

UNIVERSITY OF TORINO

Ph.D. in Modeling and Data Science

Final dissertation



**Deterioration: methods for estimation in natural and artificial systems**

Supervisor: Attilio Fiandrotti  
Co-supervisor: Marco Grangetto

Candidate: Davide Tricarico

XXXV CYCLE

# Summary

This doctoral thesis aims to address the quantitative monitoring and prediction of system deterioration through data-driven approaches. The focus is directed towards two crucial domains: *Predictive Maintenance* and *Medical Imaging Analysis*, wherein advanced methodologies and algorithms are developed to enhance the precision and effectiveness of these applications.

The initial segment of this work focuses on the exploration of *Predictive Maintenance* techniques. A comprehensive pipeline is proposed to forecast the degradation status of systems through the utilization of onboard sensor data. The proposed methodology covers the entire design process, from the data collection design to the predictive model training and deployment. Within this pipeline, a novel signal selection algorithm, namely *CORR-FS*, is introduced to reduce input dimensionality while upholding interpretability. Real-world use cases derived from the automotive sector are employed to evaluate the efficacy of the proposed methodology. The findings demonstrate its capability to surpass the requirements set by domain experts, thereby providing valuable insights into data and system behaviour. Additionally, practical considerations concerning the implementation of the proposed approach are discussed, exploring how it can be efficiently deployed across a vehicle fleet to yield benefits for both users and service providers.

The second part of the thesis focuses on the domain of *Medical Imaging Analysis*. A novel approach is introduced to evaluate degradation and predict disease severity by utilizing image similarity in the feature space. The proposed methodology harnesses the power of deep learning models for feature extraction and employs distance-based regression for prediction. Notably, the algorithm incorporates a new loss function, *DISTMAT*, which facilitates a contrastive training process based on distance matrices. Extensive experimental evaluations are conducted on diverse medical imaging

use cases, confirming the effectiveness of the approach and surpassing existing state-of-the-art methods in most instances. The performance is examined, shedding light on the strengths, limitations, and unresolved aspects of the methodology. Furthermore, this thesis delves into the integration of the proposed approach within a real-world hospital environment, aiming to optimize the diagnostic workflow during the *COVID-19* pandemic. The obtained results from this integration are presented and analyzed in detail.

In conclusion, this thesis presents future directions for the development of both the proposed methodologies in the *Predictive Maintenance* and *Medical Imaging Analysis* applications. These directions entail rigorous testing and validation in real-world scenarios, as well as their application to diverse use cases to verify the generalizability of the achieved performance. Moreover, this research proposes the integration of natural language processing techniques in conjunction with the *DISTMAT* loss function to train more precise and task-agnostic feature extraction models, capitalizing on the vast corpus of available textual medical reports. The thesis also summarizes the unresolved aspects regarding method configuration and application, offering valuable insights for future experimental investigations and refinements.

# Acknowledgements

I extend my heartfelt gratitude to all the advisors, Professors *Attilio Fian-drotti*, *Marco Grangetto*, *Mario Dante Lucio Giacobini*, and *Elvira Di Nardo*, for their invaluable support and expert guidance throughout this PhD program. And I also thank Professor *Laura Lea Sacerdote* for her support.

I am equally grateful to my company supervisor, *Massimiliano Melis*, who offered me the opportunity to embark on this remarkable journey. I also thank my company, *AITEM Solutions*, which sponsored my entire PhD program.

Additionally, I wish to express profound thanks to my beloved family, *Cristina*, *Gianluca*, and *Rachele*, who, although not directly involved in this project and have been unjustly neglected because of this, have supported and inspired me throughout, providing endless source motivation. They are my principal source of joy and gratitude in my life.

Finally, a message to my future self, who will read this PhD thesis in years to come: this work has been the outcome of numerous sacrifices and perseverance, even when faced with seemingly insurmountable challenges and insufficient time. Remember this as a powerful incentive not to succumb to adversity or despair, for you have not given up and have proven time and again that determination and resilience lead to achievements. Congratulations! Carry this accomplishment as an incentive not to break down in front of any seemingly insurmountable obstacle that may arise in the future.

# Contents

<b>List of Tables</b>	7
<b>List of Figures</b>	9
<b>1 Introduction</b>	17
1.1 Document structure . . . . .	21
<b>2 Predictive maintenance</b>	23
2.1 State of the art . . . . .	30
2.2 Specific use cases . . . . .	32
2.2.1 Oxygen sensor for Diesel engine . . . . .	34
2.2.2 Diesel engine Fuel high pressure system . . . . .	43
2.3 Proposed methodology . . . . .	50
2.3.1 Method inference . . . . .	52
2.3.2 Method settings . . . . .	56
2.4 Results . . . . .	67
2.4.1 Metrics and requirements . . . . .	67
2.4.2 Performance . . . . .	70
2.5 Discussion . . . . .	74
2.5.1 Oxygen sensor . . . . .	75
2.5.2 Fuel high pressure system . . . . .	77
<b>3 Medical imaging</b>	81
3.1 Automatic tool for medical imaging . . . . .	85
3.2 Specific use cases . . . . .	94
3.2.1 COVID-19 . . . . .	94
3.2.2 Estimation of pediatric bone age . . . . .	104
3.2.3 Quantification of calcium score . . . . .	107
3.3 Deep Learning application in Medical Imaging . . . . .	110

3.4	Proposed methodology . . . . .	118
3.4.1	Image pre-processing . . . . .	121
3.4.2	Feature extraction . . . . .	122
3.4.3	Distance based regression . . . . .	123
3.5	Method settings . . . . .	124
3.5.1	Image pre-processing procedure and parameters . . . . .	124
3.5.2	Feature extraction training process . . . . .	126
3.5.3	Distance based regression design and tuning . . . . .	130
3.6	Results . . . . .	133
3.6.1	Metrics and benchmarks . . . . .	133
3.6.2	Performance . . . . .	140
3.6.3	Preliminary “Real Life” Results . . . . .	148
3.7	Discussion . . . . .	154
3.7.1	Dataset unbalance . . . . .	155
3.7.2	Other cases in the prioritization use case . . . . .	158
3.7.3	Reference set depletion robustness . . . . .	159
3.7.4	DISTMAT loss training . . . . .	161
<b>4</b>	<b>Conclusion</b> . . . . .	<b>165</b>
4.1	Predictive Maintenance . . . . .	165
4.2	Medical Imaging . . . . .	166
4.3	Next steps . . . . .	168
<b>A</b>	<b>List of published works</b> . . . . .	<b>171</b>
A.1	Predictive Maintenance . . . . .	171
A.2	Medical Imaging . . . . .	172

# List of Tables

2.1	Comparison of Physics-based, Data-driven, and Hybrid Approaches . . . . .	28
2.2	Cardinalities for the three classes . . . . .	42
2.3	Signals overview . . . . .	48
2.4	Dataset description per driving cycle. The <i>RDE</i> and <i>WLTC</i> cycles exhibit balanced class distribution, while the <i>Artemis</i> and <i>Real</i> cycles are predominantly associated with the Green class. This variation can be attributed to differences in pedal dynamics and test execution. . . . .	50
2.5	Data transformation settings used in the experiments . . . . .	54
2.6	Signal selection settings used in the experiments . . . . .	60
2.7	Hyperparameter settings for each evaluated model . . . . .	65
2.8	Best performance obtained for both the use cases . . . . .	70
2.9	<i>Decision Tree</i> best configuration results . . . . .	73
2.10	<i>Artificial Neural Network</i> best configuration results . . . . .	74
2.11	<i>Support Vector Machine</i> best configuration results . . . . .	74
3.1	Dataset composition for Radiology diagnostic workflow optimization for <i>COVID-19</i> task . . . . .	98
3.2	Age statistics of positive and negative classes for Radiology diagnostic workflow optimization for <i>COVID-19</i> task. *Shapiro test. . . . .	99
3.3	Gender statistics of positive and negative classes for Radiology diagnostic workflow optimization for <i>COVID-19</i> task. . . . .	100
3.4	BrixIA data main characteristics . . . . .	101
3.5	Summary of the ICIAP 2022 Per-COVID-19 challenge dataset . . . . .	104
3.6	Summary of the CAC score ranges and associated risk . . . . .	109
3.7	Summary of the CAC score dataset . . . . .	110
3.8	Image preprocessing for the different use cases . . . . .	126

3.9	Image preprocessing for the different use cases. Prioritization use case has been tested with different settings that will be discussed in the next sections . . . . .	129
3.10	Configuration of distance-based regression algorithm for the different use cases . . . . .	131
3.11	Metrics estimated for performance analysis, with benchmark values and confidence interval. . . . .	136
3.12	Metrics estimated for performance analysis. . . . .	137
3.13	Metrics estimated for performance analysis. . . . .	138
3.14	Metrics estimated for performance analysis. . . . .	139
3.15	Best performance obtained for all the use cases. Underlined the results that are better than benchmark . . . . .	140
3.16	Final competition ranking for Validation dataset. In <b>bold</b> the result related to the presented approach . . . . .	145
3.17	Final competition ranking for Test dataset. In <b>bold</b> the result related to the presented approach . . . . .	145



# List of Figures

1.1	Photos of different stage of deterioration for an apple Labs [2017]. . . . .	17
1.2	Stages of Concrete Corrosion (Monib Co Ltd. [2017]). . . . .	19
1.3	In influenza virus infection, viral glycoproteins attach the virus to a host epithelial cell. As a result, the virus is engulfed. Viral RNA and viral proteins are made and assembled into new virions that are released by budding. (Microbiology [2023]). . . . .	20
1.4	DNA damage in non-replicating cells, if not repaired and accumulated can lead to aging. DNA damage in replicating cells, if not repaired can lead to either apoptosis or to cancer Wikipedia [2018]. . . . .	21
2.1	The progression of failure over time. A very early incipient failure occurs to the part, gradually degrading the component performance, until it is no more able to deliver its functionality. Further degradation of the component causes the entire system failure. The time remaining until the component failure is called <i>Residual useful life</i> — <i>RUL</i> , during which Prognostics technologies monitor the system to detect the deterioration progression. After the component failure, Diagnostics is in charge to detect and identify the failure to notify the necessity of intervention . . . . .	24
2.2	How the approach to maintenance changed over history. Prognostics stimulated a more cost effective trade-off between risk and costs . . . . .	27

2.3	The context of automotive prognostics use cases. Signals are acquired by onboard sensor and elaborated by local Electronic control Unit. Manipulated signals can be sent by remote connection to the cloud infrastructure where can be used to perform prognostic algorithm and enable predictive maintenance. The output of predictive maintenance can be notified to both the vehicle user and the maintenance organization to schedule the intervention. . . . .	33
2.4	NKG ZFAS-U2 oxygen sensor for automotive application, NGK [2023]. . . . .	34
2.5	Oxygen sensor after exposure to exhausted gas in the after-treatment pipeline. The dust contained in the flow is deposited on the sensor, clogging the small holes that allow the gas exchange between the sensitive element and the environment and deteriorating its measurement performance. . . .	35
2.6	Functional schematic of planar wideband zirconia oxygen sensor. Embedded feedback loop circuit controls the pump current to regulate oxygen ions exchange with exhaust gas and keep the monitoring chamber gas composition constant. The magnitude of pump current is directly proportional to the oxygen concentration in exhaust gas, Handrich [2010]. . . .	36
2.7	Acceleration pedal profile used in cycles. Each cycle can be divided into two parts, each recorded with a different software: Program A and Program B . . . . .	37
2.8	Example of oxygen measurement signal during the last 5 minutes of the cycle: the cut-off manoeuvre generates a step-like stimulus, visible around time 250 seconds. . . . .	38
2.9	Detail of oxygen sensor sudden change in Figure 2.8 at 250 seconds . . . . .	39
2.10	Response time measurement process . . . . .	41
2.11	Response time trend throughout the experiment. The horizontal lines show the positions of the thresholds . . . . .	41
2.12	Distribution of response times: the colors represent the assigned classes based on the defined thresholds . . . . .	43
2.13	Response time trend smoothed with different k values . . . .	44
2.14	Number of label switches occurring as k increases . . . . .	45

2.15	Architecture of a Fuel Injection System in a modern Diesel Engine. Yellow pipes indicate low pressure fuel, red lines the high pressure. Lines connecting the Control Unit to sensor and actuators represent information exchange. The arrows explain the direction of flow or exchange. This illustration has been derived from Hannu Jääskeläinen [2023]. The components on which we concentrated our study are indicated with yellow labels. . . . .	46
2.16	High-level overview of the proposed method to predict deterioration status of a system from measured signals. . . . .	51
2.17	The impact of time window length on the predictive performance of the proposed method. . . . .	53
2.18	The number of selected signals varying the value of the parameter $r_{min}$ for the oxygen sensor use case . . . . .	59
2.19	The number of selected signals varying the value of the parameter $r_{min}$ for the fuel high pressure use case . . . . .	59
2.20	Comparison of the three obtained rankings by mean of min-max normalized feature importance . . . . .	62
2.21	Performance of the best <i>SVM</i> model configuration by using a progressively increasing number of features with the different rankings . . . . .	63
2.22	Performance of the different algorithms during the grid search by means of precision and recall for <i>Red</i> class. The red solid lines indicate the minimum performance thresholds indicated by domain experts. Good performer algorithms are expected to produce results in the upper right corner of the graphs. . . . .	66
2.23	All the combinations of parameters $C$ and $\gamma$ that produce <i>SVM</i> models able to meet the minimum performance thresholds indicated by domain experts . . . . .	66
2.24	Impact of different initial random seed values on precision and recall for <i>Red</i> class, for the two different activation functions . . . . .	67
2.25	Confusion matrix for Fuel high pressure system Validation set. $Precision_{Red} : 0.824$ and $Recall_{Red} : 0.852$ . . . . .	70
2.26	Confusion matrix for Fuel high pressure system Test set. $Precision_{Red} : 0.790$ and $Recall_{Red} : 0.742$ . . . . .	71

2.27	Mismatch matrix for Fuel high pressure system Validation set. Experiment 8 exhibits many wrongly detected time windows . . . . .	72
2.28	Mismatch matrix for Fuel high pressure system Test set. . .	73
2.29	Confusion matrix of each model for Oxygen sensor Test set.	75
2.30	Distribution of the 90 <sup>th</sup> percentile of the derivative of the oxygen variable, by class. . . . .	76
2.31	Training performance for increasing training set size . . . . .	78
2.32	Bandwidth required to transmit signals and features to the cloud . . . . .	79
3.1	Examples of X-radiographs. . . . .	82
3.2	Abdomen CT scan slices and its 3D reconstruction (Medical news today [2018]) . . . . .	83
3.3	The detail from MRI scans of a patient’s head. . . . .	84
3.4	Ultrasound image of the foetus a 30 weeks of pregnancy in a sagittal scan. Measurements of fetal Crown Rump Length (CRL). (Moroder [2012]) . . . . .	85
3.5	Fluorodeoxyglucose <i>PET</i> images of the brain. Different colours indicate different functional areas of the brain. (Catarina Silva [2019]) . . . . .	86
3.6	Data flow and structure in LeNet (LeCun et al. [1998]). The input is a handwritten digit, the output a probability over 10 possible classes. (Dive into deep learning [2023]) . . . . .	91
3.7	Architecture of <i>U-Net</i> (Ronneberger et al. [2015]). Each blue box corresponds to a multi-channel feature map. The number of channels is denoted on top of the box. The x-y-size is provided at the lower left edge of the box. White boxes represent copied feature maps. The arrows denote the different operations. (University of Freiburg [2011]) . . . . .	92
3.8	World map showing the number of total COVID19 cases in each country, WHO [2023b] . . . . .	95

3.9	Representative serial chest radiography of patients with <i>COVID-19</i> infection. A, Images in a 29-year-old man with rapid respiratory deterioration after symptom onset shows the progression from lower lung-predominant interstitial and airspace opacities on day 3 to diffuse involvement with extensive airspace disease on days 4 and 6. B, Images in a 56-year-old-man with COVID-19, presenting initially with a normal chest radiograph, which then progressed to lower lung-predominant interstitial and airspace opacities at day 9, which mildly worsened by day 15 Stephanie et al. [2022]. . . . .	98
3.10	Data selection and labelling of COVID-19 images for Radiology diagnostic workflow optimization use case. Images have been assigned to different groups according to RT-PCR test (for COVID-19) and assessment by a team of radiologists (for other diseases). . . . .	99
3.11	Sample distribution over the age classes, divided by label . . .	100
3.12	Brixia score distribution with sex stratification on the Brixia <i>COVID-19</i> dataset . . . . .	102
3.13	Brixia score: (a) zone definition and (b–d) examples of annotations. Lungs are first divided into six zones on frontal chest X-rays. Line A is drawn at the level of the inferior wall of the aortic arch. Line B is drawn at the level of the inferior wall of the right inferior pulmonary vein. A and D upper zones; B and E middle zones; C and F lower zones. A score ranging from 0 (green) to 3 (black) is then assigned to each sector, based on the observed lung abnormalities. . . . .	103
3.14	Examples of hands and wrists x-radiograph used to estimate the bone age . . . . .	105
3.15	Bone age distribution and the number of images in the training set. . . . .	106
3.16	Bone age distribution and the number of images in the validation set. . . . .	107
3.17	Gender distribution and the number of images in the training, validation, and test bone age data sets. . . . .	108
3.18	Example of CT scans acquired to estimate different levels of Agatson calcium score . . . . .	109
3.19	Distribution of CAC score in the provided dataset for the Quantification of calcium score from <i>CXRs</i> . . . . .	111

3.20	Labelling hierarchy from Irvin et al. [2019] dataset . . . . .	113
3.21	A 5-layer dense block with a growth rate of $k=4$ . Each layer takes all preceding feature-maps as input. . . . .	114
3.22	Detailed scheme of <i>BS-Net</i> from Signoroni et al. [2021]. In particular, in the top-middle the <i>CXR</i> to be analyzed is fed to the network. The produced outputs are: the segmentation mask of the lungs (top-left); the aligned mask (middle-left); the <i>Brixia score</i> (top-right). . . . .	115
3.23	Preprocessing pipeline for the second-place method used to construct inputs to the neural network. The image is manually cropped and resized to a length of 560 pixels, and the contrast is enhanced; this is followed by extraction of 49 patches of $224 \times 224$ pixels . . . . .	117
3.24	Hand masking by <i>U-Net</i> , as it is adopted in the fifth-place solution in RSNA Pediatric Bone Age Machine Learning Challenge . . . . .	118
3.25	AI4CAD: Deep learning to detect severe <i>CAD</i> from <i>CXR</i> . (D’Ancona et al. [2023]) . . . . .	119
3.26	High level illustration of adopted methodology in the case of <i>CT</i> scans. <b>a)</b> <i>query</i> : original image; <b>b)</b> <i>image pre-processing</i> ; <b>c)</b> pre-processed picture; <b>d) e)</b> <i>feature extraction</i> ; <b>f) g) h)</b> <i>reference set</i> : database of labelled cases from training set with corresponding projections; <b>i) l)</b> <i>distance based regression</i> ; <b>m)</b> final prediction . . . . .	120
3.27	Optional caption for list of figures . . . . .	121
3.28	The pre-processing stage. The original image <b>a)</b> is properly resized <b>b)</b> . The resized picture <b>c)</b> is then processed by a segmentation algorithm <b>d)</b> to extract a binary mask <b>e)</b> indicating the region of interest. To define the final cut <b>g)</b> , the binary mask <b>e)</b> is framed by the smallest rectangle <b>f)</b> . The resized picture <b>c)</b> is scaled <b>h,i)</b> and finally cropped <b>j)</b> using the computed rectangular cut to obtain the final picture <b>k)</b> . . . . .	122
3.29	Some examples of transformation applied during data augmentation . . . . .	128
3.30	How the random crop is executed during data augmentation. The crop coordinates are randomly chosen with uniform distribution . . . . .	129
3.31	Distribution of logarithmic transformed CAC score . . . . .	132

3.32	Distribution of transformed CAC score . . . . .	133
3.33	Priority Matrices: a) <i>FIFO</i> policy, b) <i>ImageNet Transfer Learning</i> and c) <i>Self supervised learning</i> AutoEncoder . . . . .	142
3.34	Scatter plot of predictions for BrixIA dataset . . . . .	143
3.35	Error distribution of predictions for BrixIA dataset . . . . .	144
3.36	Scatter plot of predictions for pediatric bone age . . . . .	146
3.37	Error distribution of predictions for pediatric bone age . . . . .	147
3.38	Scatter plot of predictions for pediatric bone age . . . . .	148
3.39	Sensitivity-Specificity curve for pediatric bone age . . . . .	149
3.40	Error distribution of predictions for pediatric bone age . . . . .	150
3.41	Overview of the integration between the machine running an implementation of the proposed methodology and existing IT environment in the hospital . . . . .	151
3.42	Overview of the evaluation method for diagnostic workflow optimization trial at hospital. . . . .	153
3.43	COVID-19 contagions in Provincia di Torino district during the period of interest. . . . .	154
3.44	Concurrent 7 days rolling average trends of pandemic decreasing contagion wave and priority <i>COVID-19</i> index score elaborated by the proposed method, during the period of interest. . . . .	155
3.45	Angular similarity within and between classes: a) <i>ImageNet Transfer Learning</i> and b) <i>Self supervised learning</i> AutoEncoder . . . . .	159
3.46	TSNE dimensionality reduction of features extracted from the images in the reference set using: a) <i>ImageNet Transfer Learning</i> and b) <i>Self-supervised learning</i> AutoEncoder . . . . .	160
3.47	System prioritization performance with a reduced reference set. . . . .	161
3.48	System critical case identification performance with a reduced reference set. . . . .	162
3.49	Sensitivity-Specificity curves for identification task performance with a reduced reference set using <i>ImageNet</i> approach. . . . .	162
3.50	Sensitivity-Specificity curves for identification task performance with a reduced reference set using <i>SSL</i> approach. . . . .	163
3.51	DISTMAT loss values during training: a) <i>CAC score</i> b) <i>RSNA Pediatric Bone Age Machine Learning Challenge</i> and c) <i>BrixIA</i> . . . . .	164





# Chapter 1

## Introduction

In [Cambridge \[2023\]](#), deterioration is defined as the fact or process of becoming worse. The concept of deterioration is a common experience in people's daily life and it applies to a wide range of systems, both artificial, such as human-made artefacts, machines and structures, and biological, such as organisms, cells or tissues. In [Figure 1.1](#), a typical case of deterioration in the food industry: apple spoilage caused by several factors such as bacteria and various fungi; it has been estimated that around one-third of all the world's food produced for humans is wasted because of spoilage every year ([Gustavsson et al. \[2011\]](#)).



Figure 1.1: Photos of different stage of deterioration for an apple [Labs \[2017\]](#).

In artificial systems, the ageing process can occur due to a large number of reasons, depending on the specific case. Some of the most notable examples are the physical environmental stressors, such as wearing, temperature

stress, cracking and fatigue, the chemical agents, such as corrosive compounds, material oxidation (Figure 1.2) and microbial contamination, the mishandling, as misuse and incorrect installation, and the incorrect storage. Material degradation often occurs in mechanical equipment, materializing in unpredicted and harmful consequences (International SAE [2023]). The process decreases the component or system level of performance to a point where they fail to provide their functionality. To restore the system is usually necessary an external maintenance intervention.

In biological systems, deterioration can happen due to various factors such as pathology, environmental agents and ageing. Pathology usually involves the interaction between the organism and an external microbial agent, such as bacteria and viruses, but it can also be caused by internal dysfunctionalities as it is in the case of tumours and autoimmune diseases. Organisms exchange energy and matter with the environment to procure the elements that are needed to keep their internal state steady, in a process also referred to as *Homeostasis*. It can happen that, during this exchange, pathogenic microbes enter the organism causing damage to the vital processes (Figure 1.3). The accumulation of exposure to external and internal stressors deteriorates health resulting in both the emergence of a disease or the process of ageing. In the case of the biological systems, the effects of deterioration can be impaired functioning, reduced quality of life, and ultimately death. The organisms include a large variety of internal processes to reduce the effects of deterioration, preserving and recovering as much as possible the original status, such as the DNA damage management depicted in Figure 1.4. In medicine, restoration of health can also involve medical treatments and other external interventions by physicians that aim to improve the functioning of cells and tissues.

In all cases, deterioration refers to the gradual decline in the performance or functionality of a system over time. When the process reaches an advanced status that the system is no more able to provide its functionality or performance fails to meet some mandatory requirements, it reached its *End Of Life – EOL*. Depending on the applications and kind of systems, the original level of functionality, or part of it, can be recovered by external action or interaction. This is the case of maintenance when an operator (human or automatized) acts on the system by fixing the effects of deterioration on the materials or of medical treatment interventions when a physician attempts to restore an acceptable level of health. In some cases, the *EOL* status is not reversible and the system functionalities cannot be restored anymore




Figure 1.2: Stages of Concrete Corrosion (Monib Co Ltd. [2017]).

(e.g., organisms' death, universe thermal death), causing severe disruptions and high costs to the processes and network the system is part of.

The pace of deterioration progress can be influenced by the interaction between the environment and the system by both slowing down it and, possibly, reverting it, as it is in the case of maintenance, and by speeding it up. For example, external intervention can cause abrupt decrements in performance, inflicting severe damages to system components (e.g. damages caused by incidents, infection or intoxication), or it can increment the overall speed of the ageing process requiring a higher than acceptable output to the system (e.g. very demanding usage, bad habits, fatigue).

Because of the variable progression of the deterioration phenomena, often is not possible to predict precisely when it will reach a critical level. Also for this reason, the estimation and prediction of degradation level is important and is emerging as a relevant field of study. Monitoring and quantifying the evolution of the ageing process enable the adoption of effective actions in




Figure 1.3: In influenza virus infection, viral glycoproteins attach the virus to a host epithelial cell. As a result, the virus is engulfed. Viral RNA and viral proteins are made and assembled into new virions that are released by budding. ([Microbiology \[2023\]](#)).

a preventive instead of reactive manner. Considering the high costs associated with the failure of the system in some domains (e.g. health, public transportation, manufacturing) it is important to prevent such conditions. On the other side, too frequent or oversized adoption of corrective actions results in not optimal strategy causing the waste of resources. Precise degradation monitoring enables the adoption of optimal intervention strategies with the right balance between the prevention of failure and resource costs. It also allows performing quantitative analysis of the degradation process supporting both management, design and troubleshooting.

This is associated with the concept of predictive maintenance in artificial systems, an approach in which, starting from acquired data, the decay is




Figure 1.4: DNA damage in non-replicating cells, if not repaired and accumulated can lead to aging. DNA damage in replicating cells, if not repaired can lead to either apoptosis or to cancer [Wikipedia \[2018\]](#).

tracked and this information is used to optimize intervention improving the trade-off between system reliability and availability. It is also the case of preventive medicine, in biology, by which is possible to minimize the risk of a disease occurring by adopting corrective actions much more in advance, with potentially vast improvements in life quality and costs. Moreover, the measurement of degradation can support the design of more appropriate treatments and can provide quantitative feedback about the effects of a cure by analyzing the changes in trends.

## 1.1 Document structure

In this PhD thesis, a comprehensive study of the techniques for the estimation of the level of deterioration in different application fields is presented. To provide the reader with an overall overview, the context and the state of the art for each application field are presented. In this document, some methodologies for the estimation of ageing level are proposed and tested

in the presented applications fields, and finally, results are discussed. The thesis is structured into four chapters, each one focusing on a specific aspect of the problem and application.

Chapter 2 presents a method for the estimation of the level of degradation in the context of predictive maintenance for automotive applications. A machine learning approach based on statistical feature extraction and selection to estimate the level of degradation in diesel engines is proposed. The chapter also presents and analyzes the results obtained with this method when processing real-world data.

Chapter 3 focuses on medical imaging as a tool for estimating the level of ageing or infection. A method based on computer vision and deep learning to quantify degradation in different radiological applications is proposed and the results obtained when applying it to real tomographies and radiologies are shown.

Finally, Chapter 4 concludes the thesis by summarizing the contributions and motivating their relevance for the study of degradation monitoring. It also proposes potential next steps for further performance improvements.

## Chapter 2

# Predictive maintenance

Due to a multitude of factors, which largely depend on the specific use case, the artificial system components degrade over time. This phenomenon is ubiquitous in every person's life and is recognized as a fundamental aspect of reality. The degradation process goes through a number of states, depending on the severity level. The specific evolution is closely related to the individual case, but can generally be described as a sequence of stages, illustrated in Figure 2.1, depicted as follows:

- **New part:** At the start, the component is fully functional. All performances are in the ranges of specifics and no deterioration is present.
- **Useful life:** some defects may be present, but at the beginning, they are barely noticeable and have no measurable effects on the overall system behavior. During this phase, the magnitude of the defect increases, as does its impact on performance and behavior.
- **Component failure:** at this stage, the component's performance is so degraded that it no longer meets its requirements. This situation has an impact on the overall system behavior, which affects its performance.
- **System failure:** the component's degradation is so pervasive that the system can no longer perform its intended functionality.

During the design phase, considerable effort is invested by system engineers to mitigate or delay the impacts of such phenomena. The careful selection of suitable materials and robust mechanisms can exert a substantial influence on the system's resistance to stressors. Additionally, practices like overengineering, involving redundancies, can enhance system reliability.

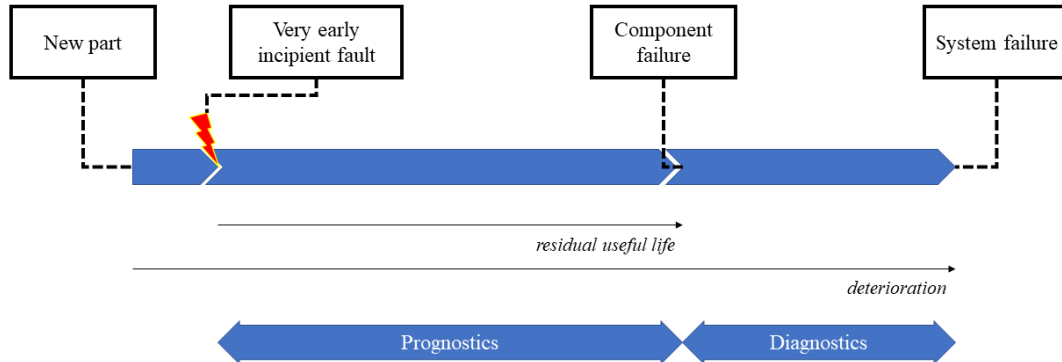


Figure 2.1: The progression of failure over time. A very early incipient failure occurs to the part, gradually degrading the component performance, until it is no more able to deliver its functionality. Further degradation of the component causes the entire system failure. The time remaining until the component failure is called *Residual useful life* — *RUL*, during which Prognostics technologies monitor the system to detect the deterioration progression. After the component failure, Diagnostics is in charge to detect and identify the failure to notify the necessity of intervention

This is achieved by multiplying the components capable of performing a particular function, ensuring that if a designated part malfunctions, the other duplicates can step in. Furthermore, such an approach provides the operator with the opportunity to address the faulty component (e.g., through repair or replacement) while the system continues to deliver its output. It should be noted that redundancy is accompanied by increased costs, which necessitate evaluation based on the unique context of the business model, application, and regulatory framework.

Component failure is typically detected through diagnostics, an engineering discipline used to detect and identify component or system malfunctions when they are present. It is usually set to raise an alert when operating requirements are no longer met. In some industries the diagnostic strategy is strongly regulated by legislation, specially if related to human or environmental safety.

Given the high costs and impacts of component and system failures, particularly in safety-critical systems such as airplanes, public transportation,



or high-output manufacturing, it is crucial to monitor the system and predict performance disruptions with sufficient margin to act. This is the goal of prognostics, an engineering discipline focused on predicting the time at which a component or system will no longer meet its intended functionality. This lack of performance is often a failure beyond which the system can no longer be used to meet desired performance. An essential concept in decision-making for contingency mitigation is the *Remaining useful life – RUL* defined as the time remaining until the failure occurs. *RUL* estimation serves as a fundamental tool for optimizing maintenance strategies by facilitating the timely replacement or repair of components, thereby minimizing downtime and operational disruptions. By accurately estimating *RUL*, maintenance actions can be scheduled proactively, maximizing resource utilization and enhancing overall system reliability.

The prognostics algorithm typically predicts the future performance of a component by evaluating the extent of deviation or degradation of a system from its expected normal operating conditions. This requires a thorough understanding of the failure mechanisms that are likely to result in the degradation leading to eventual system failures. Consequently, it is crucial to possess initial knowledge of potential failures, encompassing details such as the site, mode, cause, and mechanism of a product. Such knowledge is essential in identifying the specific system parameters that necessitate continuous monitoring.

The emergence of prognostic technology has had a transformative impact on asset management strategy, summarized in Figure 2.2. In the early days of lifecycle management history, roughly from 1930 to 1960, the predominant approach was *Corrective Maintenance*, in which a component or system is only fixed when it is broken. This approach had the advantage of reducing unnecessary part replacement, ensuring that components were used throughout their entire useful lives. However, corrective maintenance also had significant disadvantages: asset availability was generally poor or difficult to predict or control, and costs were high in the event of a system outage. With this approach, it is not possible to be prepared in advance. Furthermore, deferring maintenance activities until the final stages of degradation could lead to more significant system damages, resulting in increased costs and disruptions.

Given the pros and cons of corrective maintenance, *Preventive Maintenance* became very popular between 1960 and 1985. In this scheme, operations are scheduled in advance, and parts are replaced regardless of their

level of deterioration. This approach has the advantage of being highly predictable, minimizing system failure outages. However, the replacement of parts that were still fully functional or only minimally deteriorated made this approach wasteful and costly. Moreover, scheduled maintenance had an impact on system availability, pausing operations for the duration of the maintenance intervention.

Motivated by the need to comprehend the underlying causes for an elevated crash rate within jet aircraft, a series of comprehensive investigations on operational airplanes during the 1960s was instigated through the sponsorship of the *Federal Aviation Administration – FAA*. The findings debunked the central principle of Preventive Maintenance, which entailed defining a specific operational lifetime to individual components or subsystems, particularly in the case of complex systems.

In response, the paradigm of the *Reliability-centered Maintenance – RCM (of Automotive Engineers [2009], Nowlan and Howard [1978])*, emerged to overcome these limitations. *RCM* surpasses the anticipation of life expectancies, redirecting towards the management of the failure process itself, acknowledging that failures originate from diverse origins beyond ageing and the inability to predefine the trajectory of deterioration. This framework accentuates the divergence between user-derived asset requirements and the ones for design reliability.

Within *RCM*, the operational context of machinery is outlined through tools like *Failure Mode Effects and Criticality Analysis – FMECA*, which trace and characterize all conceivable failure modes. For each identified fault, probabilities of occurrence and criticality scores about safety, operability, and costs are computed. Subsequently, the most suitable maintenance strategy is determined based on this evaluation. In scenarios involving elevated risks and probabilities, *RCM* may even necessitate the reevaluation of system design to mitigate these factors to more tolerable levels. The prescriptions and analyses of *RCM* undergo constant review during the system's operational lifespan, remaining subject to updates based on observations.

*RCM* emphasizes the importance of managing assets based on their actual conditions, paving the road to the large adoption of *Predictive Maintenance*, between 1985 and 2005. The key concept of this approach is to use information about the system's usage, status and deterioration dynamics to modulate maintenance interventions so that they occur just before enough damage has occurred to avoid failure, but not so frequently as to be inefficient in terms of cost and time. In *Condition-based Maintenance*, physical or

empirical damage accumulation models are used to determine when a part or subsystem should be replaced. These models are updated using data from ad-hoc inspections based on risk. The *Prognostics and Health Management* approach promotes the adoption of real-time data measurement systems to be used in damage models to enable prompt intervention, fault or anomaly detection, and isolation. This approach considers the implementation of both prognostics and diagnostics techniques.

Finally, in the present days, fueled by the development of other technologies such as *Internet of Things – IoT*, *Artificial Intelligence – AI*, and *cloud computing*, the *Intelligent Health Management* paradigm is increasingly being adopted. In this approach, prognostic and diagnostic tools are integrated with asset control to provide a holistic overview capable of optimizing asset usage and maximizing its efficiency.

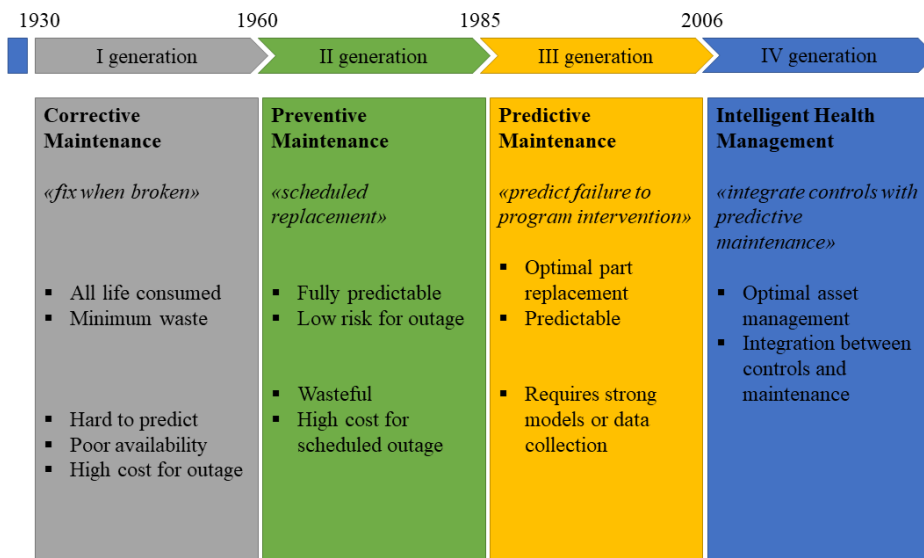


Figure 2.2: How the approach to maintenance changed over history. Prognostics stimulated a more cost effective trade-off between risk and costs

In the field of prognostics, there are three main technical approaches to building models: *physics-based*, *data-driven*, and *hybrid*. The key aspects of the different methods are summarized in Table 2.1.

