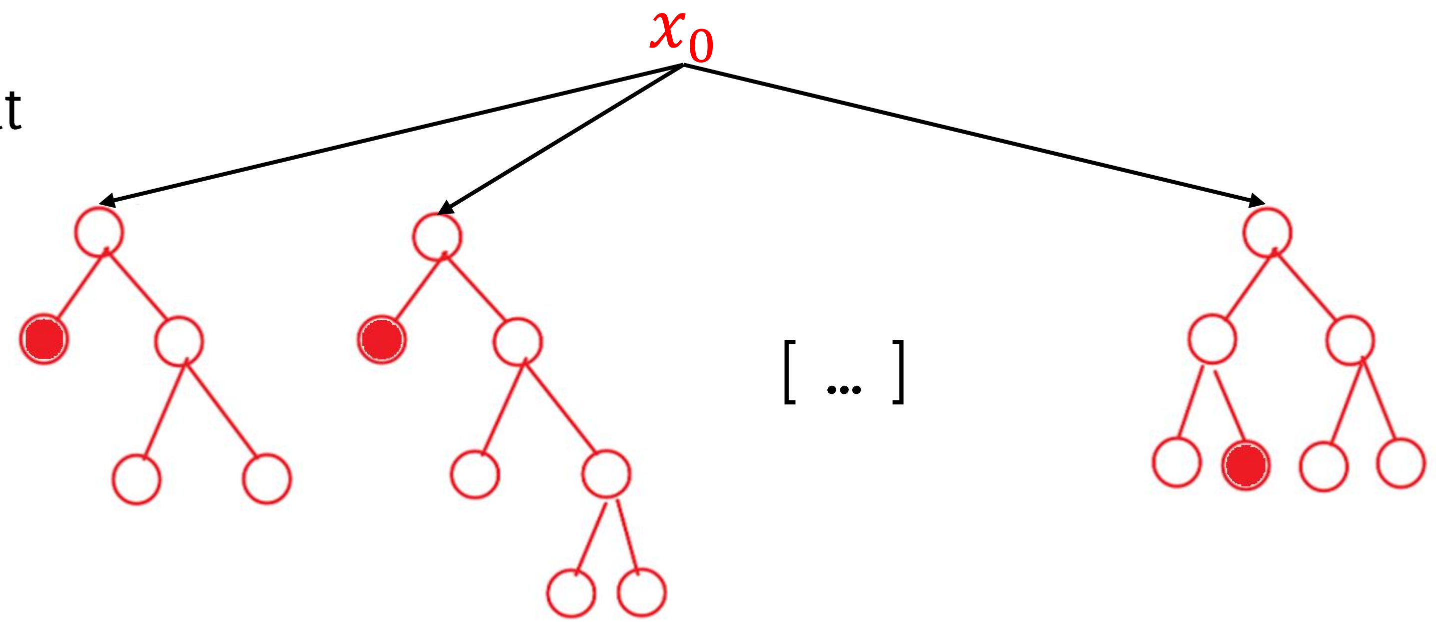

Sliding PIF

Sometimes only *local* knowledge about normal model family $\mathcal{F}$ is available.

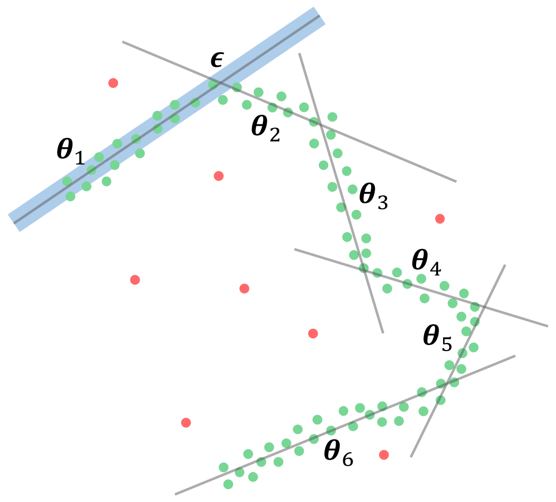

We can sample models $\boldsymbol{\theta_i}$ belonging to a *local* family by performing PIF in a sliding-window fashion.

