Thesis Opportunities AA19/20

Giacomo Boracchi, Politecnico di Milano, DEIB

http://home.deib.polimi.it/boracchi/

Updated on October 16th, 2019

Reserach topics for a thesis

Image processing and deep learning

- Anomaly detection
- Image restoration and enhancement
- Deep learning for unconventional data
- Object detection by advanced template matching
- X-ray Multispectral Image Analysis

Change/Anomaly Detection in Datastreams

- Anomaly detection based on ensemble of histograms
- Change detection in datastreams

Machine learning for health

- ML methods for Seismocardiogram and ECG signal analysis in wearable devices
- Machine Learning for diagnosing nasal pathologies through CFD
- Biomarkers for Huntington Disease

Stage Opportunities

- @STMicroelectronics (Agrate): 3 calls
- @Fondazione Don Gnocchi (Milano)

These slides…

…provide both:

- A short overview of the research background (**Background** slides)
- Possible research directions along which to develop a thesis under our supervision (**Research directions** slides)
- Research institutions / companies involved in these research activities
- Stage opportunities

The Team

You'll be supervised by myself and another experienced colleague.

Here we are:

Giacomo Boracchi

Diego Carrera

Luca Frittoli

Filippo Leveni

Luca Magri

Simone Melzi

Andrea Schillaci

Diego Stucchi

For any enquiry

Just drop me an email and we can arrange a short meeting to discuss…

giacomo.boracchi@polimi.it

Image processing and deep learning

Deep Learning for Anomaly Detection

Background

Anomaly detection is a very general problem with applications ranging from: quality inspection to health and industrial monitoring

It is typically treated as an unsupervised problem, being anomalies unknown

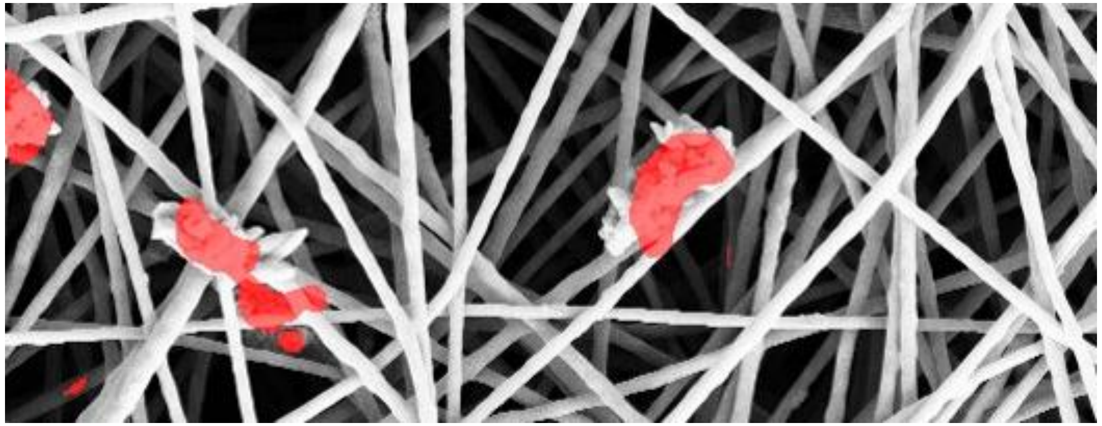

nanofiber production

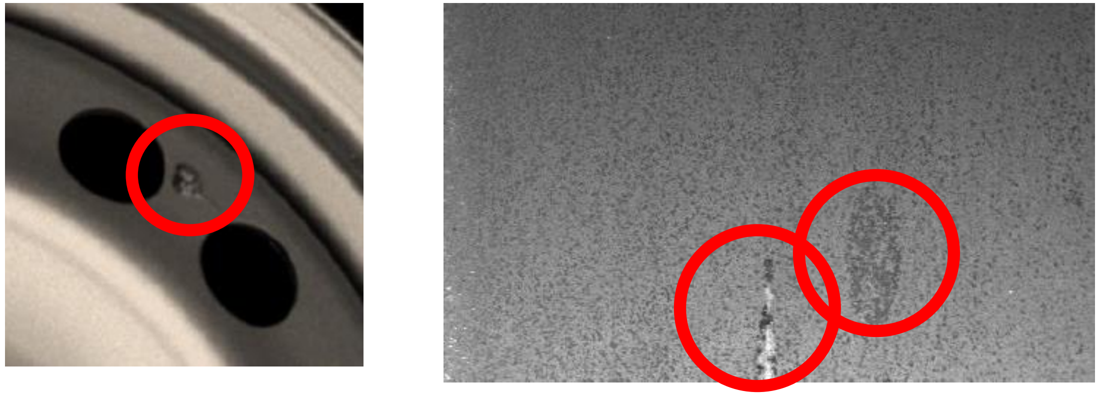

manufacturing

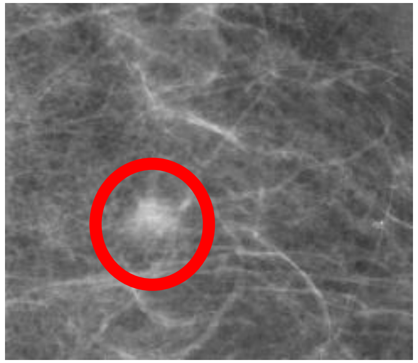

mammograms

Deep Learning for Anomaly Detection

Background

Deep Learning (DL) represents the state-of-the-art in supervised tasks such as classification and semantic segmentation

Unsupervised DL methods have been studied less, and the use of supervised techniques in industrial or health-related monitoring problems might not be straightforward

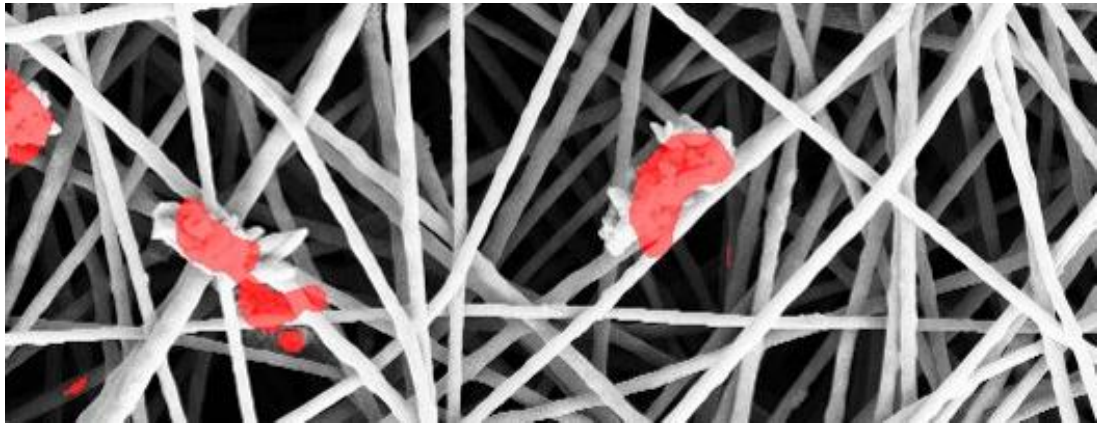

nanofiber production

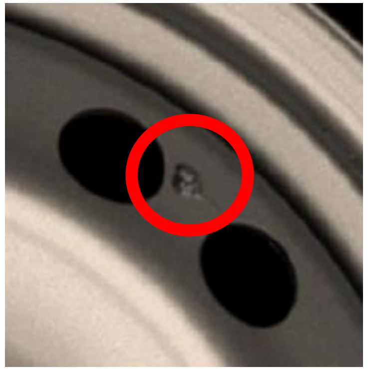

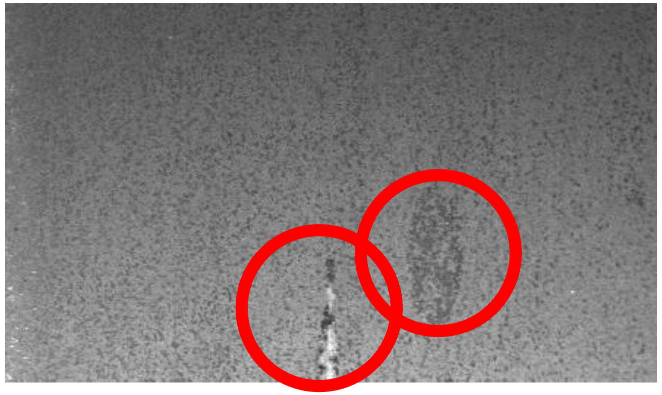

manufacturing

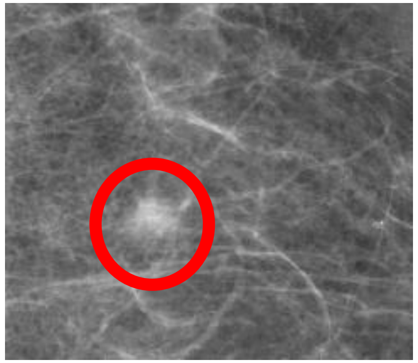

mammograms

GANs and Anomaly Detection

Background

Generative Adversarial Networks (GANs): very effective generative models for images. These have been applied in a plethora of applications.

GANs are appealing for anomaly detection as they successfuly parametrize the manifold where images live

Images generated after 5 training epochs

Image aritmetric: examples similar to word embedding

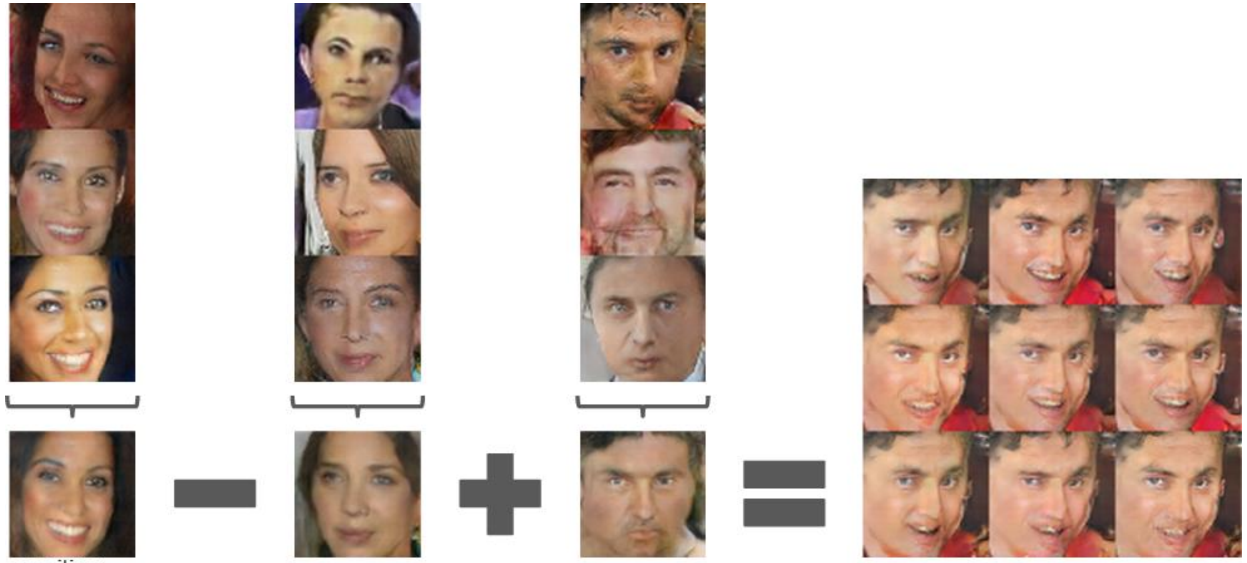

Radford, A., Metz, L., & Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. ICLR 2016

Thesis: Deep Learning for Anomaly Detection

Research Directions

Research goals for a few thesis:

- Design **new training methods for (deep) autoencoders** that are specifically designed for AD.
- Investigate how **GANs can be used to learn data representations** and to augment training data
- Investigate **sliding-window solutions** (such as Fully Convolutional CNN) in GAN-based detectors
- Develop AD methods able to deal with **inaccurate reference templates**

Materials and Methods:

- Access to a server mounting **GPUs will be provided**
- **Annotated training sets** from industrial production
- Reference algorithm based on sparse-representations

Background

Image restoration and enhancement

Noise always affects digital images and video frames

$$z(x,t) = y(x,t) + \eta(x,t)\,, \qquad x \in X\,, \eta \sim \mathcal{N}(0,\sigma^2)$$

Denoising algorithms provide $\hat{y}$, an estimate of the noise-free image y.

z

$\hat{y}$

Image restoration and enhancement

Background

Restoration algorithms have been widely investigated for their many applications in science and engineering

These often leverage **data-driven models** (including deep Convolutional Neural Networks, **CNN**) to provide effective representations of noise-free images.

Image restoration and enhancement

Artifacts to be removed by restoration algorithms include noise and blur.

Both de-noising/de-blurring algorithms and the image formation process for this kind of artifacts have been widely investigated in the literature.

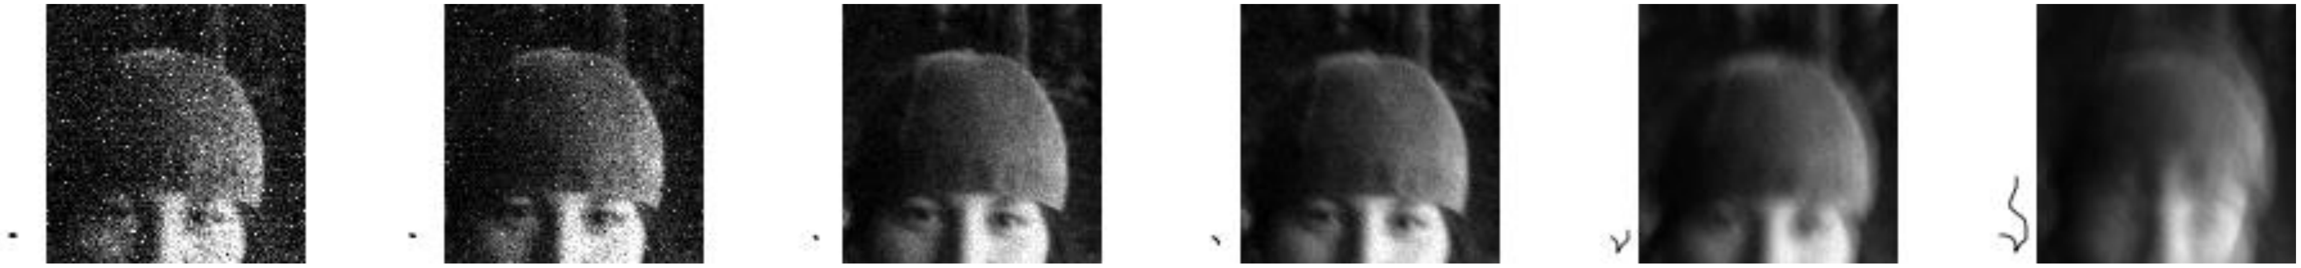

G. Boracchi and A. Foi, *Modeling the Performance of Image Restoration from Motion Blur*, IEEE Transactions on Image Processing, vol. 21, no. 8, pp. 3502 - 3517, Aug. 2012,

Image restoration and enhancement

By leveraging image formation model and clean natural images it is possible to train powerful neural networks for