*Physics-based* prognostics involves incorporating physical models of a system into the estimation of its current level of deterioration, its evolution over time, and its *Remaining Useful Life – RUL*. These models are developed based on prior knowledge of the system components, their interactions,

	<b>Physics-based</b>	<b>Data-driven</b>	<b>Hybrid</b>
<b>Fundamental aspects</b>	Incorporates physical models of a system based on prior knowledge and physics of the phenomena	Relies on pattern recognition and machine learning techniques to detect changes in system states	Merges aspects of both physics-based and data-driven approaches to achieve better performance
<b>Pros</b>	Higher accuracy for well-understood systems, provides insight into system operation, can handle small datasets.	Fast and cost-effective, provides system-wide coverage, suitable for complex systems.	Combines strengths of both approaches, can leverage available prior knowledge, suitable for complex systems with limited data.
<b>Cons</b>	Difficult and labor-intensive to develop, simplifications may reduce accuracy, requires specialized domain knowledge, uncertainty management can be complex, may not perform well on complex systems.	Wide confidence intervals, requires substantial amount of training data, poor performance with small datasets.	Requires domain knowledge for model selection, combining approaches can be complex and time-consuming.
<b>When to Use</b>	Well-understood systems with established physics-based models, systems where deeper understanding is required	Complex systems where understanding of first principles is not comprehensive, systems where run-to-failure data is challenging to obtain	When available prior knowledge can be leveraged to improve accuracy and coverage
<b>Examples</b>	fatigue life model for ball bearings, crack growth model, stochastic defect-propagation model.	Multi-step prediction for remaining useful life of bearings based on vibration data, CNN-based method for detecting and classifying faults in rotating machinery, LSTM-based model for predicting wind turbine blade failures	Usage of data-driven techniques to tune the parameters of physics models

Table 2.1: Comparison of Physics-based, Data-driven, and Hybrid Approaches

and the underlying physics, but, specially for complex ones, may represent a simplified view of the system that can reduce accuracy while increasing coverage, introducing necessary simplifications that need to be accounted for in uncertainty management. Developing these models can be time-consuming and labour-intensive, particularly for complex systems, and requires specialized domain knowledge. Micro and macro-level models are distinguished within this approach, with micro-level models relying on sensed system parameters to infer critical damage properties otherwise rarely available by direct observation. The micro-level models are often referred to as the damage propagation model. Macro-level models are mathematical representations of the system as a whole and define the relationships between input, state, and output variables. Notable examples of this approach include fatigue life model for ball bearings (Yu and Harris [2001]), crack growth model (Paris and Erdogan [1963]), and stochastic defect-propagation model (LI et al. [2000]).

*Data-driven* prognostics relies on pattern recognition and machine learning techniques to detect changes in system states (Liu et al. [2009]) and their adoption is increasing due to promising results obtained by these methods. *Data-driven* approaches are particularly suited when the understanding of the first principles of system operation is not comprehensive or when the system is too complex to develop an accurate model at a reasonable cost. Therefore, the main advantage of *data-driven* approaches is their ability to be deployed faster and cheaper than other approaches, and their higher ability to provide system-wide coverage. However, a significant disadvantage is their potential for wider confidence intervals than other approaches, and their need for a substantial amount of data for training. This last aspect is particularly relevant: in real cases obtaining run-to-failure data can be challenging, time-consuming and expensive and collecting enough data to include all possible conditions (both usage and environmental) is nearly impossible for most non-stationary systems. In fact, both the quantity and quality of data are important for the efficacy of *data-driven* methods.

In an attempt to leverage the strengths of both *physics* and *data-driven approaches*, *hybrid* methods are also adopted (Pecht and Jaai [2010], Liu et al. [2012]). Pure *data-driven* or *physics-based* solutions are rarely used in real-life situations, and aspects of each approach are often merged to achieve better performance. One typical case of this setting is the usage of data-driven techniques to tune the parameters of physics models. One example of hybrid modelling is the application of particle filter method (Del

Moral [1997], Liu and Chen [1998]) which combines a *physics-based* model of the system with a *data-driven model* of the noise in the system. The particle filter method has been successfully applied to a range of engineering applications, including gas turbine engines, bearings, and batteries. Another example of *hybrid* modelling is the use of deep learning techniques to improve the accuracy of *physics-based* models. In this approach, the deep learning algorithm is trained on a large dataset of system data, and the resulting model is used to improve the predictions of the *physics-based* model.

## 2.1 State of the art

In recent years predictive maintenance emerged as a relevant field of research. The rise of some core technologies such as the *Internet of Things paradigm – IOT* and *Artificial intelligence – AI* has enabled important successes in research and many applications in the real world. Predictive maintenance has been studied by many fields, which makes it almost impossible to address it in dept as it whole. In the following paragraphs, a curated selection of contemporary academic literature will be examined to address the following inquiries:

- *Q1*: To what extent has predictive maintenance been adopted across industrial sectors?
- *Q2* What are the principal challenges faced in its application to design processes?
- *Q3* To what extent are the various methodologies and techniques currently being applied?

The work by Pandian and Ali [2009] provides an overview of the entire topic, discussing possible methods to address the problem, such as *Artificial Neural Networks – ANN* and *Genetic Algorithms*. In Nunes et al. [2023], the authors focused on the main challenges behind the development of data-driven predictive maintenance systems: the reliability of sensor readings, the time constraints in the elaboration and transmission of collected information and the inability of prognostic algorithms to generalize over different use cases. The relevance of predictive systems is spread across multiple applications (Dalzochio et al. [2023]): it is largely adopted in aviation (Behera et al. [2019], Boller [2001], Byington et al. [2002], Cook [2007], Desell et al.

[2014], Furch et al. [2017], Hruz et al. [2021], Leao et al. [2008]) to monitor failures of aircraft and helicopters such as the one occurring in gearboxes and jet engines; prognostics is also gaining momentum in the automotive field (Chen et al. [2021], Homborg et al. [2018], Liu et al. [2005], Nixon et al. [2018], Vachtsevanos and Valavanis [2018], Jagannathan and Raju [2000], Lebold et al. [2012]) where is applied for both commercial and military applications. It is also adopted in the naval domain (Pal [2019], Cipollini et al. [2018], Liu et al. [2007], Tambe et al. [2015]) for both mechanical parts failure prediction and fleet monitoring, railways (Rabatel et al. [2011]), and in many industrial applications (de Andrade Lopes et al. [2021], Dalzochio et al. [2020], Lall et al. [2012a], Lall et al. [2012b], Lee et al. [2014], Nordal and El-Thalji [2021], Shamayleh et al. [2020]) such as the monitoring of power transformers, electronics, medical devices, and rotatory machinery.

Among the different approaches proposed by the researchers, three main approaches arise: *knowledge-based*, *physics-based* and *data-driven*. In *knowledge-based* approaches, applied in industries since the beginning of the 1990s (Montero Jimenez et al. [2020], Freyermuth [1991], Majstorović and Milačić [1990], Vingerhoeds et al. [1995], Vafaei et al. [2019], Cao et al. [2019], Baban et al. [2018], Berredjem and Benidir [2018], Tang et al. [2019], Boral et al. [2019]), domain experts' knowledge is synthesized in a set of *IF-THEN-ELSE* rules that are applied to detect faults, determine the state of degradation of equipment or, in diagnostics, determine the root cause of faults. *Physics-based* approaches exploit mathematical models to describe the deterioration phenomenon and are often applied to widely explored use cases in literature such as fatigue and crack propagation (Nasution et al. [2012], Vachtsevanos et al. [2006], Qiu et al. [2002], Tinga [2010], Tinga [2013], Tinga et al. [2014], Yiwei et al. [2019]). Under the definition of *data-driven* approach, both statistical (Kaiser and Gebraeel [2009], You et al. [2010]) and machine learning techniques (Dalzochio et al. [2020]) are adopted. Examples of statistical approach are the Markov models (Rabiner [1989]), Wiener processes (Iqbal and Aziz [2011]), *Principal Component Analysis – PCA*, (Abdi and Williams [2010], Kargupta et al. [2004]) and autoregressive-moving-average (Box et al. [2008]). Machine learning techniques comprise Decision Trees and Random Forest (Le et al. [2017], Breiman [2001]), Support Vector Machine (Nixon et al. [2018], Hsu and Lin [2002]), Artificial Neural Network (Iannace et al. [2019], Ripley [2007], Jain et al. [1996], Namuduri et al. [2020], Schmidhuber [2015]) and Deep Neural

Network (Ducoffe et al. [2019]). In some cases, data-driven and physics-based approaches have been merged get the benefits from both, as in Susto et al. [2015].

In light of the aforementioned state-of-the-art developments, several key findings can be highlighted to address the questions ( $Q1 - Q3$ ) posed at the outset of this section:

- *Finding 1*: Predictive Maintenance has been proven as a viable and efficaciously implemented approach across a diverse array of industries, including sectors such as aerospace, automotive, naval, railways, and industrial machinery.
- *Finding 2*: In the process of configuring Predictive Maintenance systems, due consideration must be accorded to constraints about execution time, communication bandwidth, and computational capabilities, which present notable challenges to deployment. Furthermore, it is imperative to acknowledge the inherent limitations in the reliability of available sensor data, which may be susceptible to substantial deviations. Additionally, the transition from one use case to another seldom entails the application of identical solutions, primarily due to the limited generalization ability of proposed systems.
- *Finding 3*: All three principal methodological categories find utilization within the domain of Predictive Maintenance. While knowledge-based and physics-based approaches have enjoyed well-established prominence for over three decades, data-driven methodologies, with a particular emphasis on *DL*, are experiencing an adoption boost in recent years across a multitude of application domains.

## 2.2 Specific use cases

This thesis will investigate the behaviour and performance of our proposed methodology by analyzing two specific use cases in the automotive domain: the Oxygen sensor for the Diesel engine and the Diesel engine Fuel high-pressure system. The Oxygen sensor case is an example of component-level prognostics, while the fuel high-pressure case is system-level. In both cases, experimental activities have generated a set of data with different levels of deterioration. The next sections will provide a more detailed description of the use cases, along with an analysis of the available data.



The application of prognostics solutions to automotive use cases poses several challenges from the implementation point of view. Signals are acquired by local sensors and processed by the onboard *Electronic Control Unit – ECU*, with limited hardware capabilities and rigid software architecture that poses significant constraints in the choice of more complex algorithms. Nowadays is possible to transfer the collected data to the cloud, enabling elaboration that requires more computational power and memory. By collecting data from the entire fleet, is it also possible in the cloud to run algorithms that compare vehicles, looking for anomalies, drifts and generalized trends that can require proactive intervention by the fleet managers. Vehicle fleet monitoring is also a valuable tool to analyze how the cars are behaving in the real life, confirming or not some of the choices made during the design phase and supporting troubleshooting. Centralized prognostic monitoring makes also it possible to optimize the scheduling of maintenance interventions, allowing the service organization to be prepared in advance to provide a flawless experience to the client, for example by procuring the necessary parts in advance. Because of hardware limitations and connection coverage, the cost of data transfer must be taken into account during the design. To achieve such a target splitting the whole algorithm into two parts, one running on the onboard *ECU* and another on the cloud, can represent an applicable tradeoff. Figure 2.3 illustrates the overall situation, showing the key elements and challenges of the automotive prognostic use case.

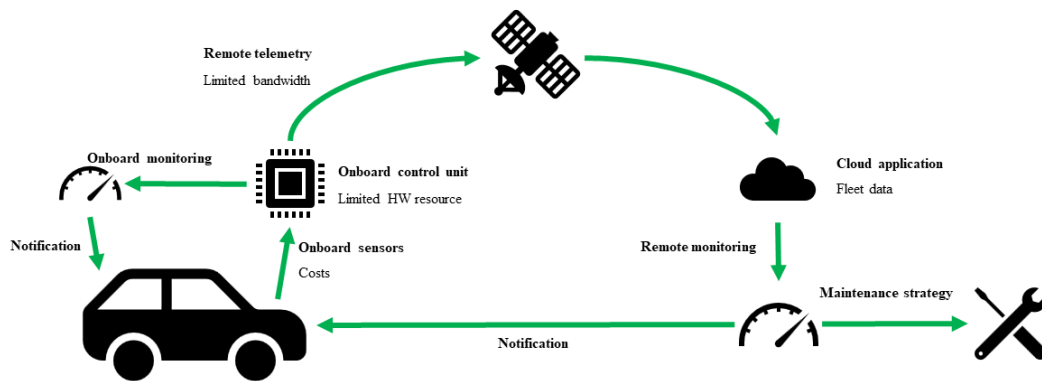


Figure 2.3: The context of automotive prognostics use cases. Signals are acquired by onboard sensor and elaborated by local Electronic control Unit. Manipulated signals can be sent by remote connection to the cloud infrastructure where can be used to perform prognostic algorithm and enable predictive maintenance. The output of predictive maintenance can be notified to both the vehicle user and the maintenance organization to schedule the intervention.

## 2.2.1 Oxygen sensor for Diesel engine

### Description

The oxygen sensor, showed in Figure 2.4, is a critical component utilized in measuring the oxygen concentration in exhaust gas of an internal combustion engine. The device examined in this study is a planar wideband zirconia sensor that incorporates two solid-state electrochemical fuel cells, namely the Nernst and pump cell. The zirconia membranes surrounding the two cells exhibit an ion permeability-gap voltage dual relation: the permeability can be controlled by regulating the voltage, and the voltage can be measured to estimate the oxygen exchange rate. In this way, the membrane surrounding the pump cell can act as an ion pump, while the one around the Nernst cell can serve as a reference value. An electronic circuit that includes a feedback loop controls the gas-pump current to maintain a constant output of the Nernst electrochemical cell. The magnitude of this current is proportional to the oxygen concentration of the exhaust gas. The oxygen in the exhaust is ionized on the sensor surface and transferred into the electrochemical cell utilizing the gas-pump current. Nevertheless, the mechanism of the sensor is highly sensitive to temperature, which is maintained by an electric heater, embedded in the device, at around 700°C. The Figure 2.6 illustrates the main functional components of the device.



Figure 2.4: NKG ZFAS-U2 oxygen sensor for automotive application, [NGK \[2023\]](#).



Figure 2.5: Oxygen sensor after exposure to exhausted gas in the after-treatment pipeline. The dust contained in the flow is deposited on the sensor, clogging the small holes that allow the gas exchange between the sensitive element and the environment and deteriorating its measurement performance.

The exhaust gas the sensor is continuously exposed contains soot particles resulting from imperfect fuel combustion. The soot can progressively accumulate in the small holes on sensor's surface, as visible in Figure 2.5. These small holes enable gas exchange between the sensitive element and the exhaust gases, depending on several conditions such as flow direction, temperature, humidity, and combustion setpoints. This accumulation of dust causes the sensor's dynamic to deteriorate, causing it to become progressively slower in detecting changes in oxygen concentration. The clogging process can also be reversed in specific conditions of high speed and hot flow that can remove or burn the accumulated matter, thus restoring some of the sensor's original performance.

In diesel engines, the measurement of oxygen concentration serves various purposes. The *ECU* is connected to the oxygen sensor and regulates emission-relevant strategies based on its readings. One typical usage of the signal is the compensation of drifts in the injection system by comparing




Figure 2.6: Functional schematic of planar wideband zirconia oxygen sensor. Embedded feedback loop circuit controls the pump current to regulate oxygen ions exchange with exhaust gas and keep the monitoring chamber gas composition constant. The magnitude of pump current is directly proportional to the oxygen concentration in exhaust gas, [Handrich \[2010\]](#).

the measured value with an expected one predicted using combustion models. The oxygen sensor measurement is also utilized in the aftertreatment system onboard diagnostic and control, in particular to regulate the combustion setpoints for control of regeneration cycles in *Lean NO<sub>x</sub> Trap – LNT*, a device used for emission reduction, and as a parameter of chemical reaction models.

Due to its importance, it is critical to monitor the clogging phenomenon correctly and trigger appropriate recovery or maintenance actions to ensure the optimal working status of the exhaust gas aftertreatment system, thereby limiting the release of potentially harmful pollutant gases. Slow readings, as in the case of heavy clogging, will result in more gas emissions into the environment and lower the efficiency of the overall system.

The aim of this investigation is to predict the onset of clogging in the oxygen sensor by leveraging data collected from a vehicle’s *ECU*. Direct measurement of clogging requires sensor removal and extensive sediment analysis, both of which are unsuitable for in-vehicle usage. To circumvent

this issue, we analyze this system using a dataset provided by an automotive manufacturer, developing a method for estimating the sensor’s clogging state by combining onboard measurements, including the one from oxygen sensor itself.

### Data

To achieve the objective of designing a system capable of predicting when the oxygen sensor is clogged and estimating the level of deterioration, a dataset has been provided by a major automaker. The dataset comprises measurements acquired during experiments, or *cycles*, which consist of readings from both onboard sensors and additional measurement equipment, integrated in the test bench, recorded while a real engine is run in a controlled environment using a predefined gas pedal profile, illustrated in Figure 2.7. The pedal sequence has been designed by engineers to reproduce the clogging phenomena and simulate on-road realistic usage. This setting minimizes the impact of external, uncontrollable factors and maximizes the repeatability of tests. The same engine has been used to produce the whole dataset.

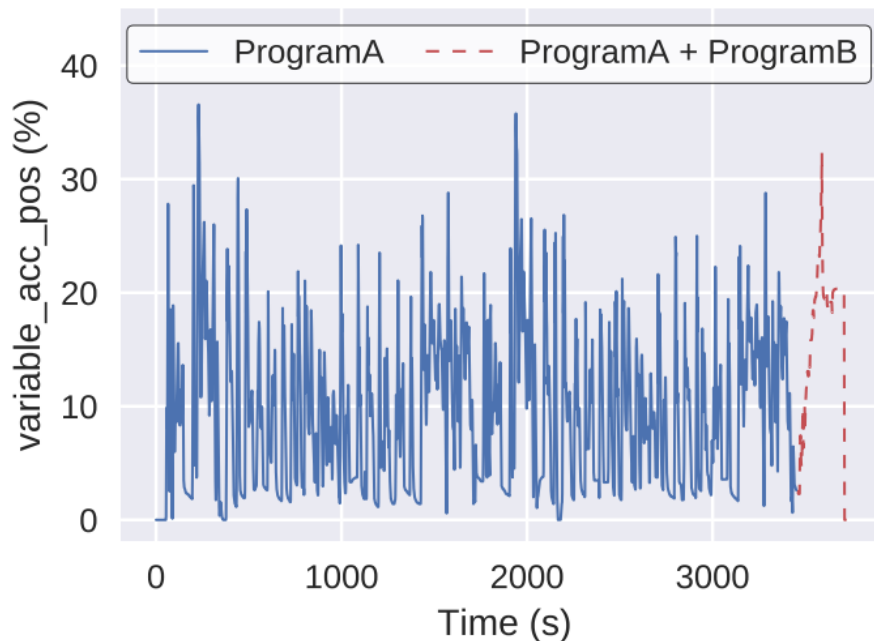


Figure 2.7: Acceleration pedal profile used in cycles. Each cycle can be divided into two parts, each recorded with a different software: Program A and Program B

Each cycle is divided into two parts with different durations: the first

part spans almost one hour (3450 seconds), while the second part covers only the last 5 minutes of each cycle. During the last 5 minutes, a specific manoeuvre, referred to as a *cut-off*, is executed to stimulate oxygen sensor dynamics and provide an indirect measurement of its clogging status, which consists of the sudden release of the gas pedal and is of particular interest in determining whether the cycles are clogged or not. The impact of the cut-off manoeuvre on the measurement of the oxygen sensor can be observed in Figure 2.8. Specifically, this figure depicts the oxygen sensor's response to the sudden release of the gas pedal, which leads to an influx of fresh air entering the exhaust line. As a result of this procedure, the oxygen sensor measurement experiences a sudden increase and reaches a value of 20.95%, as showed in Figure 2.9. This phenomenon is an important indicator of the oxygen sensor's dynamic behaviour and provides an indirect measurement of its clogging status, in fact the slope of such transition is influenced by the air flow reaching the sensitive element that is much reduced in clogged sensors. It is important to note that such a manoeuvre is only applicable as a measure of clogging in so heavily controlled and repeatable environment, as it is difficult to predict the engine usage by the driver in real-life settings, which could lead to unpredictable results.




Figure 2.8: Example of oxygen measurement signal during the last 5 minutes of the cycle: the cut-off manoeuvre generates a step-like stimulus, visible around time 250 seconds.

Two different software programs were used to collect the measurements:

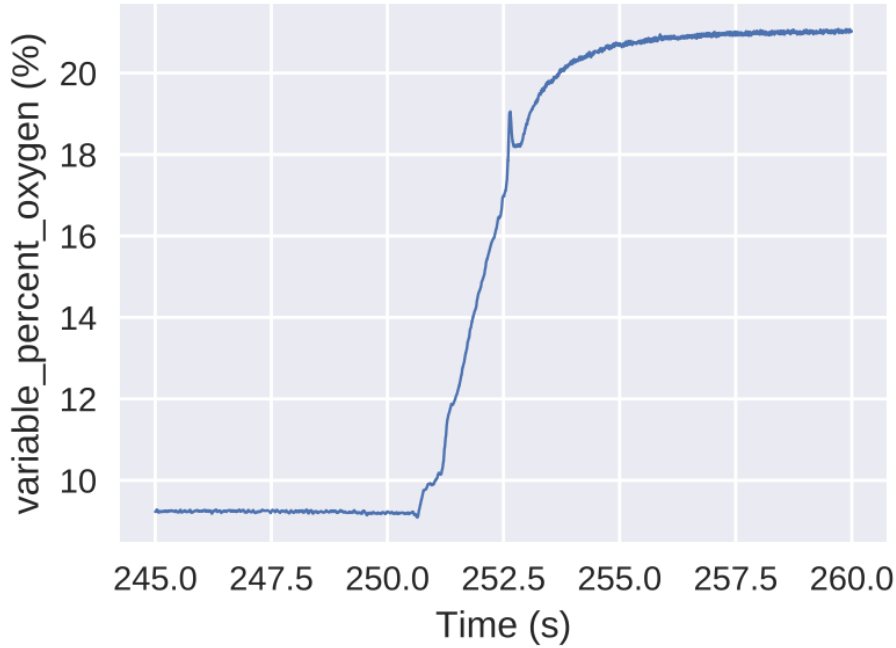


Figure 2.9: Detail of oxygen sensor sudden change in Figure 2.8 at 250 seconds

*Program A*, which covers the entire experiment, and *Program B*, which is used only in the last part of the cycle. This separation is shown in Figure 2.7. These programs have different characteristics and were used for similar but complementary purposes. *Program A* is capable of recording up to 50 different signals with a 1 Hz sampling frequency, while *Program B* can record 440 different parameters, including the ones measured by the test bench equipment, at up to 320 Hz. Although both programs theoretically can be used to record the same cycles, most probably because of tool malfunctions, the actual number of usable cycles differs between *Program A* and *Program B*. The latter is available in only 388 of the total 400 experiments. Within the signal dataset recorded by *Program B*, various measurements originating from the test bench equipment have been captured. Nevertheless, to faithfully represent solely the data typically available on board an average vehicle for this study, these variables have been excluded and dealt with using a domain-agnostic approach to avoid any context-specific bias.

The dataset provided did not include direct measurements of clogging status. To retrieve the different levels of degradation associated with each cycle, a semi-supervised labelling algorithm was designed, taking advantage of the last portion of the test recorded by *Program B*, during which a step-like command is given to the engine, stimulating the sensor dynamic.

The algorithm consists of three main steps: Response time measurement, Labelling, and Mapping.

#### *A. Response time measurement*

The oxygen level measured in the exhaust gas is directly influenced by the combustion in the diesel engine. When fuel reacts with fresh air, it burns, reducing the amount of oxygen in the exhaust gas. The more fuel burnt, the lower the resulting oxygen concentration. Because the main driver of fuel burnt amount is the torque requested at the engine by the driver, the pedal profile influences the oxygen content in the exhaust gas. During the *cut-off* manoeuvre, performed in the last 5 minutes of the cycle, the pedal is controlled so that the engine load moves from a relative amount of torque required to zero, by a sudden release of the pedal. By not burning any fuel, the exhaust air maintains the atmosphere's amount of oxygen, 20.95%, while passing through the engine, arriving in the exhaust line. The pedal transition is so rapid that the oxygen sensor reacts as fast as possible to the new measurement. If the sensor is affected by severe clogging, the exhaust gas, now full of oxygen, enters the sensor and reaches its sensitive element less quickly than when clogging is light or absent. This happens because soot mechanically blocks the holes or reduces their section, providing more resistance to gas flow. Figure 2.10 shows the typical shape of oxygen measurement response during cut-off. Starting from the initial oxygen level,  $O_{2_{start}}$ , at the pedal release, oxygen increases to the  $O_{2_{end}}$  setpoint, which is the fresh air oxygen content. To measure the response time, the time required by the sensor to increase its measurement to a target value equivalent to 63% of the total expected value range (from  $O_{2_{start}}$  to 20.95%) was counted. By applying this methodology, indicated by engineers from the automotive maker, to the provided data, the response time ranges anywhere between 1 and 2 seconds. This justified the choice of using *Program B* (up to 320 Hz) for this purpose, compared to *Program A* (1 Hz). Figure 2.11 illustrates how the calculated response time evolves along the different cycles. It can be noticed that there is a non-monotonicity of clogging phenomena and its general increasing trend, interrupted by several reset events.

#### *B. Labelling*

In this step, three categories of clogging based on the measured response time have been defined: green for clean sensors, yellow for intermediate clogging, and red for severe clogging. However, the determination of thresholds for separating these categories required expert input. Despite their efforts, domain experts were unable to provide clear-cut values for these



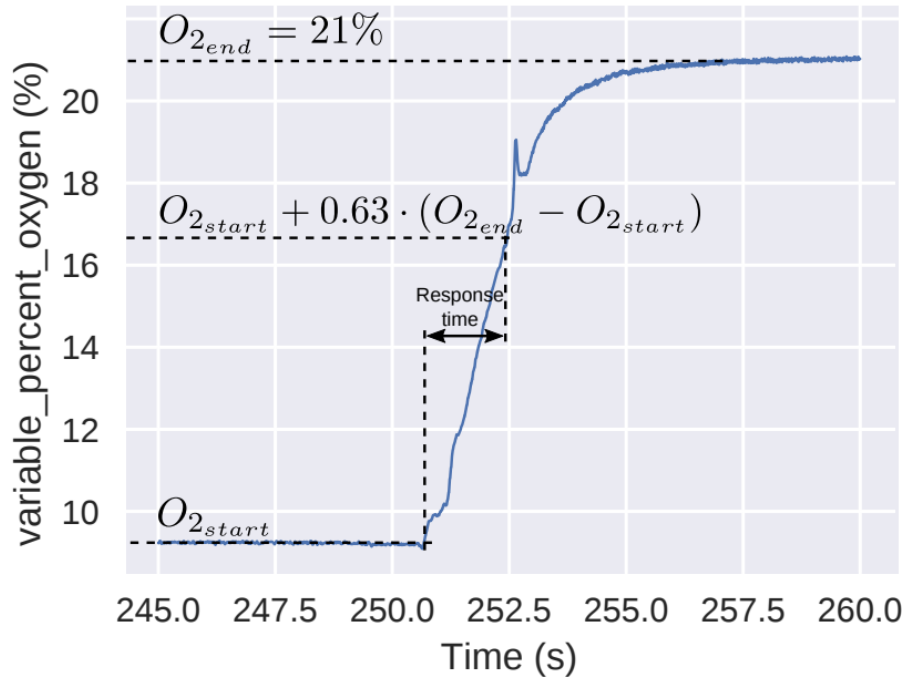


Figure 2.10: Response time measurement process




Figure 2.11: Response time trend throughout the experiment. The horizontal lines show the positions of the thresholds

thresholds. As a result, the thresholds have been established using the distribution of response times depicted in Figure 2.12. The objective was to ensure that the red class (which indicates malfunctioning) represented a minority class. Therefore, the lower bound for the red class has been set as first. Then, the samples have been split evenly between the green and yellow classes by setting the threshold that separated them. The response time evolution throughout the cycles is shown in Figure 2.11. It can be observed a jagged profile with many spikes, which is counterintuitive and undesired since soot accumulation is expected to occur slowly. To mitigate this problem, the response time trend has been smoothed using a moving average filter. Specifically, each sample in the sequence is averaged with the previous  $k$  and next  $k$  samples. The purpose of minimize jaggging and label switches in the entire labelling sequence have been aimed by analyzing the effect of different values of  $k$  on the number of label changes. Figure 2.14 shows the resulting amount of label changes for different values of  $k$ . Has been determined the final value for  $k$  using the elbow method, which yielded  $k = 2$  as an acceptable value. The effect of different values of  $k$  on the labelling is illustrated in Figure 2.13. With this setting, the green and yellow classes consisted of almost the same number of cycles, 164 and 163, and the remaining 61 cycles were assigned to the red category, as summarized in Table 2.2.

Class	Cardinality
Green	164
Yellow	163
Red	61

Table 2.2: Cardinalities for the three classes

### C. Mapping

In this final step, the labels obtained during the labelling phase are assigned to the cycles in the unlabeled dataset. Since different software programs are used to acquire the signals used to compute the label (*Program B*) compared to the one used for onboard sensors (*Program A*) that will be used in the predictive algorithm, this step requires careful attention in terms of implementation. Specifically, is needed to ensure that the mapping between the labelled and unlabeled datasets is accurate and consistent to avoid introducing errors or biases in the subsequent analysis.





Figure 2.12: Distribution of response times: the colors represent the assigned classes based on the defined thresholds

## 2.2.2 Diesel engine Fuel high pressure system

### Description

The high-pressure fuel system is an essential component of modern diesel engines for efficient operation and controlling emissions. Compared to earlier lower-pressure fuel injection, high-pressure offers power and fuel consumption benefits by injecting fuel as a larger number of smaller droplets. This increases the ratio of surface area to volume, resulting in improved vaporization from the surface of the fuel droplets and a more efficient combining of atmospheric oxygen with vaporized fuel, leading to more complete combustion ([Wikipedia \[2023\]](#)). The composition of this system, illustrated in Figure 2.15, may vary among different products, but some key components are generally recognizable at a higher scope:

- **Fuel tank:** stores the fuel at ambient pressure and temperature.
- **Fuel tank pump:** surges fuel from the tank and creates the necessary pressure to reach the high-pressure pump.
- **Fuel filter and heater:** due to impurities contained in the fuel, it is important to filter and heat it. The presence of compounds such as

Figure 2.13: Response time trend smoothed with different  $k$  values

paraffin that tend to solidify at low temperatures requires heating of the fuel to maintain fluidity.

- **High-pressure pump:** delivers fuel at a constant pressure to the common rail and is mechanically coupled with the cam.
- **Metering unit valve:** regulates the amount of fuel entering the pump to ensure that only the required amount of fuel is supplied to the common rail, avoiding compression of excess fuel to high pressure, improving hydraulic efficiency and avoiding excessive high fuel temperatures.
- **Common rail:** stores fuel at high pressure and delivers it to the injectors.




Figure 2.14: Number of label switches occurring as k increases

- **Pressure regulator valve:** allows the system to partially discharge excess pressure in the rail by opening its valve.
- **Rail pressure sensor:** provides the direct measure of fuel pressure in the common rail.
- **Fuel injectors:** regulate fuel injection in the combustion chamber.
- **Electronic Control Module:** processes signals from different sensors to control different actuators providing the required engine output.

To achieve its critical functionality, the high-pressure fuel system in modern diesel engines must maintain a narrow range of fuel pressure inside the rail. This range is necessary to guarantee the required flow dynamic to the injectors when they are opened. The *ECU* achieves this target by regulating the quantity of fuel entering and exiting the rail through the *Metering Unit Valve – MU* and the *Pressure Regulator Valve – PR*, respectively. The balance between these two actuators provides the necessary control of pressure levels. To regulate their action and compute optimal valve commands, the *ECU* reads the current measurement from the *Rail Pressure Sensor* and compares it with the target value. The performance of the high-pressure fuel system relies on the proper functioning of both the actuators and the




Figure 2.15: Architecture of a Fuel Injection System in a modern Diesel Engine. Yellow pipes indicate low pressure fuel, red lines the high pressure. Lines connecting the Control Unit to sensor and actuators represent information exchange. The arrows explain the direction of flow or exchange. This illustration has been derived from [Hannu Jääskeläinen \[2023\]](#). The components on which we concentrated our study are indicated with yellow labels.

pressure sensor. Any deterioration that impairs the ability of the *MU*, *PR*, or *Rail Pressure Sensor* to control the inlet and outlet flow will compromise the system's performance. In particular, the *MU* and the *PR* have limitations that prevent them from compensating for each other's shortcomings. Similarly, an incorrect measurement from the *Rail Pressure Sensor* will bias the entire control system and set the pressure level to an unacceptable value, either too high or too low. Therefore, the high-pressure fuel system must be properly maintained to ensure optimal performance and minimize harmful emissions.

## Data

In this study, the focus is on the deterioration of two crucial actuators: the *MU* and the *PR* in high-pressure fuel systems. Due to their robust design, observing the natural deterioration of these components in real-life scenarios is rare. However, in a large fleet of vehicles, such as the ones from

the car manufacturer that provided the data, the occurrence of this event is not negligible in absolute numbers and can lead to significant customer dissatisfaction and safety issues. Failure to provide sufficient fuel (e.g., *MU* too closed or *PR* too opened) to the injector renders the engine incapable of igniting and starting the vehicle. Similarly, overpressure (e.g., *MU* too opened or *PR* too closed) triggers engine shutdown for safety reasons.

Since this event is rare, it is challenging to acquire complete data in a timely manner at the engine test bench. Therefore, to overcome this limitation and obtain different levels of deterioration, the control algorithm parameters of the two components, namely the high-pressure fuel pump reference level and the valve timing aperture, have been manually tuned. This procedure has been validated internally and received the endorsement of domain experts. By imposing different engine conditions, these components have been forced to operate under conditions that drift from the specifications, affecting the fuel pressure and its flow, thereby simulating common system malfunction. This approach enables to precisely regulate the amount of deterioration and to emulate a faulty system up to a point where the *ECU* onboard diagnostics detects the failure and triggers engine shutdown or fuel delivery is not sufficient to sustain engine run.

To develop and validate a tool capable of predicting the deterioration status of the high-pressure system in automotive engines, a dataset has been provided by a major car manufacturer. This dataset consists of measurements acquired during experiments that were carried out using a predefined sequence of the accelerator pedal and engine loads. These predefined sequences, referred to as *driving cycles*, are designed to represent different usage profiles, such as urban, rural, and highway driving. In total, five different cycles have been tested, of which three are used during the standard homologation procedure of vehicles in different markets: *Real Driving Emissions – RDE* (Suarez-Bertoa et al. [2019]), *Worldwide Harmonized Light Vehicles Test Cycle – WLTC* (Tutuianu et al. [2015]), and *Artemis* standard test cycles (André [2004]). The remaining two cycles are obtained by recording actual driving sessions. This experimental setup and the adoption of the aforementioned procedure for inducing deterioration resulted in a dataset with over 230 experiments.

During each experiment, the measurement equipment records over 600 different signals, including some that directly monitor the fuel high pressure and injection systems, and others that provide a broader view of the engine condition (e.g., torque control, engine airflow, exhaust gas after treatment,

etc.), and finally others that are traced to check onboard diagnostics algorithms. The entire set of signals can be summarized into 13 categories, as indicated by domain experts and illustrated in Table 2.3.

Category	Number
Fuel Rail	184
Fuel injection	182
Engine airflow	33
NOx emissions	30
Oxygen levels	26
Torque control	21
Catalytic converter	15
Exhaust manifold	13
Exhaust temperature	12
Engine rotation	11
After treatment	10
Diagnostics	5
Others	74
Total	616

Table 2.3: Signals overview

To acquire the signals from the experiments, two possible sampling strategies are available: *Linear Sampling* and *Angular Sampling*. With *Linear Sampling*, the measurement is recorded at a constant frequency throughout the experiment, while with *Angular Sampling*, the system acquires the signal at a frequency that depends on the rotational speed of the engine and can vary over time. The frequency of *Linear Sampling* can range from 1Hz to 160Hz, while the frequency of *Angular Sampling* increases with the engine speed. This aspect is to be taken into account when data is manipulated, to compensate for sample misalignment between signals.

Each experiment in the provided dataset is associated with a label determined by applying a policy defined by domain experts. This labelling policy is based on monitoring the high-pressure control tracking error, denoted as  $P_{error}$ , which represents the difference between the current pressure measured in the rail by the sensor and its target value for optimal operation. The value of  $P_{error}$  is computed hundreds of times per second during the entire driving cycle and is smoothed using an iterative filter to remove spikes not related to deterioration. However, this signal exhibits strongly non-linear behaviour and cannot be used directly to infer the deterioration level in real-life situations. The policy labels the experiments with a 3 levels class etiquette, *Red*, *Yellow* and *Green*, based on  $P_{error}$  as follows:



- **Red:** The system is in a critical state, and the car must go to the service. This happens when  $P_{error} > \alpha$  for 5 seconds continuously at any time during the cycle.
- **Yellow:** The system starts drifting from the nominal behaviour, but it is not in a critical state. This happens when  $P_{error} > \beta$  for 2.5 seconds continuously at any time during the cycle.
- **Green:** The system works normally, i.e., in all other cases.

The two parameters  $\alpha$  and  $\beta$  are thresholds indicated by the domain experts representing the high-pressure fuel system requirements. The determination of these two parameters does not involve iterative procedures or other automated methods; instead, it is directly inferred from the system's design. These values signify thresholds beyond which the influence of fuel pressure deviation becomes observable in the comprehensive engine performance. The designated levels denoted as  $\alpha$  and  $\beta$ , correspond, respectively, to severe and mild impacts on emissions, fuel consumption, and torque, with their specific numeric values depending on the specific engine. If the onboard diagnostics detects a failure in the system, the experiment is discarded. This is justified by the goal of preventive maintenance, as the developed tool is meant to highlight deterioration before it creates issues in the entire system.

Because the high-pressure fuel system can still work properly, even if severely deteriorated, when not stimulated with certain manoeuvres, each experiment is labelled separately. It should be noted that different experiments related to the same deterioration can result in different labels by simply applying the policy. In such cases, the worst-case label among the cycles has been extended to others with the same deterioration. For each configuration of the deterioration parameters (i.e., high-pressure pump and valve drifts), from 2 to 5 different experiments have been obtained with different driving cycles. Label distribution is detailed in Table 2.4. The observed variations in class distribution are influenced by the specific executed *driving cycle*. A balanced distribution is notably achieved during experiments conducted using the *RDE* and *WLTC* cycles. Conversely, the *Artemis* and *Real* cycles predominantly yield instances classified as Green. These dissimilarities can be rationalized by disparities in pedal dynamics across the *driving cycles*. Furthermore, certain cycles have not been universally tested for all levels of induced deterioration due to project and technical constraints, which contributes to these observed variations.