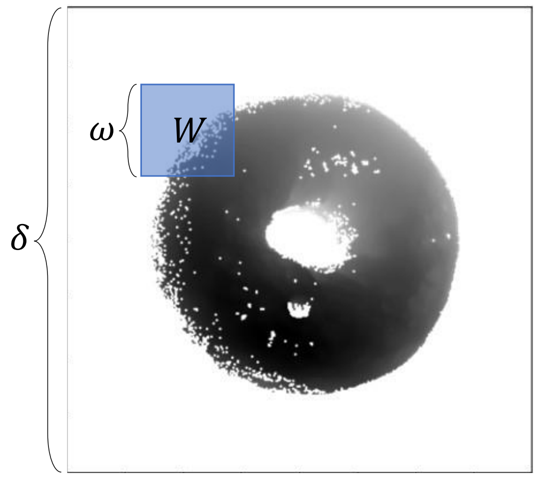

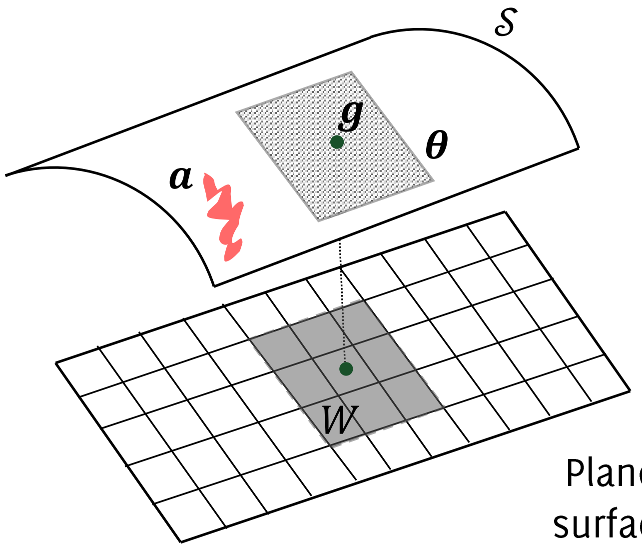

Planes to approximate a surface $\mathcal{S}$ belonging to an unknown family $\mathcal{F}$.

Sliding PIF

It is sufficient to opt for any model family $\mathcal{F}$ matching the normal *manifold's dimensionality* to estimate $\{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_m\}$ for the Preference Embedding.

We might generalize to broader families of matching dimensionality and allow for *more flexible* local approximations.