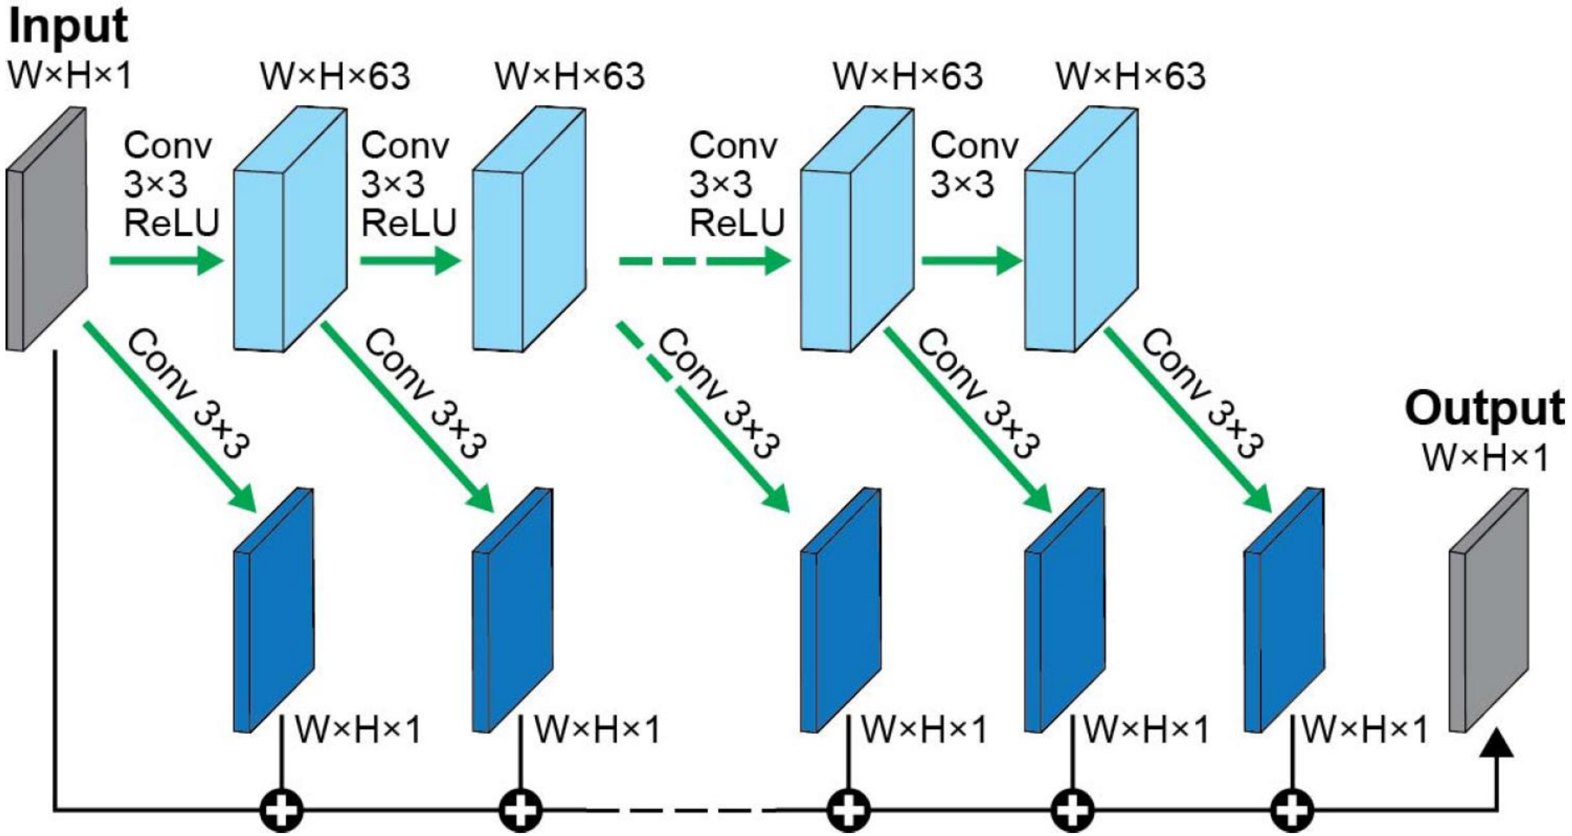

Remez, T., Litany, O., Giryes, R., & Bronstein, A. M. *Class-Aware Fully Convolutional Gaussian and Poisson Denoising*. IEEE TIP 2018

Thesis: Image Restoration by data-driven (deep) representations

Research Directions

Goals (possibly more than one thesis)

- Study **data-driven representations (CNN, autoencoders) of noise-free images** and their use in combination with **key principles of denoising** algorithms such as:
  - Multiscale processing, self-similarity
- Consider denoising as a «mapping problem» and train a network to **learn transformations** from corrupted images to «clean» ones
- Force **rotation invariance** through specific layers such as in RotEqNet

Materials and Methods:

- Libraries for adding noise and corrupt images in realistic manner
- Access to a server mounting GPUs will be provided
- Reference algorithms to compare with the proposed solution

Computer Vision-related Thesis

Our template matching algorithms in action

Our template matching algorithms in action

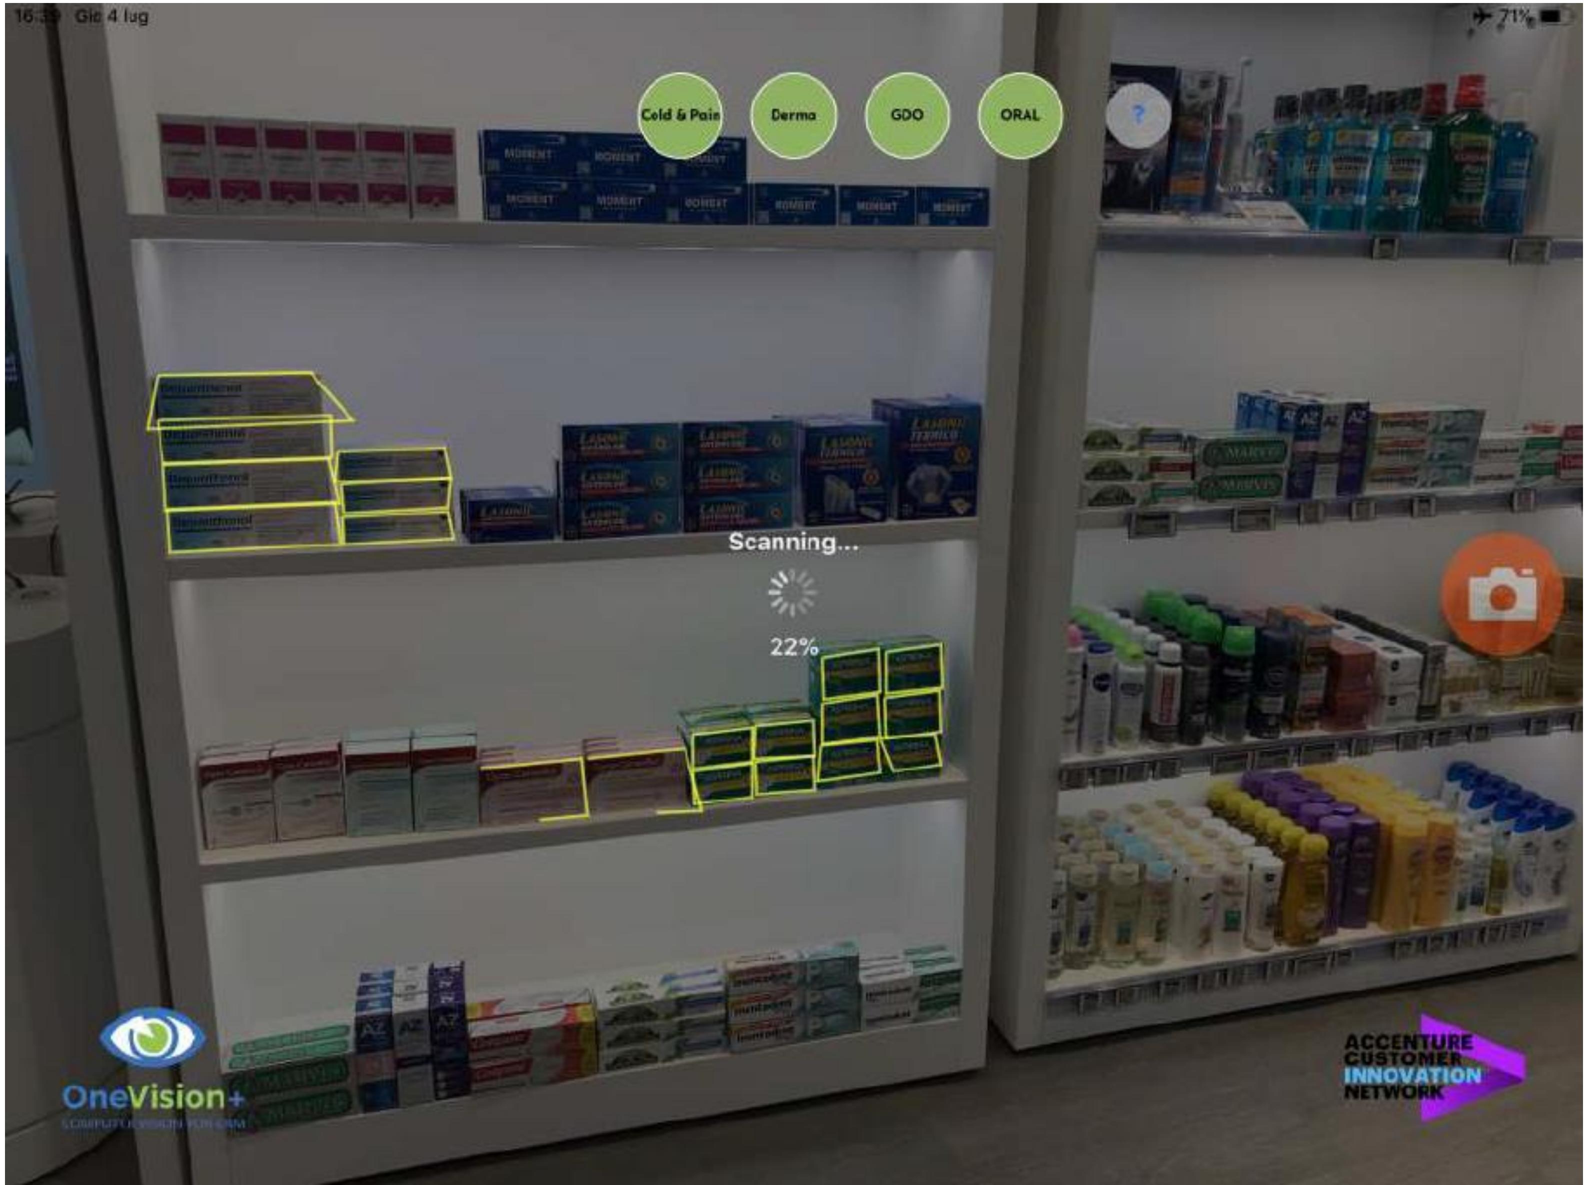

Our template matching algorithms in action

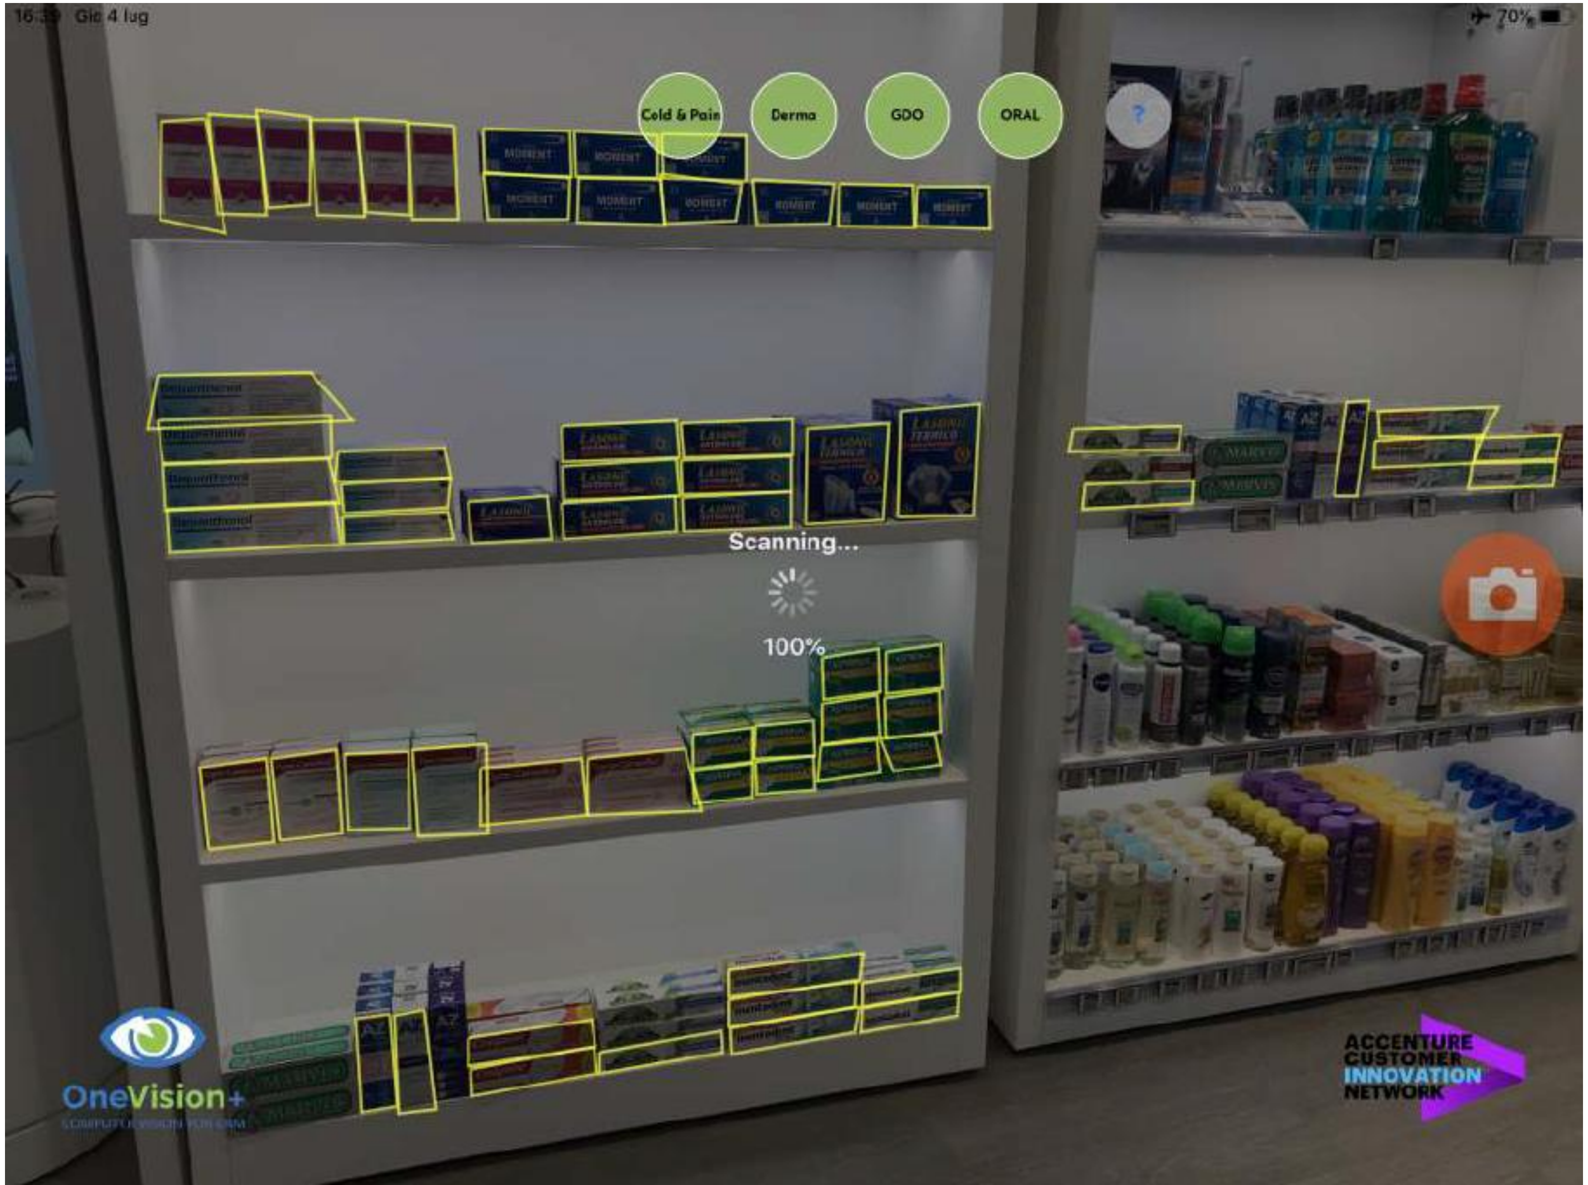

Object Detection by Advanced Template Matching

Background

Template matching is a very practical technique for finding all the occurrnces of a template image in a given query image.