The method proposed aims to ascertain the degree of deterioration of the high-pressure fuel system, classifying it into one of the three aforementioned categories: Red, Yellow, and Green.

	Cycle	Duration	Green	Yellow	Red	Total
Homologation	RDE	60	25	24	38	87
	WLTC	31	28	37	22	87
	ARTEMIS	54	15	7	4	26
Real	Driver1	43	15	7	4	26
	Driver2	66	19	1	2	22

Table 2.4: Dataset description per driving cycle. The *RDE* and *WLTC* cycles exhibit balanced class distribution, while the *Artemis* and *Real* cycles are predominantly associated with the Green class. This variation can be attributed to differences in pedal dynamics and test execution.

## 2.3 Proposed methodology

In both presented use cases, a set of signals is collected by the onboard measurement system and processed by a local *ECU*. Due to hardware capability limitations, part of the processing can be sent to a remote centralized computational infrastructure to analyze and determine the system’s deterioration level. To achieve the required levels of performance, data exchange should be carefully designed, taking into account available bandwidth and time constraints. In this context, the proposed methodology is designed to be divided into blocks that can be conveniently allocated on both local and remote units, depending on the specific use case. The methodology accepts signals in the form of time series as input and predicts the deterioration level as output, represented by discrete levels *Green*, *Yellow*, and *Red*.

The proposed approach is depicted in Figure 2.16 and can be summarized as follows. In brief, a collection of signals is recorded by the onboard measurement system and divided into time windows. For each time window, data is transformed through feature extraction for each signal and its derivatives. The extracted features from each time window are then utilized as input to the model to predict the deterioration level. During the training of this methodology, a larger set of signals has been conveniently recorded with associated deterioration labels. Signal selection is performed to discard unneeded signals, and feature selection is carried out to reduce the dimensionality of the problem. The model is trained by optimizing a loss function computed by comparing the predictions with the actual values.

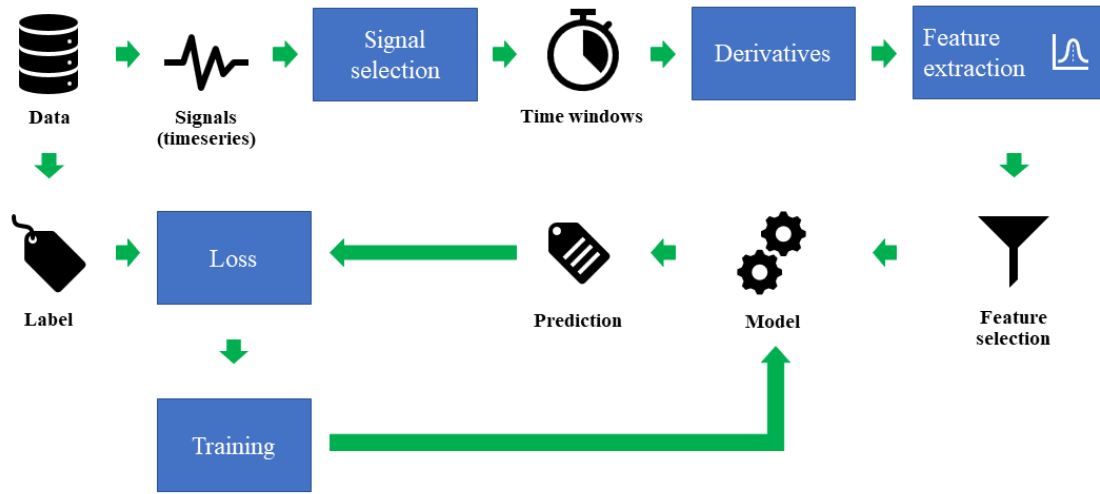


Figure 2.16: High-level overview of the proposed method to predict deterioration status of a system from measured signals.

In order to estimate the level of degradation, the proposed methodology follows a series of crucial steps, which are detailed below:

- **Data Transformation**, to represent long sequences of samples using few scalar values (e.g. statistics), reducing the problem dimensionality.
- **Model**, to predict the deterioration status of the system.

In order to attain the desired level of performance and optimize resource utilization, certain stages must be carried out during the algorithm design phase:

- **Signal selection**, to exclude redundant or irrelevant signals from the recorded corpus.
- **Feature selection**, to discard features that do not contribute to improving the predictive performance.
- **Model Training, Tuning and Selection**, to update the model parameters to improve predictive performance by optimizing the loss function.

In the following sections, each of the above stages will be detailed.

### 2.3.1 Method inference

#### Data Transformation

Based on the scenario presented in the preceding sections, it is evident that during normal operations, onboard *ECU* can record hundreds of different signals at considerably high frequency, resulting in thousands or millions of samples per acquisition. The huge volume of samples makes the problem intractable with a number of experiments that is compatible with real-world applications. Therefore, it is crucial to apply dimensionality reduction techniques to the recorded time series.

In both use cases, the computational capabilities of the *ECU* are exploited to aggregate samples into time windows  $w(k)$  and summarize each signal with statistics. Specifically, average, standard deviation and  $N_p$  equally spaced percentiles are computed for each signal. A deeper understanding of the data can be gained through the analysis of features computed within the frequency or time-frequency domains. These features offer insights into the signal's dynamics, revealing patterns that may remain concealed in the time-only domain. This additional information can potentially enhance the performance of the proposed methodology. However, these analyses involve complex mathematical operations, including fast Fourier transforms and convolution processes, which lead to a significant computational load and demand for heavy memory occupation. The onboard *ECU* utilized in the presented use cases lacks dedicated hardware, such as *Digital Signal Processing – DSP*, and sufficient computational power to accommodate the computation of such features. Consequently, due to these limitations and according to similar evaluations made by domain experts, these advanced statistics have not been employed for the oxygen sensor and high-pressure fuel system use cases.

In the experiments conducted for the oxygen sensor use case,  $N_p = 9$  percentiles between the 10<sup>th</sup> and 90<sup>th</sup> were used. For the fuel high-pressure system, the average and standard deviation were not considered, while  $N_p = 11$  percentiles were computed. Nine percentiles were computed between the 10<sup>th</sup> and 90<sup>th</sup>, and the 1<sup>st</sup> and 99<sup>th</sup> percentiles were also included.

For a signal  $X$  from the overall set of measurements  $\mathcal{X}$  and its sample at time  $t$ ,  $x[t]$  the signal value is deemed to belong to the time window  $w(k)$  if the condition in Equation 2.1 is satisfied:

$$x[t] \in w(k) \rightarrow k \cdot \Delta T \leq t < (k + 1) \cdot \Delta T \quad (2.1)$$

The length of the time window, represented by  $\Delta T$ , is one tunable parameter that can be regulated. In the context presented, since the extracted features from the samples in the time window can be sent to a remote computational node, the value of the time window length has an impact on the required bandwidth. This parameter can be chosen suitably to maximize performance by using an iterative process. For increasing values of  $\Delta T$ , the method's performance is evaluated to find the best tradeoff.

For the Fuel high pressure system use case, the time window length was evaluated in the range of values from 60 to 300 seconds, with intervals of 10 seconds. For each value, the performance was computed on a validation set using F-measure, and it was observed that the values between 110 and 160 seconds exhibited the most balanced performance, as can be observed in Figure 2.17. Increasing  $\Delta T$  reduces the availability of samples in the training set, while smaller values can cause instability in the computation of percentiles. Therefore, a value of  $\Delta T = 120s$  was chosen as it provides a good tradeoff between these opposing factors.




Figure 2.17: The impact of time window length on the predictive performance of the proposed method.

The phenomenon of Oxygen sensor clogging exhibits extremely long dynamics, and using low values of  $\Delta T$  makes it impossible to capture the slow divergence of sensor response time. Moreover, initial data explorations

have indicated that extracting statistics from only a subset of the cycles can yield misleading patterns. Specifically, the variations in pedal manoeuvres and engine setpoints during the test exert diverse influences on the oxygen sensor, inducing distinct dynamics. These distinct dynamics introduce substantial variability in data from one time window to another, potentially leading to misclassification. For instance, the dynamic response slowdown induced by clogging phenomena is more noticeable during intense transients. Therefore, time windows collected during periods without such transients, such as urban low-speed drive, may incorrectly categorize a degraded sensor as *Green* or *Yellow*. Additionally, the occurrence of these known transients, which is known as *a priori* in controlled experimental settings, remains unpredictable in real-world scenarios where it depends solely on driver actions. Hence, in the attempt to cover the longest possible time window, we arbitrarily set  $\Delta T = 3450s$ , which is the entire length of the first part of the cycle.

By following the aforementioned process, the dimensionality was effectively reduced from millions to a maximum of  $\#(\mathcal{X}) \cdot (N_p + 2)$ , where  $\mathcal{X}$  represents the set of signals, but the information about the trend and dynamics of each time series was lost completely. In an attempt to recover some of this information, the derivative of each signal with respect to time was computed, and their statistics in the time window were calculated and added to the percentiles of the original signals. Finally, the adoption of this method results in  $2 \cdot \#(\mathcal{X}) \cdot (N_p + 2)$  features for each time window. For the fuel high-pressure system, this yields  $2 \cdot 43 \cdot 11 = 946$  features, while for the oxygen sensor, it produces  $2 \cdot 14 \cdot (9 + 2) = 308$  features.

Table 2.5 summarizes the data transformation parameters setting used to conduct the experiments.

Use case	$\Delta T$	$N_p$	Average	Standard deviation
Oxygen sensor	3450s	9	Yes	Yes
Fuel high pressure system	120s	11	No	No

Table 2.5: Data transformation settings used in the experiments

## Model

In the literature, numerous classification algorithms have been proposed, each with its strengths and weaknesses that can be adapted to different use cases. However, no algorithm has emerged as a superior choice in all

fields. Therefore, in the methodology proposed in this work, the model is selected case by case by an iterative data-driven procedure based on the obtained performance. In the use cases presented in this work, several popular algorithms have been tested, including:

- **Logistic Regression – LR** (Hastie et al. [2001]): The posterior probabilities of the outcomes of a dependent variable are estimated through linear functions on multiple independent variables. This is a linear model that has been adapted to the multiclass classification task of the presented use cases by adopting multinomial logistic regression together with regularization to balance the tradeoff between bias and variance. Given its capability to detect linear relationships between variables, this model’s performance is potentially limited when applied to complex systems, such as the ones presented in the introduced use cases. Both applications contain intricate mechanisms characterized by numerous non-linearities, such as pressure pump control, oxygen release, and flow dynamics during combustion. Such complexities cannot be effectively captured by models of this nature.
- **Decision Tree – DT** (Kingsford and Salzberg [2008]): This is a supervised ML algorithm used for classification tasks. The final decision is provided by visiting a tree-like model, in which each node represents a decision based on the comparison between the value of one feature and a threshold. Starting from the root node, the tree is traversed to arrive at one of the leaves, where the final classification decision is made. The training procedure defines the feature to be analyzed and the threshold to be applied for every node, trying to maximize the purity of the resulting set split. Due to its predefined set of static rules, *DTs* can be more prone to overfitting the training data, making them less robust than other methods. However, given a decision, it can always be explained by checking the decision path in the tree.
- **Random Forest – RF** (Breiman [2001]): Decisions from a set of different *DTs*, built using a different subset of features or constraints, are ensembled to provide a final, more robust classification. This is done to overcome the robustness weakness of *DT*.
- **XGBoost** (Chen and Guestrin [2016]): A gradient boosting technique (Friedman [2001]) is applied to *DT*. Through different boosting rounds, a *DT* is trained to iteratively improve its performance on previously misclassified training points.

- **Support Vector Machines – SVM** (Cortes and Vapnik [1995]): This is a supervised *ML* algorithm used for classification and regression tasks. It constructs a hyperplane or set of hyperplanes in a high-dimensional space to separate data points into different classes (Wikipedia contributors [2023e]) maximizing the margin. *SVM* achieves a high level of performance, but its decisions are not always explainable because it is a *black box* method.
- **Artificial neural networks – ANN** (Schmidhuber [2015]): This is a family of techniques inspired by biological neural systems. *ANNs* are based on a set of connected nodes called artificial neurons. Neurons are usually structured in connected layers divided into an input layer, one or more hidden layers, and an output layer. In particular, in the proposed pipeline, the *MultiLayer Perceptron – MLP* feed-forward *ANN* is exploited. While recent years have witnessed the introduction of more intricate and optimized network architectures, these advancements have leveraged specific data characteristics. For instance, *Convolutional Neural Networks – CNNs* excel at capturing relationships between closely located data points, *Transformers* prove highly proficient at modelling sequences of symbols and capturing long-range dependencies, and *Recurrent Neural Networks – RNNs* take into account past computations when generating current outputs. The decision to focus solely on *MLP* is justified by several factors, including its relatively lower complexity, compatibility with the nature and volume of the available data, and considerations related to embedded hardware and software constraints indicated by domain experts.

### 2.3.2 Method settings

Although the steps outlined in the proposed methodology are applicable to a wide range of applications and problems, when implemented for a specific use case, each step must be tailored to select the best techniques that maximize the performance and resource usage trade-off, as well as assign appropriate values to the tunable parameters that regulate system behaviour. In light of this, the following sections provide a detailed analysis of the procedures employed for *signal selection*, *feature selection*, and *model configuration*.



### Signal selection

The onboard *ECU* is capable of collecting and analyzing numerous signals, however, not all of them are relevant for predicting degradation status as defined in the use cases. The elimination of irrelevant signals in both use cases is carried out by consulting domain experts. They use their knowledge of the system and its degradation phenomena to remove signals that are not present in the real-world scenario and those that are not linked to the system's behaviour. Furthermore, it is possible that some signals do not carry any information since their values remain constant over time. The proposed methodology analyzes the signals' variations across different experiments, thus removing the ones that do not change. After performing this *a-priori* analysis, the number of signals has been reduced for both use cases. Specifically, for the oxygen sensor use case, the number of signals has been reduced from 50 to 31. On the other hand, for the fuel high-pressure system use case, the number of signals has been reduced from 614 to 285. Additionally, some signals may introduce redundancy in the information they represent by being slight elaborations of other signals. Therefore, it is important to implement methods for reducing the number of signals. While various algorithms have been proposed in the literature for this task, some approaches based on mathematical combinations of original signals, such as the *Principal Component Analysis - PCA*, may not yield interpretable results. The use of non-interpretable signals makes it impossible to take advantage of domain experts' knowledge. To address this, an algorithm based on the Pearson correlation coefficient, named *CORR-FS*, is defined as follows:

1. For each cycle  $k$ , the Pearson correlation coefficient between the signals  $i$  and  $j$  is computed and stored as  $r_{k,ij} = r_{k,ji} = \frac{\text{cov}(i,j)}{\sigma(i)\sigma(j)}$ .
2. For each pair of variables  $i$  and  $j$ , the overall correlation coefficient  $r_{ij} = r_{ji}$  is computed as the average correlation coefficient for that pair of variables over all cycles, as defined in Equation 2.2.

$$r_{ij} = \frac{1}{n} \sum_{k=1}^n r_{k,ij} \quad (2.2)$$

3. The list of *remaining variables*  $\mathcal{L}$  is initialized with all available variables.

4. For each variable  $i \in \mathcal{L}$ , the sum of squared correlation coefficients  $s_i$  is computed as defined in Equation 2.3.

$$s_i = \sum_{j \in \mathcal{L}} r_{ij}^2 \quad (2.3)$$

5. The variable

$$b = \operatorname{argmax}_{i \in \mathcal{L}} s_i \quad (2.4)$$

is extracted from  $\mathcal{L}$  as the most representative of the remaining variables.

6. All variables  $v \in \mathcal{L}$  such that  $r_{vb} > r_{min}$  are extracted from  $\mathcal{L}$ , in that they are well represented by  $b$ .
7. If  $\mathcal{L}$  is empty, the algorithm terminates, otherwise it continues with Step 4.

The algorithm produces, as a result, a set of signals  $x \in g$  and its representative signal  $x_g$  for each group  $g \in G$ . The choice of the representative signal of each group depends on the specific use case. One strategy is to adopt Equation 2.4 to select the most representative signal. Another approach consists of analyzing the signals in the group with the support of domain experts. In both of the use cases described, the latter strategy has been adopted.

The proposed algorithm requires the tuning of only one parameter,  $r_{min}$ , which regulates the minimum threshold for correlation coefficient to consider a pair of variables strongly correlated. Its value is defined empirically by adopting an iterative procedure that evaluates how many signals are selected for increasing values of  $r_{min}$  and selects the value to be used by the elbow method.

After applying the proposed approach to the oxygen sensor use case, by testing the  $r_{min}$  values within the range between 0 and 1, as illustrated in Figure 2.18, the optimal threshold is determined to be  $r_{min} = 0.8$  resulting in the selection of 14 signals. Similarly, for the fuel high-pressure system, by analyzing the  $r_{min}$  values within the range between 0.71 and 1, as illustrated in Figure 2.19, the best threshold is determined to be  $r_{min} = 0.95$  resulting in the selection of 43 signals.

The summary of signal selection parameter setting and their results in the presented use cases is provided in Table 2.6.

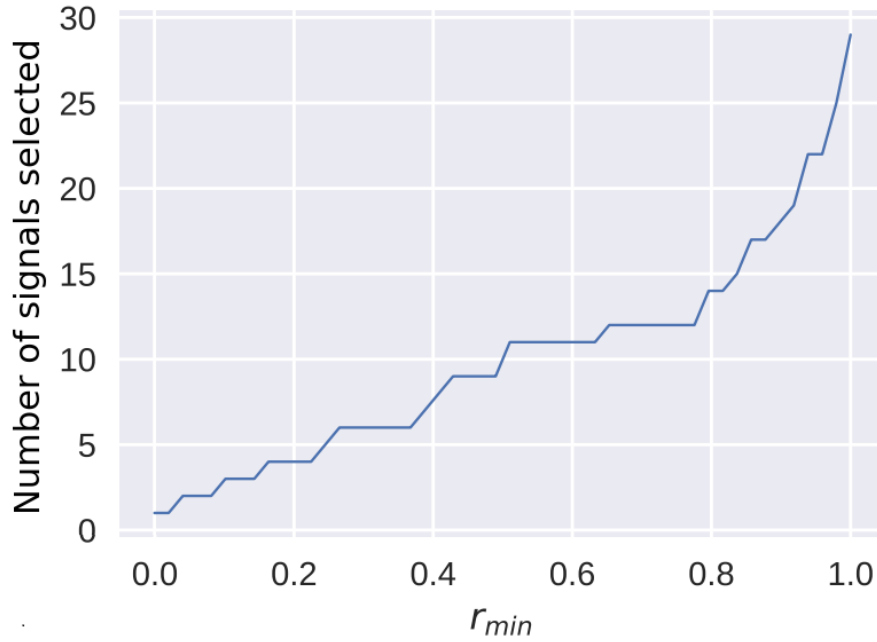


Figure 2.18: The number of selected signals varying the value of the parameter  $r_{min}$  for the oxygen sensor use case




Figure 2.19: The number of selected signals varying the value of the parameter  $r_{min}$  for the fuel high pressure use case

### Feature selection

Reducing the number of features to be utilized by the classification model represents a crucial step in the *ML* pipeline. Generally, using fewer features simplifies the model's structure (James et al. [2013]) and reduces its training time (Liu [2010]), as well as the amount of data required. Additionally, it helps in avoiding the curse of dimensionality (Kramer [1991]). The fundamental concept behind *feature selection* is that data can contain some

Use case	$r_{min}$	Signals selected		
		Original	A-priori analysis	CORR-FS
Oxygen sensor	0.8	50	31	14
Fuel high pressure system	0.95	614	285	43

Table 2.6: Signal selection settings used in the experiments

features that are redundant or irrelevant to the task under examination (Kratsios and Hyndman [2021]). Specifically, for the use cases presented, reducing the number of features has an impact on the necessary communication bandwidth and memory storage.

In the case of the oxygen sensor application, this procedural step has been omitted due to time constraints associated with the project and an evaluation by domain experts who deemed the resulting set of 308 features to be suitable within the described context of automotive prognostic use cases. Subsequent sections will expound upon the evaluation of the achieved results, revealing that certain features exhibit superior predictive performance compared to others. This observation suggests that the application of this step to the oxygen sensor use case could potentially yield substantial benefits. Conversely, in the fuel high-pressure system use case, feature selection is performed by evaluating the performance of different feature subsets using a classification algorithm (Blum and Langley [1997]). This approach directly demonstrates the predictive performance of the feature subset and highlights the combined prediction capabilities of the selected features. However, it requires an exhaustive search of all the possible feature subsets, which is infeasible. Therefore, the proposed methodology employs a sub-optimal heuristic approach to reduce the computational complexity of the task. Specifically, dedicated ranking algorithms are used to rank the features from the most relevant to the least, and different feature subsets  $S_{R_i}(j)$  are created by progressively adding the variables according to the obtained ranking  $R_i$ , as in Equation 2.5.

$$S_{R_i}(j) = \bigcup_{k=1}^j R_i(k) \quad \forall j \in (1, \dots, |R_i|) \quad (2.5)$$

Finally, classification performance is evaluated to determine the best combination of features by using a learning curve (Sammut and Webb [2011]), for each ranking  $R_i$ .

Two ranking algorithms have been investigated:

- **Minimum Redundancy Maximum Relevance – mRMR** (Ding

and Peng [2005]): the algorithm ranks the features by measuring the *Mutual Information Difference – MID* metric, which combines the importance of each feature (measured as the correlation with the target class) with the redundancy that the feature would introduce (measured respect to the other features).

- **Feature importance for the random forest** (Breiman [2001]): a *RF* classifier is trained on the target task, and the algorithm ranks the features based on the *Feature Importance – FI*, which describes how much each feature contributes to the classification process.

The training of a *RF* classifier requires tuning of some hyperparameters. Therefore, two different settings are adopted: the ones suggested in Genuer et al. [2008] and the ones obtained via coarse grid search (i.e., optimizing the F-measure performance on the Validation set). This approach results in three different feature rankings, one from *mRMR* and two from *RF*s, named, respectively, *mRMR*, *RF* and *RF-Optimized*.

In Figure 2.20, the obtained feature rankings are compared based on the change in the normalized feature importance across ranking position. It is observed that the rankings differ in terms of trend and results. The rankings based on the *RF* method tend to assign negligible importance to a larger number of features compared to the *mRMR* ranking, with the *RF* ranking exhibiting the sharpest trend. On the other hand, the *mRMR* ranking shows a different trend, with only a few features having a high score that linearly decreases in importance. The difference in trends arises from the different ways the algorithms calculate the feature importance, their distinct objectives and assessment criteria. *MRMR* prioritizes features that demonstrate a balance between informativeness (relevance) and redundancy within the dataset, aiming to select the most distinctive and non-repetitive attributes. Consequently, *MRMR* tends to produce a more linear ranking of feature importance because it is focused on optimizing the diversity of selected features. Since *RF* prioritizes features based on their direct impact on predictive performance, it may result in a ranking that exhibits more intricate or nonlinear patterns. This is because *RF* may identify complex interactions and dependencies among features, which are not necessarily linear, but crucial for accurate predictions.

To evaluate the different feature subsets, an *SVM* model is trained and tested for each of the resulting combinations. The adoption of such an algorithm represents a good trade-off between the ease of hyperparameters




Figure 2.20: Comparison of the three obtained rankings by mean of min-max normalized feature importance

tuning and the classification performance. To deal with the impact of hyperparameters on the classification performance, a lightweight grid search of 400 hyperparameter configurations is conducted, creating 400 different *SVM* models. For each ranking, 100 subsets are evaluated, resulting in a total of 120,000 different configurations. Each of these 120,000 configurations is tested on the Validation set to evaluate the performance in terms of F-measure on the *Red* class, keeping, for each feature selection, only the best performer among the 400 possible *SVM* models. Figure 2.21 illustrates the obtained results for both the Training and Validation sets. Analyzing the learning curves, *RF*-based rankings demand fewer features compared to the one based on *mRMR*. Moreover, the *RF-Optimized* ranking produced a more consistent performance, particularly between 15 and 25 features. After the first 25 features, the F-measure in Validation and Training sets diverges, suggesting overfitting on the training data. For this reason, the first 25 features from the *RF-Optimized* ranking are selected as input for the classification model. Looking back at the signals from which these 25 features are calculated, only 6 signals are used, making it possible to further reduce the number of signals monitored by the *ECU* from 43 to just 6.

### Model Training, Tuning and Selection

As there is no universal solution for classification tasks in *ML*, the best option must be selected among the available techniques, depending on the use case constraints and the performance achieved. Each method requires some hyperparameters to be tuned to regulate its behaviour and achieve optimal

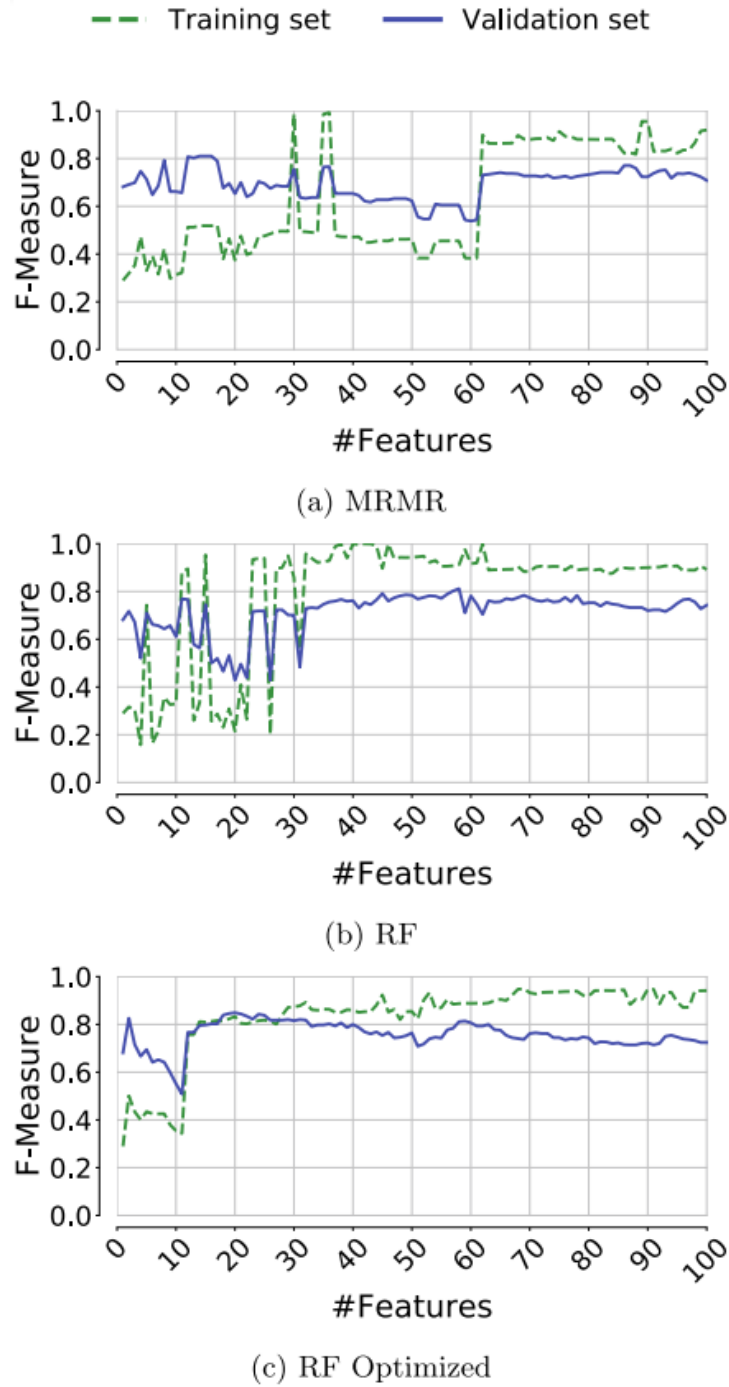


Figure 2.21: Performance of the best *SVM* model configuration by using a progressively increasing number of features with the different rankings

performance. Therefore, the proposed methodology performs ad-hoc hyperparameter tuning by grid search for each of the classification techniques

mentioned above.

For the oxygen sensor use case, *DT*, *SVM*, and *MLP* models are considered, using 10-fold cross-validation and an exhaustive grid search. Each configuration is evaluated in terms of F1-score for the *Red* class. The best performance of *DT* is obtained with a maximum depth of 6 and a minimum of 4 samples for each split. The best *SVM* configuration uses a linear kernel and  $C$  approximately 2.5. The model with two hidden layers of 500 and 1000 nodes is the best option for *MLP*.

For the fuel high-pressure system use case, the model selection is carried out in two phases.

Firstly, a grid search process, summarized in Table 2.7, is run to select the hyperparameters and observe which classifiers meet the required minimum precision and recall thresholds, indicated by the domain experts. Given the infrequent nature of system failures and the substantial consequences of unnecessary maintenance interventions on customer satisfaction and operational costs, the precision of forecasting the *Red* class assumes greater significance than its recall. However, an extremely low recall renders the tool ineffective in achieving its primary predictive maintenance objective. Specifically, experts in the domain have prescribed that solutions capable of recognizing over 50% of deteriorated instances with fewer than 30% false alarms are deemed acceptable, aligning with both customer satisfaction and business model considerations, thus facilitating efficient triggering of maintenance interventions. The examined business model consider the resulting increase or saving on costs of operations alongside the estimated impact on brand reputation and sales. Consequently, the precision of the *Red* class must surpass 0.7, while its recall should stand at no less than 0.5. All the *ML* algorithms that fail to meet these minimum requirements in any configuration are discarded. Figure 2.22 shows the results of this step. None of the configurations tested for *LR* and *RF* can meet the requirements, and the two algorithms are discarded. This result was expected, given the two methodologies' limitations discussed during their introduction. Poor performance of *LR* can be explained by the non-linearities of the system that cannot be captured effectively by such a linear model. Only a few configurations of *XGBoost* exceed the indicated thresholds, which is insufficient to consider this method robust enough to be kept.

Next, the best model for each classifier among those that meet the requirements, *SVM* and *MLP*, is selected as the candidate model, to finally choose the final model. Figure 2.23 displays the parameter combinations  $C$  and  $\gamma$



that produce *SVM* models capable of fulfilling the minimum performance thresholds set by domain experts. Concerning the *MLP*, a random seed parameter is utilized to initialize the neurons' weights and biases. For a stable algorithm, this parameter is not expected to significantly affect the performance. Nevertheless, investigating this aspect by performing 100 experiments and varying only this value, the precision and recall of the *Red* class fluctuated, and they exceeded the minimum value recommended by domain experts only in a few cases, especially for the precision metric, as shown in Figure 2.24. These outcomes suggest that *MLP* tuning is not stable, and hence, it is rejected.

Classifier	Parameter	Values
Logistic regression	Solver	newton-cg, lbfgs, sag, saga
	C	$[10^{-3}, 10^3]$ step 50 in a log scale
	penalty	l2
	multi class	multinomial
	max iter	[100, 500, 1000]
	class weight	[balanced, None]
Random Forest	Impurity Decrease	[0, 0.02] step 0.005
	Min samples leaf	[5, 35] step 5
	Estimators	10, 15, 20, 30, 50, 100, 150, 200, 250, 500
	Split Criterion	entropy
	Max features	auto, log2, None, 0.5
XGBoost	# of boosting rounds	1000
	Maximum tree depth	2, 5, 7, 10
	Learning rate	0.01, 0.05, 0.25, 0.5
	Minimum child weight	1, 0.01, 0.05, 0.25, 0.5
	$\gamma$	0, 0.5, 1, 5, 10
	subsample ratio	0.1, 0.3, 0.5, 0.7, 0.9, 1.0
	subsample ratio of columns	0.1, 0.3, 0.5, 0.7, 0.9, 1.0
SVM	Kernel	rbf
	C	$[10^{-3}, 10^3]$ step 100 in a log scale
	$\gamma$	$[10^{-3}, 10^3]$ step 100 in a log scale
MLP	1st Layer	[25, 225] step 25
	2nd Layer	[25, 225] step 25
	Activation	Logistic, tanh
	Seed	100 random values
	Solver	Adam
	Tolerance	$10^{-4}$

Table 2.7: Hyperparameter settings for each evaluated model




Figure 2.22: Performance of the different algorithms during the grid search by means of precision and recall for *Red* class. The red solid lines indicate the minimum performance thresholds indicated by domain experts. Good performer algorithms are expected to produce results in the upper right corner of the graphs.

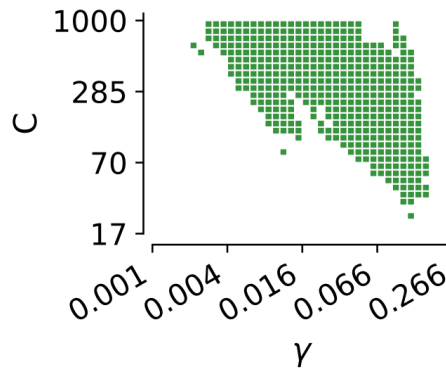


Figure 2.23: All the combinations of parameters  $C$  and  $\gamma$  that produce *SVM* models able to meet the minimum performance thresholds indicated by domain experts

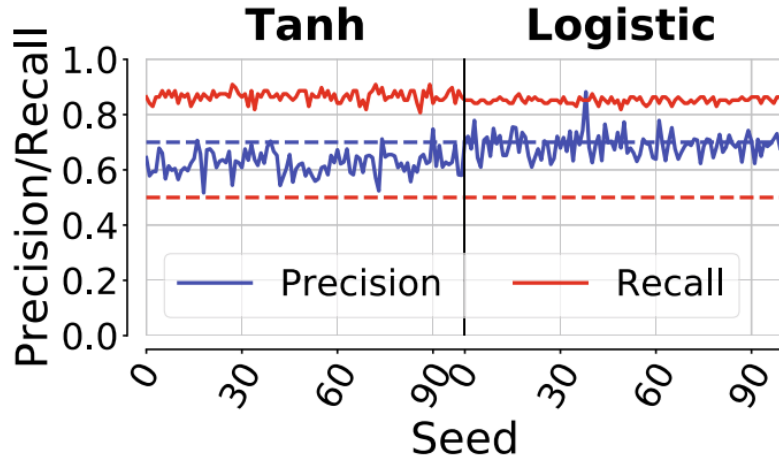


Figure 2.24: Impact of different initial random seed values on precision and recall for *Red* class, for the two different activation functions

## 2.4 Results

The methodology outlined in the preceding section has been implemented for the two use cases presented. The subsequent sections introduce the metrics employed to measure performance, along with reference values for assessing the method’s compliance with requirements. Finally, performance is demonstrated via both metric values and visualizations.

### 2.4.1 Metrics and requirements

The absence of previous research on both oxygen sensor and fuel high pressure system use cases in predictive maintenance hinders the selection of suitable benchmarks to assess the proposed method’s performance. Therefore, typical classification metrics are adopted and estimated to understand the algorithm’s behaviour.

In the domain of multiclass classification, a suite of metrics has been proposed to evaluate the efficiency and effectiveness of predictive models. The classification accuracy metric quantifies the ratio of accurately classified instances to the total, providing an overall assessment of the model’s predictive capacity, irrespective of individual classes. For each specific class, the precision metric quantifies the fraction of correctly predicted instances within those that have been predicted as belonging to that particular class. This metric highlights the model’s accuracy in classifying samples of that

specific category. Conversely, recall, often denoted as sensitivity or true positive rate, measures the equilibrium between accurately identified instances of that class and the total instances linked to that class label. This metric evaluates the model’s aptitude in correctly recognizing all instances of that class within the dataset. The F1-score, stemming from the harmonic mean of precision and recall, demonstrates the equilibrium between these two metrics and is particularly potent in scenarios characterized by class imbalances. It provides a consolidated evaluation of the model’s performance. Given a specific class of interest denoted as  $k$ , the precision  $Precision_k$  can be calculated as:

$$Precision_k = \frac{TP_k}{TP_k + FP_k} \quad (2.6)$$

the recall  $Recall_k$  as:

$$Recall_k = \frac{TP_k}{TP_k + FN_k} \quad (2.7)$$

and the F1-score  $F1_k$  as:

$$F1_k = 2 \cdot \frac{Precision_k \cdot Recall_k}{Precision_k + Recall_k} \quad (2.8)$$

Where  $TP_k$  represents the number of samples from the  $k$  class correctly identified,  $FP_k$  is the number of samples wrongly associated with the  $k$  class and  $FN_k$  the number of samples from the  $k$  class not identified by the method.

Although classification accuracy is generally adopted in the literature, it is not very representative in the case of unbalanced datasets, as it is for the two introduced use cases. Domain experts have identified the *Red* class as the most relevant one when considering real-world applications. This class indicates the system stage before it fails, and the diagnostic running on the local *ECU* starts to flag the damage, asking for immediate intervention. Correctly predicting the *Red* class brings significant benefits, enabling preventive maintenance and optimal aftermarket service operations.

To delve into the rationale behind this determination, consider the situation where the system erroneously predicts an instance as *Red*. In the context of the automotive prognostic scenario discussed earlier, this will prompt users to perform maintenance on their vehicles and trigger preparatory measures within the service organization, like procuring replacement parts in advance. The repercussions of such unnecessary activities are twofold: customers experience discontent due to disruptions caused by having their cars in the workshop, while the company incurs financial losses from the operation.

Furthermore, the premature replacement of functional components leads to increased wastage, contributing to environmental concerns.

When a *Red* case goes undetected, the system degradation will progress and its failure will eventually be identified by onboard diagnostics, resulting also in this case in customer dissatisfaction and the implementation of necessary maintenance procedures without the benefit of anticipation. Unlike other contexts such as aerospace or nuclear facilities, the failures pertinent to the cases presented here lack safety implications that necessitate absolute prevention. This latter scenario aligns with current practices, accepted by the customer community. Hence, the integration of *Predictive Maintenance* in the automotive sector mandates meticulous evaluation, considering the trade-off between the benefits coming from the anticipation and the risk of unnecessary operations, with the precise prediction of the class that prompts user notification, specifically, the *Red* class, being of the highest importance. Therefore, precision and recall for this class have been selected as performance metrics. Precision is computed as the ratio between the correct *Red* class predictions ( $TP_{Red}$ ) and the total ones, while recall is computed as the ratio between the correct *Red* class predictions and the total number of actual *Red* class samples, as shown in the following equations:

$$Precision_{Red} = \frac{TP_{Red}}{TP_{Red} + FP_{Red}} \quad (2.9)$$

$$Recall_{Red} = \frac{TP_{Red}}{TP_{Red} + FN_{Red}} \quad (2.10)$$

Here,  $FP_{Red}$  indicates the number of false alarms, and  $FN_{Red}$  indicates the number of positive cases missed by the algorithm.