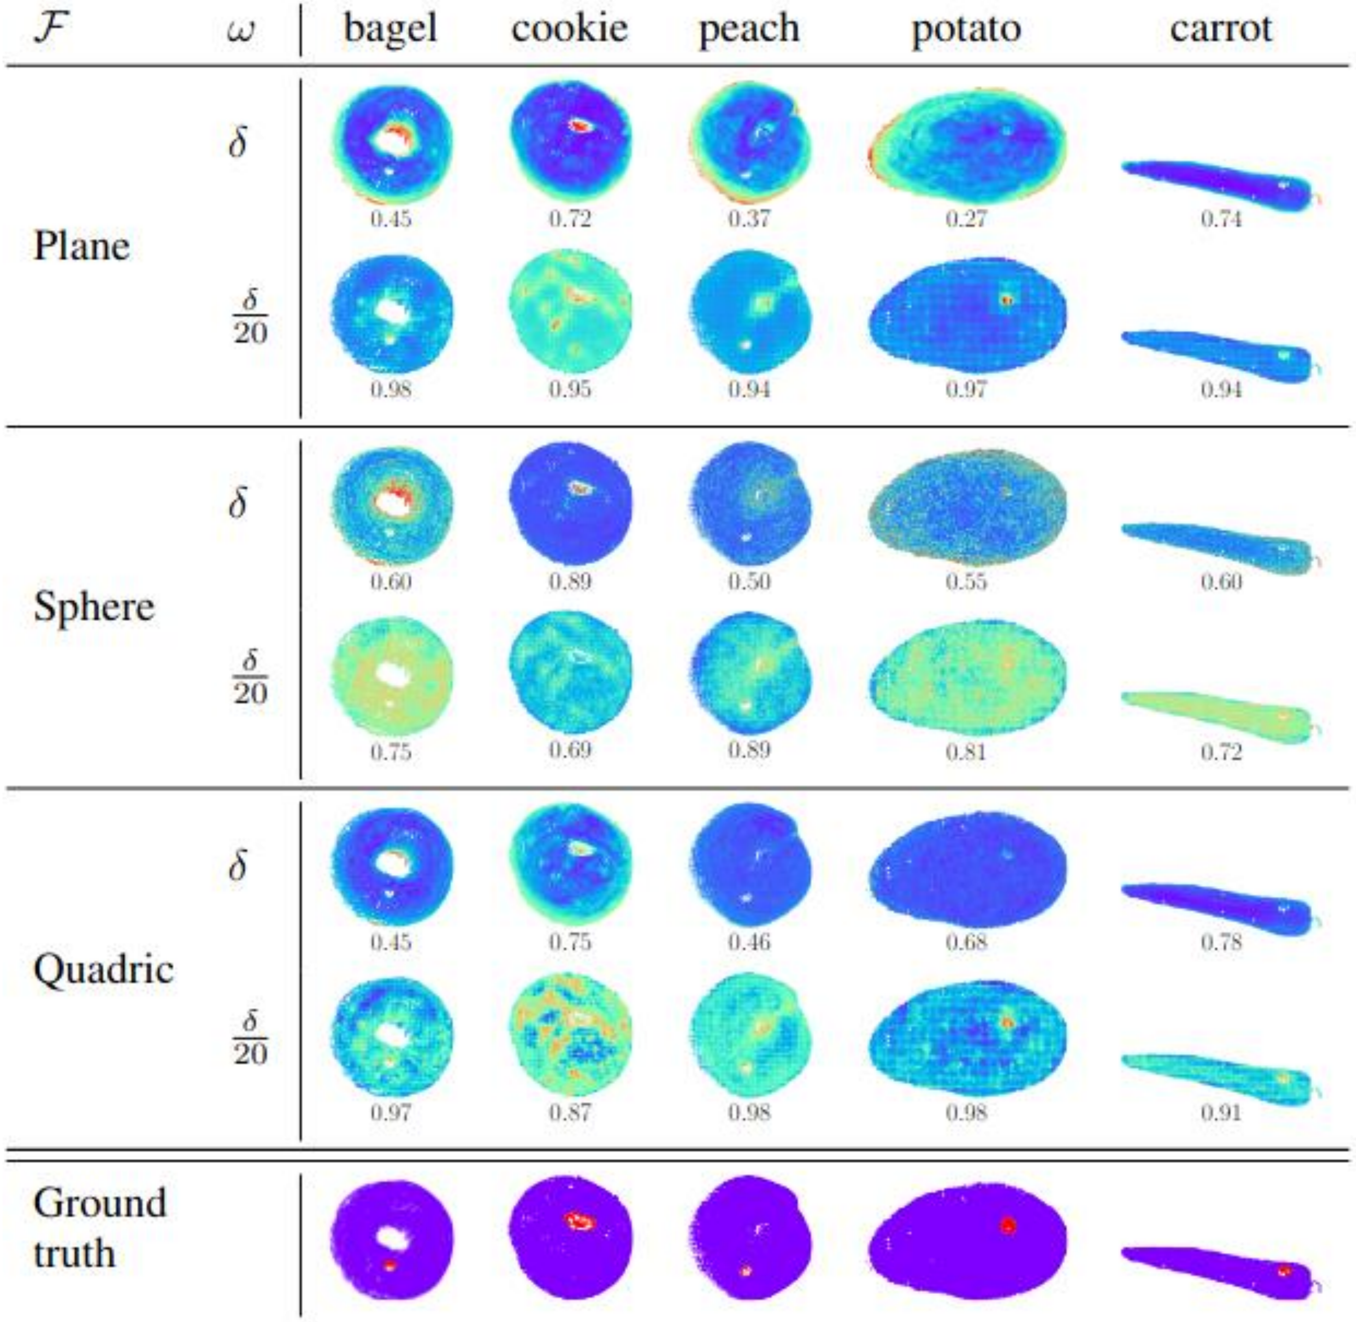

Concluding Remarks

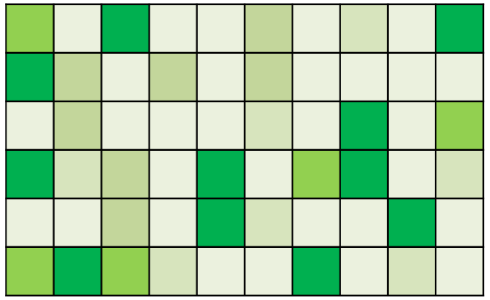

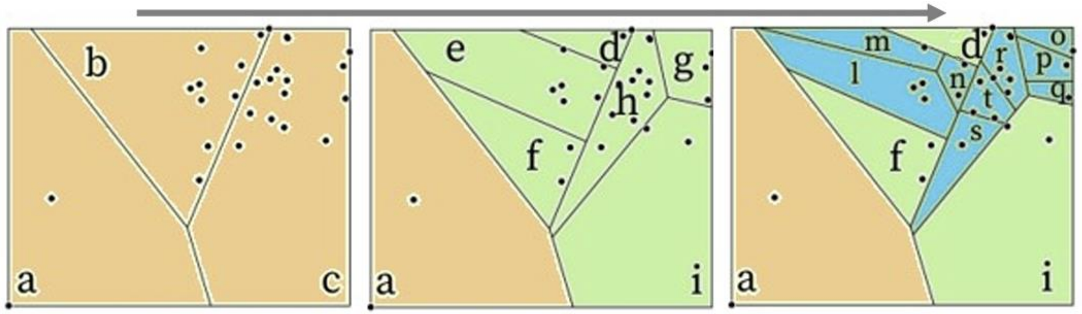

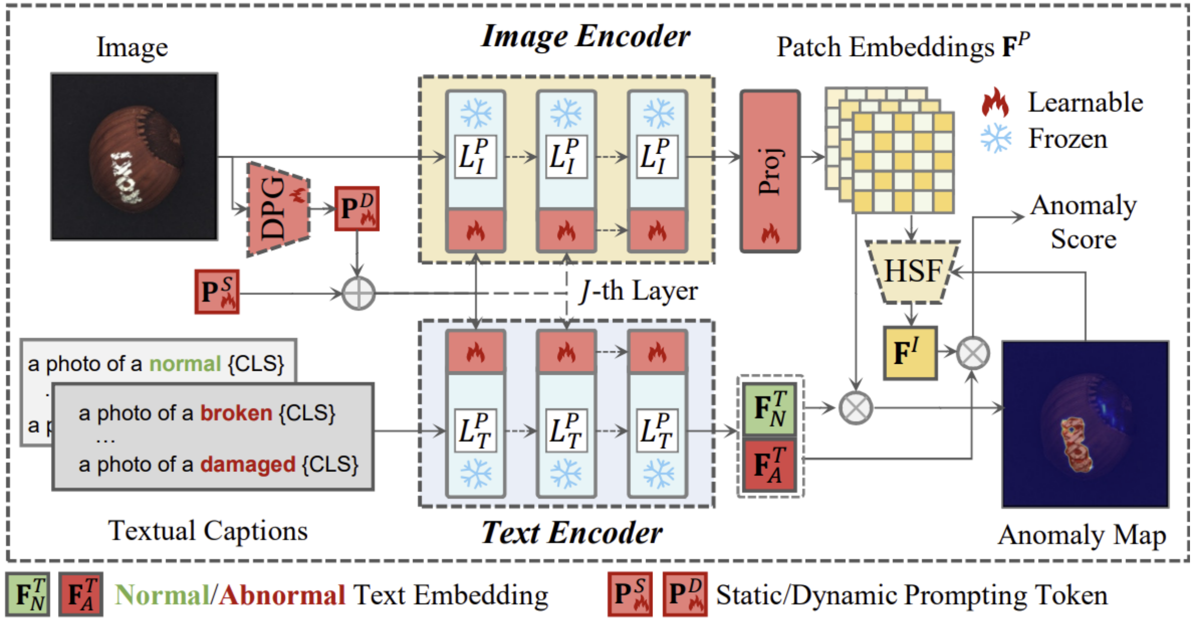

Anomaly detection is a very hot research problem with multiple applications

The field underwent the «deep learning tsunami», but still faces a dual nature

- Model based
- Statistical

That need to be taken into account when designing novel solutions.