Deep-learning techniques are not very flexible, requiring re-training for new templates.

Traditional computer vision techniques (e.g. SIFT + RANSAC) are more practical but require an underlying model describing the transformation between the template and its occurrences in the query image.

Research Directions

Thesis: Template matching for non-rigid distortions

Non-rigid distortions can not be modeled to enable RANSAC.

Goals: Design efficient clustering algorithms for feature matches that identify multiple, distorted, instances of the template in the query image:

- Studying state-of-the-art **template matching techniques** based on **image features**.
- Design of **learning-based** feature matching schemes.
- The design of an innovative technique to **group feature matches** that refer to the same template instance in the query image.

Materials and Methods:

- Access to a server mounting GPUs.
- Dataset from our industrial partner.

Research Directions

Thesis: Efficient template matching in large scale

Goals: Dynamically identify which template to search for first when the number of templates is very large (e.g. a few hundreds):

- Studying state-of-the-art **template matching techniques** based on **image features.**
- Implementation of learning-based criterion to **identify the template** that will be **most likely found** in a region of the query image.
- **Transfer learning** of pre-trained models with very little samples in segmentation.
- **Chromatic transformation** to improve matching.
- Definition of **new templates** from a single image.

Materials and Methods:

- Access to a server mounting GPUs.

Model fitting for computer vision

Background

Goals: Fitting one or multiple instances of a geometric model – also called "structures" – to measured data, which is invariably contaminated by noise and outliers.

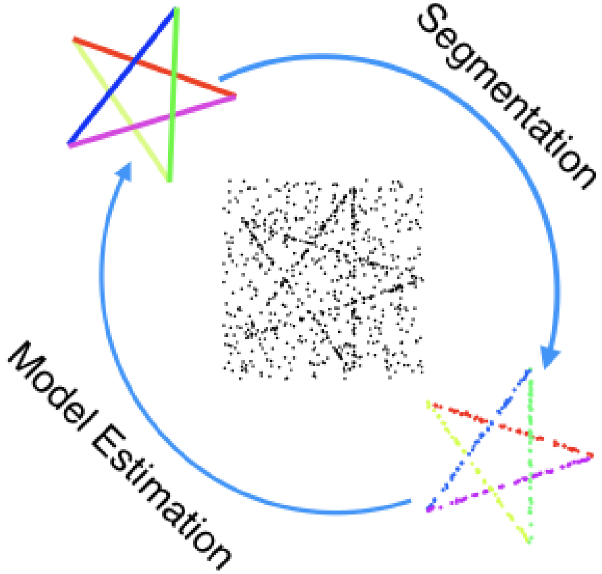

Research Challenges:

- outliers
- pseudo-outliers
- chicken-&-egg-dilemma
- ill-posed problem
- scale estimation
- model selection

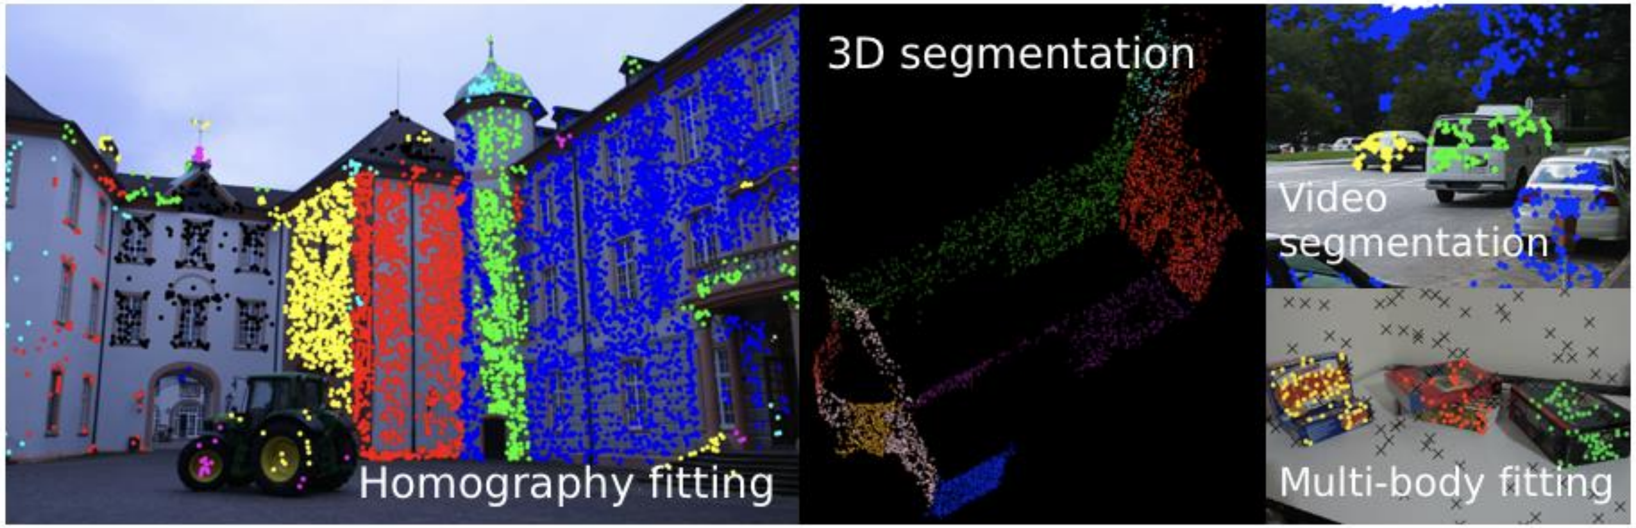

Model fitting for computer vision

Background

Possible apporaches (among others)

Consensus analysis: (e.g. Ransac, Hough transfrom)

1. represent models (columns of the matrix)
2. count the number of points explained by each model (consensus)
3. maximize consensus

Preference analysis: (e.g. Residual Histogram Analysis, T-Linkage, Higher order clustering)

1. represent points (rows of the matrix)
2. each point "votes" for the preferred models
3. clustering of points preferences

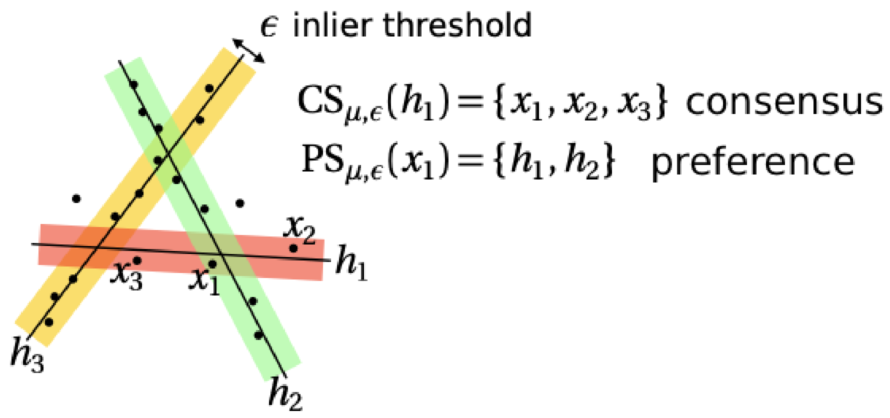

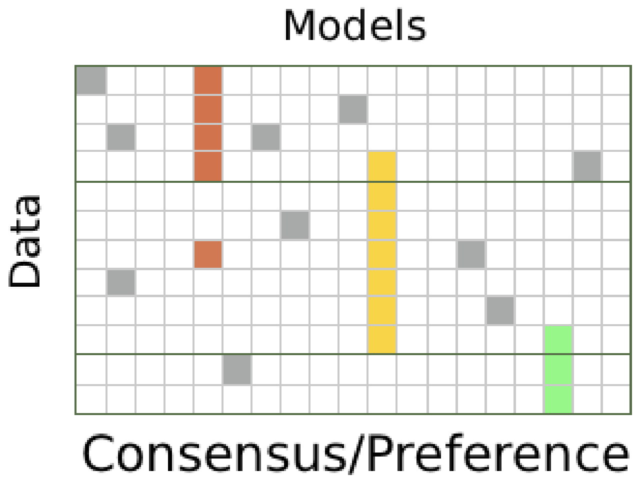

Thesis: Dealing with outliers

Research goals for a few thesis:

- Develop methods able to **accurately remove outliers** for geometric computer vision tasks
- Study and compare existing **outlier rejeciton techniques** (robust statistic, low rank decomposition, NFA, topological)
- Investigate how **the sparsity of data/model** representations can be exploit to detect outliers.
- In situations where there are many outlier points, any wrong model would achieve similar consensus. Correct models can be detected as models have anomalous consensus

Materials and Methods:

- Annotated datasets for two-view body segmentation and video segmentation.

Some reference to start with:

- Tepper and Sapiro, "Fast L1-NMF for Multiple Parametric Model Estimation", 2016

Thesis: Robust estimation in an uncertain word

Research Directions

Research goals for a few thesis:

- Investigate how the inherent uncertainty in data location, can be integrated in the estimation of geometric structures to obtain a more efficient robust fitting algorithm.
- Study the use of randomized model verification to characterize the 'non-randomness' of a solution.
- Explore the extension of this result to the case of multiple models.

Materials and Methods:

- Annotated datasets for two-view body segmentation and video segmentation.
- Develop a robust feature matching application

Some references to start with:

- Tordoff and Cipolla, "Uncertain RanSaC", Conference on Machine Vision Applications 2005.
- Raguram, Frahm, and Pollefeys. "Exploiting uncertainty in random sample consensus." ICCV 2009.

Nonconventioal imaging modalities

Deep Learning on unstructured data:

Background

Most DL algorithms are based on **discrete convolution**, an operation defined for data with an underlying **grid structure** (e.g. images)

Industrial data might lack such a structure. For instance, data might consist of a set of **3D points or measurements in scattered locations** (e.g. point clouds, defect maps, sensor networks acquisitions)

A popular solution is to **discretize** the data points and force them into a grid structure. However, this procedure might cause a **loss of information**

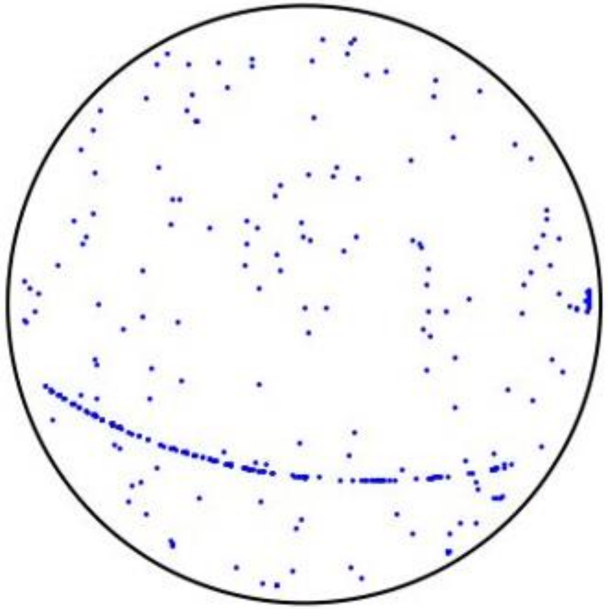

A Wafer Defect Map (WDM) is a list containing the coordinates of the defects on a wafer. Defects are displayed on a huge grid (20,000 x 20,000) that would be impossible to handle as an image

Point clouds are measurements scattered in the 3D space from LIDAR or RGB-D cameras. Acquisitions from sensor networks can also be conveniently modeled like that

Thesis: Continuous Convolutions to learn from unstructured data

Research Directions

Goals: develop new DL methods to handle and learn directly from sets of points. In particular, this project will address:

- The study of **continuous convolutions**, where the kernel is a **continuous function** instead of a discrete matrix
- The design of a **continuous convolutional** layer
- The development of a DL architecture based on continuous convolutions and validation of the solution

Materials and methods:

- Access to a server mounting GPUs will be provided
- Annotated from sparse wafer defect maps, public datasets point clouds
- This thesis can be possibly addressed during an internship at STMicroelectronics

Background

Multispectral Image Analysis

Multispectral and hyperspectral images gather data from multiple spectral bands. Each pixel, is thus associated to a spectrum of a few tens / few hundreds of intensity values, each representing a certain wavelength.