Domain experts, considering the impact on customers and maintenance operations, have indicated that, for both oxygen sensor and fuel high-pressure system use cases, the minimum acceptable level of performance is reached if more than half of *Red* cases are correctly detected and if more than 70% of alarms raised by the method are correct. This translates into the requirement of obtaining precision  $> 70\%$  and recall  $> 50\%$ . Moreover, to provide a synthetic indication of both metrics, the F1-score is computed as:

$$F1score_{Red} = 2 \cdot \frac{Precision_{Red} \cdot Recall_{Red}}{Precision_{Red} + Recall_{Red}} \quad (2.11)$$

## 2.4.2 Performance

The optimal performance for both use cases is consolidated in Table 2.8. Notably, the requirements have been successfully met in both instances.

Use case	$Precision_{Red}$	$Recal_{Red}$
Oxygen sensor	0.902	0.902
Fuel high pressure system	0.790	0.742
<b>Requirement</b>	0.700	0.500

Table 2.8: Best performance obtained for both the use cases

By applying the proposed methodology to the Fuel High-pressure system use case, the requirements for both  $Precision_{Red}$  and  $Recall_{Red}$  metrics are successfully met by the best configuration, achieving values of 0.790 and 0.742, respectively, when evaluated on the Test set. The confusion matrices for the Validation and Test sets are depicted in Figure 2.25 and Figure 2.26. It can be observed that only a few instances classified as *Green* or *Yellow* are misclassified as *Red*. This result is of great importance considering the goal of minimizing unnecessary interventions on vehicles, as indicated by domain experts. The consistent performance across both the Validation and Test sets suggests the generalizability of the obtained classifier.

Actual Label	Predicted Label		
	G	Y	R
G	286	35	9
Y	36	111	7
R	9	4	75

Figure 2.25: Confusion matrix for Fuel high pressure system Validation set.  $Precision_{Red}$  : 0.824 and  $Recall_{Red}$  : 0.852

To enhance the visualization of algorithm results and analyze instances where the prediction of deterioration level fails, a visualization technique




Figure 2.26: Confusion matrix for Fuel high pressure system Test set.  $Precision_{Red} : 0.790$  and  $Recall_{Red} : 0.742$

called the *Mismatch matrix* has been introduced. This matrix represents time windows, with each cell corresponding to a specific one, and each column representing the cycle from which they are extracted. The columns are grouped based on the cycle label, with *Red* cycles on the left, *Yellow* cycles in the middle, and *Green* cycles on the right. Solid vertical lines are used to separate the different classes. The cells are vertically stacked according to the sequence of their associated time windows, so the higher a cell is placed in the column, the later it occurs within the associated cycle. Due to variations in cycle lengths, some cells may be empty for certain columns, and they are coloured in grey. A blank cell indicates the correct classification of the time window, while the colour of a cell indicates the wrongly assigned class when the model fails to attribute the correct class. For example, if a cell appears in yellow within a *Green* cycle, it means that the associated time window has been incorrectly classified as *Yellow* instead of *Green*.

The *Mismatch matrices* for the Validation and Test sets are presented in Figure 2.27 and Figure 2.28, respectively. It is worth noting that in some cases, the engine behaviour differs from the assigned label throughout the entire cycle. One such instance is experiment 8 in the Validation set, where the majority of cases are predicted as *Green*, while the labels indicate *Yellow* for the system conditions. Focusing specifically on the *Red* experiments, only a few time windows have been misclassified, with one critical case

being experiment 1 in the Test set, where 17 out of 33 time windows were not correctly associated with the *Red* class. The other experiments also have some time windows classified as *Yellow* or *Green*, indicating that a malfunctioning Fuel High-pressure system can exhibit normal behaviour during certain parts of the cycle. To address this aspect, it is important to establish a decision-making process that considers multiple samples rather than relying solely on the result of a single time window. Several options for handling this aspect will be discussed in the next section.




Figure 2.27: Mismatch matrix for Fuel high pressure system Validation set. Experiment 8 exhibits many wrongly detected time windows

The performance achieved by the proposed methodology in the Oxygen sensor use case meets the required level across multiple configurations. The metrics corresponding to the optimal configurations of *DT*, *SVM*, and *ANN* are summarized in Table 2.9, Table 2.11, and Table 2.10, respectively.

The *DT* model exhibits the poorest performance, as expected due to its simplicity. However, this drawback is compensated by the interpretability of its internal representation. To gain further insights into the results, confusion matrices are presented in Figure 2.29 for all the different models. The majority of the results distribute along the diagonals of these matrices,






Figure 2.28: Mismatch matrix for Fuel high pressure system Test set.

	<i>Green</i>	<i>Yellow</i>	<i>Red</i>
Precision	0.8365	0.7529	0.8475
Recall	0.8110	0.7853	0.8197
$F_1$ score	0.8235	0.7688	0.8333
Accuracy	0.8015		

Table 2.9: *Decision Tree* best configuration results

aligning with the high values obtained for the performance metrics. Interestingly, most classification errors occur between *adjacent* classes, indicating that the models have learned the sequential progression of the degradation process. In only one instance, the *SVM* model mistakenly identified a *Green* cycle as *Red* (two classes ahead). This observation suggests that the models perceive *Green* cycles to be more similar to *Yellow* cycles rather than *Red* cycles.

	<i>Green</i>	<i>Yellow</i>	<i>Red</i>
Precision	0.8869	0.8679	0.9016
Recall	0.9085	0.8466	0.9016
$F_1$ score	0.8916	0.8571	0.9016
Accuracy	0.8814		

Table 2.10: *Artificial Neural Network* best configuration results

	<i>Green</i>	<i>Yellow</i>	<i>Red</i>
Precision	0.9026	0.8303	0.8261
Recall	0.8476	0.8405	0.9344
$F_1$ score	0.8742	0.8354	0.8769
Accuracy	0.8582		

Table 2.11: *Support Vector Machine* best configuration results

## 2.5 Discussion

This section provides a detailed analysis of the results obtained from the described use cases. It presents various findings that offer deeper insights into the data, the system’s deterioration process, and the resulting model. Furthermore, practical aspects of applying the proposed methodology to the predictive maintenance context are examined, including the required amount of training data, bandwidth considerations, memory requirements, and the potential for distributing the algorithm between local and remote nodes. Lastly, the crucial aspects of the decision-making process for initiating maintenance procedures are discussed.

A prevailing commonality observed in both of these use cases pertains to the issue of data imbalance. In the high-pressure fuel system scenario, this imbalance predominantly manifests within the test dataset. Conversely, in the context of the oxygen sensor application, the imbalance is more pronounced, with the Red class containing only 15.7% of the entire sample set. These imbalances, characterized by an uneven distribution of classes, possess the capacity to introduce biases into the machine learning models deployed, thereby detrimentally affecting their predictive precision and ability to generalize. As a proactive measure aimed at mitigating the adverse ramifications of this data imbalance, the literature has proposed an array of techniques, including resampling, synthetic data generation, and the adoption of class-weighted learning strategies. The potential application of such methodologies to the presented use cases represents a future step for the presented research. Furthermore, given that time constraints




Figure 2.29: Confusion matrix of each model for Oxygen sensor Test set.

during the data acquisition phase have contributed to a portion of this data imbalance, an expanded experimental phase holds promise in providing a more natural equilibrium by augmenting the Red class data, thus rectifying the existing dissimilarities.

### 2.5.1 Oxygen sensor

#### Importance of interpretable signal selection

In the proposed methodology, the signal set is selected during the signal selection step, where the interpretable algorithm *FS-corr* is employed. This algorithm ensures interpretability in all its aspects, as it does not involve the combination of signals to create new ones or alter the original information content. Moreover, the parameter  $r_{min}$ , which represents the minimum correlation threshold for identifying redundant signal pairs, is also interpretable. These characteristics enable domain experts to make corrections to the selected signals by adding or removing them and effectively validate the correlation analysis and algorithm results based on their experience and system knowledge. As a result, this feature reduces the amount of data required for configuring the methodology, thereby facilitating its industrial deployment.

#### Feature importance

Among the various methods employed, the *DT* algorithm offers an interpretable model that presents a hierarchical structure of features visited to reach the final prediction. Of particular importance is the root node of the

$DT$ , which represents the most crucial feature for discriminating the level of degradation. Identifying this feature can be regarded as one of the key findings of this analysis.

In the best-performing  $DT$  configuration, the root node is associated with the 90<sup>th</sup> percentile of the oxygen value derivative. This outcome aligns with expectations since the clogging phenomena impact the dynamics of sensor readings, resulting in a slower response and consequently lower derivative values compared to *Green* class cycles. Additionally, the labelling process is based on the response time duration when the sensor is stimulated with a step-like impulse. While expected, this result is nontrivial as it demonstrates the ability to identify clogged sensor events without high-frequency signal sampling during the vehicle's typical long-term usage.

The distribution of the 90<sup>th</sup> percentile of the oxygen value is illustrated in Figure 2.30, clearly depicting how the three classes are separated based on this variable. Specifically, the *Green* and *Red* classes exhibit distinct distributions (the value 0.8 can be identified as threshold), with *Yellow* samples dispersed between the other two classes. This result implies that the application of the feature selection step to oxygen sensor use case may wield a positive influence on both performance and the necessary number of features.

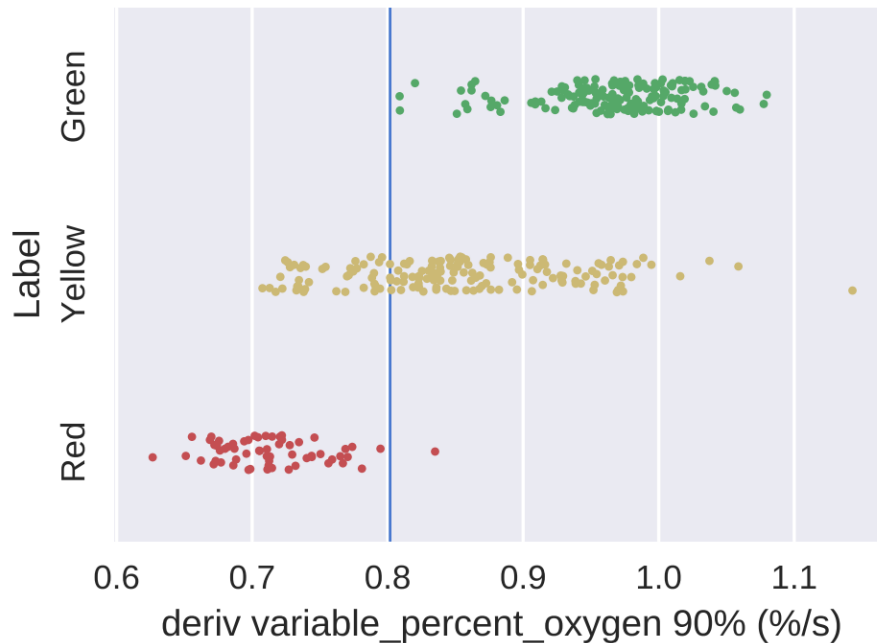


Figure 2.30: Distribution of the 90<sup>th</sup> percentile of the derivative of the oxygen variable, by class.

## 2.5.2 Fuel high pressure system

### Sensitivity to training size

When developing a data-driven solution for real-world applications, the process of data collection assumes critical significance as it influences all subsequent steps. In the context of the Fuel High-Pressure system, data collection necessitates conducting multiple experiments at an engine test bench to record system signals at varying levels of deterioration. The associated costs of this activity escalate linearly with the number of cycles to be conducted. Therefore, it is crucial to consider approaches that can achieve satisfactory performance with a relatively smaller number of samples. To assess this aspect, a learning curve was constructed to evaluate the performance of the *SVM* model when trained with an increasing amount of data. The learning curve methodology involved incrementally expanding the training set size by one unit and selecting 100 different subsets of that size of samples as the training set. Performance was then computed on both the Training and Validation sets. Intuitively, an increase in the training set size should lead to improved model performance. However, it is important to monitor the concurrent evolution of metrics in both the Training and Validation sets to observe if overfitting occurs and to assess the generalization capability of the model.

The resulting visualization illustrates the  $Precision_{Red}$  metric for increasing training set sizes and is depicted in Figure 2.31. The overlaid heatmaps provide insights into the distribution of the metric across the 100 different subsets tested. The intensity of the red colour indicates the number of samples present. The black curves represent the average performance for the Training (dashed line) and Validation (solid line) sets. The requirement of  $Precision_{Red} > 0.7$  is reported as a reference with a red horizontal solid line. It is evident that the  $Precision_{Red}$  exhibits a decreasing trend when tested on the training set, whereas it shows an increasing trend when tested on the validation set. These trends converge to similar values when the entire dataset is utilized. Additionally, as the training set size increases, the spread of the performance metric decreases. This suggests that it is possible to develop a reasonably generalized model by employing a suitable variety of data, as demonstrated in this particular use case.




Figure 2.31: Training performance for increasing training set size

### Bandwidth requirements

In scenarios where the available onboard hardware cannot execute the deterioration detection algorithm, the structural characteristics of the algorithm enable the allocation of the predictive component to a more powerful remote computational unit, such as a cloud computing platform. This division necessitates the transfer of acquired data from the local *ECU* to the remote unit. However, the deployment of the proposed methodology in real-world applications requires careful consideration of the required bandwidth due to connectivity limitations and the associated data transfer costs.

To quantify the bandwidth requirements, the number of samples that need to be collected and transferred by the *Electronic Control Unit* per second has been computed in different scenarios: (i) all signals, which amounts to 614 signals; (ii) the subset of signals after the a-priori signal selection, resulting in 285 signals; (iii) the subset of signals remaining after the complete signal selection process, yielding 43 signals; and (iv) the set of computed features after feature selection, comprising 25 features. As the efficient compression of the *ECU* data is out of scope for the present work, raw data samples delivery to the server is assumed over some robust network channel (e.g., GSM, HSDPA, LTE). In all cases, the bandwidth is estimated based on each sample being encoded as a 4-byte floating-point number. For scenario (iv), the feature values are transmitted once every 120 seconds.

The resulting visualization is presented in Figure 2.32, depicting the number of bits per second required for transmission over one hour. As expected, when all signals are sent to the cloud, the *ECU* would need to transfer the highest amount of information, over 1 Mbps, a solution deemed infeasible by

domain experts. Intuitively, reducing the number of signals through selection methods reduces this amount by half in scenario (ii) and to 110 kbps in scenario (iii). While the latter value is more aligned with application constraints, it still entails significant costs for onboard and cloud connectivity. Finally, when only the selected features are transmitted, the required bandwidth drops to 100 bytes every 120 seconds, a much more affordable value for the automotive scenario with unstable connections.




Figure 2.32: Bandwidth required to transmit signals and features to the cloud

### Memory demand

The proposed methodology incorporates the computation of percentiles for the selected signals by the onboard *ECU*. Traditionally, calculating these values requires storing the samples of each signal in memory. In the experimental setup used to achieve the presented results, the feature selection stage limited the number of signals to be processed to only 6. While this is a small number, storing all the samples for an entire 120 seconds of the time window can have a significant impact when the available hardware has limited memory capacity. Assuming each sample is encoded as a 4-byte floating-point number and considering a sampling rate of 6.25 ms per signal, this results in 20,000 samples per signal and 120,000 samples in total within the time window, equivalent to 480 kB of memory.

To overcome this limitation, several alternative algorithms have been proposed in the literature. [Jain and Chlamtac \[1985\]](#) introduced a heuristic algorithm for estimating percentiles on streaming data, eliminating the need to store all values. [Greenwald and Khanna \[2001\]](#) presented a method for accurately computing percentile estimates with an occupancy of  $O(\frac{1}{\epsilon} \log(\epsilon N))$  in memory, where  $\epsilon N$  represents the precision and  $N$  denotes the number

of samples. Cormode and Vesely [2020] further demonstrated the tightness of these results. Depending on the specific application, these percentile calculation methods can be adopted to reduce the required memory, thus enabling the implementation of the ECU.

### Decision-making strategy

Despite the provided methodology offering deterioration evaluation of the High pressure fuel system for each time window, the car manufacturer’s ultimate objective is to accurately recall vehicles for service interventions. To align the proposed algorithm’s target with the carmaker’s goal, it is necessary to establish a decision-making policy for car recalls. This becomes especially important upon examining the obtained results, as the predictions are not always precise. Relying solely on a single time window to trigger maintenance actions can result in significant customer dissatisfaction and high service costs.

To address this aspect, a preliminary analysis was conducted to evaluate the outcomes of employing a voting strategy on a sequence of consecutive time windows. Domain experts considered the results satisfactory when applying a majority voting approach to the time windows within a one-hour interval. This strategy accurately identified and recalled only the *Red* cars for service in both the validation and test sets.

Due to the limited number of available experiments, it has not been possible to fine-tune this approach and explore more complex solutions. To overcome this limitation in the real-world application, it is possible to release an initial silent version of the method that does not trigger any service actions. During this initial phase, data is collected and stored in the cloud from the vehicles, along with relevant information about the status of deterioration for the High pressure fuel system acquired during periodic interventions of Preventive maintenance. These data can then be analyzed to design the optimal decision strategy based on a data-driven calibration process. Once the policy is defined, a hard version can be released that actively triggers service operations based on model predictions of High pressure fuel system deterioration.



## Chapter 3

# Medical imaging

Medical imaging is a field of medicine that concerns itself with the methodologies and procedures employed for visualizing the internal structures and functions of the body ([Wikipedia contributors \[2023c\]](#), [FDA \[2018\]](#)). In medicine, these images are analyzed by physicians to extract valuable information for the diagnosis, treatment, and monitoring of various diseases. The most commonly used medical imaging techniques include X-rays – *XR*, computed tomography – *CT*, magnetic resonance imaging – *MRI*, ultrasound – *ULS*, and positron emission tomography – *PET* ([WHO \[2023c\]](#)).

The basic principle behind the *XR* and *CT* is the use of ionizing radiation ([Envision Radiology \[2023\]](#)). A beam of radiation is passed through the body and it is absorbed differently by different tissues, depending on the density and material. A detector, placed on the other side of the scanned body, receives the residual fraction of the beam and, by measuring its energy, can estimate the density of traversed tissues. In *XR*, the beam is directed to reconstruct a two-dimensional projection of the internal structures that can be visualized as an image. Some examples of X-rays images are illustrated in Figure 3.1. *CT* scans, on the other hand, rotate X-ray beam around the body to produce cross-sectional images, or *slices*, which are then processed to create a 3D image. Examples of slice pictures and 3D reconstruction are illustrated in Figure 3.2. Compared with *XR*, *CT* produces more detailed representations of internal organs, blood vessels, soft tissues, and bones ([NIH \[2022\]](#)). The usage of ionizing radiation must be considered when using these modalities because the accumulation of high doses can be harmful ([CDC \[2021\]](#)).

Differently from *XR* and *CT* that use ionizing radiations, the *MRI* uses strong magnetic fields and radio waves to produce detailed reconstructions



Figure 3.1: Examples of X-radiographs.

of the examined body (NIH [2023]). This makes this modality a safer option for patients. The strong magnetic field causes hydrogen atoms to align in a certain way, and then radio waves make them emit a signal that is captured by the *MRI* machine and elaborated to create the final picture (Radiological Society of North America [2022]), as the one in Figure 3.3. Also, this modality can be used to produce three-dimensional representations including soft tissues, organs and bones. *MRI* is a very expensive and time-consuming investigation compared to other imaging methods, as *XR* and *CT* (myVMC [2018]). Because of its strong magnetic field, this modality cannot be used with patients with medical devices such as pacemakers.

With *ULS*, high-frequency sound waves are used to scan the body (NIH [2016]). An emitter probe called *transducer* produces the sound waves and detects their echo to produce the image. *ULS* is a safer alternative for the other modalities because it does not use ionizing radiation or strong magnetic fields (Mayo clinic [2022]), but it has some limitations that can affect its diagnostic accuracy, such as the inability to penetrate through bones or air, the limited penetration in adipose tissue (RF Wireless World 2012 [2012]). Moreover, because the *transducer* is manipulated by an operator, *ULS* is operator-dependent, meaning that the quality of the images can vary depending on the skills and experience of the technician performing




Figure 3.2: Abdomen CT scan slices and its 3D reconstruction (Medical news today [2018])

the exam (Herment et al. [1987]). Figure 3.4 shows an example of reconstruction from *ULS*.

*PET* is a type of nuclear medicine procedure that measures the metabolic activity of the cells of body tissues detecting the concentration of a radioactive substance, radiopharmaceutical, that is injected in the body and is adsorbed by the target tissue or organ (The Johns Hopkins University [2023]). This substance emits positrons that are detected by the *PET* machine and used to reconstruct the image. One example is presented in Figure 3.5. The differences in the absorption of radiopharmaceuticals provide information about the cellular changes in organs and tissues earlier than *CT* and *MRI* scans (Cleveland Clinic [2022]). Because of the injection of a radioactive tracer, which can be harmful in high doses, *PET* is not suitable for patients with certain diseases, such as kidney ones, and is much more expensive than



Figure 3.3: The detail from MRI scans of a patient's head.

other modalities.

Medical imaging plays a crucial role in the diagnosis and treatment of various diseases and conditions. Different imaging modalities are used for different purposes, depending on the type of condition being diagnosed or monitored (Fayad [2023]). *XR* are often used to diagnose bone fractures, dental problems, and chest abnormalities. *CT* scans are used to diagnose and monitor a variety of conditions, including cancer, heart disease, and neurological disorders. Also *MRI* is often used to diagnose and monitor cancer, heart disease, and neurological disorders, as well as to visualize soft tissues, such as muscles and organs. *ULS* is often used to diagnose and monitor pregnancy, as well as to diagnose and treat a variety of medical conditions, including heart disease, liver disease, and cancer. *PET* scans are often used to diagnose and monitor cancer, heart disease, and brain disorders, and can detect cellular changes in organs and tissues earlier than *CT* and *MRI* scans. In addition to diagnosis and treatment, medical imaging is also used for disease monitoring and progression assessment. By tracking changes in the body over time through imaging, physicians can monitor the effectiveness of treatments and adjust them as necessary. For example, *MRI* can be used to monitor the progression of multiple sclerosis, while *PET* imaging can be used to monitor the response to cancer treatment.




Figure 3.4: Ultrasound image of the foetus a 30 weeks of pregnancy in a sagittal scan. Measurements of fetal Crown Rump Length (CRL). (Moroder [2012])

Each imaging modality has its strengths and limitations, and the choice of modality depends on the specific condition being diagnosed or monitored.

### 3.1 Automatic tool for medical imaging

With the increase in the worldwide population and its ageing, more and more people have access to healthcare in recent years. In this context, Medical Imaging, a crucial tool in modern medicine, has seen tremendous advancements and became one of the most relevant areas of study. In the coming paragraphs, a curated selection of contemporary literature will be analysed to respond to the following questions:

- *Q1*: To what extent is the adoption of automated tools in the realm of Medical Imaging significant and what drives its relevance?
- *Q2*: What form the primary challenges encountered in the conception




Figure 3.5: Fluorodeoxyglucose *PET* images of the brain. Different colours indicate different functional areas of the brain. (Catarina Silva [2019])

of these tools?

- *Q3*: Which methodologies or techniques are currently in use within this domain?

The global rising demand for medical imaging (Smith-Bindman et al. [2008]) has created an increasing need for accurate and efficient interpretation of the produced data (Grossi [2021]). By contrast, the analysis of the generated information is a complex task that requires the involvement of highly skilled specialists, who are difficult and expensive to train and hard to find on the work market (Rimmer [2017]). Consequently, the workload for radiologists and other medical imaging professionals has increased, leading to longer waiting times for patients and potential burnout for practitioners.

Medical Imaging modalities, such as *CT*, *MRI*, and *ULS*, generate high-dimensional and complex data. Traditional human interpretation of these images is time-consuming and prone to errors, especially when subtle changes or in presence of rare pathologies (Krupinski [2010]). Moreover, human experts' inter- and intra-observer variability is a known issue in radiology (Covert et al. [2022]). Studies have been conducted to explore the variability among contours drawn by the same radiologist (intra-observer) and among different radiologists (inter-observer). In Woo et al. [2020] cancer mass size measurements from 13 board-certified radiologists on a set of 10 *CT* scans have been compared, highlighting always having at least 9% average difference (Intraclass Correlation Coefficient – *ICC* score = 0.91) to each other when reviewing the same *CT* image sets; this level of disagreement is not low enough to achieve clinically acceptable performances, since consistency and objectivity are essential attributes of effective medical imaging interpretation.

The increasing complexity and volume of imaging data, the demand for time-critical decision-making and the shortage of highly skilled personnel make Medical Imaging a prime candidate for AI-driven automation. The integration of AI-driven medical imaging analysis with electronic health record – *EHR* systems has the potential to revolutionize patient care. By combining imaging data with relevant clinical information, AI algorithms can support more comprehensive and personalized diagnostic and treatment recommendations. AI-driven tools can provide more information to physicians, enabling the detection of subtle patterns and features in the images that may be challenging to identify by human observers, improving diagnostic accuracy. They can also provide consistent and objective analyses, reducing variability and potentially improving patient outcomes. This integration can also streamline workflows, minimize data input errors, and enhance communication between healthcare providers, ultimately contributing to improved patient outcomes and more efficient healthcare delivery. Moreover, AI algorithms can rapidly process and analyze large amounts of data, providing radiologists with more information and significantly speeding up the analysis of each case, increasing workflow efficiency and minimizing the waiting time for the patients.

Despite the potential benefits of AI adoption within medical imaging, it is accompanied by some challenges necessitating careful consideration (Kelly et al. [2019], Saw and Ng [2022], Waller et al. [2022], Bi et al. [2019]). Foremost among these challenges is the imperative for suitably extensive

datasets packed with accurate annotations for robust AI algorithm training, supporting its reliability and precision. Additionally, the insufficiency of clinician input during AI algorithm development poses a substantial difficulty. Clinician engagement in the developmental process is imperative to ensure clinical relevance and utility. Data access limitations and regulatory constraints arise as significant obstacles. Healthcare data frequently resides within a mosaic of disparate medical imaging archival systems, rendering access and utilization in machine learning endeavours arduous. Regulatory limitations, encompassing data privacy, medical device, and reimbursement regulations, may impede AI's assimilation into healthcare (Avi Goldfarb [2022]). AI algorithms demand reliability and transparency, with a comprehensive comprehension of their functioning. Misaligned incentives, such as misincentives for healthcare providers, have the potential to hinder AI adoption. Lastly, the matter of patient acceptance must be duly confronted. Patient comfort and comprehension concerning AI in healthcare, including its benefits and risks, necessitate meticulous communication and education.

The adoption of AI-driven methods in Medical Imaging had a transformative effect on the field, with various applications proposed over the years, following the key progresses in the AI research field.

The first integration of AI-driven methodologies in Medical Imaging can be traced back to the 1970s - 1980s, when Expert Systems – *ESs* and Rule-based systems – *RBs* emerged as the first approaches applied. These computer systems are designed to solve complex tasks by researching through a set of predefined *IF-THEN-ELSE* rules. This set of rules is defined based on the knowledge of the problem and requires extensive manual effort to be created and maintained. This approach has been successfully applied to perform many tasks such as image segmentation (Stansfield [1986], Matesin et al. [2001], Costin [2013]), feature extraction (Zarandi et al. [2011]) and diagnosis (Chabat et al. [2000], Al-Ani and Rawi [2014], Buchanan and Duda [1983]). Because of the rule set static structure, *ESs* and *RBs* are hard to maintain, are constrained by the availability of expert knowledge and could not easily scale to accommodate the increasing complexity and diversity of medical imaging data. These limitations put severe obstacles to the large-scale adoption of such systems.

Consequently, to the successes obtained in other fields, Machine Learning – *ML* techniques have been applied to Medical Imaging since the 1990s (Erickson et al. [2017]). Differently from *ES* and *RL*, *ML* algorithms allow for the automatic extraction of features and patterns from imaging data,



overcoming the limitations of rule-based systems. Some of the most common techniques adopted are:

- **Support Vector Machines – SVMs:** *SVM* is a supervised *ML* algorithm used for classification and regression tasks. It constructs a hyperplane or set of hyperplanes in a high-dimensional space to separate data points into different classes ([Wikipedia contributors \[2023e\]](#)). The algorithm chooses support vectors, which are data points closest to the hyperplane, to influence the position and orientation of the hyperplane. It has been proposed by Vapnik in 1992 ([Cortes and Vapnik \[1995\]](#)) and adopted successfully in many Medical Imaging applications, such as MRI classification ([Othman et al. \[2011\]](#), [Abdullah et al. \[2011\]](#)), tumour detection and segmentation ([Wu et al. \[2012\]](#), [Zhang et al. \[2004\]](#)) and microcalcifications identification ([El-Naqa et al. \[2002\]](#)). *SVM* achieves a high level of performance, but because of its internal structure and the adoption of kernel-based mathematical transformations, the decisions it makes are not always explainable in a straightforward way.
- **Decision Trees – DTs and Random Forests – RFs:** *DT* is a supervised *ML* algorithm used for classification task. The algorithm provides the final decision by visiting a tree-like model, in which each node represents a decision based on the comparison between the value of one feature and a threshold. Starting from the root node, the method visits the tree arriving at one of the leaves, where the final classification decision is made. Similarly to the rule-based methods, the task is accomplished by a sequence of *IF-THEN-ELSE* rules, but in the case of *DTs* the generation of such sequence is algorithmically derived from data. The training procedure defines the feature to be analyzed and the threshold to be applied for every node, trying to maximize the purity of the resulting set split. Because of its predefined set of static rules, *DTs* can be more prone to overfitting the training data, being less robust than other methods, but, given a decision, it can be always explained by checking the decision path in the tree. In *RFs*, to overcome the robustness weakness of *DT*, decisions from a set of different *DTs*, built using a different subset of features or constraints, are ensembled to provide a final, more robust classification. Despite their simplicity, *DTs* and *RF* have been successfully applied to different tasks in Medical Imaging

(Criminisi and Shotton [2013]), such as emphysema medical image classification (Narayanan et al. [2019]), cell and retina vessels segmentation (Hartmann et al. [2021]) and *ULS* images of metacarpophalangeal joint classification for Rheumatoid Arthritis (Min and Haijiang [2020]).

- **K-nearest neighbors – KNN:** *KNN* is a simple, easy-to-implement, non-parametric, and supervised *ML* algorithm used for both classification and regression problems. The basic principle of *KNN* is that similar samples (similar in the sense of class membership or regression value) exist in close proximity in the space of the features. By measuring the geometric distance of an unlabelled object and all the labelled ones in the space defined by the feature values, *KNN* assigns the label by merging the ones of the  $k$  nearest. *KNN* have been adopted in segmentation of echocardiographic images (Heena et al. [2022]), brain cancer image filtering (Florimbi et al. [2018]), brain tumor and breast cancer classification (Nair and Kashyap [2020]).

Despite the promising results achieved by *ML* techniques, they still require manual feature engineering and are limited by the quality and quantity of labelled data available for training.

The introduction of Deep Learning – *DL* techniques in the early 2010s marked a pivotal moment in AI-driven medical imaging. By automatically extracting the features that are more meaningful for the task to be solved, *DL* overcomes some of the limitations of *ML* methods. This ability eliminates the need for manual feature engineering and enables a more effective and efficient feature extraction. This benefit comes with a price: the automatically determined features are no more directly interpretable based on domain knowledge, and so the decision taken by the *DL* algorithm is not explainable. Some key contributions can be identified as widely adopted as a base of a lot of *DL* applications in Medical Imaging:

- **Convolutional Neural Networks – CNNs:** The *CNN* is a class of artificial neural network that makes use of the mathematical operation called *cross-correlation* to learn spatial hierarchies of features through backpropagation by using multiple building blocks from raw image data (LeCun et al. [1995]). The successive layers of the network combine the features computed by the previous one to compute more complex ones, in the first layer the features are computed directly from input image pixels. Because of the planar shape of the input, features are organized in a matrix also called *feature maps*. The convolution operation

consists of the weighted sum of a set of adjacent features (or pixels in the case of input layer) in the *feature maps*, called the *receptive field* of the feature. Convolutional layers can be interleaved by *Pooling layers*, which reduce the dimension of each *feature map* by merging near features (e.g. mean, minimum, maximum). After a certain number of convolutional and pooling layers, the feature maps are linearized and fed to a sequence of *Fully connected layers*, where every neuron is connected with everyone in the previous and next layer. This last stage is the one dedicated to the compute the decision of the algorithm. The previous stage, composed of the sequence of Convolutional and Pooling layers is referred to as *Feature extraction* part. *CNNs* they are suitable for all those fields in which the locality of the features is a fundamental aspect of the data structure, such as in computer vision, in time series, in problems related to three-dimensional objects, to videos and the recognition of audio tracks. Example of *CNN* architecture is illustrated in Figure 3.6.




Figure 3.6: Data flow and structure in LeNet (LeCun et al. [1998]). The input is a handwritten digit, the output a probability over 10 possible classes. (Dive into deep learning [2023])

- U-Net:** Introduced in 2015, *U-Net* is a specialized *CNN* architecture designed for biomedical image segmentation. *U-Net* has since become one of the most widely used deep learning models in medical imaging, with numerous adaptations and extensions for various imaging modalities and tasks (Ronneberger et al. [2015]). The *U-Net* architecture, as depicted in Figure 3.7, comprises a *contractive* path and an *expansive* path, forming a U-shaped structure. The *contractive* path consists of a series of *convolutional* layers, followed by *rectified linear unit – ReLU* and *max pooling operation* operations. These operations reduce the spatial information while increasing the feature information, with a

high number of feature channels. On the other side, the *expansive* path combines the feature and spatial information through a sequence of *up-convolutions* and concatenations with high-resolution features from the contracting path, by incorporating *skip connections*. The latter represents one significant aspect of the *U-Net*, establishing direct links between corresponding layers in the *contractive* and *expansive* sections so that can be effectively propagated both local and global contextual information enhancing its ability to perform accurate segmentation tasks.




Figure 3.7: Architecture of *U-Net* (Ronneberger et al. [2015]). Each blue box corresponds to a multi-channel feature map. The number of channels is denoted on top of the box. The x-y-size is provided at the lower left edge of the box. White boxes represent copied feature maps. The arrows denote the different operations. (University of Freiburg [2011])

- **Transfer Learning – TL:** *TL* is a *ML* technique where the knowledge, in form of parameters, of a model already trained on a task, is applied to another model, but to solve a different, but related task (Wikipedia contributors [2023f]). Training large *DL* models can be a resource-intensive activity and may require a massive amount of data to converge to an acceptable level of performance. One of the most notable examples in the computer vision is the usage of weights values

from a model pre-trained on a massive public generic dataset of pictures for classification tasks (e.g. [Deng et al. \[2009\]](#), [Lin et al. \[2014\]](#)) as initialization for a task-specific model to be *fine-tuned* on a much smaller dataset. For this reason, *TL* has been particularly beneficial for medical imaging applications, as it allows researchers to overcome the challenges associated with limited and imbalanced datasets, initializing the model parameters with values coming from pre-trained ones.

- **Self-supervised learning – SSL:** *SSL* is a *ML* paradigm that utilizes unlabeled samples to obtain useful representations of the data ([Wikipedia contributors \[2023d\]](#), [Assran et al. \[2023\]](#)). The crucial benefit of *SSL* is that it doesn't require human-annotated labels solving one of the main issues in many Medical Imaging applications. The typical *SSL* pipeline consists of computing a supervisory signal that is then used as a pseudo-label to train the model. The task of fitting the supervisory signal is called *Pretext* or *Auxiliary task*. In computer vision, some examples of pseudo-labels are the belonging or not of two image portions to the same original image, the belonging or not of two geometric transformations to the same original picture, or the reciprocal positioning of various portions of the same original image. The obtained model can be used as a feature extractor for a downstream task-specific model or as initialization in a *TL* framework.

Considering the aforementioned advancements in the state of the art, several observations can be emphasized to address the questions (*Q1 - Q3*) raised at the start of this section:

- *Finding 1:* Escalating demand for the analysis of Medical Imaging and its escalating intricacy exercise considerable pressure on specialized practitioners, significantly impacting their workloads. Furthermore, recent advances in medical imaging devices generate increasingly meaningful, intricate data, necessitating greater effort from operators. Additionally, partly driven by this workload pressure, it is common to observe substantial variability among physicians when interpreting the same image. Moreover, error rates can vary considerably depending on the level of expertise and the psychophysical condition of the specialist. In this context, automatic tools serve as valuable aids to physicians, alleviating their workload.
- *Finding 2:* The principal challenges entailed in the adoption of automatic tools for Medical Imaging encompass ethical considerations, the

absence of clinician input, constraints related to data availability, obstacles associated with data access, regulatory impediments, challenges in extending applicability to diverse populations and settings, and patient acceptance.

- *Finding 3*: The adopted approaches can be classified into three primary categories: Expert and/or Rule-based systems, Machine Learning, and Deep Learning. While the former has been present for over three decades, the latter has gained substantial momentum more recently. Particularly, DL is progressively emerging as the predominant methodology in numerous Medical Imaging applications.

## 3.2 Specific use cases

This chapter aims to present a methodology for detecting and estimating deterioration in biological systems, with a specific focus on the human body, through the analysis of several use cases in the field of Medical Imaging. The use cases that will be discussed include *COVID-19* for multiple tasks, the estimation of pediatric bone age, and the evaluation of coronary artery calcium score. The subsequent sections will provide a more detailed description of these use cases, including an analysis of the available data and the current state of the art.

### 3.2.1 COVID-19

*COVID-19* is a contagious disease mainly affecting the respiratory system, caused by the SARS-CoV-2 virus (WHO [2023a]). The virus spreads mainly between people close together transported by the small droplets and aerosol in the breath of an infected subject. Is also possible to get infected by the virus when a contaminated surface is touched, but this can be considered a secondary route of infection. A portion of people report asymptomatic evolution, but, when present, reported symptoms varying from person to person, often including fever, dry cough, breathing difficulties, and loss of smell and taste (WHO [2023a]). Even if most of the symptomatic subjects, with high dependence on age, develops only mild to moderate symptoms, *COVID-19* can also evolve from severe manifestations to death. During the press conference of the WHO's general director Tedros Adhanom Ghebreyesus on 6<sup>th</sup> April 2023 (WHO [2020]), *COVID-19* has been officially characterized as a world pandemic. At the date of 7<sup>th</sup> April 2023 have been

confirmed more than 762.200.000 cases of *COVID-19* infection causing more than 6.889.000 deaths worldwide. In Italy, the first Western country to experience diffuse infection, the total number of cases has been more than 25.695.000 with more than 189.000 death (WHO [2023b]). Figure 3.8 provides an overview of worldwide infection.




Figure 3.8: World map showing the number of total COVID19 cases in each country, WHO [2023b]

The initial surge of positive cases increased the pressure on healthcare systems to a critical point and imposed strict rules on the whole society to contain the diffusion of the virus, with inevitable effects on the economy and people's wellness. The high number of subjects to be managed, produced a shortage of available beds in hospitals, especially in Intensive Care Units – *ICUs*, medical equipment, and medical staff all over the world. In an attempt to reduce the criticality, governments imposed restrictions on the mobility and activities of citizens, like small gathering cancellations, closure of educational institutions, border restrictions, individual movement restrictions, and national or regional lockdowns. Although these countermeasures have been proven effective in contagion containment (Haug et al. [2020]), produced severe side effects on the society and economics of the countries.

In the first phases, the presence of *COVID-19* infection was diagnosed by RT-PCR test, consisting of the analysis of nasopharyngeal and throat swabs using chemical reactions to obtain and replicate hundreds of thousands of times a segment of genetic material to be extracted. If the virus is present in the collected sample the amplification increases its content enough to be detected and measured. RT-PCR sensitivity ranged from 32% for the

pharyngeal swab to 63% for the nasal swab (Ghoshal et al. [2020]), while its specificity ranged between 54% and 73% (Loeffelholz and Tang [2020]). The result of the test required hours to be ready, impacting the management of logistic and diagnostic workflow in the emergency departments (Long et al. [2020]).

During this thesis, the proposed methodology has been applied on two different tasks for COVID19:

- **Radiology diagnostic workflow optimization:** to rearrange the diagnostic worklist prioritizing the cases most probable to be infected by COVID19, based on CXR analysis.
- **Monitoring of COVID19 infection severity:** to quantify the severity of infection from CT scan or CXR images.

For each of the tasks, the next sections will be provided with a brief description and a summary of the available data.

### **Radiology diagnostic workflow optimization**

#### *Description*

During the first phases of the *COVID-19* pandemic, the medical operators were not equipped with an accurate diagnostic tool that can provide results promptly. In this scenario, the chest radiography *CXR* emerged as an alternative tool for *COVID-19* in the diagnostic work-up of symptomatic patients suspected of being infected. Recognition of some peculiar radiological findings in *CXRs*, such as multi focal and bilateral ground glass opacities and consolidations, suggests the infection, as illustrated in Figure 3.9. The main radiological societies endorse this approach (ACR [2020], Akl et al. [2021]), also in consideration of the simplicity and speed of cleaning operation, crucial for the infection containment, and the large availability of radiological equipment, also in a portable format. The main limitation of *CXR* usage as a diagnostic tool for *COVID-19* is due to its limited performance in the detection of disease, in the early stages, affecting its effectiveness in clinical practice. Has been estimated that *CXRs* analysis by radiologists provides 61.1% sensitivity to *COVID-19* with a specificity around 63% (Gatti et al. [2020]). Moreover, the analysis of *CXRs* requires the involvement of highly specialized personnel. Because of the shortage of such skilled operators, when dealing with a large number of radiographs, the analysis can produce a bottleneck in the workflow. For these reasons,



the availability of a tool able to automatically pre-analyse the CXRs, and prioritise the most suspicious cases during the workflow will be highly beneficial for the doctors. Opposed to the most frequently used diagnostic queue policy, *first in first out*, *FIFO*, this tool can help to promptly identify the infected case and adopt all the required containment actions to reduce the contagion. Given the prioritization elaborated from the tool, a subject identified as a potential infected can be separated from the other people in the waiting room, and necessary care can be provided faster. By giving precedence to the examination of cases with a higher likelihood of being infected, a reduction in the average waiting time for patients testing positive for *COVID-19* is achieved, as they are no longer required to wait for *COVID-19*-negative cases to be attended to first. Furthermore, the prioritization ranking can serve as one of the factors considered by physicians during the diagnostic process, promoting more precise diagnostic outcomes.

#### *Data*

To conduct the experiments described here and evaluate the performance of our proposed approach a dataset of CXRs collected by AOU San Luigi Gonzaga ([Alesina et al. \[2023\]](#)) has been used. The images have been proposed in anonymized DICOM format and the study has been approved by the IRB, according to the Helsinki Declaration. The positive group of 234 subjects has been determined by the presence of a positive nasopharyngeal swab RT-PCR test in the 24 hours with respect to the CXR, and these pictures have been collected during the virus outbreak in the first semester of 2020. The negative group of 300 people contains CXRs produced before the emergence of the virus in Italy. The second group has been retrospectively reviewed by a team of radiologists with 20, 10, and 3 years of experience to distinguish between healthy cases and subjects affected by disease, different from *COVID-19* (e.g., bacteria pneumonia, lung cancer). Both the positive and negative CXRs are produced in the same departments of the same hospital. The dataset is composed of 728 pictures, of which 250 are from the *COVID-19*-positive group. Because of the imbalance between positive and negative cases in the lateral view CXRs (170 out of 186 are from the negative group), these have been excluded from this study, keeping only antero-posterior and posterior-anterior views. The dataset used in this study is composed of 542 pictures about 536 subjects from 3 categories: *COVID-19*, *No Finding* and *Other*. [Figure 3.10](#) and [Table 3.1](#) summarize the data selection and the classes.

The age distribution of subjects in the positive and negative groups is




Figure 3.9: Representative serial chest radiography of patients with *COVID-19* infection. A, Images in a 29-year-old man with rapid respiratory deterioration after symptom onset shows the progression from lower lung-predominant interstitial and airspace opacities on day 3 to diffuse involvement with extensive airspace disease on days 4 and 6. B, Images in a 56-year-old-man with *COVID-19*, presenting initially with a normal chest radiograph, which then progressed to lower lung-predominant interstitial and airspace opacities at day 9, which mildly worsened by day 15 [Stephanie et al. \[2022\]](#).

Class	Description	Number of images
No Finding	Healthy subjects	146
COVID-19	COVID-19 positive case	234
Other	People affected by disease different from <i>COVID-19</i>	162

Table 3.1: Dataset composition for Radiology diagnostic workflow optimization for *COVID-19* task




Figure 3.10: Data selection and labelling of COVID-19 images for Radiology diagnostic workflow optimization use case. Images have been assigned to different groups according to RT-PCR test (for COVID-19) and assessment by a team of radiologists (for other diseases).

illustrated in Figure 3.11 and summarized by statistics in Table 3.2 and Table 3.3. The two cohorts do not differ significantly in age average (positive: 67.2, negative: 66.7, *Kruskal-Wallis* test p-value: 0.792, alpha: 0.05), variance (positive: 253.7, negative: 281.6, *Bartlett* test p-value: 0.153, alpha: 0.05) and distribution (*Mann-Whitney* test p-value: 0.389, alpha: 0.05). They do not differ significantly also in gender distribution (*Fisher* exact test p-value: 0.539, alpha: 0.05).

Label	Average	Variance	Median	Quantile				p-value*
				2.5 <sup>th</sup>	25 <sup>th</sup>	75 <sup>th</sup>	97.5 <sup>th</sup>	
COVID-19	67.2	235.7	69.5	37.9	56.0	79.0	92.0	0.0012
Negative	66.7	281.6	70.0	26.0	57.0	78.0	91.0	<1e-7

Table 3.2: Age statistics of positive and negative classes for Radiology diagnostic workflow optimization for *COVID-19* task. \*Shapiro test.




Figure 3.11: Sample distribution over the age classes, divided by label

Label	Total	Female	Male	Female %	Male %
COVID-19	234	97	137	41.5%	58.5%
Negative	300	133	167	44.3%	55.7%

Table 3.3: Gender statistics of positive and negative classes for Radiology diagnostic workflow optimization for *COVID-19* task.

### Monitoring of COVID19 infection severity

#### *Description*

RT-PCR is a diagnostic tool but does not provide any quantification of the severity of infection. Under the stress conditions of the *COVID-19* pandemic peaks, it is extremely important not just to know if an incoming patient is positive or not at the triage, but also to have accurate monitoring of the disease evolution to adopt the most appropriate healthcare procedures for each patient. In this context, thoracic imaging, specifically *CXR* and *CT*, is playing an essential role in the management of patients, especially those evidencing risk factors (from triage phases) or moderate to severe *COVID-19* signs of pulmonary disease (Rubin et al. [2020]). *Computed tomography* is considered to be the primary diagnostic modality for examining patients with *COVID-19* (Aljondi and Alghamdi [2020]), because of its better performance. *CXR* has limited sensitivity for the early stage of infection but is the most commonly used diagnostic imaging modality because it is a widespread, relatively cheap, fast, and accessible tool, which may be easily

brought to the patient’s bed, even in the emergency departments. Monitoring of affected patients requires serial image acquisitions, often on daily basis, which can rise concerns about the X-ray dose and, since it requires a lower amount of radiation, also in this case *CXR* is preferred to *CT* in clinical practice.

The task of quantifying the infection severity requires highly skilled physicians and can take a significant amount of time, which is not compatible with the needs of the hospital departments. Also in this case, the shortage of such qualified operators can be a bottleneck in the healthcare process, preventing the medical equipe from promptly addressing criticalities and providing the most effective treatments. In this scenario, the adoption of an automatic tool that scans the produced thoracic images, both *CXR* and *CT*, suggesting to the doctor a measure of infection severity can speed up the operations and improve operational efficiency.

*Data*

To develop and test the proposed methodology when applied to the task of *COVID-19* severity estimation, two datasets have been used.

One dataset referred to as *BrixIA* (Signoroni et al. [2021]), is composed of 4703 *antero-posterior* – *AP* and *posteroanterior* – *PA CXRs*, collected from the ASST Spedali Civili di Brescia, one of the biggest hospitals in the north side of Italy, in the period between March 4<sup>th</sup> and April 4<sup>th</sup> 2020, the first pandemic wave peak. The main characteristics of the dataset are summarised in Table 3.4. The images are publicly available at [Università degli Studi di Brescia \[2022\]](#). All the images are annotated with a severity score in the range of 0-18, where 0 indicates the absence of infection signs, and 18 the maximum severity. The score distribution is shown in Figure 3.12. Pictures are provided in full resolution DICOM format.

Modality	CR (62%) - DX (38%)
View position	AP (87%) - PA (13%)
Manufacturers	Carestream, Siemens
Image size	(1056-3050) x (1186-3376)
Total samples	4703 images
Training set	3311 images
Validation set	945 images
Test set	447 images

Table 3.4: BrixIA data main characteristics

This score is referred to as *BrixIA score*, (Borghesi and Maroldi [2020], Borghesi et al. [2020]) and is determined as follows: 1) the CXR image is




Figure 3.12: Brixia score distribution with sex stratification on the Brixia *COVID-19* dataset

divided between the two lungs, 2) each lung is divided, from top to bottom, in three zones, 3) for each zone, a score is indicated in range 0-3 and 4) the local scores are summed to provide the overall evaluation, the *Global score*. If the zone under evaluation does not present abnormalities, it is scored 0, while if interstitial infiltrates, or interstitial and alveolar infiltrates with interstitial dominant, or interstitial and alveolar infiltrates with alveolar dominant, it is associated respectively with score equal to 1, 2 or 3. The zonal segmentation of the lungs and score examples are illustrated in Figure 3.13. The labelling of this dataset is a collective effort of a team of about 30 specialists from the radiology staff at ASST Spedali Civili di Brescia. A smaller group of them, composed of 4 radiologists with different levels of seniority, has been asked to rate individually a subset of 150 images from the test set. For each picture in this subset, a *Gold Standard score* has been indicated based on majority voting, considering seniority in case of a tie. This *Gold Standard score* is inherently considered more reliable than the one in the other pictures.

Another dataset has been used to conduct the experiments for the *CT* modality. This dataset is provided by the ICIAP 2021 conference organizer




Figure 3.13: Brixia score: (a) zone definition and (b–d) examples of annotations. Lungs are first divided into six zones on frontal chest X-rays. Line A is drawn at the level of the inferior wall of the aortic arch. Line B is drawn at the level of the inferior wall of the right inferior pulmonary vein. A and D upper zones; B and E middle zones; C and F lower zones. A score ranging from 0 (green) to 3 (black) is then assigned to each sector, based on the observed lung abnormalities.

for the *COVID-19* infection percentage estimation competition (Bougourzi et al. [2021]), the *Per-COVID-19 challenge*. The dataset includes *CT* scan slices (simply slices from now on) of patients either infected with *COVID-19* or healthy and the problem statement is to estimate the percentage of *COVID-19* infection rate. The dataset is divided into three sets, Train, Test and Validation. The Train set includes 3054 *CT* scan slices from 132 patients, out of which 128 are diagnosed as infected by *COVID-19*. The method of diagnosis is a *CT* scan prognosis by expert thoracic radiologists and the positive *reverse transcription polymerase chain reaction – RT-PCR*. The remaining 4 patients are healthy i.e. not infected by *COVID-19*. In the Train set, *CT* scan slices are provided with their respective estimated *COVID-19* infection rate, calculated by two expert radiologists. The infection rate is a figure between 0% and 100% and is determined as the ratio between the picture area of the infected regions and the one occupied by the lungs. The Validation set includes 1301 slices from 57 patients, among which 55 are *COVID-19* positive. This set is blinded, i.e. the *COVID-19* ground truth infection rate is not known. The Train and Validation sets have been collected from different patients including both Male and Female patients. The age of patients ranges from 27 to 70 years old. The data collection was done between June to December 2020, from two hospitals: Hakim Saidane, in Biskra (Algeria), and Ziouch Mohamed, in Tolga (Algeria) (Bougourzi et al. [2021]). Finally, the Test set includes 4449 *CT* Scan slices from 130 patients, all tested positive for *COVID-19* using both diagnosis methods, *RT-PCR* and thoracic *CT* scan prognosis. The dimensions of these slices are 630x630. The test set is also blind, i.e. the ground truth labels are not known. Table 3.5 summarizes the main characteristics of the three sets. All the pictures are provided in *PNG* format.

Set	Patients	Positive patients	Total slices
Train	132	128	3054
Val	57	55	1301
Test	130	130	4449

Table 3.5: Summary of the ICIAP 2022 Per-COVID-19 challenge dataset

### 3.2.2 Estimation of pediatric bone age

#### *Description*

Metabolic and endocrine disorders can affect skeletal maturation, making it faster or slower than expected. In particular, the wrists and hands are sensitive to hormonal anomalies and can be examined by physicians for discrepancies concerning the expected bone maturity. In paediatrics, the bone age assessment is routinely used to compare skeletal maturation with chronological age; huge differences represent warning signs suggesting further evaluation for possible endocrine and metabolic diseases (Zerin and Hernandez [1991]).

Despite the importance of bone age assessment in the monitoring of hormonal disorders, the process has not been updated since the mid of previous century. The two main approaches are the comparison of the patient’s left-hand radiograph (Figure 3.14) with an atlas of reference cases (Greulich and Pyle [1959]) or the elaboration of a complex scoring that takes into account 20 specific bones in the hand (Tanner [1983]). In both cases, the procedure is tedious and requires a lot of time from the specialists. Concerns regarding inter-observer variability in manual bone age estimation have risen since 1992 (Berst et al. [2001], Heyworth et al. [2011]). All these aspects, limit the application of bone age estimation in clinical practice, especially in time-critical contexts. Hence, this procedure is an ideal candidate for the development of an automatic image evaluation method that can estimate bone age. This value can be then compared with the chronological one to monitor the presence of growth disorders and their evolution.

#### *Data*

A public dataset from RSNA has been used to conduct the experiments described in this document. The dataset has been published in the context of the *RSNA Pediatric Bone Age Machine Learning Challenge* in 2019 (Halabi et al. [2019]), created to show and application of machine learning and artificial intelligence in medical imaging. The curation and use of the dataset for this competition have been approved by the Stanford University






Figure 3.14: Examples of hands and wrists x-radiograph used to estimate the bone age

and University of Colorado institutional review boards.

The dataset consists of 14236 hand radiographs obtained from Children’s Hospital Colorado (Aurora, Colo), and Lucile Packard Children’s Hospital at Stanford. Each image has been labelled with the patient’s gender and bone age estimation, elaborated from the original clinical radiology report. The dataset has been split into 3 subsets: 12611 samples have been selected for Training set, 1425 images for Validation set and the remaining 200 for Test one. The latter only contains pictures from the Lucile Packard

Children’s Hospital. The gender distribution across the various datasets is depicted in Figure 3.17, and it demonstrates a balance between the two genders. Examining the distribution of bone age within the Training and Validation sets, as shown in Figure 3.15 and Figure 3.16, reveals non-uniform labelling distribution across the age spectrum, that is consistent in the two datasets. Specifically, there is an underrepresentation of extreme age values for newborns and young adults, while the majority of samples fall within the age range of 9 to 15 years old.




Figure 3.15: Bone age distribution and the number of images in the training set.

Training and Validation sets labelling has been obtained by merging different estimations:

- **Clinical report:** extracted from the review of the original clinical record.
- **Independent review:** four pediatric radiologists who reviewed the cases independently (two pediatric radiologists from each institution).
- **Second review:** second review by one of the pediatric radiologists who reviewed the cases approximately 1 year after the first review.




Figure 3.16: Bone age distribution and the number of images in the validation set.

The reviewers adopted the approach from Greulich and Pyle standard ([Greulich and Pyle \[1959\]](#)) during their analysis for bone age estimation. For the definition of ground truth in *Test* set, the same steps have been applied, but, to ensure a more reliable estimation, the different estimations have been merged by applying different weights to the contribution, based on reviewers' deviation from the group ([Larson et al. \[2018\]](#)). 6.1 months *Mean Absolute Error* – *MAE* has been measured for human evaluations.

### 3.2.3 Quantification of calcium score

#### *Description*

Cardiovascular events remain one of the most frequent causes of mortality and morbidity worldwide, despite a significant decrease over the past decades. In most cases, the event occurs in subjects without past known cardiac diseases. Thus, the adoption of preventive measures to identify those who have a high probability of manifesting cardiovascular events in the future is crucial to decrease the number of these kinds of fatalities. Analyzing the risk factors of each patient, such as the age, gender, diet, lifestyle and family history, is one of the ways to evaluate the a priori probability of




Figure 3.17: Gender distribution and the number of images in the training, validation, and test bone age data sets.

experiencing cardiac events, but does not provide an assessment of the real status of the cardiovascular system.

The Agatston calcium score has been proposed as a method to quantify the severity of atherosclerotic coronary calcium, *CAC*, and has been demonstrated valuable in the prediction of cardiovascular events (Parikh et al. [2018]). It is computed as a summed score based on calcified plaque area and the maximum measured density among all the identified calcified lesions (Agatston et al. [1990]). The absence of coronary calcium in cardiac computed tomography is associated with a very low rate of major cardiac events in the next 3-5 years. The obtained score is then compared with percentiles based on age, gender and ethnicity (McClelland et al. [2006]). Traditionally, a *CAC* score equal to zero is associated with very low risk, with values between 1 and 99 with mildly increased risk, 100 to 299 with moderately increased, and over 300 to severely increased risk. Risk categories are summarized in Table 3.6.

To measure the *CAC* levels is used the non-contrast material-enhanced electrocardiographically gated cardiac *CT* scans. In the produced image the area of calcifications on cross-sectional sections weighted by *CT* attenuation

CAC value range	Risk
0	very low
[1 – 99]	mildly increased
[100 – 299]	moderately increased
> 300	severely increased

Table 3.6: Summary of the CAC score ranges and associated risk

is measured to calculate the final score. Figure 3.18 shows some examples of acquired images and associated scores. Also, non-contrast-enhanced cardiac fluoroscopy and radiography are adopted to compute the risk level.




Figure 3.18: Example of CT scans acquired to estimate different levels of Agatston calcium score

As already discussed in the previous sections, *Chest radiographs* are more commonly adopted than *CT* scans and are preferred in several cases because of simplicity, lower x-ray dose and availability. In fact, *CXR* modality can also be acquired with portable devices, better fitting the routine clinical practice in hospitals worldwide. In this context, having an automatic tool able to analyze the *CXRs* to estimate the *CAC score* can bring massive benefits to healthcare, dramatically simplifying the cardiovascular checkup and enabling massive screening in the population, with impact on overall mortality.

### Data

To test our proposed methodology for this task, a dataset of *CXR*s has been used. The dataset contains a total of 505 radiographs coming from 505 different patients, collected specifically for the problem of interest by a group of radiologists of the "A.O.U. Citta della Salute e della Scienza di Torino", during the period from 2009 and 2022. Images are anonymized and saved in *Digital Imaging and Communications in Medicine – DICOM* format. The main aspects of the provided dataset are summarized in Table 3.7.

Property	Value
Total radiographic images	505
Total patients	505
- Women	206
- Men	299
Average age	59
Maximum age	94
Minimum age	3
Patients with CAC = 0	186
Average CAC	1192
Standard deviation of CAC	2135
Median	155
Maximum value	13049
Minimum value	0

Table 3.7: Summary of the CAC score dataset

The same group of specialists also labelled each image, providing its CAC score. Among all 505 patients, 186 are healthy with a CAC score equal to 0, while the rest are distributed over a wide range that reaches a maximum of 13049. Figure 3.19 illustrate the distribution of the risk score, it is possible to observe the peak at CAC equal to 0 coinciding with healthy patients. Most of the remaining values are distributed in the range between 1 and 2000.

### 3.3 Deep Learning application in Medical Imaging

Over the past decade, deep learning has emerged as one of the most prominent areas of research in computer science, demonstrating successful applications in a wide range of fields including computer vision (Chai et al. [2021]), timeseries analysis (Hahn et al. [2023], Mahmoud and Mohammed




Figure 3.19: Distribution of CAC score in the provided dataset for the Quantification of calcium score from *CXRs*

[2021], Torres et al. [2021]), natural language processing (Otter et al. [2020]), audio analysis (Liu et al. [2022], Roger et al. [2022], Zhu et al. [2021]), and content generation (Harshvardhan et al. [2020], Oussidi and Elhassouny [2018]). However, due to the vast amount of literature on this topic, providing a comprehensive review in this document is not feasible. Therefore, the scope of this study is limited to a set of relevant and significant findings in medical imaging, which are pertinent to the use cases presented.

### Chest radiography

With the advent of the machine and deep learning technologies (LeCun et al. [1998], LeCun et al. [2015]), numerous methodologies have been proposed to address various tasks in the analysis of *CXRs* and *CTs* (Bhattacharya et al. [2021]). The availability of large public datasets (Irvin et al. [2019], Soda et al. [2021], Cohen et al. [2020b], Bougourzi et al. [2021]) and the emergence of the *COVID-19* pandemic have stimulated research on *convolutional neural networks* – *CNNs* for the early detection of SARS-CoV-2 positive cases, resulting in the introduction of several novel approaches in

the past year. A significant majority of the proposed solutions have focused on disease identification using deep learning algorithms (VJ et al. [2021], Brunese et al. [2020], Haghanifar et al. [2022], Ozturk et al. [2020], Afshar et al. [2020], Wang et al. [2021b], Li et al. [2020], Hu et al. [2020]), with some level of interpretability (Panwar et al. [2020], Karim et al. [2020], Wang et al. [2021a]). Other studies have tackled different tasks, such as quantifying infection severity (Cohen et al. [2020a], Aboutaleb et al. [2022]), image segmentation (Alom et al. [2004], Amyar et al. [2020]), predicting disease progression (Sriram et al. [2021]), and image synthesis (Zunair and Hamza [2021]).

The article by Irvin et al. [2019] is a seminal work in medical imaging, specifically in the area of *CXR* analysis. The paper introduces a large-scale dataset of *CXRs* containing approximately 224,316 images from over 65,000 patients. The dataset is unique in that it includes radiologist-labelled annotations for 14 different pathologies, including pneumonia, pleural effusion, and atelectasis, with a hierarchical labelling structure illustrated in Figure 3.20. The authors propose a novel approach for addressing missing labels, a common issue in medical datasets, which involves introducing a labeller disagreement among radiologists and a multi-label model that accounts for these disagreements. By training a model based on *Densenet-121* (Huang et al. [2017], 3.21), the authors achieve high *Area Under Curve* – *AUC* values for all of the diseases covered, ranging between 0.97 (Pleural Effusion) and 0.85 (Atelectasis). These impressive results demonstrate the efficacy of the proposed approach and have significant implications for the development of automated systems for *CXR* analysis.

In Signoroni et al. [2021], a novel *DL* architecture called *BS-Net* is proposed for estimating the *BrixIA score* from *CXRs*. Figure 3.22 illustrates its structure. To address the issue of patient pose variability, the proposed model performs image alignment using the *Spatial Transformation Network* – *STN* (Jaderberg et al. [2015]) and eliminates irrelevant portions of the image (i.e., turning the pixels to black) by employing a nested version of *U-Net*, also known as *U-Net++* (Zhou et al. [2018]), which is trained to perform lung segmentation. *BS-Net* provides both *BrixIA score* estimation and disease segmentation, highlighting the severity of infection in different lung zones. The performance of *BS-Net* was evaluated on the *Gold Standard score* subset of the *BrixIA* dataset (150 images), resulting in a *Mean Absolute Error* – *MAE* of 1.787.

The task of diagnostic workflow optimization by using automatic *CXR*






Figure 3.20: Labelling hierarchy from Irvin et al. [2019] dataset

analysis has been studied in Baltruschat et al. [2021] and Annarumma et al. [2019]. In both cases, the adoption of an automatic system reduced the reporting delay for critical and urgent imaging findings when compared with the standard worklist processing *FIFO*, demonstrating the potential of automated systems to improve the efficiency and accuracy of *CXR* analysis, which can have significant implications for patient outcomes and healthcare costs.

Several *DL* methods have been proposed and evaluated on the publicly available *Per-COVID-19 challenge* dataset with the purpose of estimating the severity of *COVID-19* infection by analyzing *CT* scans. In Chaudhary et al. [2022], the authors propose using the Swin-L transformer (Liu et al. [2021]) as a feature extractor, and a two-layer MLP is trained to estimate the infection severity. In the study presented by Anwar [2022], an ensemble of six *Resnet* models is used to calculate the severity. The features




Figure 3.21: A 5-layer dense block with a growth rate of  $k=4$ . Each layer takes all preceding feature-maps as input.

computed by the networks are stacked, and a *Squeeze and Excitation* – *SE* block is added to weight the features of the backbone *CNNs*. Finally, a decision fully-connected layer is attached to the *SE* block to obtain the estimation. In Miron and Breaban [2022], the authors propose a solution based on a modified version of *ResNeSt* (Zhang et al. [2022]) complemented by two fully-connected layers. The first layer consists of 101 neurons corresponding to discrete percentage scores (0-100), and the second consists of one neuron for the final prediction. Three loss functions are adopted: (i) Smooth L1 loss function between the predicted value and the expected one, (ii) Smooth L1 loss was between the sum of first layer and the real infection score, (iii) the Kullback–Leibler divergence loss between the softmax of the output of the first layer and the probability distributions of the ground truth. The overall loss is the sum of the three losses. In the article presented by Napoli Spatafora et al. [2022], a mixup data technique is used to augment the dataset, significantly increasing the diversity of the training data. To compute the prediction, a model based on *Inception-v3* (Szegedy et al. [2016]) is used. Most of the proposed approaches adopt *TL* to initialize their model, using the network weights from the *ImageNet* classification task. All of the presented methods provide promising performance, with *MAE* in the range of 4.17-4.99 for the validation set and 3.55-6.53 for the




Figure 3.22: Detailed scheme of *BS-Net* from [Signoroni et al. \[2021\]](#). In particular, in the top-middle the *CXR* to be analyzed is fed to the network. The produced outputs are: the segmentation mask of the lungs (top-left); the aligned mask (middle-left); the *Brixia score* (top-right).

testing set. The *Pearson Correlation – PC* ([Cohen et al. \[2009\]](#)) is also high, ranging from 0.949-0.936 for the validation set and 0.855-0.709 for the testing set. These results demonstrate the potential of *DL* methods for analyzing *CT* scans and estimating the severity of *COVID-19* infection, which can have significant implications for patient outcomes and healthcare resource allocation.

However, methodological flaws and underlying bias are pervasive in many of these studies, as highlighted in the study by [Roberts et al. \[2021\]](#). After reviewing 62 studies proposed during 2020 for the diagnosis or prognosis of *COVID-19* from *CXR* or *CT* images using *ML* methods, none of them is of potential clinical use due to methodological flaws and/or underlying biases, mostly because of the data used in the studies. The article identifies that the usage of public datasets increases the results’ reproducibility, but because the labelling process and subject selection are not always described with sufficient details, the results are not reliable enough to be applied to the real-life context. The adoption of image exclusion criteria that are not

enough detailed, is another aspect that makes it difficult to apply the proposed methodology in the clinical setting. Finally, in the cases the samples for *COVID-19* and the control group are from very different production processes or very different in subject composition (e.g. paediatric patients as controls), the risk is to overfit and having the model learning these differences, instead of the visual patterns of the disease.

A deep learning study has been carried out using lung ultrasound videos for the detection of *COVID-19* (Roy et al. [2020]). Besides detection, it also performs localization on each lung ultrasound frame to provide better estimate of the disease.

### Estimation of pediatric bone age

Several studies have been conducted to automate the task of estimating pediatric bone age from *XR* images, and some of them have been tested on the *RSNA Pediatric Bone Age Machine Learning Challenge* dataset during the competition (Halabi et al. [2019]). The method that performed the best (Alexander Bilbily [2022]) utilizes both image and gender information. Specifically, the image is analyzed by a *CNN* model based on the *Inception-v3* architecture to extract features, and the gender, coded with values 0 or 1, is fed to a fully connected layer of 32 neurons. The features from gender and picture are then concatenated and fed into an *MLP* model with two layers of 1000 neurons, followed by an output layer to compute the final estimation. To achieve the claimed performance, multiple high-performance models are ensembled. In the work presented by Ian Pan [2020], the gender is also taken into account by training specific models for each sex. In this approach, image patches of 224x224 pixels are analyzed by a *CNN* model based on the *ResNet-50* architecture, as shown in Figure 3.23, and the overall prediction is computed as the median of the estimations from the 49 patches. The methodology that reached the third-place in the competition is composed of a novel-designed *CNN* structure, namely the *ICE module*, considerably smaller than *Inception*-based models (approximately 1% of parameters). This module is composed of a *Transposed convolution* stage followed by a *Convolutional* layer, and finally, a *Pooling* layer computes the final assessment. Each of the model parts is trained separately, and the best ones are merged to be used in the test set. Ensemble of different *CNN* architectures have also been adopted in the solution that reached the fifth place in the competition. Moreover, this methodology introduces a segmentation stage to mask the area of the *XR* image that is not related to the hand, illustrated

in Figure 3.24. In all presented solutions, very reduced  $MAE$  in the range of 4.2-4.55 months were observed. In an attempt to achieve better results by combining the strengths and limitations of the proposed solutions, a merge of all of them has been proposed in Pan et al. [2019], obtaining 3.79 months  $MAE$ . These studies demonstrate the potential of  $DL$  methods for automating the estimation of pediatric bone age.

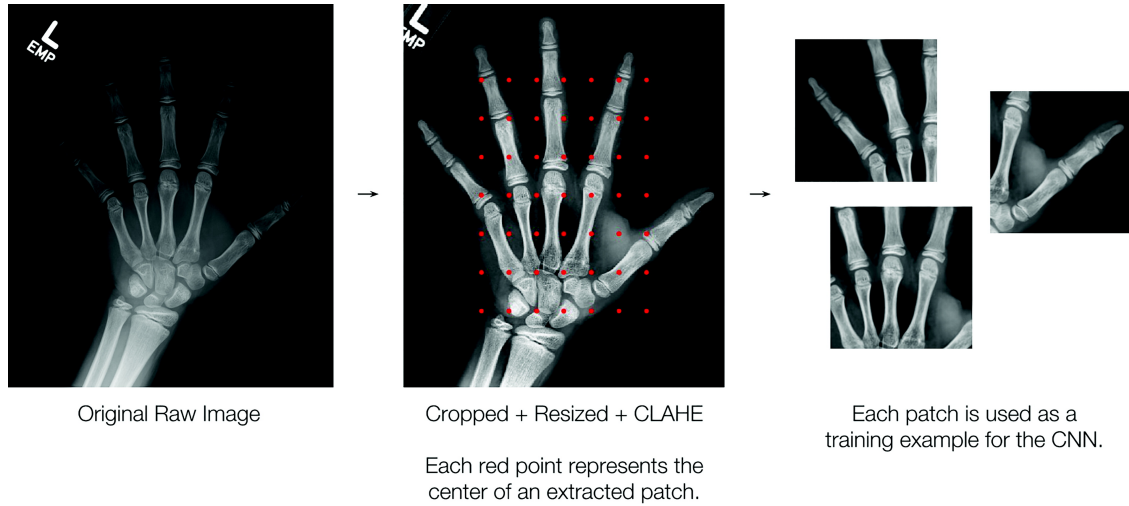



Figure 3.23: Preprocessing pipeline for the second-place method used to construct inputs to the neural network. The image is manually cropped and resized to a length of 560 pixels, and the contrast is enhanced; this is followed by extraction of 49 patches of  $224 \times 224$  pixels


### Quantification of calcium score

Due to the importance of evaluating the *Coronary Artery Calcium* –  $CAC$  score accurately and efficiently, several approaches based on applying  $DL$  algorithms to  $CXR$  have been proposed in recent years (Feng et al. [2019]). In Kamel et al. [2021], different models are trained to classify each  $CXR$  into the risk categories defined by the  $CAC$  value. The models are derived from the  $VGG-16$  architecture (Simonyan and Zisserman [2014]). In D’Ancona et al. [2023], a  $DL$   $CNN$  model is tested as a supportive tool to estimate patient risk for  $CAC$  score. In both cases, the results demonstrate the potential of  $DL$ -based methods to improve the estimation of the  $CAC$  score. In the first study, the presented model achieves an  $AUC$  of 0.73 for the detection of cases with an index higher than zero. In the second study, when integrated into a *Logistic regression* –  $LR$  binary classifier, the model prediction emerged as the strongest severe predictor for *Coronary Artery Disease* –  $CAD$ , as visible in the plot in Figure 3.25. These studies

### 1. Manual Mask Annotations



### 2. Dilated Convolutional U-Net



### 3. Convolutional Neural Network Ensemble

ResNet-50 with global average pooling  
 Inception-V3 with global average pooling  
 Xception with global average pooling  
 Xception with global max pooling  
 Inception-ResNet-V2 with global average pooling  
 Inception-ResNet-V2 with global max pooling

M/F Embedding Layer

Bone Age Regression Layer

Figure 3.24: Hand masking by *U-Net*, as it is adopted in the fifth-place solution in RSNA Pediatric Bone Age Machine Learning Challenge

demonstrate the potential of *DL*-based methods as a supportive tool for *CAC* score estimation.

## 3.4 Proposed methodology

In all of the use cases presented, a medical image, specifically *XR* and *CT*, is employed to estimate a variety of quantities, including but not limited to infection severity, diagnostic priority, bone age, and coronary artery calcium score. These use cases exhibit both commonalities and distinctions that can either facilitate or hinder the widespread application of the same algorithm across all instances. Notable differentiators include the involvement of distinct anatomical regions (e.g., chest versus hand), diverse imaging modalities (e.g., *CT* versus *XR*), and the targeting of varying clinical aspects (e.g., infection severity versus bone maturity). In an hypothetical industrial context, the capacity to employ a consolidated methodology adaptable to diverse scenarios with minimal adjustments would grant significant advantages to a company, decreasing the costs, the time-to-market and facilitating user adoption. It's worth noting that the datasets at this stage suffers from a significant class imbalance issue, which is prevalent in




Figure 3.25: AI4CAD: Deep learning to detect severe *CAD* from *CXR*. (D’Ancona et al. [2023])

almost all use cases. Medical applications are particularly affected by class imbalance, as disease prevalence can vary significantly, making it difficult to gather a sufficient number of patients for a specific application. Conversely, retrospective studies, like those described in this context, may have an overabundance of positive cases. Additionally, it’s important to note that extreme values at the dataset boundaries, such as the *CAC* score, *BrixIA* score, and Bone age, occur less frequently than those within the central range. If left unaddressed, these imbalances can result in biased models that significantly undermine their predictive accuracy and generalization capacity, particularly for underrepresented classes or values. In scientific literature, many techniques have been proposed to mitigate this problem. While some data augmentation techniques will be used, a comprehensive exploration of this aspect will be examined in the dedicated section within the Discussion. Bearing this context in mind, the proposed methodology has been tailored to accommodate one or more images as input and subsequently generate one or more continuous values as output.

The approach proposed is illustrated in Figure 3.26 for the case of *CT*

scans, and can be summarized as follows. In a nutshell, each image is projected in a well-behaved lower dimensional *target feature space* such that nearby points in this feature space correspond to images with a similar label, while distant points correspond to different values. Whenever the target quantity shall be predicted for a query slice, it is projected in the target feature space. The query image projection is compared with the projections of the labelled pictures from the training set, i.e. the reference images. The query picture target value is predicted by interpolating the ones of the nearest reference samples in the feature space.

The proposed method is implemented as a 3-stages pipeline as follows:

- **Image pre-processing**, to standardize the images pixel intensity and resolution across the different acquisition settings
- **Feature extraction**, to project the picture into the above-mentioned well-behaved feature space, and
- **Distance based regression**, to predict the value looking at the query image neighbors in the feature space from reference set.

In the following, each of the above stages are detailed.