Relevant challenges in this field are classification, recognition and unmixing i.e. separating an acquired spectrum as the sum of few basic spectra (the endmembers)

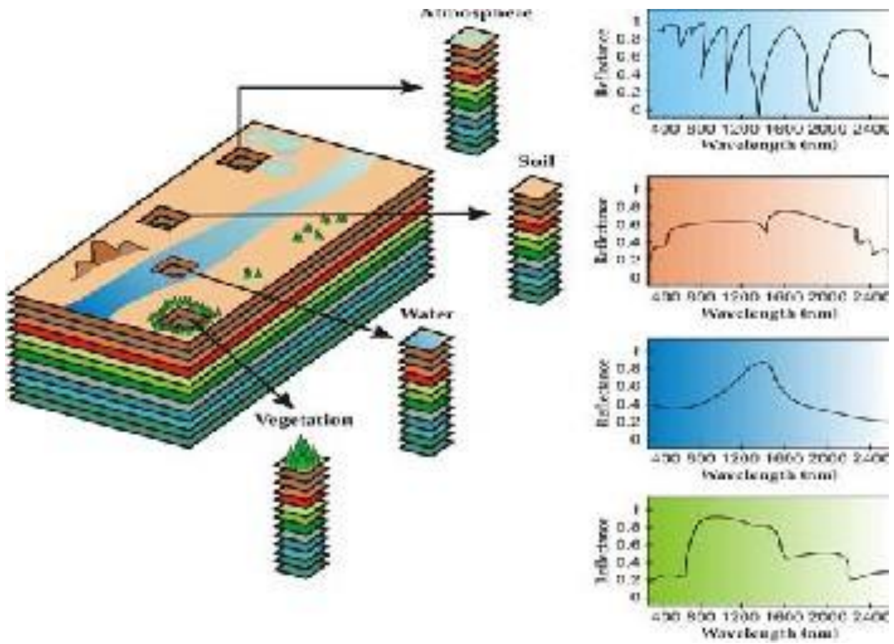

Spectral Unmixing for X-ray NDT

Research Directions

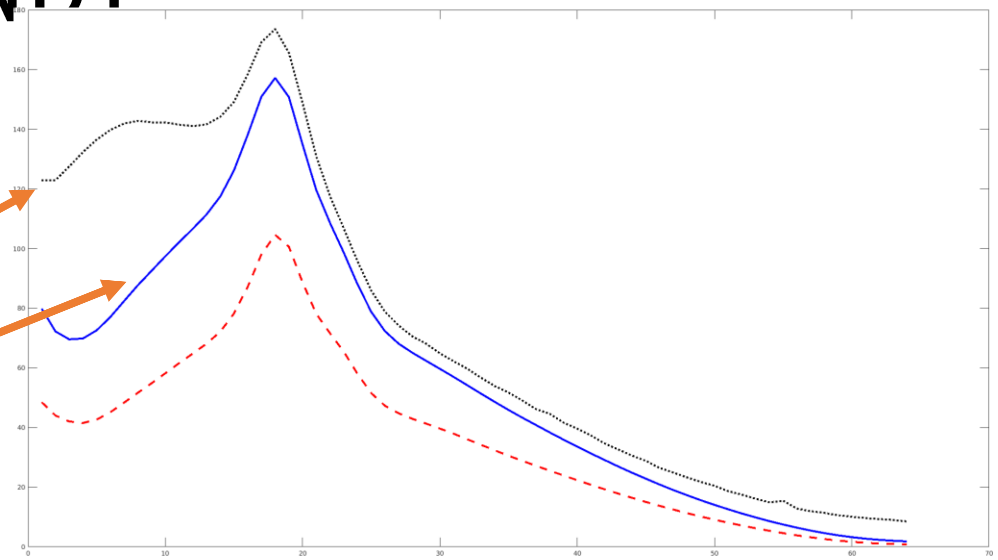

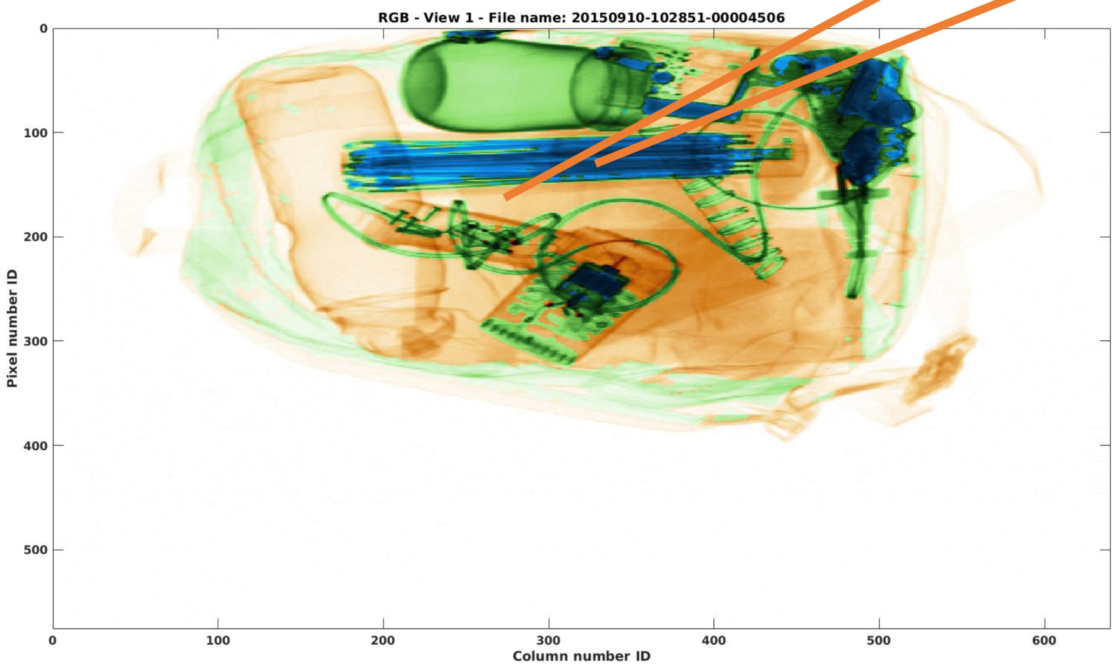

Collaboration with

GILARDONI

Thesis: Multispectral Image Analysis

Research Directions

Goals:

- Develop advanced classification and recognition algorithms for satellite imaging that combine additional information like GIS data and proximity constraints.
- Design spectral unmixing algorithms able to identify known materials in x-ray multispectral images acquired from baggage inspection systems.

Materials and Methods: multispectral images from opendigital and from X-ray machinery are available for testing unmixing and classification algorithms

Machine Learning for Health

Online and long-term ECG monitoring in wearables

Background

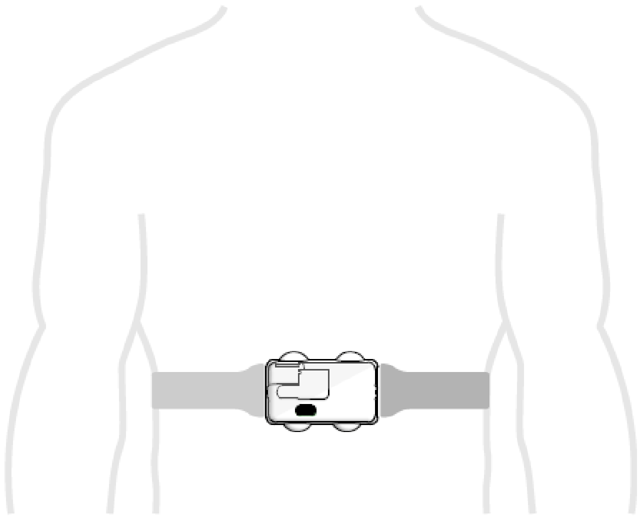

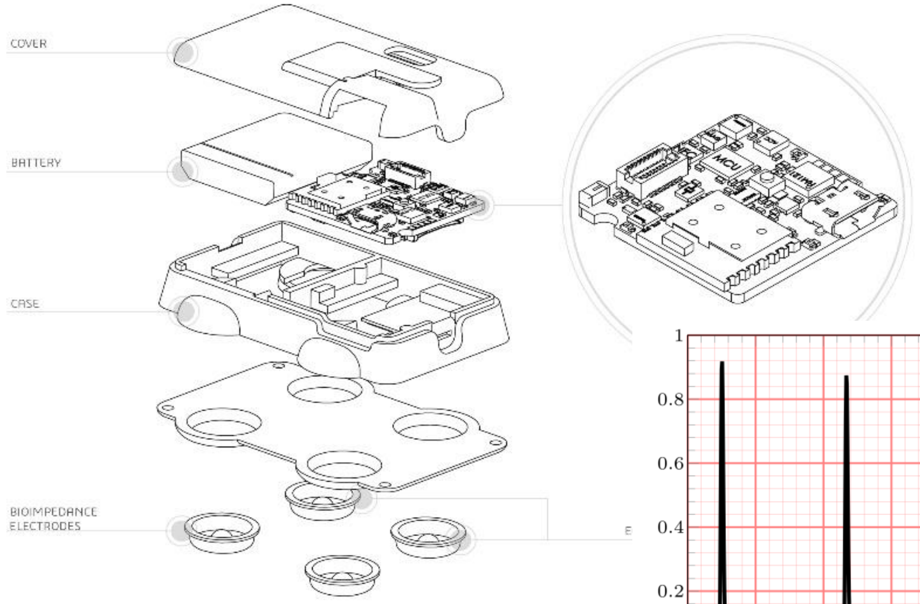

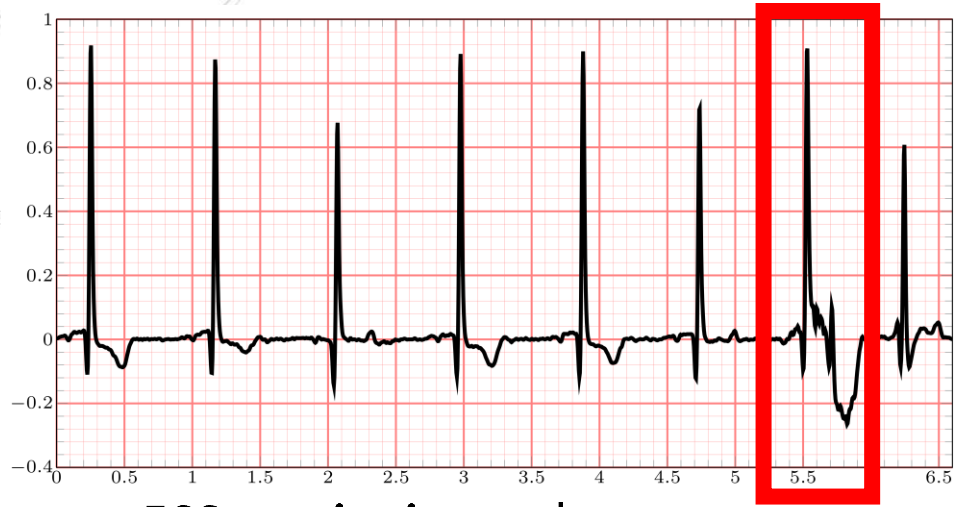

We have been designing algorithms to enable **online and long term ECG monitoring** on low-power wearable devices. Our algorithms are based on **data-driven models** to automatically **detect anomalies** (e.g., arrhythmias), these models undergo a **domain adaptation** procedure to track heart rate variations. The model is based on **a learned and user-specific model** of heartbeats.

IJCAI Demo webpage: http://home.deib.polimi.it/carrerad/IJCAI_2018_Demo.html

life.augmented

Thesis proposal on ECG Monitoring

Research Directions

Goals:

- Train 1d-CNN / deep RNN to classify heartbeats and identify arrhythmias
- Design efficient algorithms to enable online monitoring
- Design easy-to-train transfer learning algorithm that require few training data. In particular, adapt a pre-trained general-purpose models to:
  - analyze ECG signals from each specific user
  - Operate at different heart-rates
  - Operate in presence of motion artIfacts

Materials and Methods: We currently have a **prototypal wearable device** equipped with ECG and Bioimpedence leads as well as MEMS accelerometers. Annotated datasets from multiple users have been prepared.

Machine Learning for Seismocariogram analysis

Background

- Seismocardiogram (SCG), records micro-accelerations of the chest wall due to the heart movements, and it's a very informative signal since, w.r.t. traditional ECG:
    - SCG provides a direct measure of the **heart mechanical activity**, and not just the heart electrical triggers (assessed by the ECG).
    - The analysis of SCG is a quite new research, and efficient **automatic algorithms for segmenting, classifying and identifying anomalies** in this signal have still to be investigated.

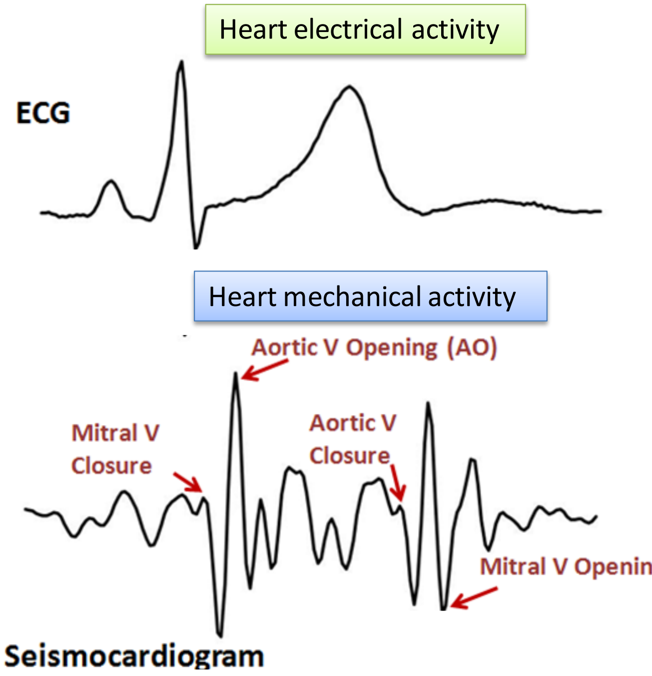

- Recently, Fondazione Don Gnocchi developed a wearable system for SCG monitoring, which have been used to monitor patients in clinics, and also in the International Space Station to monitor astronauts during their sleep.
- See the video at http://ow.ly/7z7Q30mgI1K

Research Directions

Thesis proposal: ML methods for Seismocariogram analysis

Goals:

- Develop data-driven algorithms for the automatic localization of patterns (associated with salient moments of the cardiac cycle, such as the opening and closure of valve) within the SCG waveform
- Develop patient-specific ML models to detect anomalies in the SCG signal.

Materials and Methods:

- The activity will be carried out in collaboration with the Wearable Sensor and Telemedicine Laboratory coordinated by Ing. Marco Di Rienzo, of the IRCCS Fondazione Don Carlo Gnocchi Hospital in Milano.
- The SCG recording device (SeisMote) as well as annotated datasets from patients will be available for the thesis.

Collaboration with Fondazione Don Carlo Gnocchi Onlus

Sleep monitoring and analysis

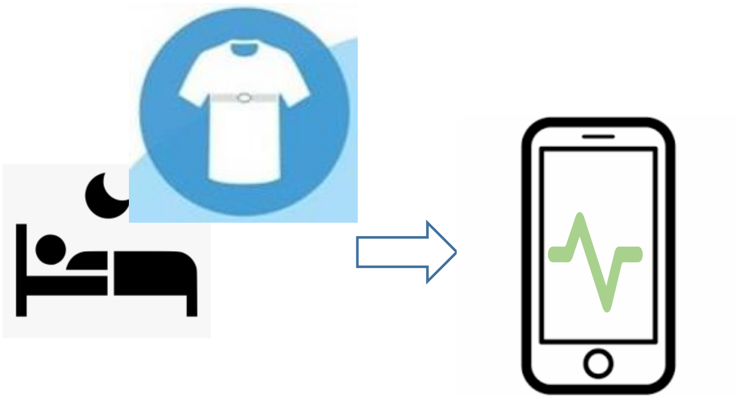

Goal: defining an indicator of **sleep quality** starting from data related to users' vital signs gathered by means of a smart t-shirt.