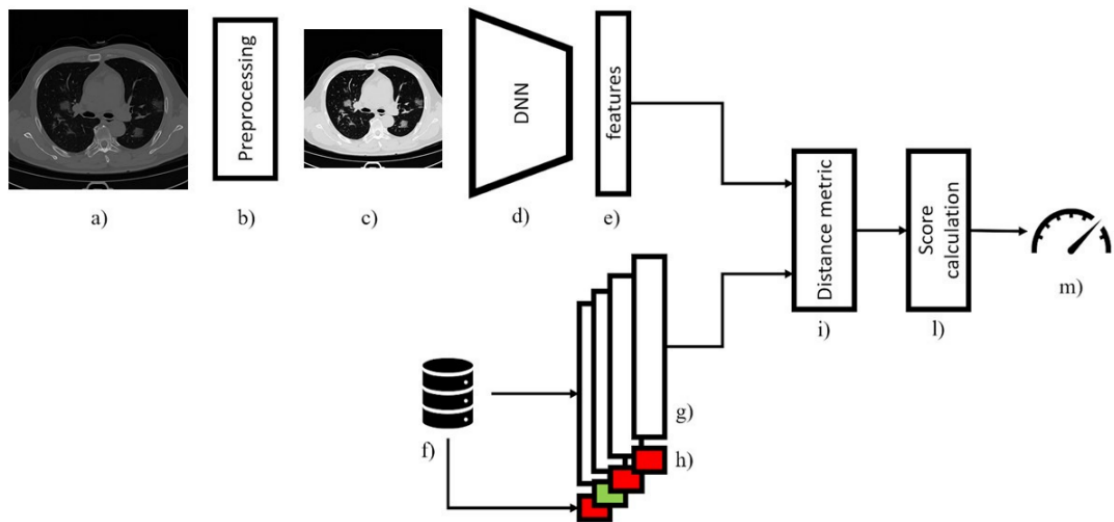


Figure 3.26: High level illustration of adopted methodology in the case of *CT* scans. **a)** *query*: original image; **b)** *image pre-processing*; **c)** pre-processed picture; **d)** **e)** *feature extraction*; **f)** **g)** **h)** *reference set*: database of labelled cases from training set with corresponding projections; **i)** **l)** *distance based regression*; **m)** final prediction



### 3.4.1 Image pre-processing

Image appearance can be strongly affected by the adopted acquisition process and technology. In particular the contrast, the brightness and the pixel intensity histogram can change when equipment from different vendors are used, or different settings are applied. This variability is evident in Figure 3.27, where differences in pixel intensity are present among the samples. It can be also observed, that the image size and scale can differ significantly between pictures.



Figure 3.27: Samples from dataset provided by ICIAP 2021 Challenge organizers during testing phase. Depending on the production process, slices appearance can strongly differ in term of colour distribution, scale and size. In terms of colour distribution, the slices in the first row are significantly brighter than the last ones. Image size of third slice is different from the others. The sixth slice has a different scale.

Because these sharp differences can lead to poor performance, a pre-processing strategy is put in place, illustrated in Figure 3.28 for the case of *CXRs* and composed of the following steps:

- **Image resize:** In this step, slices are resized to a common pixel format. The optimal size is selected based on the specific application requirements, taking into account the trade-off between performance and computational resource load.
- **Region of interest cropping:** This step involves identifying the meaningful portion of the picture and discarding the rest through a cropping operation.

- **Pixel intensity scaling:** In this step, pixel values are re-scaled to a limited range. Depending on the specific case, different scaling strategies can be used, such as normalization, min-max scaling, clipping, or a combination of these techniques.




Figure 3.28: The pre-processing stage. The original image **a**) is properly resized **b**). The resized picture **c**) is then processed by a segmentation algorithm **d**) to extract a binary mask **e**) indicating the region of interest. To define the final cut **g**), the binary mask **e**) is framed by the smallest rectangle **f**). The resized picture **c**) is scaled **h**,**i**) and finally cropped **j**) using the computed rectangular cut to obtain the final picture **k**).

### 3.4.2 Feature extraction

The preprocessed images are projected into the target feature space using the deep convolutional neural network *DenseNet-121*. This architecture is characterized by a sequence of *dense blocks* (concatenation of non-linear operators with skip connections to prevent gradient vanishing effect), connected by *transition layers* (concatenation of convolution and pooling layers), and finally, a *linear layer* is responsible for the class prediction. In this case, the original architecture has been truncated before the output classification layer, resulting in a convolutional feature extractor that yields a 1024-dimensional projection of the image provided as input.

Consider a pre-processed picture  $x$  from the set of all possible pictures  $X$ , and a deep neural network model  $M(\cdot)$ , we have:

$$\forall x \in X \rightarrow M(x) = z \in \mathcal{F} \subseteq \mathbb{R}^{1024}, \quad (3.1)$$

where  $z$  is the feature vector corresponding to the image  $x$  and belongs to the target feature space  $\mathcal{F}$ . With the notation  $x_i$  the  $i$ -th sample in  $X$  and  $z_i = M(x_i)$ , consider  $y_i \in \mathcal{S} \subseteq \mathbb{R}$ , the label associated with  $x_i$ , the goal is to design  $M(\cdot)$  such that:

$$\forall i, j, k \in \mathbb{N} : \{\|y_i, y_j\| < \|y_i, y_k\|\} \rightarrow \{\|z_i, z_j\| < \|z_i, z_k\|\} \quad (3.2)$$

where  $i, j, k$  indicate three samples in set  $X$  and  $\|a, b\|$  is the distance between points  $a$  and  $b$ . This distance metric can be chosen conveniently to maximize the performance, using an iterative process, among the ones presented in the literature.

### 3.4.3 Distance based regression

To compute the final prediction, the method exploits the distance in the target feature space between the query image and the reference pictures for which the label is known.

Consider a set  $X^{ref}$  of  $n^{ref}$  references with relative label  $Y^{ref}$ , such that given a  $x_i^{ref} \in X^{ref}$  its quantification is  $y_i^{ref}$ ,  $\forall i = [1, 2, \dots, n_{ref}]$ . The methodology computes *feature reference set*  $Z^{ref}$  as:

$$Z^{ref} := \{M(x_i^{ref}) \quad \forall i = [1, 2, \dots, n_{ref}]\} \quad (3.3)$$

Whenever it is requested to predict the output for a query picture  $x^{query}$ , its projection in the target feature space is computed as:

$$z^{query} = M(x^{query}) \quad (3.4)$$

The distance between the resulting 1024-dimensional vector  $z^{query}$  and the set of reference vectors  $Z^{ref}$  is then computed.

The elements in the feature reference set are sorted from the nearest to the most distant and related elements in  $Y^{ref}$  are stored in the sorted list  $Y^{ref,sorted}$ .

To compute the final prediction  $\widehat{y^{query}}$ , the proposed algorithm computes:

$$\widehat{y^{query}} = \frac{1}{\sum_{i=1}^m w_i} \sum_{i=1}^m w_i \cdot y_i^{ref,sorted}, \quad y_i^{ref,sorted} \in Y^{ref,sorted} \quad \forall i = 1, 2, \dots, n^{ref} \quad (3.5)$$

where  $m \leq n^{ref}$  is a parameter of the algorithm, regulating the number of nearest neighbours considered in the score estimation, and  $w_i$  is a weight coefficient used to regulate the contribution of each sample. To tune the first, an iterative process has to be put in place: for increasing values of  $m$ , the performance of the method is evaluated, looking for the best trade-off. The definition of the weight coefficients can be performed empirically, depending on their distance in the target feature space or their position in the ordered list  $Y^{ref,sorted}$ .

### 3.5 Method settings

The definition of the proposed methodology is both general enough to accommodate various use cases and specific enough to effectively address the *Deep Regression* task for medical imaging. However, achieving optimal performance for each use case requires careful selection of the techniques to apply at each step, as well as the assignment of appropriate values to the tunable parameters that regulate system behaviour. To this end, the next sections provide a detailed explanation of the design choices made for each of the three main steps of the methodology: *Image pre-processing*, *Feature extraction*, and *Distance based regression*.

#### 3.5.1 Image pre-processing procedure and parameters

Given that the presented use cases involve different types of medical images, the pre-processing step must be tailored to each specific use case. This step is critical because the appearance of images can vary significantly across different subsets, depending on the acquisition process, settings, and technology used. Differences in image contrast, brightness, and pixel intensity distribution are some of the most visible effects. Additionally, image size can vary from sample to sample, affecting the pixel scale of details and the model's ability to generalize. Therefore, it is essential to carefully consider and adjust the pre-processing step for each use case to optimize the performance of the proposed methodology.

Image size plays a crucial role in the model's performance. Using high-resolution images can help the model better recognize certain patterns that can be exploited to accomplish its task, but it is necessary to carefully choose the feature scale that matches to not negatively impact predictive performance. Furthermore, the use of large images has a substantial impact on computational power demand, making the overall pipeline slower or even intractable for certain computation platforms.

In the use case related to Radiology worklist prioritization, resizing has been applied to every image such that the biggest dimension (width or height) results in 1000 pixels. This operation is performed to preserve the original proportion to avoid the introduction of unnatural distortion. The same approach has been applied to all other use cases, with different picture sizes: 624 pixels has been used as the value for *CAC* score, while 578 pixels for the *COVID-19* infection score from *CXRs* and Bone Age estimation, and 512 for the *COVID-19* infection score from *CT* scans.

In some cases, accomplishing a specific task can be achieved by only looking at a particular portion of the whole image. In the case of *COVID-19* infection signs in *CXR*, for example, not only are the lungs visible in the picture, but also the abdomen, arms, and part of the head. Removing these areas not only reduces the amount of data to process, making the overall system less intensive in computational resource demand but also has the effect of producing better convergence in the model training, avoiding it from focusing on background clutter in other regions that are not relevant for the task. In the diagnostic workflow optimization and pediatric bone age estimation use cases, the application of *Region of interest – ROI* masking is beneficial in obtaining better performance. In the first case, a pre-trained *U-Net* model has been applied to identify the lungs, and a proper rectangular section is defined to contain them. This will remove the abdomen, arms, and head, which can all be sources of distracting details and limit the model to focus only on the chest portion of the images. Testing its performance on the dataset, has been discovered that it encounters some issues in capturing one or both lungs when the subject is affected by severe opacity. In a significant number of these cases, the *U-Net* recognizes only one lung in case of unilateral opacity or none in case of bilateral. To reduce the occurrence of these cases, some checks have been introduced in the resize algorithm that: i) force the horizontal symmetry in the cropping coordinates to take advantage of the recognition of one lung to detect the other, ii) limit the amount of area that can be cropped out so that resulting images cannot be

too small, and iii) enlarge the cropping area of some pixels to avoid that near-border details are eliminated.

In the case of bone age estimation, a simple algorithm has been implemented that scans the picture to find the smallest rectangle to exclude the background pixels, exploiting the fact that they are significantly darker than the *ROI*.

When different radiological machines or settings are used, this can significantly affect the characteristics of the produced images. In addition to differences in image dimensions, the pixel intensity scale is also affected by the variability of the production process. Therefore, procedures must be put in place to standardize the pixel intensity distribution, which helps to minimize the influence of specific setups on the images fed into the model. In some cases, the settings used can produce images with very low contrast and brightness, making it difficult to discern any patterns.

To maximize the performance of the proposed methodology during experiments in the *COVID-19* severity and *CAC* scores from *CXR*, as well as the bone age estimation from the *XR*, histogram adjustment using the CLAHE algorithm (Pizer et al. [1987]) has been performed. The adoption of this procedure enhances the contrast in the images, potentially making the presence of patterns more evident. This approach has been shown to be beneficial in Tjoa et al. [2022] for *XR* scans.

Table 3.8 summarizes the image pre-processing settings used for the different use cases.

Parameter	COVID-19			Bones	Heart
	Prioritization CXR	Severity score CXR	CT	Age XR	CAC score CXR
Image size	1000 × 1000	578 × 578	512 × 512	578 × 578	624 × 624
ROI	Lung	-	-	Hand	-
Pixel	normalization	percentiles	percentiles	percentiles	percentiles
- range	-	2 <sup>nd</sup> - 98 <sup>th</sup>	min - 90 <sup>th</sup>	2 <sup>nd</sup> - 98 <sup>th</sup>	2 <sup>nd</sup> - 98 <sup>th</sup>
- colour	-	CLAHE	-	CLAHE	CLAHE

Table 3.8: Image preprocessing for the different use cases

### 3.5.2 Feature extraction training process

At the heart of the proposed methodology lies the feature extraction model  $M(\cdot)$ , defined as *DenseNet-121* deep learning model, whose behaviour significantly impacts the overall system performance. As highlighted in Equation

3.2, it is crucial that the computed projections in the target feature space accurately reflect the similarity between the corresponding labels of two samples. To this end, selecting the most appropriate training procedure is crucial, with several alternatives requiring a thorough evaluation to ensure optimal performance.

During the investigation of the *Radiology diagnostic workflow optimization* use case, various experiments were conducted to measure the influence of different training procedures on system performance. Specifically, the following cases were tested: (i) adoption of *TL* without any fine-tuning, using *ImageNet* classification task weights, (ii) initializing the model parameters with *TL* as in (i), and fine-tuning on *CXR* by leveraging *SSL* techniques (i.e., Autoencoder (Kramer [1991], Kramer [1992])) on the *BrixIA* dataset, (iii) additional fine-tuning starting from (ii), by training  $M(\cdot)$  for binary classification between healthy and infected cases, and (iv) additional fine-tuning starting from (ii), by training  $M(\cdot)$  with triplet loss contrastive learning (Chechik et al. [2010]).

The performance and structure of the induced feature space were studied in each case and presented in the results section of this work. Based on the findings of this analysis, a new loss function was proposed and tested on all other use cases, with promising results. This loss function, named *DISTMAT*, is designed to regulate the relative distance among the sample projections such that the proximity is related to label similarity, as in Equation 3.2. In other words, the loss function encourages the feature space to be such that the distance between two samples is as close as possible to their label similarity. This helps ensure that samples with similar labels are mapped to nearby regions in the feature space, while samples with dissimilar labels are mapped to distant regions. The result is a feature space that is optimized for the *Distance based regression* procedure. The details of this method are described in the next section.

Regulating the other parameters is also crucial to improve the effectiveness of the training procedure in machine learning. The batch size, which represents the number of samples taken into account during each learning step, is often limited by available computational resources. The number of epochs has a significant impact on the success of the training process; if the number is too low, the model will not have enough learning steps to converge, but if it is too high, the risk of overfitting to the training set increases. The learning rate, which regulates the influence of each learning step in the model parameters update, is also an important parameter to

regulate. Increasing this parameter could make the model converge faster, but can also cause instability and poor results. Lower values, on the other hand, make the training procedure smoother but could prevent the model from converging.

Variability in the training set is an important characteristic to improve the model performance and generalization. If the set is too small or has low variability, this can lead to poor generalization because it is not representative of the sample variation that can be found in the real world. To mitigate this issue, a set of transformations can be applied to the sample during the training procedure (some examples are illustrated in Figures 3.29 and 3.30). This technique, called *Data augmentation*, has been applied to all the use cases. The transformations consist of geometrical modifications of the image that preserve its overall structure.

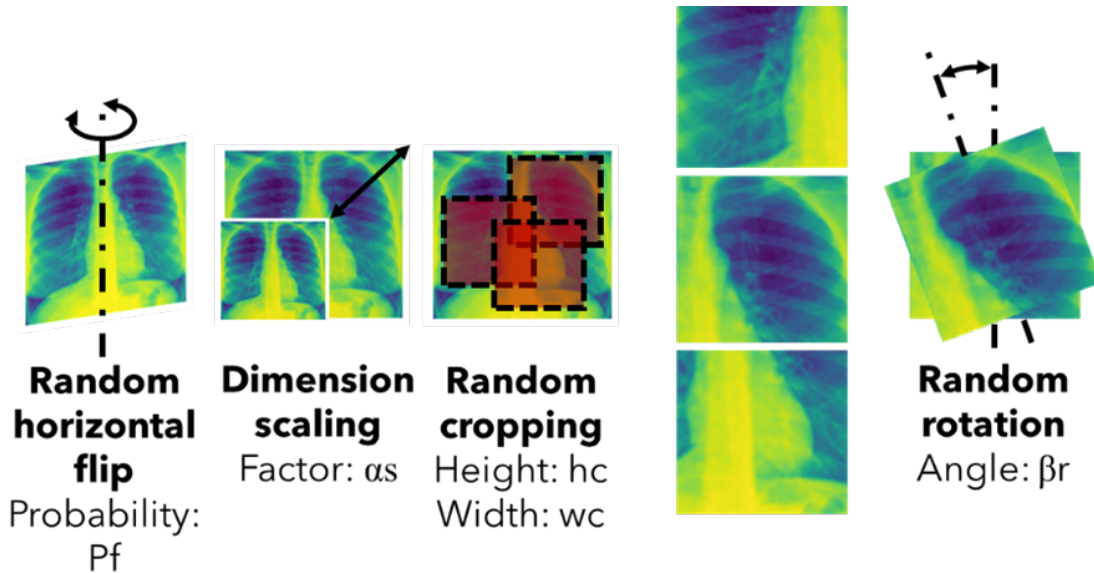


Figure 3.29: Some examples of transformation applied during data augmentation

Table 3.9 describes the training procedure settings used for the different use cases.

#### Distance matrix loss – DISTMAT

To train the model  $M(\cdot)$  achieving the sought property in Eq. 3.2, has been defined a novel loss function, namely *DISTMAT*, based on the relative distance of the samples in the target feature space  $\mathcal{F}$  and prediction space  $\mathcal{S}$ .






Figure 3.30: How the random crop is executed during data augmentation. The crop coordinates are randomly chosen with uniform distribution

Parameter	Prioritization CXR	COVID-19		Bones	Heart
		Severity score CXR	Severity score CT	Age XR	CAC score CXR
Loss function	various	DISTMAT	DISTMAT	DISTMAT	DISTMAT
Batch size	various	14	12	12	12
Learning rate	various	0.05, 0.005	0.05	0.05, 0.005	0.01, 0.001
Epochs	various	100	30	100	500
Augmentation					
- Flip	Horizontal	Horizontal	Horizontal	Horizontal	Horizontal
- Rotation		25°		25°	25°
- Shift		0.1		0.1	0.1
- Crop	512 × 512	512 × 512	448 × 448	512 × 512	512 × 512
- Resize			512 × 512		
- Scale		0.1		0.1	0.1
- Distortion		Elastic: $\alpha=60, \sigma=12$ Grid: limit=0.3 Optical: distort=0.2, shift =0.2		Elastic: $\alpha=60, \sigma=12$ Grid: limit=0.3 Optical: distort=0.2, shift =0.2	Elastic: $\alpha=60, \sigma=12$ Grid: limit=0.3 Optical: distort=0.2, shift =0.2

Table 3.9: Image preprocessing for the different use cases. Prioritization use case has been tested with different settings that will be discussed in the next sections

Considering a random subset of  $n_b$  pictures  $X^{batch}$ , related labels  $Y^{batch}$ , and computed projections in the feature space  $Z^{batch}$ , the proposed method

defines the following distance matrices:

$$D_z \in \mathbb{R}^{n_b \times n_b} \rightarrow D_z[i, j] = \|z_i, z_j\| \quad z_i, z_j \in Z^{batch} \forall i, j \in [1, 2, \dots, n_b] \quad (3.6)$$

$$D_y \in \mathbb{R}^{n_b \times n_b} \rightarrow D_y[i, j] = \|y_i, y_j\| \quad y_i, y_j \in Y^{batch} \forall i, j \in [1, 2, \dots, n_b] \quad (3.7)$$

To achieve faster convergence of the training procedure, because of the different magnitude of  $D_x$  and  $D_y$ , is beneficial to scale each of the matrices dividing by its maximum value:

$$\widehat{D}_z \in [0, 1]^{n_b \times n_b} \rightarrow \widehat{D}_z[i, j] = \frac{D_z[i, j]}{\max(D_z)}, \forall i, j \in [1, 2, \dots, n_b] \quad (3.8)$$

$$\widehat{D}_y \in [0, 1]^{n_b \times n_b} \rightarrow \widehat{D}_y[i, j] = \frac{D_y[i, j]}{\max(D_y)}, \forall i, j \in [1, 2, \dots, n_b] \quad (3.9)$$

The *loss function*  $L_D$  is defined as the *Mean Absolute Error* – *MAE* between  $\widehat{D}_z$  and  $\widehat{D}_y$ :

$$L_{DM} = \frac{1}{n_b^2} \sum_{i=1}^{n_b} \sum_{j=1}^{n_b} |\widehat{D}_z[i, j] - \widehat{D}_y[i, j]| \quad (3.10)$$

### 3.5.3 Distance based regression design and tuning

The final prediction computation is governed by Equation 3.5, which contains various parameters and functions that regulate the system’s behaviour. These settings must be adjusted through an iterative procedure to identify the optimal combination in terms of method performance. To determine the appropriate value of the parameter  $m$ , which represents the number of nearest samples in the reference set considered when computing the weighted average, the procedure evaluates the method’s performance for increasing values of  $m$  and selects the optimal value that produces the best results. The coefficient  $w_i$  can be used to incorporate distance into the final score computation. The fundamental concept behind this approach is that closer samples should have a greater weight in the average computation. While this feature was not utilized in all use cases, it was employed in the experiments conducted during the optimization phase, where a logarithmic decay function produced the most favourable outcomes. This function is outlined in Equation 3.11.

$$w_i = \left( \frac{1}{\log_2(i+1)} - \alpha \right) \cdot \beta, \quad \alpha = \frac{1}{\log_2(N+1)}, \quad \beta = \frac{1}{1-\alpha} \quad (3.11)$$

The Equation 3.5 includes the term  $Y^{ref,sorted}$ , which contains the label associated with the samples in the reference set sorted according to the distance between the query projection  $z^{query}$  and the set of reference projections  $Z^{ref}$ . This formulation necessitates the selection of a distance metric. In all use cases in which the *DISTMAT* loss function was utilized, the Euclidean distance was adopted due to the loss definition. In the remaining case, radiology workflow optimization, the Cosine similarity produced the best results.

The table 3.10 summarizes the configuration of the distance-based regression step used for the different use cases.

Parameter	COVID-19			Bones	Heart
	Prioritization CXR	Severity score CXR	CT	Age XR	CAC score CXR
$m$	10	2	21	27	5
Distance metric	Cosine	Euclidean	Euclidean	Euclidean	Euclidean
Weights	Eq. 3.11	1	1	1	1

Table 3.10: Configuration of distance-based regression algorithm for the different use cases

The available data for the *Coronary Artery Calcium* – CAC score estimation use case exhibits a highly imbalanced distribution, with a majority of labels concentrated at lower values. Specifically, out of the 505 patients, 186 are associated with a score of 0. To address this skewed distribution, the proposed methodology incorporates several mathematical steps to maximize performance when computing the final score. This transformation comprises the following sequential procedures:

1. The CAC score  $CAC_i$  associated with the  $i$ -th patient is used to compute the quantized CAC score  $CACq_i$ , according to Equation 3.12:

$$CACq_i = \begin{cases} 0 & \text{if } CAC_i \in [0,10] \\ 10 & \text{if } CAC_i \in (10, 200] \\ 200 & \text{if } CAC_i \in (200, 500] \\ 500 & \text{otherwise} \end{cases} \quad (3.12)$$

2. The transformed quantized CAC score  $\widehat{CACq}_i$  is computed using Equation 3.13:

$$\widehat{CACq}_i = \log(CACq_i + 1) \quad (3.13)$$

3. The transformed quantized CAC scores of selected neighbours are combined using Equation 3.5 to obtain  $\widehat{yq}^{query}$ .
4. The final prediction  $\widehat{y}^{query}$  is obtained by applying the function in Equation 3.14:

$$\widehat{y}^{query} = \exp(\widehat{yq}^{query}) - 1 \quad (3.14)$$

The impact of log transformation on the distribution of the CAC scores becomes evident upon examination of Figure 3.31. A comparison with Figure 3.19 reveals that, aside from values at zero, the samples exhibit a more even dispersion across the entire range. The ultimate distribution, including quantization, is depicted in Figure 3.32, wherein the values are now partitioned into four distinct setpoints.




Figure 3.31: Distribution of logarithmic transformed CAC score




Figure 3.32: Distribution of transformed CAC score

## 3.6 Results

The implemented and applied methodology for deterioration estimation from medical imaging has been extensively utilized across all the introduced use cases in the preceding sections. Subsequently, the subsequent sections delve into the metrics adopted for assessing performance and the corresponding benchmarks. Finally, the achieved results are presented through the utilization of metrics and visualizations, facilitating a comprehensive comparison with the defined benchmarks.

### 3.6.1 Metrics and benchmarks

In this section, for each use case, performance metrics and corresponding benchmark values will be defined to facilitate the interpretation of results.

#### Radiology diagnostic workflow optimization

The proposed methodology is implemented to analyze a set of *CXR*s to prioritize cases that are more likely to be from *COVID-19* positive patients.

The tool evaluates each image, computing a *COVID-19* risk indicator, which is then used to sort the *CXRs* and optimize the diagnostic workflow. The evaluation of the methodology’s performance utilizes specific metrics and related benchmarks, which are summarized, with confidence intervals, in Table 3.11.

The evaluation of the proposed algorithm’s performance in sorting cases involves the computation of the *Mean Average Precision – MAP* (Zhu [2004], Wikipedia contributors [2023b]). The *MAP* is a metric employed in information retrieval to assess the effectiveness of ranking algorithms. It quantifies the precision and relevance of retrieved documents in a ranked list by calculating the average precision for each query and subsequently computing the mean across all queries to obtain a single evaluation score. The precision value for a query is determined by the ratio of relevant documents retrieved to the total number of retrieved documents. This evaluation process ensures that higher precision scores are achieved when relevant documents are ranked higher, and it assumes values within the range of 0 to 100%, where the latter indicates that all relevant results are positioned at the top. A higher *MAP* score indicates a superior ranking algorithm that effectively retrieves relevant documents.

In essence, the computation of the *MAP* metric follows these steps:

1. The sorted list of cases, prioritized using the proposed methodology, is defined as  $\mathcal{R} := \{c_1, c_2, \dots, c_n\}$ , where  $n$  represents the total number of cases.
2. For each case  $c_i$  ranked at the  $i$ -th position, if it pertains to a *COVID-19* positive patient, the sorted list of preceding cases, denoted as  $\mathcal{R}_i := \{c_j \in \mathcal{R}\}$  with  $j < i$ , is determined. This list represents the cases prioritized relative to  $c_i$ .
3. For each list  $\mathcal{R}_i$ , the precision of  $c_i$  ranking is computed as the ratio between the number of *COVID-19* positive cases and the total number of cases in  $\mathcal{R}_i$ .
4. Finally, the *MAP* is computed as the average precision across all positive cases.

The *MAP* is a valuable metric as it takes into account both the order and relevance of retrieved documents, providing a more comprehensive evaluation compared to traditional precision and recall measures. Its widespread adoption in information retrieval research enables effective comparison and

benchmarking of various ranking algorithms, thereby facilitating advancements in the field.

The utilization of the *COVID-19* risk indicator extends to the identification of critical cases through the establishment of a threshold, beyond which an alarm can be triggered for prompt action. To assess the performance of the *COVID-19* risk indicator in this regard, standard classification metrics such as Sensitivity and Specificity can be employed. Sensitivity and Specificity serve as evaluation metrics for classification tasks, quantifying the algorithm’s effectiveness in determining the class membership of a given query image. These metrics respectively indicate the proportion of correctly identified actual *COVID-19* positive and negative cases. Their values range between 0 and 1, with 1 signifying that all cases were correctly assigned to their respective classes.

To estimate each metric, 6 times iterated 5-fold stratified cross-validation has been used. At each iteration, cases are randomly shuffled, and 5 folds were extracted. Four folds out of five are used as the reference set, while the remaining ones are used as queries. This procedure results in 30 different evaluations for each metric from which the average and 95% confidence interval are computed based on the Gaussian distribution hypothesis (Shapiro test  $p\text{-value} > 0.05$ ).

To assess the potential benefits of the proposed method, reference benchmarks have been established for each of the selected metrics. Regarding the prioritization of the diagnostic workflow, except for specific cases identified during triage and hospital access, the queue is predominantly managed according to a *First-In-First-Out – FIFO* policy. This policy can be regarded as a form of random sorting. When applied to the provided data, it results in a *MAP* score of 47.79% (40–48%). This value serves as a reference for the prioritization task.

To gauge Sensitivity and Specificity, the performance level of physicians can be adopted as the benchmark. Literature data (Gatti et al. [2020], Stephanie et al. [2020], Islam et al. [2021], Cozzi et al. [2020], Borakati et al. [2020]) indicates that radiologists achieve a Sensitivity of 61% (55–67%) and a Specificity of 63% (40–89%) in *COVID-19* diagnosis from *CXR*s. The relatively low performance of physicians in diagnosing *COVID-19* is due to the wide range of clinical manifestations, symptom overlap with other infections and individual variability in disease presentation. Imaging tools as *CXR*s may not capture early-stage lung abnormalities, while delays in testing can

impact sensitivity. These values are utilized to comprehend the metrics obtained through the proposed methodology and to establish the threshold for triggering alarms. Specifically, the system has been fine-tuned to attain a level of Specificity similar to the benchmark when evaluating Sensitivity, and vice versa.

Metric	Description	Interpretation	Benchmark
Mean Average Precision – MAP	Given the sorting induced by the COVID-19 indicator, this metric indicates, on average, the portion of positive cases placed before another positive one	Higher values indicate the positive cases are placed on top of the diagnostic queue	47.8% (40–48%)
Sensitivity	Portion of COVID-19 positive individuals correctly identified by the system (COVID-19 risk indicator higher than a threshold)	Higher values indicate the method can identify a larger majority of positive cases	61.1% (55–67%)
Specificity	Portion of COVID-19 negative individuals correctly identified by the system (COVID-19 risk indicator lower than a threshold)	Higher values indicate the method can identify a larger majority of negative cases	63.0% (40–89%)

Table 3.11: Metrics estimated for performance analysis, with benchmark values and confidence interval.

### Monitoring of COVID19 infection severity

The availability of existing literature on the quantification of *COVID-19* infection severity through *CXR*s and *CT* scans, on the same datasets introduced in previous sections, allows for a straightforward establishment of metrics and benchmark values. The metrics and their corresponding benchmarks are presented in Table 3.12.

In Signoroni et al. [2021], the authors conducted a study to evaluate the algorithm’s performance in assessing the severity of *COVID-19* infection using the *BrixIA* dataset. The evaluation involved measuring the *Mean Absolute Error – MAE* between the predicted and actual values of the *BrixIA* score. Specifically, a predefined subset of 150 samples, referred to as the *Gold Standard score* subset, was utilized for this assessment. The computed *MAE* for this evaluation is 1.787.

During the *Per-COVID-19* challenge, various teams presented diverse



Metric	Description	Interpretation	Benchmark		
			BrixIA	Per-COVID-19 Validation	Testing
Mean Absolute Error – MAE	Mean absolute difference between the estimated value and the real one	Lower values indicate better performance	1.787	4.17-4.99	3.55-6.53
Root Mean Square Error – RMSE	Square root of the mean squared difference between the estimated value and the real one	Lower values indicate better performance	N.A.	8.359-9.081	7.510-10.275
Pearson Correlation – PC	Pearson correlation index computed between the estimated value and the real one	Higher values indicate better performance	N.A.	0.936-0.949	0.709-0.855

Table 3.12: Metrics estimated for performance analysis.

approaches to address the task of predicting the severity of infection based on *CT* scans. To determine the challenge winner, the submitted solutions were assessed on two distinct subsets: the validation set and the testing set. For each method, performance metrics such as *MAE*, *Root Mean Square Error – RMSE*, and *Pearson Correlation – PC* were computed. The *RMSE* is calculated as the square root of the average of the squared differences between the predicted and actual scores. Similarly, the *PC* is determined as the Pearson correlation coefficient between these quantities.

To establish the rankings of the solutions, the *MAE* was prioritized first, followed by the *RMSE*, and finally the *PC*, for each subset. On the validation set, the top-performing teams achieved *MAE* scores ranging from 4.317 to 4.993, *RMSE* scores between 8.359 and 9.081, and *PC* values from 0.936 to 0.949. On the testing set, these teams achieved *MAE* scores ranging from 3.557 to 6.536, *RMSE* scores between 7.510 and 10.275, and *PC* values from 0.709 to 0.855.

### Estimation of pediatric bone age

The metrics and benchmarks for applying the proposed methodology to the pediatric bone age estimation use case can be derived directly from the publicly available data of the *RSNA Pediatric Bone Age Machine Learning Challenge*. In this challenge, multiple teams competed to provide the most

accurate estimation of bone age by employing machine learning algorithms on the wrist or *XR* images. The metrics under analysis are summarized in Table 3.13.

Metric	Description	Interpretation	Benchmark
Mean Absolute Error (MAE)	Mean absolute difference between the estimated value and the real one	Lower values indicate better performance	3.79-4.55 months

Table 3.13: Metrics estimated for performance analysis.

To determine the final ranking, the various solutions have been assessed based on the *MAE* between the estimated and actual values. In the ultimate ranking, the top-performing teams achieved *MAE* values ranging from 4.2 to 4.55 months, while an ensemble of all the solutions achieved a *MAE* of 3.79 months.

#### Quantification of calcium score

The assessment of the *CAC* score is a method employed to quantitatively measure the severity of atherosclerotic plaque deposition in the coronary arteries. In clinical practice, this score serves as a valuable tool for monitoring disease progression and distinguishing between individuals with varying levels of cardiovascular risk. Specifically, the score is utilized as a discriminative factor by comparing it to a predetermined threshold, enabling the classification of patients into binary categories. Typically, a threshold value of 0 is employed, wherein patients with no detectable calcium are categorized as healthy, while those with positive scores are deemed at risk. Notably, positive scores are associated with a low prevalence of obstructive *Coronary Artery Disease – CAD*, nonobstructive *CAD*, and a relatively low annualized risk of major adverse cardiac events (Agha et al. [2022]). The proprietary nature of the dataset employed in this analysis, coupled with the scarcity of publicly available datasets, presents a challenge in conducting a comprehensive comparison between the proposed methodology and existing literature solutions. However, by employing similar evaluation metrics, it is possible to partially bridge this gap and establish a common framework for analysis. To compute these metrics, a separate testing set consisting of 90 samples randomly selected using stratified sampling has been utilized. The metrics and corresponding benchmark values adopted in this particular use case are succinctly summarized in Table 3.14.

Metric	Description	Interpretation	Benchmark
Area Under Curve – AUC	The integral of the sensitivity-specificity curve across the entire range of decision thresholds	Higher values indicate superior classification performance	73.0% (69.0-76.0)%
Sensitivity	The proportion of patients with non-zero <i>CAC</i> scores correctly identified	Higher values signify the method’s ability to detect a larger portion of positive cases	71.6% (62.5-77.9)%
Specificity	The proportion of patients with zero <i>CAC</i> scores correctly identified	Higher values indicate the method’s ability to identify a larger majority of negative cases	63.0% (55.7-69.9)%
Median Error – MdE	The median difference between actual and predicted <i>CAC</i> scores, divided into different value ranges	Lower values indicate more accurate estimations provided by the method	N.A.

Table 3.14: Metrics estimated for performance analysis.

Among the existing literature, the works of [Kamel et al. \[2021\]](#) and [D’Ancona et al. \[2023\]](#) examine the performance of their respective solutions in discriminating between zero and non-zero cases. Both studies consider the *Area Under Curve – AUC* as a classification performance metric, while the latter also evaluates sensitivity and specificity. The obtained *AUC* value of 0.73 is consistent across both works, with a 95% confidence interval of (0.69-0.76). The other metrics, sensitivity and specificity, result in 0.716 (0.625-0.779) and 0.630 (0.557-0.699), respectively.

Due to the lack of previous studies analyzing the regression task of predicting the exact value of the *CAC* score, it is challenging to establish metrics and reference values directly. Furthermore, the dataset’s high imbalance makes it difficult to apply common metrics such as *MAE* or *PC* due to the presence of a few values with very high *CAC* scores and the majority of values concentrated near zero, which can yield unreliable results. To address this issue and facilitate discussion on the proposed methodology’s performance in the regression task, the test set samples have been divided into subsets based on the following ranges: [0,10], (10,200], (200, 500], and > 500. These ranges were determined by physicians considering different levels of cardiovascular risk and the available data distribution. For each

range, a separate *Median Error – MdE* is computed to provide a robust indication of the predictive performance within each risk class.

### 3.6.2 Performance

The performance achieved by applying the proposed methodology to the presented use cases is briefly summarized in Table 3.15. It is worth noting that the method consistently outperformed the benchmark in the majority of use cases while demonstrating competitive performance in the remaining ones.

Use case	Metric	Result	Benchmark
Radiology diagnostic workflow optimization	<i>MAP</i>	ImageNet: <u>71.8%</u> <i>SSL</i> : <u>89.3%</u>	47.8% (40–48%)
	Sensitivity	ImageNet: <u>78.2%</u> <i>SSL</i> : <u>92.9%</u>	61.1% (55–67%)
	Specificity	ImageNet: <u>77.2%</u> <i>SSL</i> : <u>94.9%</u>	63.0% (40–89%)
Monitoring of COVID19 infection severity, <i>XR</i>	<i>MAE</i>	<u>1.550</u>	1.787
Monitoring of COVID19 infection severity, <i>CT</i> , Validation	<i>MAE</i>	4.912	4.17-4.99
	<i>RMSE</i>	8.700	8.359-9.081
	<i>PC</i>	0.943	0.936-0.949
Monitoring of COVID19 infection severity, <i>CT</i> , Testing	<i>MAE</i>	5.020	3.56-6.53
	<i>RMSE</i>	9.006	7.510-10.275
	<i>PC</i>	0.798	0.709-0.855
Estimation of pediatric bone age	<i>MAE</i>	5.291	3.79-4.55
Quantification of calcium score, zero score classification	<i>AUC</i>	<u>81.4%</u>	73.0% (69.0-76.0)%
	Sensitivity	<u>89.3%</u>	71.6% (62.5-77.9)%
	Specificity	<u>67.6%</u>	63.0% (55.7-69.9)%
Quantification of calcium score, score estimation	<i>MdE</i>	[0,10]: 2.6 (10,200]: 103.6 (200,500]: 306.3 (500, +inf): 769.6	N.A.

Table 3.15: Best performance obtained for all the use cases. Underlined the results that are better than benchmark

The performance evaluation of the proposed methodology in optimizing the diagnostic workflow for *COVID-19* involves two distinct scenarios. In

the first scenario, the feature extraction model is initialized with weights from Imagenet using *Transfer Learning*, without further training. In the second scenario, starting from the same initialization, the model undergoes *Self-Supervised Learning* as part of an AutoEncoder.

In both cases, the sorting of cases produced by the methodology exhibits significant improvements compared to the *FIFO* policy, with a substantial margin. As expected, the model trained with AutoEncoder achieves superior performance compared to the Imagenet-initialized model. This notable advantage is also observed when employing the priority index to trigger alarms in the context of binary classification. Notably, both sensitivity and specificity can be optimized to surpass the benchmark performance when setting an appropriate threshold.

To assess the effectiveness of the prioritization induced by the proposed methodology, a dedicated visualization technique has been introduced. This visualization is referred to as the *Priority Matrix*, where the generated queues are represented as columns, and each cell corresponds to a patient. The colour of each cell indicates the patient's class membership: black represents positive cases for *COVID-19*, grey represents patients affected by other diseases, and white represents healthy individuals. As a result of the six iterations of stratified five-fold cross-validation, a total of 30 different queues are generated, corresponding to the 30 columns in the matrix. An optimized diagnostic queue is expected to exhibit a *Priority Matrix* where the majority of black cells are positioned at the top, indicating that *COVID-19* positive patients are prioritized over others. Conversely, randomly distributed colours in the matrix indicate poor queue optimization. Figure 3.33a illustrates the starting point using the benchmark method, *FIFO* policy, where the colours are randomly distributed and no discernible patterns are visible. By applying the prioritization from the proposed methodology in the scenario utilizing ImageNet *Transfer Learning*, as shown in Figure 3.33b, it becomes evident that black cells tend to occupy the upper portion of the matrix, indicating the prioritization of *COVID-19* positive patients. Finally, in the scenario employing *Self-Supervised learning*, depicted in Figure 3.33c, a much higher number of black cells are positioned in the upper part of the matrix, clearly demonstrating the effective sorting induced by the proposed methodology. In line with expectations and *MAP* values, the prioritization performance of the scenario utilizing *Self-Supervised learning* surpasses that of the ImageNet *Transfer Learning* approach.




Figure 3.33: Priority Matrices: a) *FIFO* policy, b) ImageNet *Transfer Learning* and c) *Self supervised learning* AutoEncoder

The performance evaluation of estimating the severity of *COVID-19* infection through the analysis of *CXR* images has been conducted on the subset of the *BrixIA* dataset known as the *Global consensus score*. This evaluation compares the predicted label scores generated by the proposed methodology with the actual label scores. The *MAE* is computed to quantify the disparity between the predicted and actual values. The proposed methodology exhibits a significant improvement over the benchmark method, demonstrating an *MAE* of 1.550 compared to the benchmark's *MAE* of 1.787. This

substantial margin suggests an enhanced capability in accurately quantifying the severity of the infection.

To facilitate a comprehensive interpretation of the results, visualizations have been generated. A scatter plot, depicted in Figure 3.34, illustrates the correlation between the predicted values and the actual labels. The predictions display a strong positive correlation with the true values, with a distinct separation between low and high scores. However, for extremely high values, the system tends to underestimate the severity of the infection. This discrepancy may be attributed to the under-representation of extreme scores in the training set. Specifically, out of the total 3642 samples in the training set, only 286 samples, accounting for less than 7.9%, have labels equal to or higher than 16.

Further analysis of the error distribution, as depicted in Figure 3.35, reveals a concentration of errors around zero. The majority of errors fall within the range of -2 to +2, aligning with the  $MAE$  value of 1.550 obtained.




Figure 3.34: Scatter plot of predictions for BrixIA dataset

The *Per-COVID-19* challenge entails the development of computer vision solutions for estimating the severity of *COVID-19* infection from horizontal *CT* scans.

In the initial phase of the challenge, the participating teams were provided with training and validation sets. The training set contained labelled slices, which the teams utilized to design and train their methods. Conversely, the




Figure 3.35: Error distribution of predictions for BrixIA dataset

validation set was released without labels. Therefore, evaluating the performance necessitated the following steps: i) executing the proposed solution on the images in the validation set to generate infection estimations for each scan, ii) compiling a CSV file containing the produced estimations and iii) submitting the CSV file to an online automated platform, which computes performance metrics and generates a leaderboard of the teams. During this first phase, teams had the opportunity to submit estimations multiple times, enabling them to experiment with different approaches and identify the most effective one. After the initial phase, the proposed methodology achieved a notable rank of 5<sup>th</sup> among the participating teams, demonstrating competitive results across all three metrics. Notably, the *PC* exhibited a high value of 0.943, closely approaching the performance of the top-performing team (0.949).

The favourable ranking achieved during the first phase enables the proposed methodology to undergo evaluation on a substantially larger and more diverse testing set in the second phase. Analogous to the validation set, the testing set was made available with only the images, devoid of any labels. However, unlike the first phase, only a single submission was permitted,



without access to the automated online platform. The computation of metrics was carried out by the organizers, and once again, the proposed methodology attained the 5<sup>th</sup> position among all participating teams, demonstrating competitive performance.

The rankings from the first and second phases can be observed in Table 3.16 and Table 3.17 respectively. Notably, there is a change in the performance of all teams, with a significant drop in terms of *PC* between the two phases. Additionally, it is significant to observe that among all the teams, the proposed methodology demonstrates the smallest disparity in *MAE*, implying a commendable level of generalization, particularly considering the more diverse composition of the testing set in comparison to the validation set. Unfortunately, due to the absence of ground truth in both the validation and testing phases, further exploration of our results was limited.

Position	Team	MAE	PC	RMSE
1	SenticLab.UAIC	4.317	0.947	8.359
2	Captain-CSgroup	4.479	0.947	8.406
3	TAC	4.484	0.946	8.547
4	Taiyuan_university_lab713	4.504	0.949	8.097
<b>5</b>	<b>EIDOSlab_Unito (ours)</b>	<b>4.912</b>	<b>0.943</b>	<b>8.700</b>
6	IPLab	4.953	0.944	8.604
7	ACVLab	4.993	0.936	9.081

Table 3.16: Final competition ranking for Validation dataset. In **bold** the result related to the presented approach

Position	Team	MAE	PC	RMSE
1	Taiyuan_university_lab713	3.557	0.855	7.510
2	TAC	3.645	0.802	8.571
3	SenticLab.UAIC	4.617	0.763	9.100
4	ACVLab	4.866	0.729	10.275
<b>5</b>	<b>EIDOSlab_Unito (ours)</b>	<b>5.020</b>	<b>0.798</b>	<b>9.006</b>
6	Captain-CSgroup	5.168	0.772	8.392
7	IPLab	6.536	0.709	9.976

Table 3.17: Final competition ranking for Test dataset. In **bold** the result related to the presented approach

The proposed methodology is employed to estimate pediatric bone age through the analysis of wrist *XRs*. The system is trained to provide skeletal status estimations for each image. To assess the accuracy of the predictions,

the proposed methodology is evaluated using data from the *RSNA Pediatric Bone Age Machine Learning* challenge.

The results obtained from the proposed methodology demonstrate competitiveness when compared to other methods proposed by participating teams in the challenge. Although the achieved *MAE* value is lower than that of the top performers, the difference from the best solution is within 1.5 months. The effectiveness of the proposed methodology is evident in the scatter plot depicted in Figure 3.36, where a strong positive correlation between the predicted and actual bone age is observed. The Pearson correlation index between the predicted and actual values is calculated to be 0.9878. The distribution of errors, as illustrated in Figure 3.37, is concentrated around zero, with the majority of errors falling within the range of -10 to +10, consistent with an *MAE* of 5.291. Additionally, it is noteworthy that the distribution is skewed towards negative values, indicating a tendency of the system to overestimate age.




Figure 3.36: Scatter plot of predictions for pediatric bone age

The proposed methodology has been trained to provide estimations of the *CAC* score from each *CXR* image. In particular, the best performance has been obtained by initializing the feature extraction model with the weights from the *Self-supervised* AutoEncoder adopted for the *COVID-19* diagnostic workflow optimization use case. The predicted value serves two purposes: distinguishing between patients with a low risk of cardiovascular




Figure 3.37: Error distribution of predictions for pediatric bone age

issues and those with a higher risk, as well as accurately monitoring the deterioration of heart conditions. To assess the performance of the proposed methodology, it has been tested on a randomly selected subset of samples, ensuring stratification during the sampling process.

In the case of discrimination, the benchmarks are surpassed by a significant margin through the implementation of the proposed methodology. The resulting confusion matrix, depicted in Figure 3.38, reveals a noticeable 25% increase in Sensitivity compared to the reference. The considerably higher *AUC* serves as an indication of the proposed methodology's ability to be finely tuned for Sensitivity-Specificity operating points that outperform the benchmarks. The Sensitivity-Specificity curve, which depicts the variation of the decision threshold, is presented in Figure 3.39.

When the proposed methodology is employed to track deterioration or conduct quantitative analysis and sorting, it exhibits a low *MdE* of 2.6 for samples with *CAC* scores in the low-risk range (below 10). However, in the other score groups, the *MdE* is significantly higher, particularly when compared to the group boundaries. This occurrence can be attributed to the dominant prevalence of low values within the provided dataset. Specifically, out of the total dataset, 186 out of 505 samples possess a *CAC* score equal to zero, while the remaining values are concentrated within the <1000 range, as depicted in Figure 3.19. A more comprehensive understanding of

estimation errors is offered by Figure 3.40. This visualization confirms that a majority of the errors between medical annotations and predictions from the proposed method lie close to zero, primarily between 0 and 300. Additionally, there are noticeable outliers with extremely high values in the thousands range. These outliers contribute to the observed higher  $MdE$  for groups with elevated  $CAC$  scores.




Figure 3.38: Scatter plot of predictions for pediatric bone age

### 3.6.3 Preliminary “Real Life” Results

The performance and potential benefits of the proposed methodology, applied to optimize the diagnostic workflow for *COVID-19*, have been examined in a local hospital during the period between 9<sup>th</sup> April 2021 and 4<sup>th</sup> June 2021, amidst the ongoing pandemic contagion wave. This section presents and discusses the aspects of the integration of such a system into the hospital’s IT environment, as well as the resulting performance.




Figure 3.39: Sensitivity-Specificity curve for pediatric bone age

### System Integration and Installation

To establish a realistic and representative scenario for a real-life setting, various factors related to data protection, system interfaces, and user ergonomics must be taken into account. These aspects have been addressed during the trial conducted at the local hospital, with the active involvement and guidance of physicians.

In April 2021, a prototype version of a machine named *Aippo* was installed within the hospital network at *AOU San Luigi Gonzaga* hospital ([Università degli studi di Torino \[2020\]](#)). A high-level representation of the system is depicted in Figure 3.41. Simplifying the hospital’s IT structure for radiology imaging management, can be summarized as follows:

1. Radiology units produce the images using imaging machines.




Figure 3.40: Error distribution of predictions for pediatric bone age

2. These machines are configured to transmit the acquired scans to a centralized data collection system known as the *Picture Archiving and Communication System – PACS*. The *PACS* system provides storage and access services to users, utilizing the *Digital Imaging and Communication in Medicine – DICOM* format.
3. Radiologists access the diagnostic worklist through the *Radiological Information System – RIS*, which supports scheduling, reporting, and tracking of examinations. The worklist contains a list of *XR* scans that need to be analyzed and interpreted to generate the final report.
4. A direct link between *RIS* and *PACS* enables physicians to access the images directly from the *RIS* platform.

To integrate the proposed methodology into the aforementioned environment, it has been implemented as a composition of services and deployed on a machine within the hospital’s *Local Area Network – LAN*, accessible by other network nodes. Initially, an inbound service was created to function as a *PACS*, facilitating the reception of *DICOM* modalities from radiology units. Three machines, produced from two different vendors and utilized during the *COVID-19* pandemic, were configured to send all generated samples to both the hospital’s *PACS* and the implemented service. Upon receipt, the inbound service examines the received images to determine their relevance for the experimental purpose. Specifically, it excludes

images that fail to meet the specified criteria for body part (CHEST) and position (AP/PA). For images satisfying these criteria, they are transmitted to the examining service using specialized *Application Programming Interfaces – APIs*. The examining service implements the proposed methodology, leveraging the received images to predict the prioritization *COVID-19* risk indicator for each provided picture. Subsequently, the results are communicated to the outbound service, which, through dedicated *APIs*, transmits them to the *RIS* system. This crucial step was made possible through collaboration with the *RIS* provider, who developed the necessary *APIs* to receive the system’s predictions.

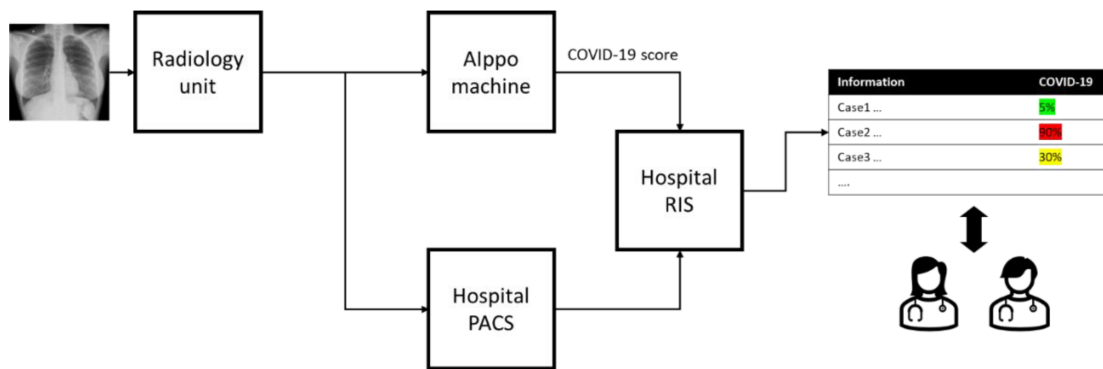


Figure 3.41: Overview of the integration between the machine running an implementation of the proposed methodology and existing IT environment in the hospital

The task of radiologists requires a high level of concentration, and the complexity of managing multiple systems and applications can significantly increase their cognitive load, potentially leading to a higher number of errors that can impact patient health. To mitigate this effect and integrate the proposed methodology into the diagnostic workflow, the prediction is incorporated into the *RIS* dashboard as a column. This enables users to easily sort the worklist based on the ranking provided by the proposed methodology with just one click. To assist physicians in quickly identifying the most critical cases, a colour scheme is associated with different ranges of predictions. The highest values are represented by the colour red, the middle values by yellow, and the lower range by no colour. This approach seamlessly integrates the prediction without introducing any new windows or applications to the existing system. Importantly, this integration has been approved by the medical professionals at the hospital, ensuring its acceptance and usability.

Unlike the cloud-based service paradigm, the integration of the system into the hospital LAN, hosted on a physically located machine within the facility, eliminates the need to transfer sensitive data over the internet or store it on external servers not managed by the hospital IT team. This approach minimizes the risk of data leakage and ensures that information is not unintentionally shared with third parties. By keeping the system within the hospital's local network, data security and privacy are enhanced, providing greater control and mitigating potential vulnerabilities associated with external cloud services.

### Evaluation method

Due to time and organizational limitations, it has not been possible to continuously monitor the system's performance by comparing its output with real antigen test outcomes. This constraint arises from the lack of integration between the system managing antigen test information and the radiological system, which necessitates manual matching of antigens to *CXR* images. Furthermore, access to antigen results is limited to sanitary operators, precluding the possibility of automating or delegating the task to others. Considering the substantial number of scans generated during the observation period, amounting to 2942, and the complexity of the task, it has proven unmanageable.

To address this limitation and gain insights into the system's performance in a real-life setting, the priority *COVID-19* index score was compared to the local infection incidence. This comparison aimed to assess whether the tool followed the trend observed during the study period, as summarized in Figure 3.42. The internal database of the Aippo machine, containing prediction values from 9<sup>th</sup> April 2021 to 4<sup>th</sup> June 2021, was analyzed alongside publicly available and daily updated data on *COVID-19* incidence in the *Provincia di Torino* (Ministero della Salute [2023]), the local district where the *AOU San Luigi Gonzaga* of Orbassano is located. The *COVID-19* incidence trend is depicted in Figure 3.43. Both trends were subjected to a 7-day rolling average filter and subsequently compared quantitatively using the *Pearson correlation* – *PC* coefficient and qualitatively through plots.






Figure 3.42: Overview of the evaluation method for diagnostic workflow optimization trial at hospital.

### Performance

During the observation period, a total of 2942 *CXR*s from the three X-ray units were subjected to analysis using the proposed method, resulting in the generation of priority *COVID-19* index score for each *CXR*. The correlation coefficient between the 7-day rolling averages of the priority *COVID-19* index score and the local pandemic trend is found to be 0.873 (p-value <  $1 \times 10^{-5}$ ).

Moreover, it is observed that the incidence curve of *COVID-19* and the average prediction generated by the proposed methodology exhibit a concurrent decrease throughout the evaluation period, as depicted in Figure 3.44. This




Figure 3.43: COVID-19 contagions in Provincia di Torino district during the period of interest.

observation provides further confirmation of the previously identified high correlation.

### 3.7 Discussion

The subsequent sections delve into various facets concerning the results and application of the proposed methodology to different use cases. The dataset imbalances about the various use cases will be examined, and potential remedies for mitigating their impact on performance will be discussed. In-depth insights will be provided regarding its application in the prioritization use case. This includes an analysis of the impact of different training strategies on the structure of the feature space. Furthermore, the system's performance in scenarios with limited labelled data will be presented. Additionally, a comprehensive examination of the results will shed light on the role of *Other* samples in instances where the method encounters difficulties. As for the other use cases, the evolution during training will be showcased, emphasizing areas that still offer room for improvement.

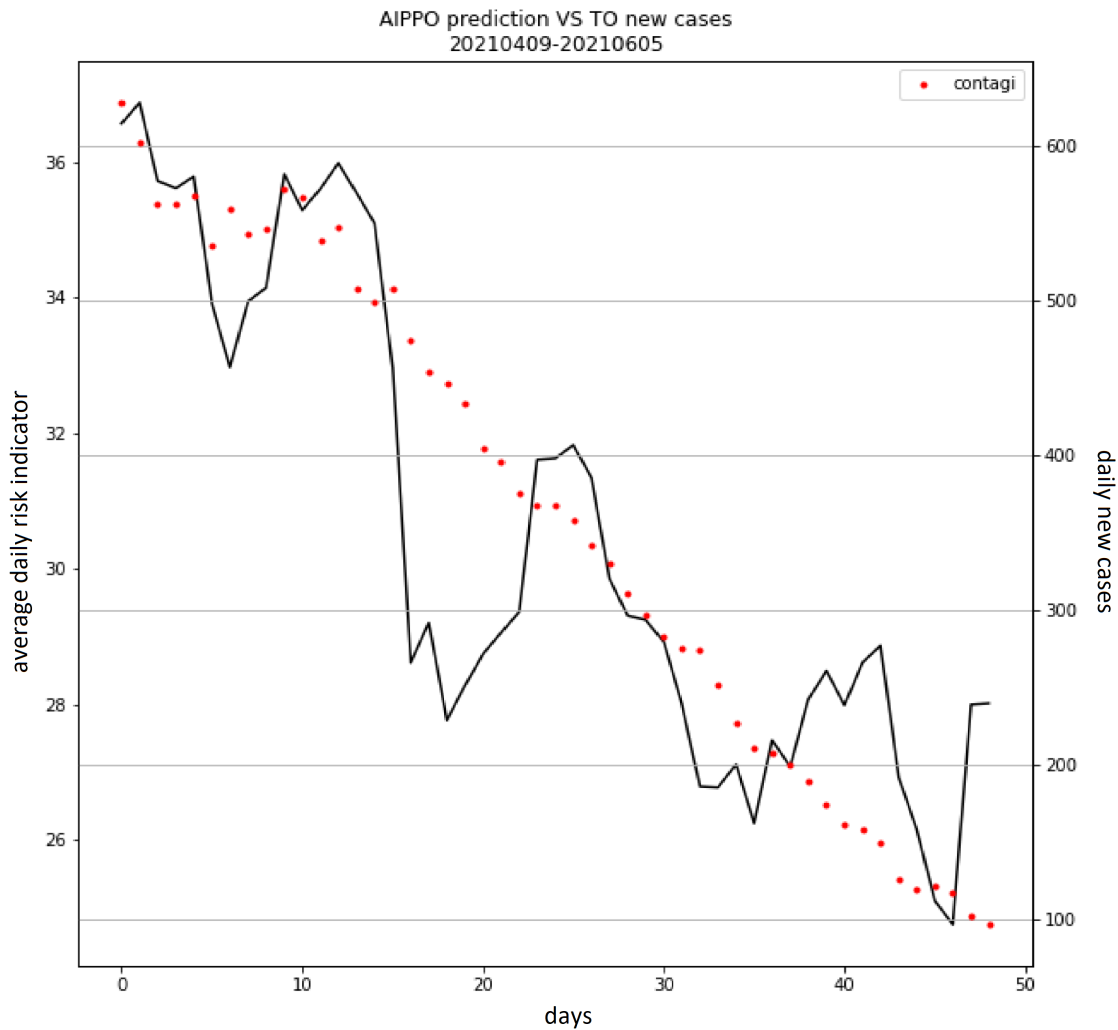


Figure 3.44: Concurrent 7 days rolling average trends of pandemic decreasing contagion wave and priority *COVID-19* index score elaborated by the proposed method, during the period of interest.

### 3.7.1 Dataset unbalance

Dataset unbalance is a prevalent issue in the development of automatic tools for the analysis of Medical Imaging. This imbalance can arise due to two primary reasons: the low prevalence of a disease or extreme conditions within the population, resulting in a scarcity of positive cases, or the sampling strategy, which may prioritize the retrieval of positive cases, leading to an excess of such cases within the dataset.

Among the use cases presented in this study, only the one related to the prioritization of *COVID-19* positive cases in the diagnostic workflow

demonstrates a relatively balanced distribution between affected and non-affected cases, with positive samples accounting for 43.3% of the total. When considering the additional division between *Other* and *Healthy* cases, they constitute 30.0% and 26.7% of the dataset, respectively. This acceptable class balance supports the proposed methodology in achieving higher performance, resulting in balanced Sensitivity and Specificity.

In contrast, the two use cases for *COVID-19* infection severity quantification from *CXR* and Bone age estimation from hand *XR* reveal the effects of underrepresented extreme values, as evident in the scatter plots in Figure 3.36 and Figure 3.38. In the *BrixIA* dataset, values rated with scores lower than 2 and higher than 15 are underrepresented, while in the *RSNA Pediatric Bone Age Machine Learning challenge* dataset, a similar pattern is observed for samples with labels below 9 years old and above 15 years old. In both cases, the system tends to overestimate samples with the lowest values and underestimate those with the highest. The lack of extreme values in the training set makes the model less aware of patterns associated with such cases. Furthermore, during the distance-based regression step, fewer high and low-scored samples are available in the reference set, compromising the ability of the proposed methodology to predict extreme values.

In the case of *CAC* score prediction, dataset imbalance is even more pronounced than in other use cases, with 36.8% of all values equal to zero, while almost all the others are lower than a score of 1000. Although the resulting performance in identifying high-risk patients surpassed the benchmark with promising results, an analysis of regression estimation errors across different score ranges reveals the disturbing impact of skewed data distribution. Within the range of scores between 0 and 10, where the majority of data resides, the *MdE* is only 2.6, which is compatible with the range limits. However, when assessing performance in higher score ranges, this metric increases to unacceptable values, reaching up to 769.6 in the set of samples with scores exceeding 500.

The impact of data imbalance on the performance of the proposed methodology varies in both its nature and extent across different use cases. The scientific literature has presented a plethora of techniques aimed at mitigating this well-recognized issue.

One approach to address the uneven distribution of data involves direct intervention, achieved by undersampling the majority class or oversampling the minority class, respectively (Koçak [2022], El Naqa et al. [2021], Wernick et al. [2010]). These two simple techniques can yield favourable outcomes,

but they may also have a negative influence on overall performance. In the case of undersampling, it entails discarding a portion of the data (Brownlee [2020]), resulting in an associated loss of valuable information and introducing arbitrary elements, such as the selection of the discarded samples. On the other hand, oversampling may lead to overfitting (Ellis [2023]) as it reintroduces the same samples multiple times during the training process. To mitigate these limitations, the *Synthetic Minority Oversampling TEchnique* — *SMOTE* has been proposed (Chawla et al. [2002]). *SMOTE* creates synthetic data points based on the original data points by interpolating between two samples. While this effectively augments the number of minority class samples, which can enhance the performance of machine learning models, it may inadvertently oversample uninformative or noisy data points (Jiang et al. [2021]) and generate artificial data points that might not accurately represent the true distribution of the minority class (Wikipedia contributors [2023a]), especially when dealing with non-uniform and discontinuous distributions across the sample space, as it is the case of images.

A more sophisticated approach to generating synthetic data for underrepresented classes involves the usage of *DL* models to create artificial medical images tailored to address class imbalances. *Generative Adversarial Networks* – *GANs* (Goodfellow et al. [2020]) represent a prominent category of deep learning algorithms frequently employed for this purpose within the domain of Medical Imaging. Nevertheless, this technique has its limitations, including the critical requirement to ensure that the synthetic data closely mirrors real-world data and the potential risk of overfitting the model to the synthetic dataset (Ricci Lara et al. [2022]).

An alternative strategy for mitigating data imbalances, without directly modifying the available samples, lies in adjusting the training procedure. One viable approach involves the adaptation of the loss function to compensate for class imbalances when computing the overall loss value. This adaptation causes the loss function’s gradient to assign greater significance to samples from the minority class, thus rendering the model more sensitive to such instances of misprediction. In the context of the proposed *DISTMAT* loss function, this compensation can be realized through the introduction of a weighted *MAE* computation between the two distance matrices, accounting for label distribution.

### 3.7.2 Other cases in the prioritization use case

Upon observing the results obtained from the application of the proposed methodology to optimize the diagnostic workflow for *COVID-19* in Figure 3.33, it becomes apparent that the system-induced automatic prioritization places the black cells, associated with positive patients, predominantly at the top of the matrices. This positioning indicates the prioritization of these cases over the healthy ones and those affected by other conditions. Additionally, in the *Self-supervised learning – SSL* approach, it is also notable that the queues tend to prioritize Other cases over *Healthy* cases. Differently, in the *ImageNet* scenario, this distinction is less evident, with *Other* samples distributed throughout the worklists. While it may be considered acceptable in certain cases to prioritize patients affected by diseases other than *COVID-19*, the selection of performance metrics and the definition of the use case penalize and adversely affect the performance values when patients with other conditions are prioritized alongside *COVID-19* cases.

From a study perspective, this behaviour is of particular interest because neither the feature extraction model nor the distance-based regression steps have been specifically trained, designed or configured to manage *Other* class differently from the *Healthy* one.

To gain further insights into this aspect, the angular similarity within and between classes has been computed and visualized in Figure 3.45a and Figure 3.45b for the *ImageNet* and the *SSL* scenarios, respectively. The impact of training can be observed in the case of the *SSL* approach: there is a decrease in the overall similarity between samples, and the distribution of *Healthy* and *Other* cases (inter-class similarity for positive cases) becomes more distinct from the *COVID-19* cases (intra-class similarity).

However, the observed phenomena cannot be solely explained by the differences in the distribution of *Healthy* and *Other* case distances. To further investigate this, the features extracted from the images in the reference set have been visualized using TSNE dimensionality reduction. The resulting plots are presented in Figure 3.46a and Figure 3.46b for the *ImageNet* and *SSL* scenarios, respectively. In these visualizations, *COVID-19* cases are represented by a red cross, *Healthy* cases by a blue dot, and *Other* cases by a green plus sign. The background colour segments the plane into areas of influence between the different classes, indicating that a sample in that area is more similar to other samples from one specific class compared to the others. Specifically, green zones represent the *COVID-19* class, yellow zones




Figure 3.45: Angular similarity within and between classes: a) ImageNet *Transfer Learning* and b) *Self supervised learning* AutoEncoder

represent *Other*, and red zones represent *Healthy*. Analyzing Figure 3.46a, it can be observed that features from *Other* images are sparsely distributed across the entire space with no compact dense zones of influence. The separation between *COVID-19* and *Healthy* cases is evident, but many cases are still located in the zone of influence of the wrong class. On the other hand, with the *SSL* approach, the distribution of features from *Other* images is more compact and predominantly situated in the areas between *COVID-19* and *Healthy* groups. Additionally, the separation between *COVID-19* and *Healthy* samples is more pronounced. These visualizations align with the observed prioritization in the *Priority Matrices*. In the *SSL* approach, by placing *Other* case features in the zones between the other two classes, they appear more similar to *COVID-19* cases than *Healthy* cases, resulting in higher scores. Differently, in the ImageNet scenario, where *Other* case features are sparser across the space, they are more uniformly distributed across the rankings.

### 3.7.3 Reference set depletion robustness

In the context of applying data-driven methods to medical imaging, one of the major challenges is the limited availability of well-curated datasets. Specifically, certain modalities such as *XRs* and *CT* scans may have a large number of images, but the process of labelling them is time-consuming and requires highly skilled personnel. Additionally, in many real-world cases, the lack of IT automation in hospitals makes it difficult to efficiently link images with their associated diagnostic reports, which are often written in a free-text format and contain medical terminology. Furthermore, disease




Figure 3.46: TSNE dimensionality reduction of features extracted from the images in the reference set using: a) ImageNet *Transfer Learning* and b) *Self-supervised learning* AutoEncoder

cases represent a minority within the dataset, resulting in imbalanced distributions. Consequently, researchers developing automatic solutions in this field often have access to only small datasets. To address this limitation, the adoption of the *ImageNet* and *SSL* approaches proves beneficial in mitigating the need for large labelled datasets to train the feature extraction model in the proposed methodology. In the *ImageNet* approach, no training is performed, while in the *SSL* approach, model training does not rely on labels but instead leverages a substantial corpus of unlabeled images. However, labelled data remains necessary for the distance-based regression step, as part of the reference set.

To assess the resilience of the proposed methodology to data depletion in the reference set, an analysis has been conducted to evaluate the performance degradation when the system is tested with progressively reduced fractions of labels. Specifically, the performance has been measured in terms of *MAP* for the prioritization task and *AUC* for identifying critical cases. To perform this analysis, ten different subsets, each representing a defined fraction size (100%, 50%, 25%, 10%, and 5%), have been randomly sampled from the original reference set. These subsets are then utilized to make predictions using the proposed methodology. The final performance is computed by averaging the results obtained from the ten subsets.

The results are presented in Figure 3.47 for the *MAP* and in Figure 3.48 for



the *AUC*. Additionally, Figure 3.49 and Figure 3.50 illustrate the *AUC* results obtained with reduced datasets for the *ImageNet* and *SSL* approaches, respectively. As expected, the performance shows a decline as the available data is reduced. However, both approaches demonstrate limited degradation when the data size is reduced to 50% and 25%. The decay in performance becomes more pronounced when only 10% and 5% of the original labels are provided to the system. Notably, the *SSL* approach consistently outperforms the *ImageNet* approach, with a significant observation where *SSL* with only 10% of labelled data performs better than *ImageNet* with the entire original reference set. This finding further confirms the advantages of the *SSL* approach in this context.




Figure 3.47: System prioritization performance with a reduced reference set.

### 3.7.4 DISTMAT loss training

In the preceding sections, the *DISTMAT* loss function was introduced and its application in the proposed methodology for various use cases was explained. This technique has consistently exhibited remarkable performance across all use cases, surpassing existing state-of-the-art approaches or yielding competitive outcomes.

Incidentally, certain aspects of the training process justify further investigation. Figure 3.51 depicts the distinct patterns of loss trends observed




Figure 3.48: System critical case identification performance with a reduced reference set.




Figure 3.49: Sensitivity-Specificity curves for identification task performance with a reduced reference set using *ImageNet* approach.

across different use cases. Notably, when training on the *BrixIA* dataset, the model exhibits a significant disparity between the training and validation sets, typically indicative of overfitting. A similar trend is observed with the *CAC score* data. Moreover, the loss progression displays considerable noise, likely attributable to the limited data availability (only 505 samples) and the imbalanced labels. Interestingly, despite the widening gap between validation and training sets over time, both cases demonstrate improved

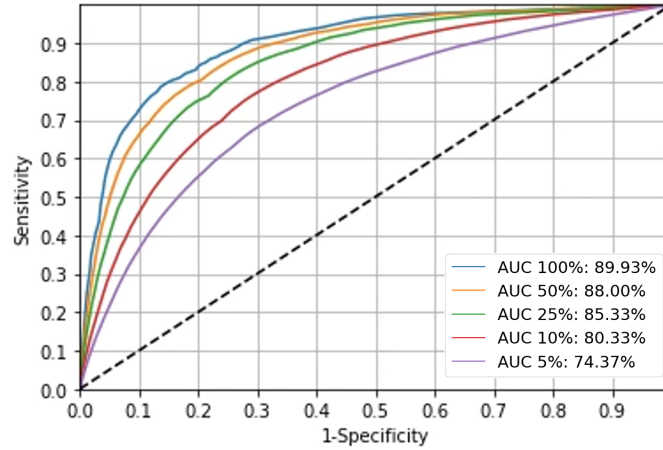




Figure 3.50: Sensitivity-Specificity curves for identification task performance with a reduced reference set using *SSL* approach.


performance on the test set during training. In fact, the optimal overall performance of the proposed methodology on the test set is achieved when the model reaches the later stages of training, coinciding with the largest disparity. One plausible explanation for this phenomenon is that a robust feature extraction on the training set, which is also employed as the reference set in the distance-based regression step, may produce additional benefits for final predictions. This particular aspect remains an open question and presents potential opportunities for improvement, to be discussed in the following chapter. On the other hand, when training on the *RSNA Pediatric Bone Age Machine Learning Challenge* dataset, the loss demonstrates a smooth progression, with a significantly reduced gap between the training and validation sets. However, due to the absence of labels for the validation and test sets, a similar analysis cannot be performed on the *Per-COVID-19 Challenge* data.



(a)



(b)



(c)

Figure 3.51: DISTMAT loss values during training: a) *CAC score* b) *RSNA Pediatric Bone Age Machine Learning Challenge* and c) *BrixIA*

## Chapter 4

# Conclusion

The degradation of systems is a pervasive phenomenon that is widely recognized as a fundamental aspect of global dynamics. Assessing the status of system degradation offers significant advantages, as it enables the prevention of failures and mitigates the occurrence of severe damages.

This chapter provides a comprehensive summary of the entire document, covering the topics discussed and highlighting the contributions made in both the *Predictive Maintenance* and *Medical Imaging* applications.

Lastly, an outlook is provided on the potential future developments of the proposed methodologies and the areas that deserve further investigation.

### 4.1 Predictive Maintenance

In Chapter 2, the degradation of artificial systems is described, and Predictive Maintenance is introduced as a technique for preventing system failures by scheduling maintenance actions based on system status and projected evolution. The advantages of this approach are discussed, along with various adopted methodologies. Furthermore, the state of the art in this field is presented.

As a contribution to the field of predictive maintenance, this thesis proposes a complete pipeline for predicting the degradation status of a system by analyzing measurements from onboard sensors. The proposed approach consists of several data-driven steps. Notably, a novel algorithm for signal selection, referred to as *CORR-FS*, is introduced to reduce input dimensionality. All stages of the proposed pipeline prioritize interpretability, enabling domain experts to understand the algorithm's logic and make corrections, or adjustments as necessary.

The proposed approach is evaluated on real-world automotive use cases, considering both component and system-level degradation estimation. In both scenarios, the proposed approach significantly exceeds the requirements specified by domain experts, demonstrating its considerable potential.

To facilitate the implementation of the proposed approach in vehicles, various implementation aspects are discussed, including memory and bandwidth requirements, and strategies to minimize them. Additionally, the impact of the proposed method on customer satisfaction for large vehicle fleets is analyzed, with an optimal strategy proposed for achieving favourable results.

In summary, the contributions to *Predictive Maintenance* made in this work are as follows:

- A data-driven pipeline for constructing a system to predict the degradation status of a system based on measured signals.
- Introduction of a novel, fully interpretable signal selection algorithm, *CORR-FS*, adjustable with a single parameter.
- Detailed explanation of the pipeline setup and its application to specific use cases.
- Experimental results on real-world automotive data, addressing both component and system-level degradation estimation.
- In-depth analysis and effective solutions for the practical implementation of the proposed method in large vehicle fleets.

## 4.2 Medical Imaging

Chapter 3 introduces the degradation of biological systems, highlighting the role of *Medical Imaging* as a valuable tool for assessing health status and investigating the internal structures of the human body. The fundamental principles and various modalities within *Medical Imaging* are described, accompanied by a brief overview of their applications.

As a significant contribution to this field, this thesis proposes a novel approach for degradation estimation from imaging data. The approach is derived from content-based retrieval and contains feature extraction using deep learning models, followed by distance-based regression that leverages a

reference set of labelled cases. Notably, a novel loss function, *DISTMAT*, is introduced to train the feature extraction model by incorporating distance matrices among computed features.

The proposed methodology undergoes testing on diverse use cases, addressing tasks such as diagnostic workflow prioritization, disease severity quantification, and bone age estimation. In most cases, the proposed algorithm outperforms existing state-of-the-art methods on the same tasks or datasets, demonstrating its significant potential. It is worth noting that while competitors rely on ad hoc solutions specific to their use cases, the proposed method remains unchanged, indicating its versatility across a wide range of applications.

Another notable aspect of the proposed solution is its interpretability. By examining the most similar images selected from the reference set, physicians can gain insight into the algorithm’s decision-making process and validate or correct the predictions. This establishes a user-machine interaction, wherein radiologists can assess and refine the selection of similar images, allowing the system to adapt its predictions accordingly.

Additionally, preliminary results from a real-life trial conducted at a local hospital during the *COVID-19* pandemic are presented. The adoption of the proposed methodology for prioritizing positive cases in the radiological diagnostic workflow demonstrates consistent performance, exhibiting a high correlation of 0.873 with the local infection trend, thus confirming the method’s potential. The practical aspects of the integration of the system within the hospital’s IT infrastructure and diagnostic process are comprehensively analyzed, and effective solutions are proposed.