Research challenges:

- Reconstruction of the hypnogram (sleep graph) by using vital signs
- Data quality
- Adaptation: sleep quality depends on users' profile and habits

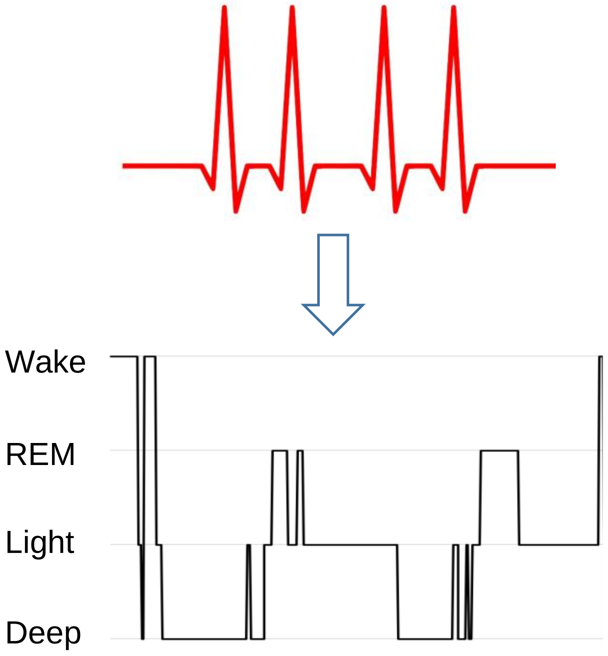

Materials and Methods: We currently have a smart t-shirt equipped with sensors for sleep monitoring and public annotated datasets

Research in collaboration with Prof. Cinzia Cappiello (DEIB)

Biomarkers Huntington's Disease (HD)

Background

HD is a **genetic, neurodegenerative condition** causing loss of volumes and atrophy in the brain.

Symptoms include motor, cognitive and emotional disorder.

MRI as well as fMRI and DTI have a great potential in the design of new data-driven and user-specific biomarkers.

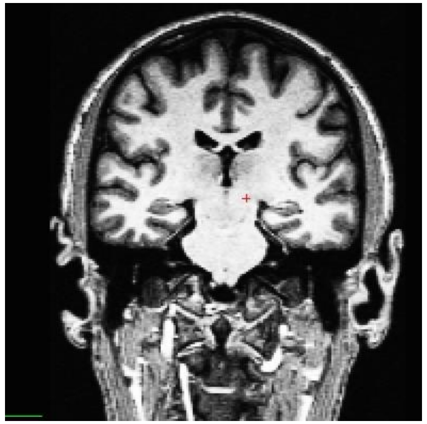

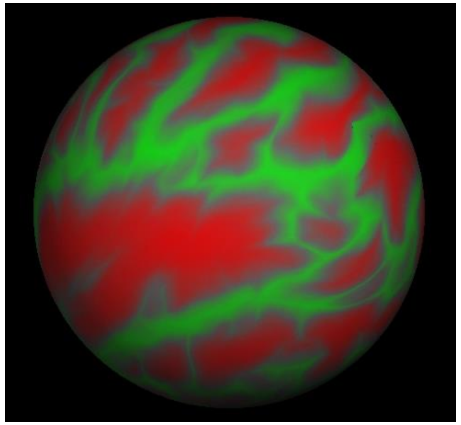

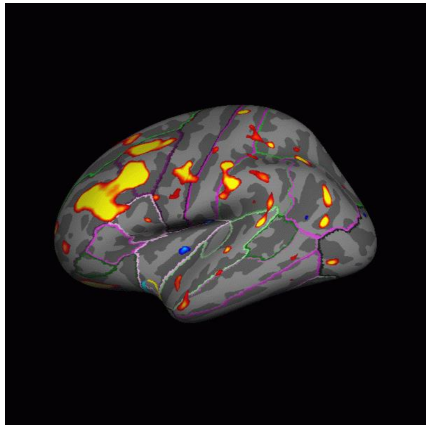

Biomarkers Huntington’s Disease (HD)

Background

MRI shows evidence of the disease **when it is already manifest**

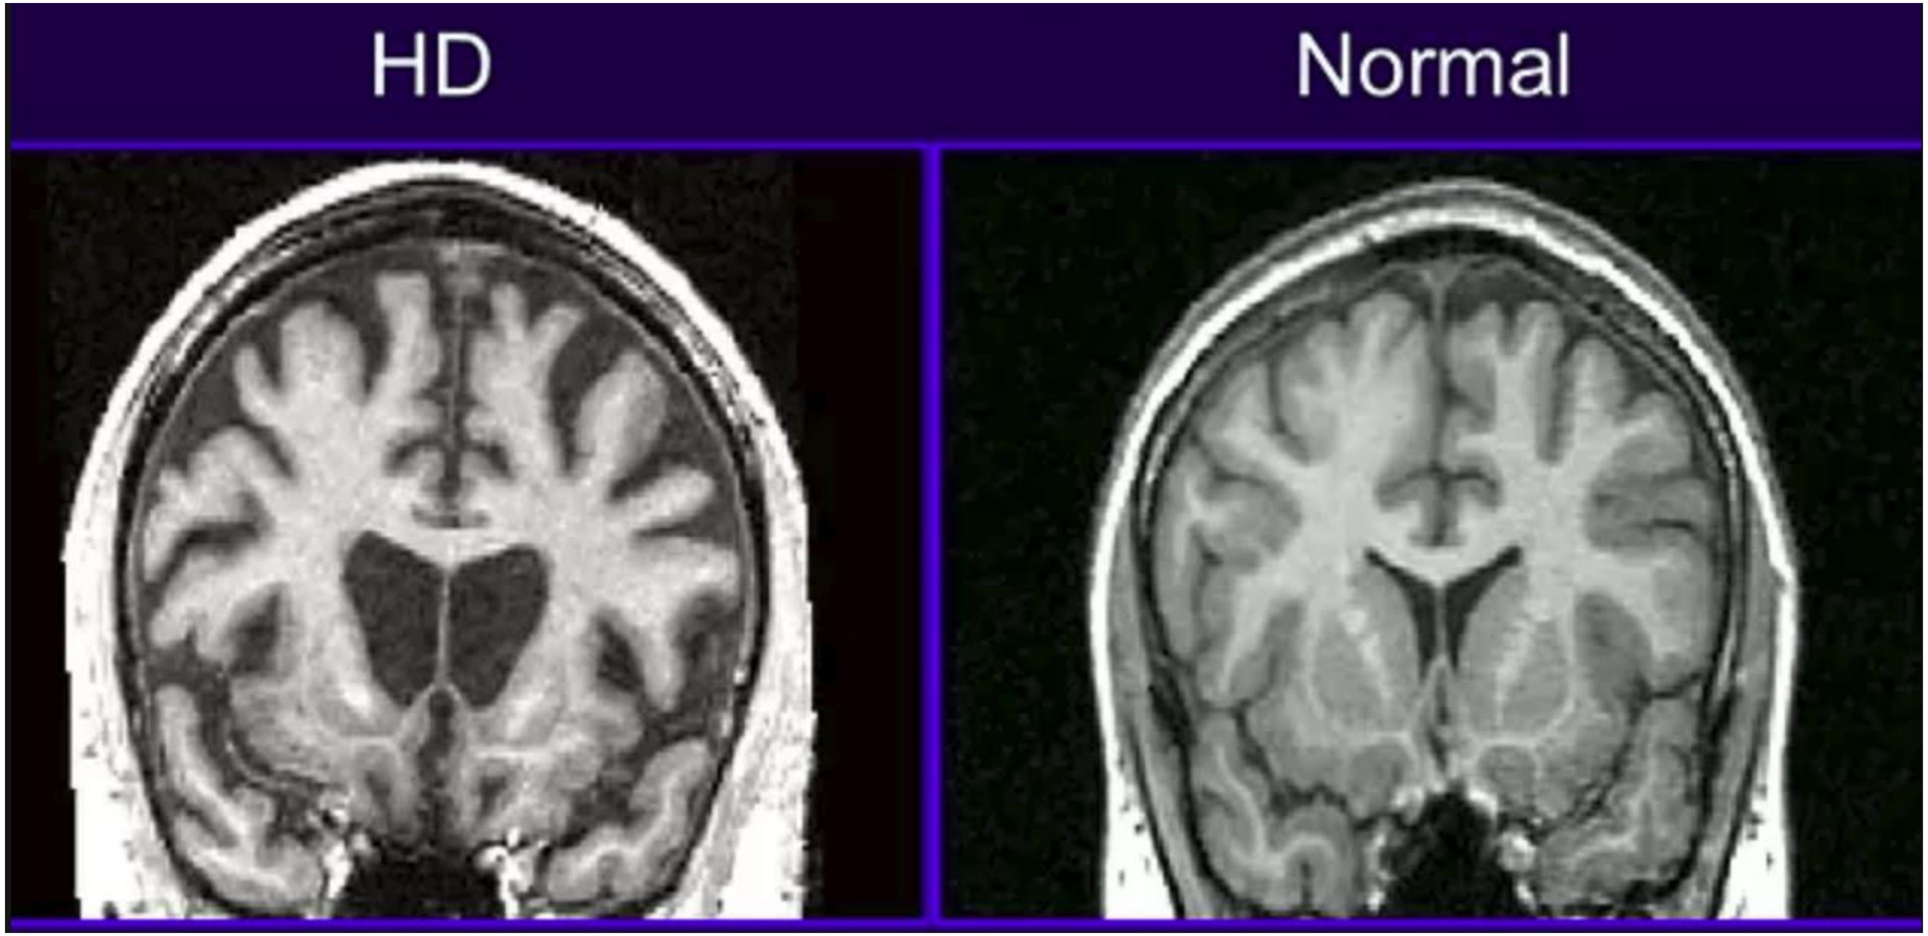

IMAGE CREDITS UC San Diego School of Medicine

A data-driven biomarker for Huntington Disease (HD)

We have developed a biomarker for quantitatively ***assessing the progression of HD****. The biomarker* ***monitors changes in shape and thickness of cortical surfaces*** *by means of* ***user-specific and data-driven models****.*

Background

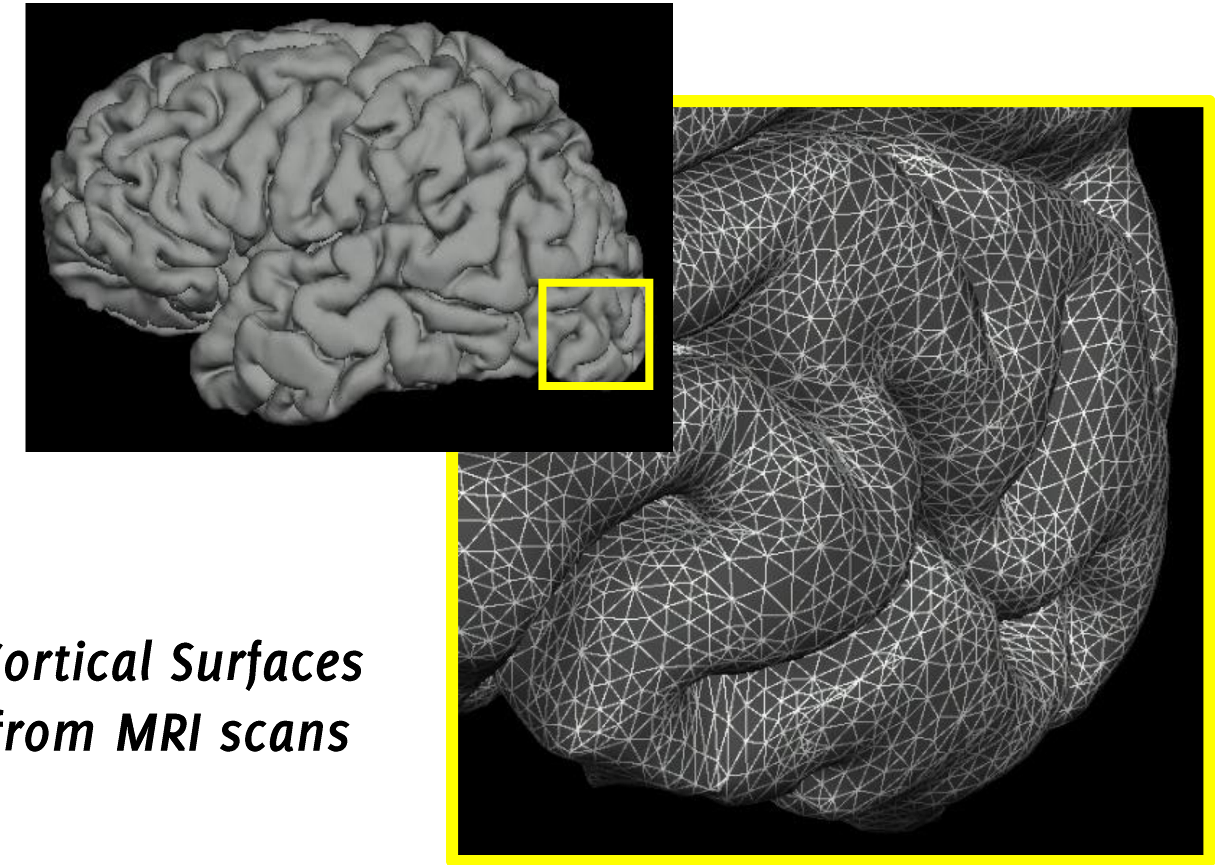

Cortical Surfaces from MRI scans

Atoms of dictionaries learned from brain regions

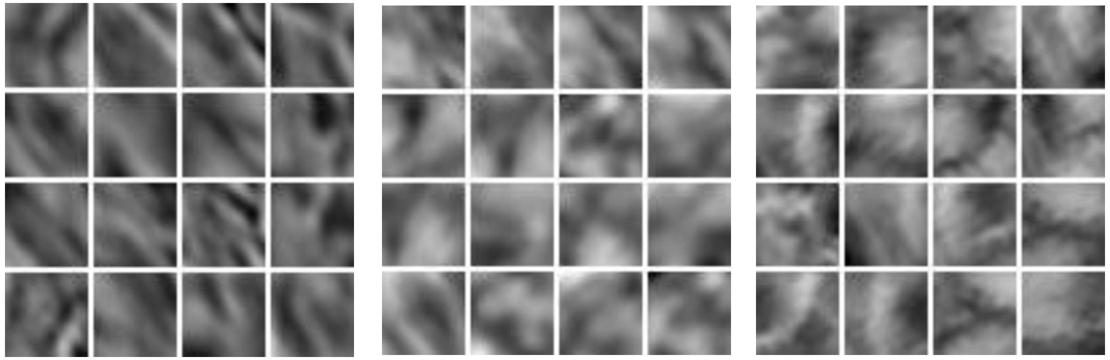

Bankssts | Precuneus | Parsorbitalis

Collaboration with IBM

Challenges of handling 3D sufaces

Background

- Measurements are not set over a regular grid
- Different number of samples in different portions

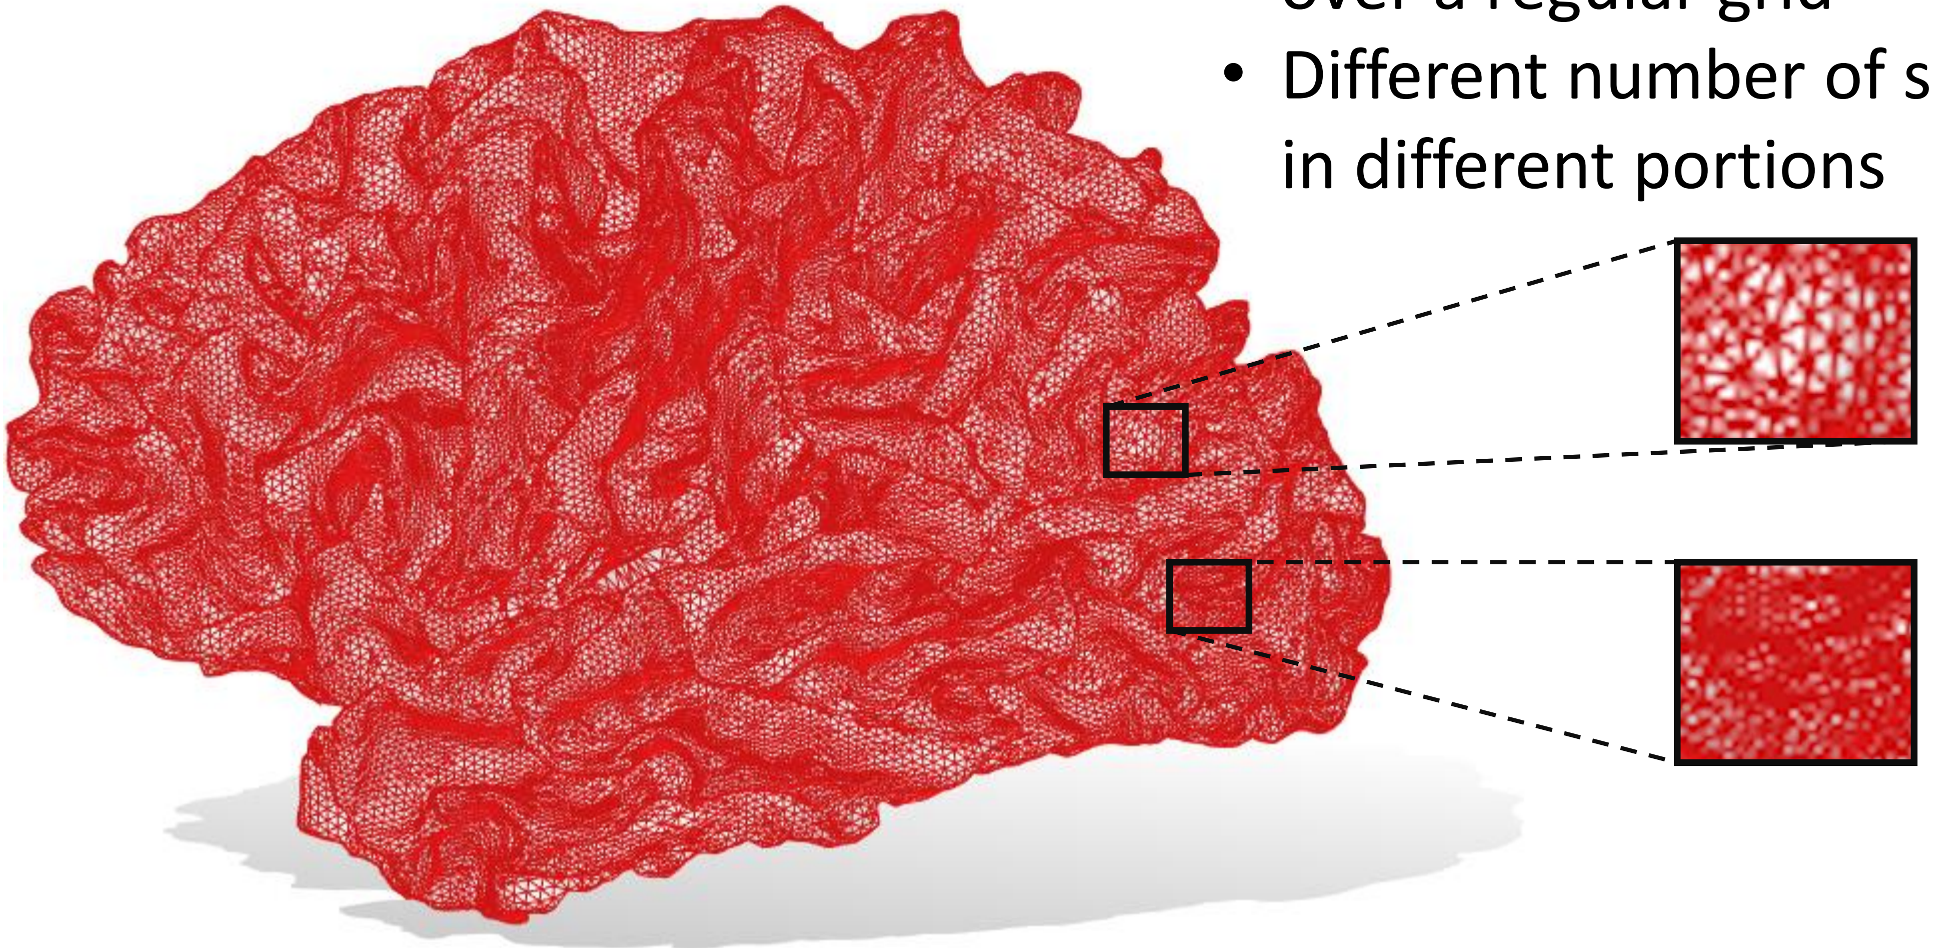

Thesis: Biomarker for HD

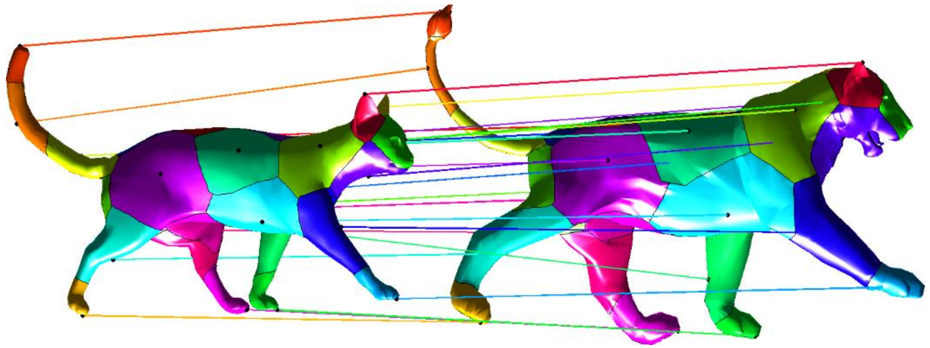

Research Directions

Goals:

- Implement a biomarker that directly operates on 3D data points, without mapping these into 2D image representations
- Directly compare cortical surfaces of the same patient as deformable surfaces. Test whether differences/deformations are meaningful for predicting HD progression

Materials and Methods:

- Longitudinal scans of more than 80 HD patients and controls.
- Well defined testing procedures.
- Functional maps software for 3d shape registration

Ovsjanikov, M., et al. *Functional maps: a flexible representation of maps between shapes*. ACM TOG 2012

Functional maps

An efficient and effective framework to estimate maps between 2D surfaces embedded in 3D.

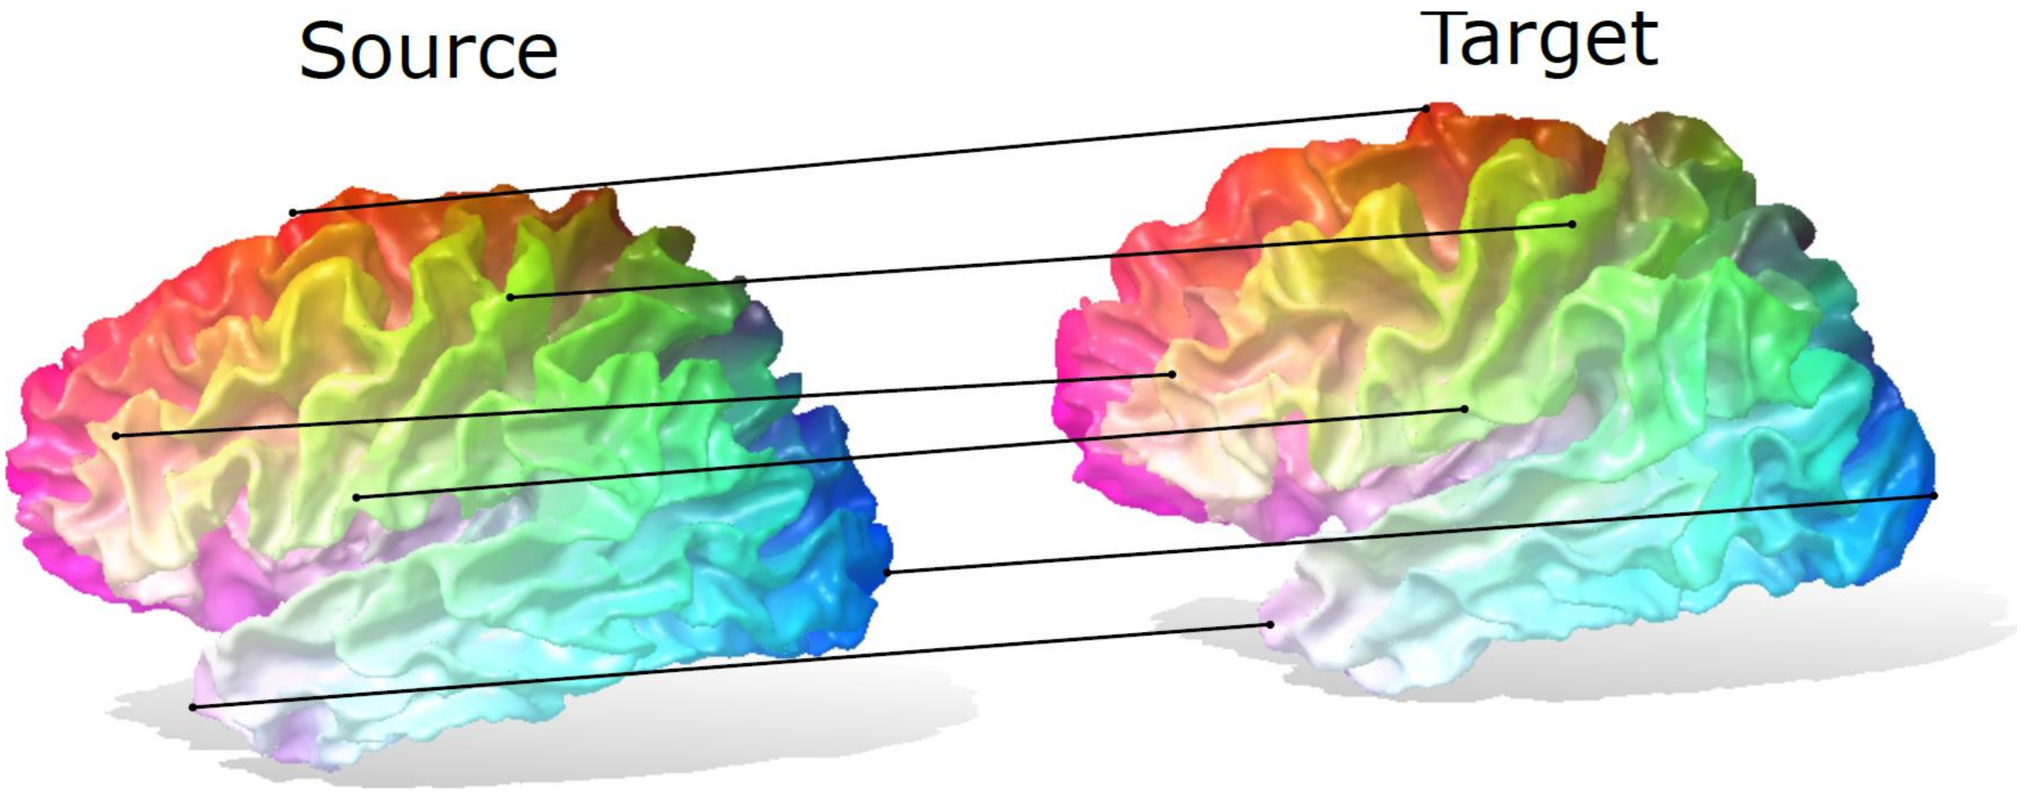

Ovsjanikov et al., Functional maps: a flexible representation of maps between shapes, ACM TOG, 2012.

ML and CFD for diagnosing nasal pathologies

Background

These thesis will be part of an established research collaboration with the aerospace department.

Goal: Design new **machine learning methods that**

- **analyze CFD simulation of the human upper airways**
- **help surgeon to make the most informed choices**

Why ML and CFD?

- Large failure rate of surgical corrections
- Lack of reliable diganostic tools
- CT scan only gives information about shape of the nasal apparatus
- CFD can give information about the functional properties during breathing

Output of a CT scan of the human upper airways

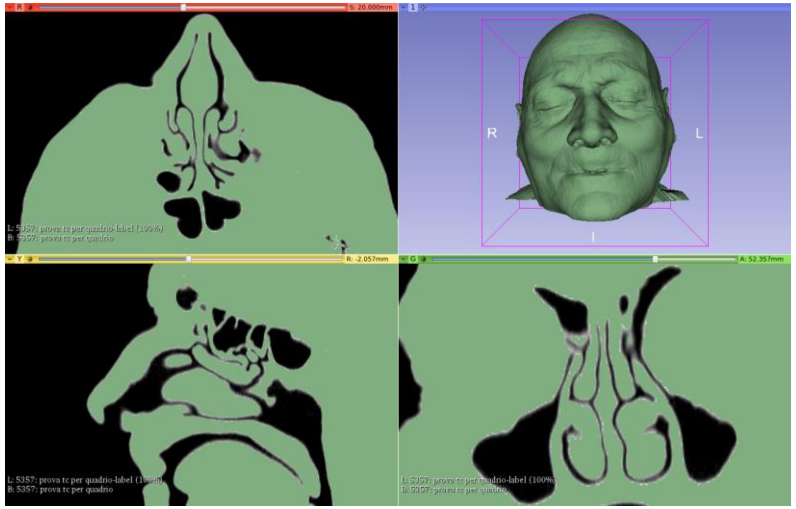

An illustration of CFD outcomes on a CT scan

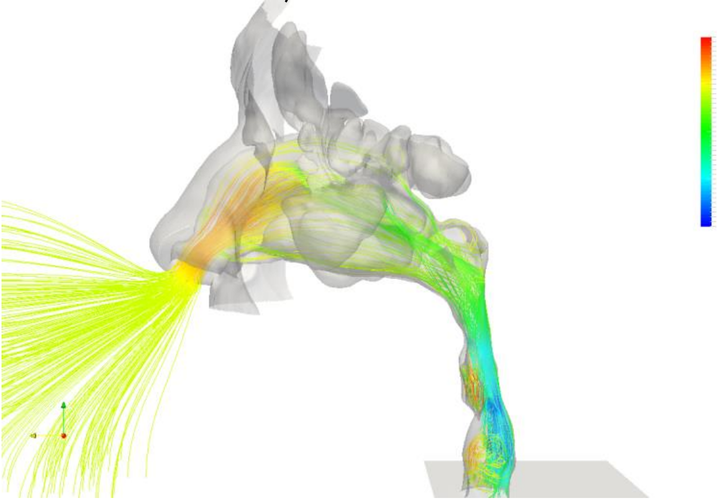

ML and CFD for diagnosing nasal pathologies

Background

Challenges:

- **Dimensionality,** CFD's output is **an high dimensional 3d (or 6d) -vector field,** it has to be **suitably transformed** to be handled by ML models. On a disk it might take a few GB!
- **Registration,** CT scan provides very different shapes from different subjects which needs to be registered.
- **Very little supervision,** compared to the input dimension. Many different pathologies to diagnose

Another illustration of CFD outcomes on a CT scan

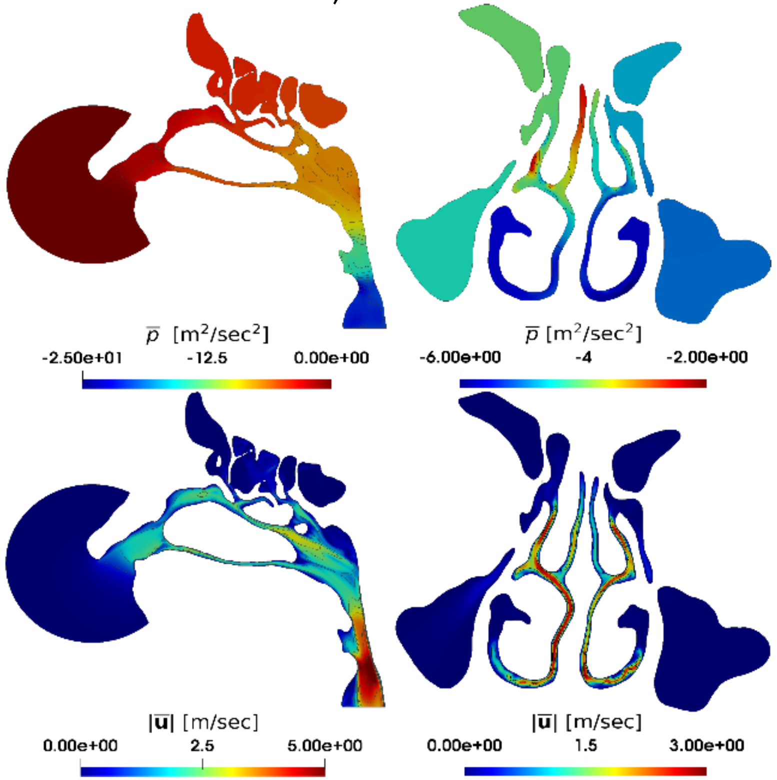

Thesis Opporunities in ML and CFD

Research Directions

Research directions for multiple thesis:

- **Dimensionality reduction and feataure design:** design the most representative features / **representation learning algorithms** to handle CFD output. Automatic selection of the most influential ones for medical diagnosis
- Investigate both grey-box and black-box approaches to perform dimensionality reduction on CFD.
- **CT registration by functional maps:** use thse geometric that allows efficient inference and manipulation to "transport" and compare properties
- **Advanced functional maps:** formulate the most suitable constraints to our problem at hand.
- **Augmentation** design data-augmentation procedures to introduce realistic CT-scan manipulation to mimic pathological effects

Example of Registration via functional maps

Research Directions

Noses from two subjects that are matched through functional maps. Colors encode estimated correspondences

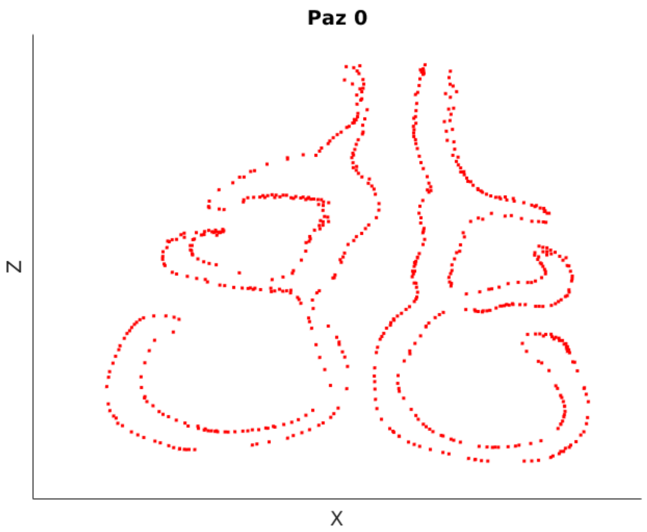

A reference section extracted from the first subject

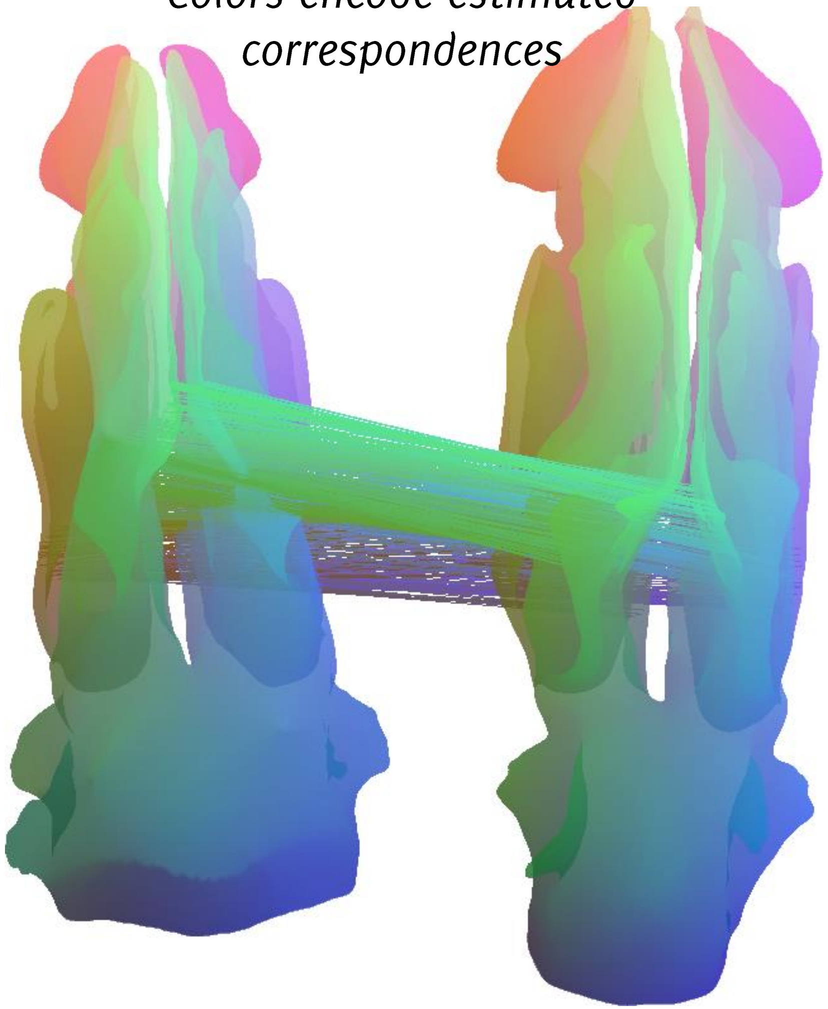

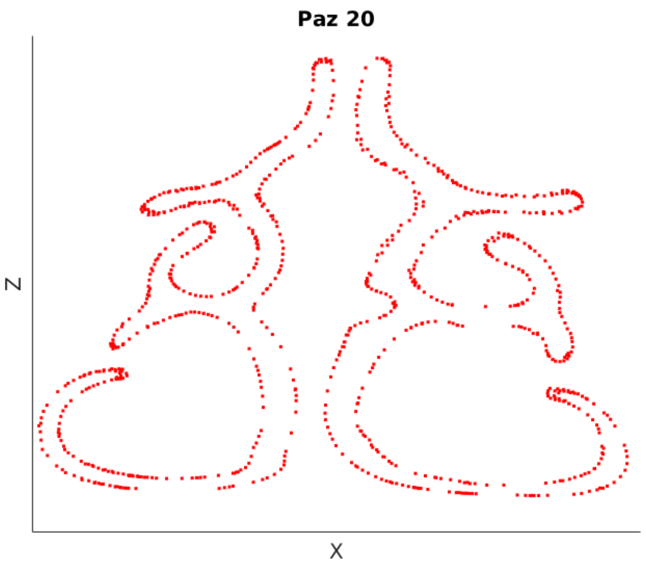

A nasal section extracted in sthe second subject, estimated through the correspondences

Ovsjanikov, M., et al. *Functional maps: a flexible representation of maps between shapes*. ACM TOG 2012

Thesis Opporunities in ML and CFD

Research Directions

Very Multidisciplinary environment:

This project brings together doctors, CFD expert, Data scientist, experts of geometric computer vision

Methods and Materials:

- Features inspired to image analysis techniques.
- State-of-the-art feature selection methods from ML literature
- Many CFD simulations of patients provided with the medical diagnosis

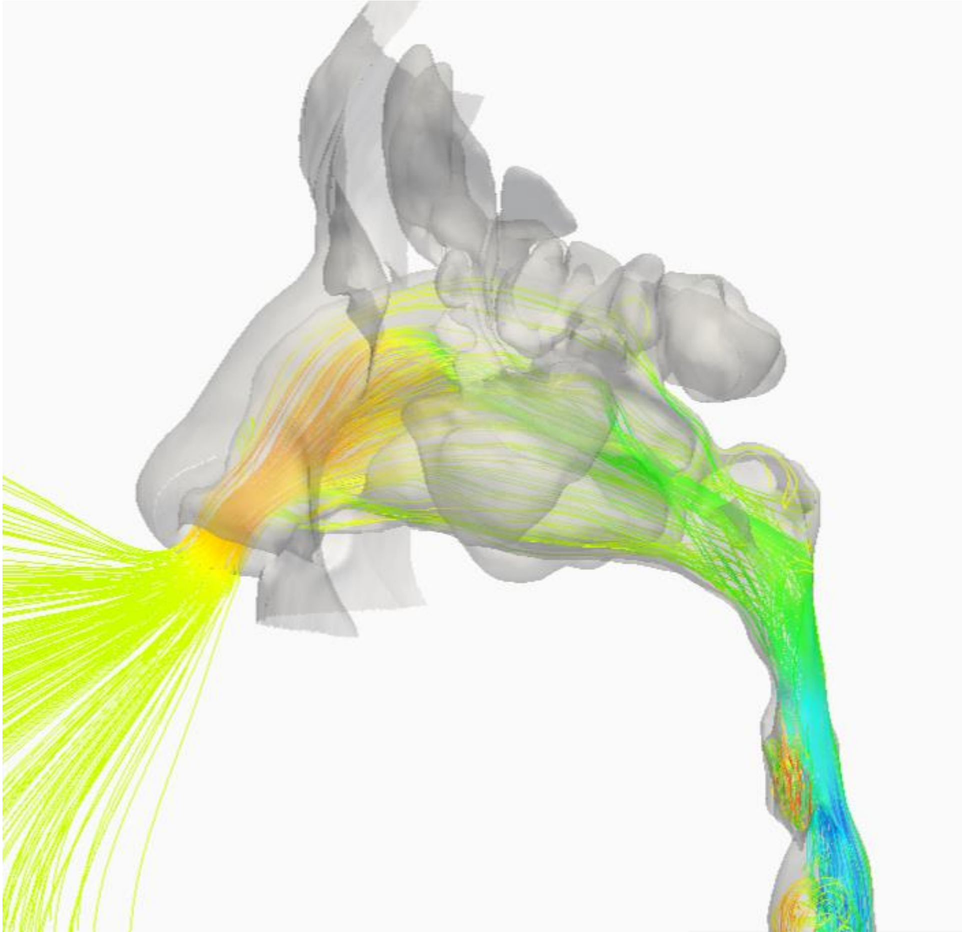

Change/Anomaly Detection in Datastreams

Anomaly detection in datastreams

Background

Usually anomaly detection is carried out **modeling** the **density of normal data** and identifying **anomalies** as samples that fall in **low density regions**.

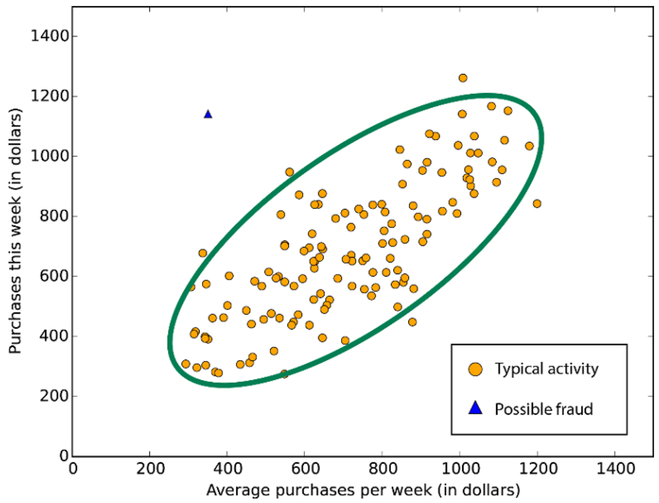

Non-parametric density models are much more flexible, but **extremely inefficient** (e.g. KDE). This is specially evident when the **data dimensionality** and the **number of samples** increases.

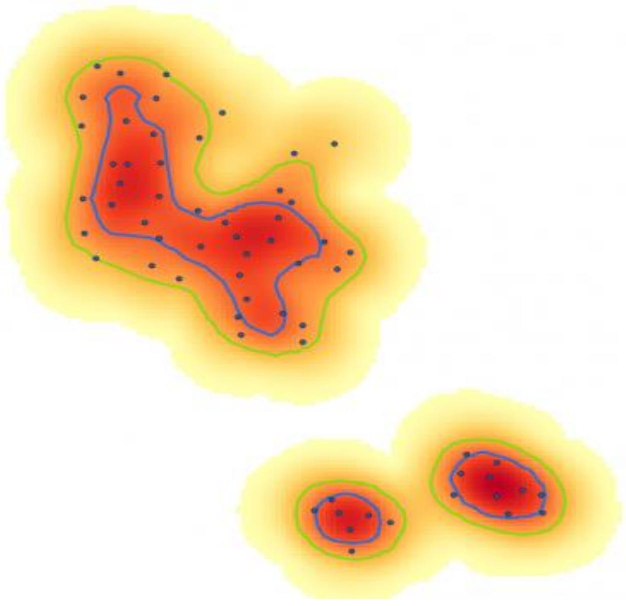

Anomaly detection in datastreams

Background

Anomalous samples often falls in top leaves

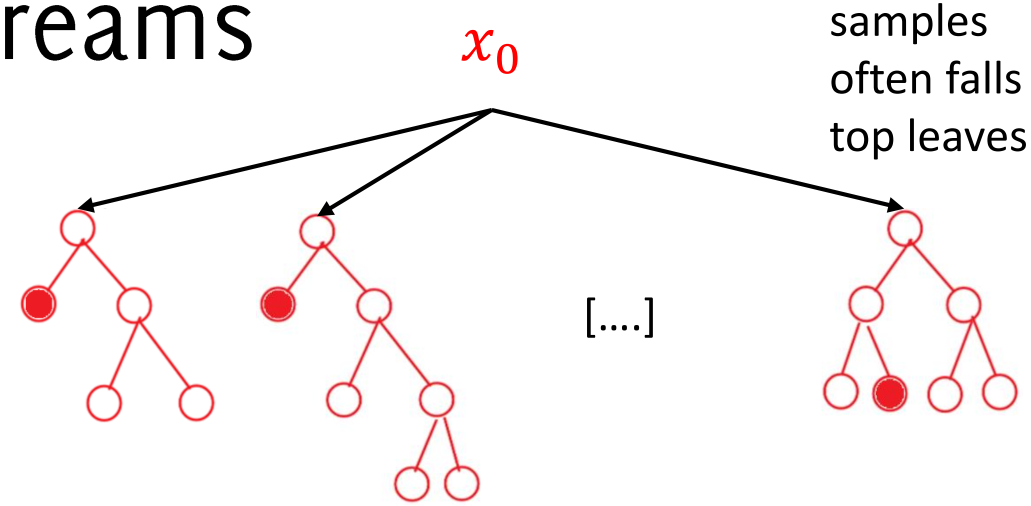

Isolation Forest (IFOR): Introduced a new paradigm in anomaly detection.

Instead of learning the density of normal data, IFOR models how **difficult is to isolate** each sample ***x*** **from the others.**

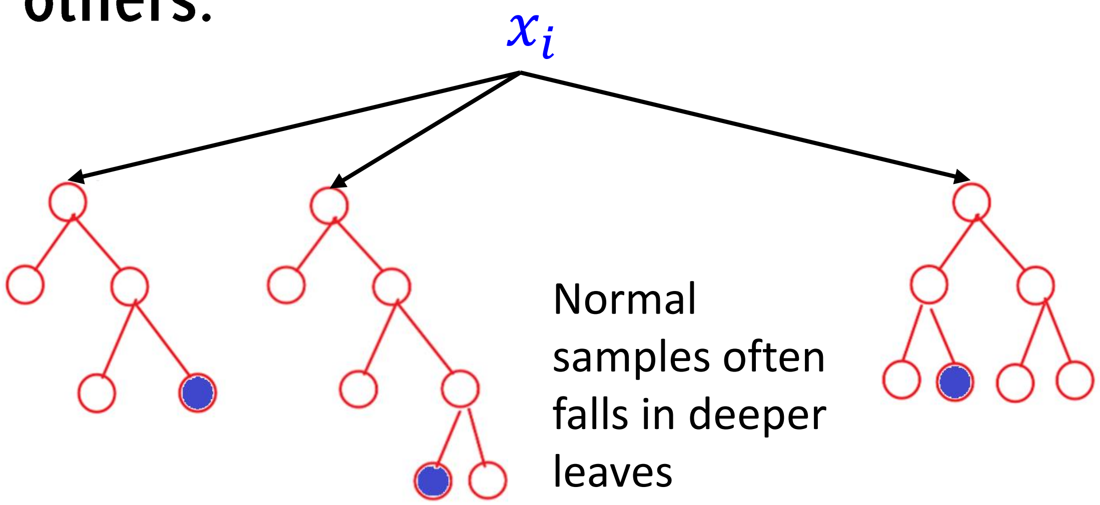

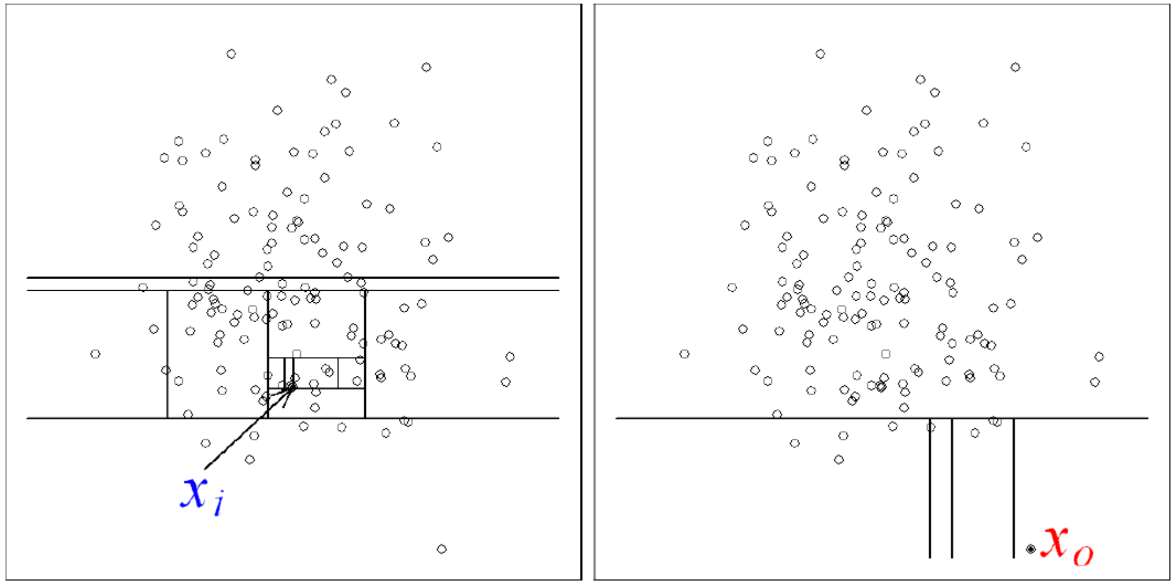

(a) Isolating x_i (b) Isolating x_o

Fei Tony Liu, Kai Ming Ting and Zhi-Hua Zhou, "Isolation Forest", *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*

Thesis: Beyond IFOR

Research Directions

Goal: investigate improvements for IFOR based on:

- Design **novel decision rule** for IFOR, possibly based on **unsupervised loss**.
- Design **adaptation mechanisms** for IFOR to handle **evolving data-streams.**
- Include **expert-driven** information in the IFOR construction.
- **Optimized**, **parallel implementation** of IFOR on GPUs.
- Define **interpretability criteria** for the most apparent anomalies.
- Investigate extensions to **categorical data**.
- Adopt **copula theory** to factorize distributions.

Materials and Methods:

- Access to a server mounting GPUs will be provided.
- Dataset from our industrial partner.

Research Directions

Thesis: Isolation Based Clustering

Goal: exploit the potential of an isolation-based approach for clustering:

- Study the state-of-the-art of **clustering techniques**.
- Model the concept of "**isolation**" in IFOR.
- Design an innovative technique to **cluster data** taking advantage of "**isolation**".

Materials and Methods:

- Access to a server mounting GPUs will be provided.
- Dataset from our industrial partner.

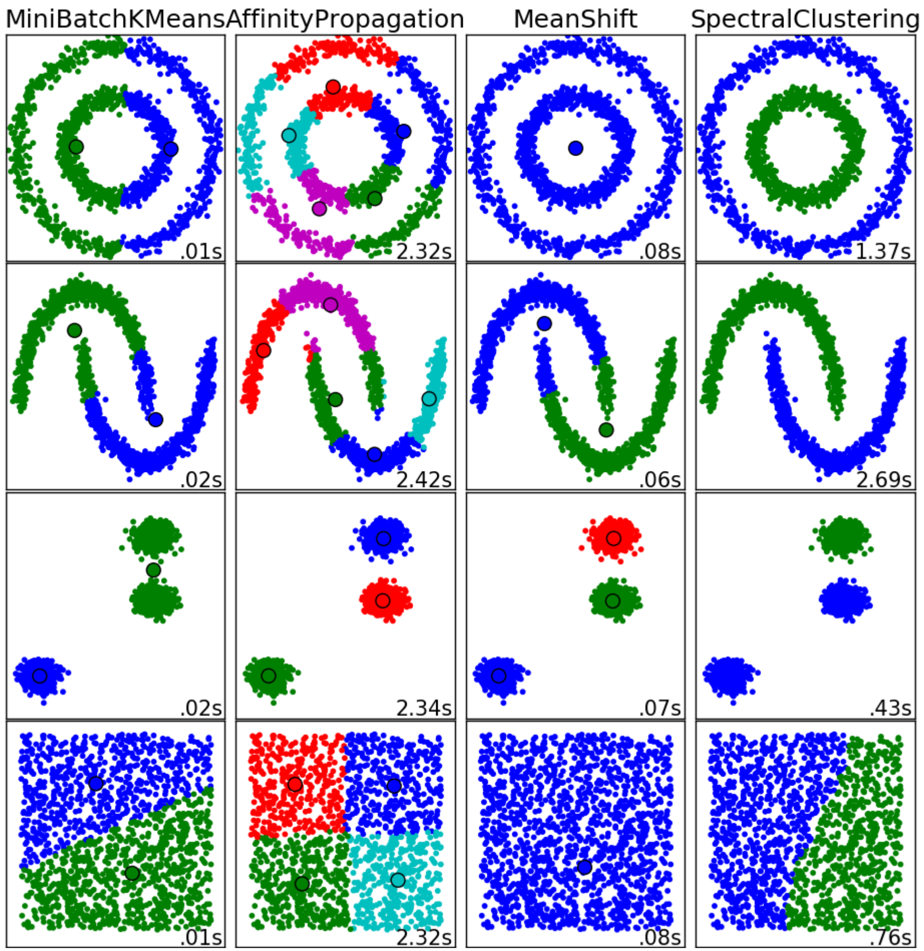

Fei Tony Liu, Kai Ming Ting and Zhi-Hua Zhou, "Isolation Forest", *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*

Change Detection in Datastreams

Efficient density models for **detecting changes in high-dimensional datatstreams**, yet controlling false positives

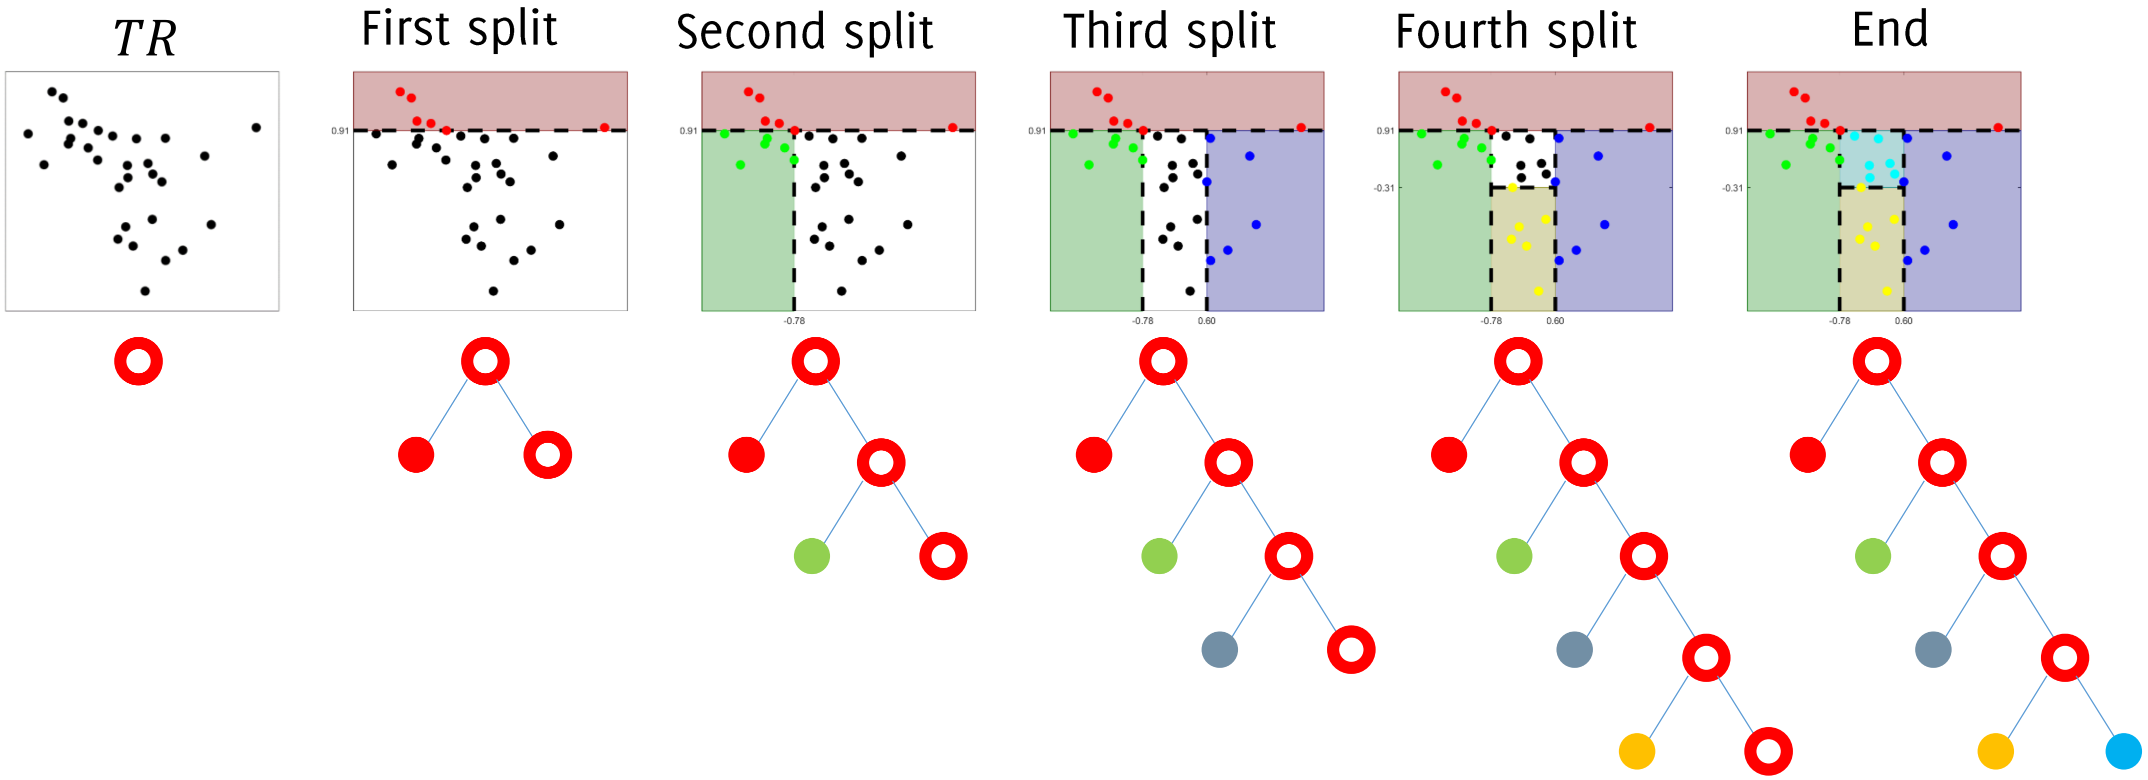

Background

QuantTree: Boracchi, Carrera, Cervellera, Macciò, “Histograms for Change Detection in Multivariate Data Streams”, ICML 2018

Change Detection in Datastreams

Research Directions

Goals:

- Design a **sequential monitoring scheme** (like CPM) **based on QuantTrees** to enable change detection in multivariate data streams
- **Study the probabilistic properties** of the QuantTree test statistics and use them to state parametric tests for change-detection purposes
- Investigate **incremental learning procedures** and their convergence guarantees for learning trees on very large training sets
- Investigate **new partitioning schemes** that are better suited for change-detection
- Design **parallel implementation** of an ensemble of histograms
- Design incremental/sequential procedures to construct QuantTrees

Materials and Methods: Experimental testbed on synthetic and real data, reference codes for QuantTrees to be used as a comparison.

Sparsity and Convolutional Sparsity

Background

Convolutional sparse models are a recent development of sparse representations

$$\boldsymbol{s} \approx \sum_{i=1}^{n} \boldsymbol{d_i} \circledast \boldsymbol{\alpha_i}\,, \qquad \text{s.t. } \boldsymbol{\alpha_i} \text{ is sparse}$$

where a signal $\boldsymbol{s}$ is **entirely encoded** as the sum of n convolutions between a filter $\boldsymbol{d_i}$ and a coefficient map $\boldsymbol{\alpha_i}$

Pros:

- Translation invariant representation
- Few small filters are typically required
- Filters exhibit very specific image structures
- Easy to use filters having different size

Sparsity and Convolutional Sparsity

There are two major problems to be addressed when adopting models based on sparse-representations:

- Dictionary learning
- Sparse coding, namely computing the representation of an input signal w.r.t. the learned dictionary

Thesis: Convolutional Sparsity

Research Directions

Design new sparse coding and dictionary learning schemes for improving image restoration performance