Other aspects, such as the demand for big data, are also discussed, analyzing how the method’s performance varies with dataset size. Remarkably, the proposed methodology shows promising results, achieving significant performance even with small fractions of data. The novel *DISTMAT* loss function works state-of-the-art performance for both the *BrixIA* and *CAC score* use cases, while delivering competitive results for the *Per-COVID-19* and *RSNA Pediatric Bone Age Machine Learning* challenges.

Furthermore, an examination of the loss progression during training reveals a significant gap between the training and validation sets, indicating potential overfitting. However, this overfitting does not impact the performance on the test set, as the best results are consistently achieved at the end of the training process when the gap is the widest. One possible explanation for this phenomenon is that having a more robust feature extraction on the

training set benefits overall performance, as the same set is utilized as a reference for the distance-based regression step. The resolution of this open point, along with the explanation of the rationale behind the necessity for a relatively high learning rate ( $1 - 5 \times 10^{-2}$ ) to reach training convergence, holds the potential to augment the performance of the proposed method.

In summary, the contributions to the field of *Medical Imaging* presented in this work include:

- A novel approach for degradation estimation in images based on similarity in the feature space, yielding state-of-the-art performance across various use cases without requiring any ad-hoc core method modifications.
- The introduction of a novel loss function, *DISTMAT*, which incorporates Euclidean distance regularization of the feature space.
- A detailed description of how the method can be set up and applied to specific use cases.
- Preliminary results from a real-life trial conducted at a local hospital, showcasing the optimization of the diagnostic workflow.
- A proposal for user-machine interaction in prediction interpretation and correction.

### 4.3 Next steps

The methodologies proposed for both the *Predictive Maintenance* and *Medical Imaging* applications have exhibited remarkable performance, surpassing requirements and outperforming state-of-the-art approaches in most use cases. These compelling results encourage further investigation and advancement of both methods, intending to enhance performance and expand their applicability.

In the realm of *Predictive Maintenance*, while implementation aspects and solutions have been discussed, the performance evaluation has been conducted on data obtained from real systems within a controlled environment. As a future step in this direction, there is the consideration for applying and testing the methodology in an "on-the-road" scenario, possibly leveraging remote telemetry systems. Given the higher variability inherent in road environments compared to controlled settings, an additional context identification system may be necessary to transition between models



that better suit different situations (e.g., urban versus rural, high loads). Furthermore, validating the approach in non-automotive domains will be crucial to ascertain its versatility across diverse systems. To complete the analysis, it is necessary to integrate a more precise and detailed economic model into the decision-making strategy, enabling a finely tailored policy design that maximizes benefits for both the carmaker and the owner.

To further advance the proposed methodology in *Medical Imaging*, it may be important to evaluate its performance across a wider range of use cases, including both medical and non-medical domains. Achieving successful results in non-medical applications would demonstrate the extensive applicability of the proposed approach to a broader array of vision regression tasks. Additionally, exploring alternative feature extraction models beyond the current *DenseNet-121* should be considered to identify the most suitable solution.

As indicated to in previous sections, certain aspects of the *DISTMAT* loss function training necessitate more in-depth investigation to unlock additional improvements to the already relevant results. Among these aspects is the integration of a weighted *MAE* to address the class imbalance present in the dataset.

The formulation of the *DISTMAT* loss function enables cross-modality training, as evidenced by the definition of  $D_y$  in Equation 3.7, which accommodates labels expressed as vectors. Consequently, the loss function can be applied to cases where  $y$  is or can be represented as, a numerical vector (e.g., tabular data, text, sound, other images). For instance, it may be feasible to directly train the *CXR* feature extraction model using radiologists' text reports, converting them into numerical vectors through text embedding extraction techniques from *Natural Language Processing – NLP* models. Recent advancements in *NLP* and the availability of specialized *Large Language Models – LLMs* trained with medical data offer promising opportunities for conducting such experiments. As a future step, there is consideration for conducting a text-image experiment using a substantial dataset comprising 377,110 *CXR* images paired with 227,835 *Radiological reports* (Johnson et al. [2019]), and leveraging specialized *LLMs* (Boecking et al. [2022]) to extract the embeddings. The accomplishment of this experiment will address the requirement for extensive labelling annotations in a stringent format, which is more compatible with machines than with humans. This will be achieved by directly utilizing the existing textual reports authored by physicians daily.



# Appendix A

## List of published works

Some of the contributions presented in this doctoral thesis have been published in dedicated articles and patents authored by the same researcher. This section provides a concise overview of these works.

### A.1 Predictive Maintenance

The initial formulation of the predictive maintenance methodology was presented in the article "*Mining sensor data for predictive maintenance in the automotive industry*" (Giobergia et al. [2018]), where the first use of the *CORR-FS* algorithm for signal selection was proposed. This article provides an early version of the methodology and applies it to the oxygen sensor use case.

A more comprehensive and refined version of the methodology, along with the results achieved for the fuel high-pressure system use case, is presented in the article titled "*Dissecting a data-driven prognostic pipeline: A powertrain use case*" (Giordano et al. [2021]). This article offers detailed insights into the implementation specifics and necessary adjustments for practical deployment.

Similarly, another article entitled "*Data-driven strategies for predictive maintenance: Lessons learned from an automotive use case*" (Giordano et al. [2022]) focuses on providing a comprehensive description and analysis of the results obtained for the oxygen sensor use case, while employing the same methodology.

Furthermore, the intellectual property rights of the proposed methodology are protected by the US patent 11423321 titled "*Method and system for predicting system status*" (Neri et al. [2022]).

## A.2 Medical Imaging

The initial introduction of the degradation estimation methodology based on medical imaging is documented in the article titled "*Convolutional Neural Network-Based Automatic Analysis of Chest Radiographs for the Detection of COVID-19 Pneumonia: A Prioritizing Tool in the Emergency Department, Phase I Study and Preliminary "Real Life" Results*" (Tricarico et al. [2022a]). This publication presents the first proposal of the methodology and focuses on its application for prioritizing COVID-19 cases. Moreover, it demonstrates the performance achieved through a retrospective analysis of chest radiographs and the preliminary results obtained from a trial conducted in a real local hospital setting.

The article titled "*Deep regression by feature regularization for COVID-19 severity prediction*" (Tricarico et al. [2022b]) introduces the *DISTMAT* loss function. This novel loss function is applied in the study to estimate the severity of COVID-19 based on computed tomography scans within the context of the *Per-COVID-19 challenge*.

The proposed methodology has also been safeguarded through intellectual property protection. For its medical application, it is covered by European patent EP3901962A1 titled "*Tool and method to analyse medical images of patients*" (Tricarico et al. [2021]). Furthermore, for its industrial applications in component fault analysis, it is protected by US patent 11,354,796 titled "*Image identification and retrieval for component fault analysis*" (Tricarico et al. [2022c]).

# Bibliography

H. Abdi and L. J. Williams. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459, 2010.

N. Abdullah, U. K. Ngah, and S. A. Aziz. Image classification of brain mri using support vector machine. In *2011 IEEE international conference on imaging systems and techniques*, pages 242–247. IEEE, 2011.

H. Aboutaleb, M. Pavlova, M. J. Shafiee, A. Sabri, A. Alaref, and A. Wong. Covid-net cxr-s: Deep convolutional neural network for severity assessment of covid-19 cases from chest x-ray images. *Diagnostics*, 12(1):25, 2022.

ACR. Acr recommendations for the use of chest radiography and computed tomography (ct) for suspected covid-19 infection, 2020. URL <https://www.acr.org/Advocacy-and-Economics/ACR-Position-Statements/Recommendations-for-Chest-Radiography-and-CT-for-Suspected-COVID19-Infect>

P. Afshar, S. Heidarian, F. Naderkhani, A. Oikonomou, K. N. Plataniotis, and A. Mohammadi. Covid-caps: A capsule network-based framework for identification of covid-19 cases from x-ray images. *Pattern Recognition Letters*, 138:638–643, 2020.

A. S. Agatston, W. R. Janowitz, F. J. Hildner, N. R. Zusmer, M. Viamonte Jr, and R. Detrano. Quantification of coronary artery calcium using ultrafast computed tomography. *Journal of the American college of cardiology*, 15(4):827–832, 1990.

A. M. Agha, J. Pacor, G. R. Grandhi, R. Mszar, S. U. Khan, R. Parikh, T. Agrawal, J. Burt, R. Blankstein, M. J. Blaha, et al. The prognostic value of cac zero among individuals presenting with chest pain: a meta-analysis. *Cardiovascular Imaging*, 15(10):1745–1757, 2022.

- E. A. Akl, I. Blažić, S. Yaacoub, G. Frija, R. Chou, J. A. Appiah, M. Fatehi, N. Flor, E. Hitti, H. Jafri, et al. Use of chest imaging in the diagnosis and management of covid-19: a who rapid advice guide. *Radiology*, 298(2):E63–E69, 2021.
- M. Al-Ani and A. A. Rawi. A rule-based expert system for automated ecg diagnosis. *International Journal of Advances in Engineering and Technology*, 6(4):1480–1492, 2014.
- M. Alesina, C. A. Barbano, C. Berzovini, M. Busso, M. Calandri, A. D. Pascale, A. Fiandrotti, P. Fonio, M. Grangetto, M. Grosso, H. A. H. Chaudhry, M. E. Mancini, T. M. Gallo, C. Martini, G. Pontone, R. Renzulli, C. Saviolo, F. Signoretta, S. Tibaldi, A. Veltri, and M. Zaffaroni. Corda dataset, Jan. 2023. URL <https://doi.org/10.5281/zenodo.7821611>.
- M. Alexander Bilbily, MD; Mark Cicero. 16bit.ai | physistm, 2022. URL <https://www.16bit.ai/physics>.
- R. Aljondi and S. Alghamdi. Diagnostic value of imaging modalities for covid-19: scoping review. *Journal of medical Internet research*, 22(8):e19673, 2020.
- Z. Alom, M. Rahman, M. Nasrin, T. Taha, and V. Asari. Covid-19 detection with multi-task deep learning approaches, 2004.
- A. Amyar, R. Modzelewski, H. Li, and S. Ruan. Multi-task deep learning based ct imaging analysis for covid-19 pneumonia: Classification and segmentation. *Computers in Biology and Medicine*, 126:104037, 2020.
- M. André. The artemis european driving cycles for measuring car pollutant emissions. *Science of The Total Environment*, 334-335:73–84, 2004. ISSN 0048-9697. doi: <https://doi.org/10.1016/j.scitotenv.2004.04.070>. URL <https://www.sciencedirect.com/science/article/pii/S0048969704003584>. Highway and Urban Pollution.
- M. Annarumma, S. J. Withey, R. J. Bakewell, E. Pesce, V. Goh, and G. Montana. Automated triaging of adult chest radiographs with deep artificial neural networks. *Radiology*, 291(1):196–202, 2019.
- T. Anwar. Sensembenet: A squeeze and excitation based ensemble network for covid-19 infection percentage estimation from ct-scans. 2022.

- M. Assran, Q. Duval, I. Misra, P. Bojanowski, P. Vincent, M. Rabbat, Y. LeCun, and N. Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023.
- F. T. Avi Goldfarb. Why is ai adoption in health care lagging?, 2022. URL <https://www.brookings.edu/articles/why-is-ai-adoption-in-health-care-lagging/>.
- M. Baban, C. F. Baban, and B. Moisi. A fuzzy logic-based approach for predictive maintenance of grinding wheels of automated grinding lines. In *2018 23rd International Conference on Methods & Models in Automation & Robotics (MMAR)*, pages 483–486, 2018. doi: 10.1109/MMAR.2018.8486144.
- I. Baltruschat, L. Steinmeister, H. Nickisch, A. Saalbach, M. Grass, G. Adam, T. Knopp, and H. Ittrich. Smart chest x-ray worklist prioritization using artificial intelligence: a clinical workflow simulation. *European radiology*, 31:3837–3845, 2021.
- S. Behera, A. Choubey, C. Kanani, Y. Patel, R. Misra, and A. Sillitti. Ensemble trees learning based improved predictive maintenance using iiot for turbofan engines. pages 842–850, 04 2019. doi: 10.1145/3297280.3297363.
- T. Berredjem and M. Benidir. Bearing faults diagnosis using fuzzy expert system relying on an improved range overlaps and similarity method. *Expert Syst. Appl.*, 108:134–142, 2018.
- M. J. Berst, L. Dolan, M. M. Bogdanowicz, M. A. Stevens, S. Chow, and E. A. Brandser. Effect of knowledge of chronologic age on the variability of pediatric bone age determined using the greulich and pyle standards. *American Journal of Roentgenology*, 176(2):507–510, 2001.
- S. Bhattacharya, P. K. R. Maddikunta, Q.-V. Pham, T. R. Gadekallu, C. L. Chowdhary, M. Alazab, M. J. Piran, et al. Deep learning and medical image processing for coronavirus (covid-19) pandemic: A survey. *Sustainable cities and society*, 65:102589, 2021.
- W. L. Bi, A. Hosny, M. B. Schabath, M. L. Giger, N. J. Birkbak,

- A. Mehrtash, T. Allison, O. Arnaout, C. Abbosh, I. F. Dunn, et al. Artificial intelligence in cancer imaging: clinical challenges and applications. *CA: a cancer journal for clinicians*, 69(2):127–157, 2019.
- A. L. Blum and P. Langley. Selection of relevant features and examples in machine learning. *Artificial intelligence*, 97(1-2):245–271, 1997.
- B. Boecking, N. Usuyama, S. Bannur, D. Coelho de Castro, A. Schwaighofer, S. Hyland, M. T. Wetscherek, T. Naumann, A. Nori, J. Alvarez-Valle, H. Poon, and O. Oktay. Making the most of text semantics to improve biomedical vision-language processing. In *The European Conference on Computer Vision (ECCV)*, October 2022.
- C. Boller. Aging aircraft, health monitoring and maintenance. 02 2001.
- A. Borakati, A. Perera, J. Johnson, and T. Sood. Diagnostic accuracy of x-ray versus ct in covid-19: a propensity-matched database study. *BMJ open*, 10(11):e042946, 2020.
- S. Boral, S. K. Chaturvedi, and V. Naikan. A case-based reasoning system for fault detection and isolation: a case study on complex gearboxes. *Journal of Quality in Maintenance Engineering*, 2019.
- A. Borghesi and R. Maroldi. Covid-19 outbreak in italy: experimental chest x-ray scoring system for quantifying and monitoring disease progression. *La radiologia medica*, 125(5):509–513, 2020.
- A. Borghesi, A. Zigliani, R. Masciullo, S. Golemi, P. Maculotti, D. Farina, and R. Maroldi. Radiographic severity index in covid-19 pneumonia: relationship to age and sex in 783 italian patients. *La radiologia medica*, 125:461–464, 2020.
- F. Bougourzi, C. Distanto, A. Ouafi, F. Dornaika, A. Hadid, and A. Taleb-Ahmed. Per-covid-19: A benchmark dataset for covid-19 percentage estimation from ct-scans. *Journal of Imaging*, 7(9), 2021. ISSN 2313-433X. doi: 10.3390/jimaging7090189. URL <https://www.mdpi.com/2313-433X/7/9/189>.
- G. Box, G. Jenkins, and G. Reisel. Time series analysis-wiley series in probability and statistics, 2008.
- L. Breiman. Random forests. *Machine learning*, 45:5–32, 2001.



- J. Brownlee. Random oversampling and undersampling for imbalanced classification, 2020. URL <https://machinelearningmastery.com/random-oversampling-and-undersampling-for-imbalanced-classification/>.
- L. Brunese, F. Mercaldo, A. Reginelli, and A. Santone. Explainable deep learning for pulmonary disease and coronavirus covid-19 detection from x-rays. *Computer Methods and Programs in Biomedicine*, 196:105608, 2020.
- B. G. Buchanan and R. O. Duda. Principles of rule-based expert systems. In *Advances in computers*, volume 22, pages 163–216. Elsevier, 1983.
- C. S. Byington, M. J. Roemer, and T. R. Galie. Prognostic enhancements to diagnostic systems for improved condition-based maintenance [military aircraft]. *Proceedings, IEEE Aerospace Conference*, 6:6–6, 2002.
- Cambridge. Deterioration - cambridge dictionary, 2023. URL <https://dictionary.cambridge.org/dictionary/english/deterioration>.
- Q. Cao, A. Samet, C. Zanni-Merk, F. de Bertrand de Beuvron, and C. Reich. An ontology-based approach for failure classification in predictive maintenance using fuzzy c-means and swrl rules. In *International Conference on Knowledge-Based Intelligent Information & Engineering Systems*, 2019.
- M. Catarina Silva. Fdg pet scan more accurately assesses severity of cognitive decline in alzheimer’s, study finds, 2019. URL <https://alzheimersnewstoday.com/news/fdg-pet-scan-accurately-assesses-cognitive-decline-severity-alzheimers-di>
- CDC. Computed tomography (ct) scans, 2021. URL [https://www.cdc.gov/nceh/radiation/ct\\_scans.html](https://www.cdc.gov/nceh/radiation/ct_scans.html).
- F. Chabat, D. M. Hansell, and G.-Z. Yang. Computerized decision support in medical imaging. *IEEE Engineering in medicine and Biology Magazine*, 19(5):89–96, 2000.
- J. Chai, H. Zeng, A. Li, and E. W. Ngai. Deep learning in computer vision: A critical review of emerging techniques and application scenarios. *Machine Learning with Applications*, 6:100134, 2021.
- S. Chaudhary, W. Yang, and Y. Qiang. Swin transformer for covid-19 infection percentage estimation from ct-scans. In *Image Analysis and Processing. ICIAP 2022 Workshops: ICIAP International Workshops, Lecce*,

- Italy, May 23–27, 2022, Revised Selected Papers, Part II*, pages 520–528. Springer, 2022.
- N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- G. Chechik, V. Sharma, U. Shalit, and S. Bengio. Large scale online learning of image similarity through ranking. *Journal of Machine Learning Research*, 11(3), 2010.
- C. Chen, Y. Liu, X. Sun, C. D. Cairano-Gilfedder, and S. Titmus. An integrated deep learning-based approach for automobile maintenance prediction with gis data. *Reliability Engineering & System Safety*, 216:107919, 2021. ISSN 0951-8320. doi: <https://doi.org/10.1016/j.res.2021.107919>. URL <https://www.sciencedirect.com/science/article/pii/S095183202100435X>.
- T. Chen and C. Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- F. Cipollini, L. Oneto, A. Coraddu, A. J. Murphy, and D. Anguita. Condition-based maintenance of naval propulsion systems with supervised data analysis. *Ocean Engineering*, 149:268–278, 2018. ISSN 0029-8018. doi: <https://doi.org/10.1016/j.oceaneng.2017.12.002>. URL <https://www.sciencedirect.com/science/article/pii/S0029801817307242>.
- Cleveland Clinic. Pet scan, 2022. URL <https://my.clevelandclinic.org/health/diagnostics/10123-pet-scan>.
- I. Cohen, Y. Huang, J. Chen, J. Benesty, J. Benesty, J. Chen, Y. Huang, and I. Cohen. Pearson correlation coefficient. *Noise reduction in speech processing*, pages 1–4, 2009.
- J. P. Cohen, L. Dao, K. Roth, P. Morrison, Y. Bengio, A. F. Abbasi, B. Shen, H. K. Mahsa, M. Ghassemi, H. Li, et al. Predicting covid-19 pneumonia severity on chest x-ray with deep learning. *Cureus*, 12(7), 2020a.
- J. P. Cohen, P. Morrison, and L. Dao. Covid-19 image data collection. *arXiv preprint arXiv:2003.11597*, 2020b.

- J. Cook. Reducing military helicopter maintenance through prognostics. In *2007 IEEE Aerospace Conference*, pages 1–7, 2007. doi: 10.1109/AERO.2007.352830.
- G. Cormode and P. Veselý. A tight lower bound for comparison-based quantile summaries. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 81–93, 2020.
- C. Cortes and V. Vapnik. Support vector machine. *Machine learning*, 20(3):273–297, 1995.
- H. Costin. A fuzzy rules-based segmentation method for medical images analysis. *International Journal of Computers Communications & Control*, 8(2):196–205, 2013.
- E. C. Covert, K. Fitzpatrick, J. Mikell, R. K. Kaza, J. D. Millet, D. Barkmeier, J. Gemmete, J. Christensen, M. J. Schipper, and Y. K. Dewaraja. Intra-and inter-operator variability in mri-based manual segmentation of hcc lesions and its impact on dosimetry. *EJNMMI physics*, 9(1):1–16, 2022.
- A. Cozzi, S. Schiaffino, F. Arpaia, G. Della Pepa, S. Tritella, P. Bertolotti, L. Menicagli, C. G. Monaco, L. A. Carbonaro, R. Spairani, et al. Chest x-ray in the covid-19 pandemic: Radiologists’ real-world reader performance. *European journal of radiology*, 132:109272, 2020.
- A. Criminisi and J. Shotton. *Decision forests for computer vision and medical image analysis*. Springer Science & Business Media, 2013.
- J. Dalzochio, R. Kunst, E. Pignaton, A. Binotto, S. Sanyal, J. Favilla, and J. Barbosa. Machine learning and reasoning for predictive maintenance in industry 4.0: Current status and challenges. *Computers in Industry*, 123:103298, 2020. ISSN 0166-3615. doi: <https://doi.org/10.1016/j.compind.2020.103298>. URL <https://www.sciencedirect.com/science/article/pii/S0166361520305327>.
- J. Dalzochio, R. Kunst, J. L. V. Barbosa, P. C. d. S. Neto, E. Pignaton, C. S. t. Caten, and A. d. L. T. da Penha. Predictive maintenance in the military domain: A systematic review of the literature. *ACM Comput. Surv.*, mar 2023. ISSN 0360-0300. doi: 10.1145/3586100. URL <https://doi.org/10.1145/3586100>.

- G. D’Ancona, M. Massussi, M. Savardi, A. Signoroni, L. Di Bacco, D. Farina, M. Metra, R. Maroldi, C. Muneretto, H. Ince, et al. Deep learning to detect significant coronary artery disease from plain chest radiographs ai4cad. *International Journal of Cardiology*, 370:435–441, 2023.
- S. M. de Andrade Lopes, R. A. Flauzino, and R. A. C. Altafim. Incipient fault diagnosis in power transformers by data-driven models with over-sampled dataset. *Electric Power Systems Research*, 201:107519, 2021. ISSN 0378-7796. doi: <https://doi.org/10.1016/j.epsr.2021.107519>. URL <https://www.sciencedirect.com/science/article/pii/S0378779621005009>.
- P. Del Moral. Nonlinear filtering: Interacting particle resolution. *Comptes Rendus de l’Académie des Sciences - Series I - Mathematics*, 325(6):653–658, 1997. ISSN 0764-4442. doi: [https://doi.org/10.1016/S0764-4442\(97\)84778-7](https://doi.org/10.1016/S0764-4442(97)84778-7). URL <https://www.sciencedirect.com/science/article/pii/S0764444297847787>.
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- T. Desell, S. Clachar, J. Higgins, and B. Wild. Evolving neural network weights for time-series prediction of general aviation flight data. pages 771–781, 09 2014. ISBN 978-3-319-10761-5. doi: 10.1007/978-3-319-10762-2\_76.
- C. Ding and H. Peng. Minimum redundancy feature selection from microarray gene expression data. *Journal of bioinformatics and computational biology*, 3(02):185–205, 2005.
- Dive into deep learning. Convolutional neural networks (lenet), 2023. URL [https://d2l.ai/chapter\\_convolutional-neural-networks/lenet.html](https://d2l.ai/chapter_convolutional-neural-networks/lenet.html).
- M. Ducoffe, I. Haloui, and J. S. Gupta. Anomaly detection on time series with wasserstein gan applied to phm. *International Journal of Prognostics and Health Management*, 10(4), 2019.
- I. El-Naqa, Y. Yang, M. N. Wernick, N. P. Galatsanos, and R. M. Nishikawa. A support vector machine approach for detection of microcalcifications. *IEEE transactions on medical imaging*, 21(12):1552–1563, 2002.

- I. El Naqa, H. Li, J. Fuhrman, Q. Hu, N. Gorre, W. Chen, and M. L. Giger. Lessons learned in transitioning to ai in the medical imaging of covid-19. *Journal of Medical Imaging*, 8(S1):010902–010902, 2021.
- C. Ellis. Oversampling vs undersampling for machine learning, 2023. URL <https://crunchingthedata.com/oversampling-vs-undersampling/>.
- Envision Radiology. X-ray vs. ct vs. mri, 2023. URL <https://www.envrad.com/difference-between-x-ray-ct-scan-and-mri/>.
- B. J. Erickson, P. Korfiatis, Z. Akkus, and T. L. Kline. Machine learning for medical imaging. *Radiographics*, 37(2):505–515, 2017.
- L. M. Fayad. Ct scan versus mri versus x-ray: What type of imaging do i need?, 2023. URL <https://www.hopkinsmedicine.org/health/treatment-tests-and-therapies/ct-vs-mri-vs-xray>.
- FDA. Medical imaging, 2018. URL <https://www.fda.gov/radiation-emitting-products/radiation-emitting-products-and-procedures/medical-imaging>.
- Y. Feng, H. S. Teh, and Y. Cai. Deep learning for chest radiology: a review. *Current Radiology Reports*, 7:1–9, 2019.
- G. Florimbi, H. Fabelo, E. Torti, R. Lazcano, D. Madroñal, S. Ortega, R. Salvador, F. Leporati, G. Danese, A. Báez-Quevedo, et al. Accelerating the k-nearest neighbors filtering algorithm to optimize the real-time classification of human brain tumor in hyperspectral images. *Sensors*, 18(7):2314, 2018.
- B. Freyermuth. Knowledge based incipient fault diagnosis of industrial robots. *IFAC Proceedings Volumes*, 24(6):369–375, 1991. ISSN 1474-6670. doi: [https://doi.org/10.1016/S1474-6670\(17\)51169-6](https://doi.org/10.1016/S1474-6670(17)51169-6). URL <https://www.sciencedirect.com/science/article/pii/S1474667017511696>. IFAC/IMACS Symposium on Fault Detection, Supervision and Safety for Technical Processes (SAFEPROCESS’91), Baden-Baden, Germany, 10-13 September 1991.
- J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.

- J. Furch, T. Tin Nguyen, and J. Glos. Diagnostics of gear fault in four-speed gearbox using vibration signal. In *2017 International Conference on Military Technologies (ICMT)*, pages 87–92, 2017. doi: 10.1109/MILTECHS.2017.7988736.
- M. Gatti, M. Calandri, M. Barba, A. Biondo, C. Geninatti, S. Gentile, M. Greco, V. Morrone, C. Piatti, A. Santonocito, et al. Baseline chest x-ray in coronavirus disease 19 (covid-19) patients: association with clinical and laboratory data. *La radiologia medica*, 125:1271–1279, 2020.
- R. Genuer, J.-M. Poggi, and C. Tuleau. Random forests: some methodological insights. *arXiv preprint arXiv:0811.3619*, 2008.
- U. Ghoshal, S. Vasanth, and N. Tejan. A guide to laboratory diagnosis of corona virus disease-19 for the gastroenterologists. *Indian Journal of Gastroenterology*, 39:236 – 242, 2020.
- F. Giobergia, E. Baralis, M. Camuglia, T. Cerquitelli, M. Mellia, A. Neri, D. Tricarico, and A. Tuninetti. Mining sensor data for predictive maintenance in the automotive industry. In *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 351–360. IEEE, 2018.
- D. Giordano, E. Pastor, F. Giobergia, T. Cerquitelli, E. Baralis, M. Mellia, A. Neri, and D. Tricarico. Dissecting a data-driven prognostic pipeline: A powertrain use case. *Expert Systems with Applications*, 180:115109, 2021.
- D. Giordano, F. Giobergia, E. Pastor, A. La Macchia, T. Cerquitelli, E. Baralis, M. Mellia, and D. Tricarico. Data-driven strategies for predictive maintenance: Lesson learned from an automotive use case. *Computers in Industry*, 134:103554, 2022.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- M. Greenwald and S. Khanna. Space-efficient online computation of quantile summaries. *ACM SIGMOD Record*, 30(2):58–66, 2001.
- W. W. Greulich and S. I. Pyle. *Radiographic atlas of skeletal development of the hand and wrist*. Stanford university press, 1959.
- R. M. Grossi. The importance of medical imaging and nuclear medicine in universal health coverage. *The Lancet. Oncology*, 22(4):423–424, 2021.

- J. Gustavsson, C. Cederberg, U. Sonesson, R. Van Otterdijk, and A. Meybeck. Global food losses and food waste, 2011.
- A. Haghanifar, M. M. Majdabadi, Y. Choi, S. Deivalakshmi, and S. Ko. Covid-cxnet: Detecting covid-19 in frontal chest x-ray images using deep learning. *Multimedia Tools and Applications*, 81(21):30615–30645, 2022.
- Y. Hahn, T. Langer, R. Meyes, and T. Meisen. Time series dataset survey for forecasting with deep learning. *Forecasting*, 5(1):315–335, 2023.
- S. S. Halabi, L. M. Prevedello, J. Kalpathy-Cramer, A. B. Mamonov, A. Bilbily, M. Cicero, I. Pan, L. A. Pereira, R. T. Sousa, N. Abdala, et al. The rsna pediatric bone age machine learning challenge. *Radiology*, 290(2):498–503, 2019.
- M. Handrich. Wideband zirconia sensor, 2010. URL [https://en.wikipedia.org/wiki/Oxygen\\_sensor#/media/File:WidebandZirconiaSensor.svg](https://en.wikipedia.org/wiki/Oxygen_sensor#/media/File:WidebandZirconiaSensor.svg).
- A. F. Hannu Jääskeläinen. Common rail injection system pressure control, 2023. URL [https://dieselnet.com/tech/diesel\\_fi\\_common-rail\\_control.php](https://dieselnet.com/tech/diesel_fi_common-rail_control.php).
- G. Harshvardhan, M. K. Gourisaria, M. Pandey, and S. S. Rautaray. A comprehensive survey and analysis of generative models in machine learning. *Computer Science Review*, 38:100285, 2020.
- D. Hartmann, D. Müller, I. Soto-Rey, and F. Kramer. Assessing the role of random forests in medical image segmentation. *arXiv preprint arXiv:2103.16492*, 2021.
- T. Hastie, R. Tibshirani, and J. Friedman. The elements of statistical learning. springer series in statistics. *New York, NY, USA*, 2001.
- N. Haug, L. Geyrhofer, A. Londei, E. Dervic, A. Desvars-Larrive, V. Loreto, B. Piniør, S. Thurner, and P. Klimek. Ranking the effectiveness of worldwide covid-19 government interventions. *Nature human behaviour*, 4(12):1303–1312, 2020.
- A. Heena, N. Biradar, N. M. Maroof, S. Bhatia, R. Agarwal, and K. Prasad. Machine learning based biomedical image processing for echocardiographic images. *Multimedia Tools and Applications*, pages 1–16, 2022.

- A. Herment, J. Guglielmi, P. Dumée, P. Peronneau, and P. Delouche. Limitations of ultrasound imaging and image restoration. *Ultrasonics*, 25(5): 267–273, 1987.
- B. Heyworth, D. Osei, P. Fabricant, D. Green, S. Doyle, R. Widmann, D. Scher, R. Schneider, P. Cahill, and M. Herman. A new, validated shorthand method for determining bone age. In *Annual Meeting of the American Academy of Orthopaedic Surgeons*, pages 16–19, 2011.
- A. Homborg, T. Tinga, and J. Mol. Listening to corrosion. In *AVT-305 Research Specialists Meeting on Sensing Systems for Integrated Vehicle Health Management for Military Vehicles*, pages 1–14, Belgium, 2018. NATO Science & Technology Organization. ISBN 978-92-837-2205-2. doi: 10.14339/STO-MP-AVT-305-12-PDF.
- M. Hruz, M. Bugaj, A. Novák, B. Kandra, and B. Badánik. The use of uav with infrared camera and rfid for airframe condition monitoring. *Applied Sciences*, 11(9), 2021. URL <https://www.mdpi.com/2076-3417/11/9/3737>.
- C.-W. Hsu and C.-J. Lin. A comparison of methods for multiclass support vector machines. *IEEE transactions on Neural Networks*, 13(2):415–425, 2002.
- S. Hu, Y. Gao, Z. Niu, Y. Jiang, L. Li, X. Xiao, M. Wang, E. F. Fang, W. Menpes-Smith, J. Xia, et al. Weakly supervised deep learning for covid-19 infection detection and classification from ct images. *IEEE Access*, 8:118869–118883, 2020.
- G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, 2017. doi: 10.1109/CVPR.2017.243.
- M. Ian Pan. Deep learning for pediatric bone age assessment, 2020. URL <https://github.com/i-pan/boneage>.
- G. Iannace, G. Ciaburro, and A. Trematerra. Fault diagnosis for uav blades using artificial neural network. *Robotics*, 8(3):59, 2019.
- International SAE. Materials degradation in mechanical design: Wear, corrosion, fatigue, and their interactions web seminar replay, 2023. URL <https://www.sae.org/learn/content/pd331722/>.



- I. M. Iqbal and N. Aziz. Comparison of various wiener model identification approach in modelling nonlinear process. In *2011 3rd Conference on Data Mining and Optimization (DMO)*, pages 134–140, 2011. doi: 10.1109/DMO.2011.5976517.
- J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Illcus, C. Chute, H. Marklund, B. Haghighi, R. Ball, K. Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 590–597, 2019.
- N. Islam, S. Ebrahimzadeh, J.-P. Salameh, S. Kazi, N. Fabiano, L. Treanor, M. Absi, Z. Hallgrimson, M. M. Leeftang, L. Hooft, et al. Thoracic imaging tests for the diagnosis of covid-19. *Cochrane Database of Systematic Reviews*, (3), 2021.
- M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. *Advances in neural information processing systems*, 28, 2015.
- S. Jagannathan and G. Raju. Remaining useful life prediction of automotive engine oils using mems technologies. In *Proceedings of the 2000 American Control Conference. ACC (IEEE Cat. No.00CH36334)*, volume 5, pages 3511–3512 vol.5, 2000. doi: 10.1109/ACC.2000.879222.
- A. K. Jain, J. Mao, and K. M. Mohiuddin. Artificial neural networks: A tutorial. *Computer*, 29(3):31–44, 1996.
- R. Jain and I. Chlamtac. The p2 algorithm for dynamic calculation of quantiles and histograms without storing observations. *Communications of the ACM*, 28(10):1076–1085, 1985.
- G. James, D. Witten, T. Hastie, and R. Tibshirani. *An introduction to statistical learning*, volume 112. Springer, 2013.
- Z. Jiang, T. Pan, C. Zhang, and J. Yang. A new oversampling method based on the classification contribution degree. *Symmetry*, 13(2):194, 2021.
- A. E. Johnson, T. J. Pollard, S. J. Berkowitz, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, R. G. Mark, and S. Horng. Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1):317, 2019.

- K. A. Kaiser and N. Z. Gebraeel. Predictive maintenance management using sensor-based degradation models. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 39(4):840–849, 2009.
- P. I. Kamel, P. H. Yi, H. I. Sair, and C. T. Lin. Prediction of coronary artery calcium and cardiovascular risk on chest radiographs using deep learning. *Radiology: Cardiothoracic Imaging*, 3(3):e200486, 2021.
- H. Kargupta, R. Bhargava, K. Liu, M. Powers, P. Blair, S. Bushra, J. Dull, K. Sarkar, M. Klein, M. Vasa, and D. Handy. Vedas: A mobile and distributed data stream mining system for real-time vehicle monitoring. In *SDM*, 2004.
- M. R. Karim, T. Döhmen, M. Cochez, O. Beyan, D. Rebholz-Schuhmann, and S. Decker. Deepcovidexplainer: explainable covid-19 diagnosis from chest x-ray images. In *2020 IEEE international conference on bioinformatics and biomedicine (BIBM)*, pages 1034–1037. IEEE, 2020.
- C. J. Kelly, A. Karthikesalingam, M. Suleyman, G. Corrado, and D. King. Key challenges for delivering clinical impact with artificial intelligence. *BMC medicine*, 17:1–9, 2019.
- C. Kingsford and S. L. Salzberg. What are decision trees? *Nature biotechnology*, 26(9):1011–1013, 2008.
- B. Koçak. Key concepts, common pitfalls, and best practices in artificial intelligence and machine learning: focus on radiomics. *Diagnostic and Interventional Radiology*, 28(5):450, 2022.
- M. A. Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE journal*, 37(2):233–243, 1991.
- M. A. Kramer. Autoassociative neural networks. *Computers & chemical engineering*, 16(4):313–328, 1992.
- A. Kratsios and C. Hyndman. Neu: A meta-algorithm for universal uap-invariant feature representation. *The Journal of Machine Learning Research*, 22(1):4102–4152, 2021.
- E. A. Krupinski. Current perspectives in medical image perception. *Attention, Perception, & Psychophysics*, 72(5):1205–1217, 2010.
- M. Labs. Apple deterioration, 2017. URL <https://www.medallionlabs.com/blog/food-manufacturers-shelf-life-testing/>.

- P. Lall, R. Lowe, and K. Goebel. Particle swarm optimization with extended kalman filter for prognostication of accrued damage in electronics under temperature and vibration. In *2012 IEEE Conference on Prognostics and Health Management*, pages 1–13, 2012a. doi: 10.1109/ICPHM.2012.6299536.
- P. Lall, R. Lowe, and K. Goebel. Prognostication of accrued damage in board assemblies under thermal and mechanical stresses. In *2012 IEEE 62nd Electronic Components and Technology Conference*, pages 1475–1487, 2012b. doi: 10.1109/ECTC.2012.6249031.
- D. B. Larson, M. C. Chen, M. P. Lungren, S. S. Halabi, N. V. Stence, and C. P. Langlotz. Performance of a deep-learning neural network model in assessing skeletal maturity on pediatric hand radiographs. *Radiology*, 287(1):313–322, 2018.
- V. T. Le, C. P. Lim, S. Mohamed, S. Nahavandi, L. Yen, G. E. Gallasch, S. Baker, D. Ludovici, N. Draper, V. Wickramanayake, et al. Condition monitoring of engine lubrication oil of military vehicles: A machine learning approach. *Proc. 17th Austral. Int. Aerosp. Congr.(AIAC)*, 2017.
- B. P. Leao, K. T. Fitzgibbon, L. C. Puttini, and G. P. B. de Melo. Cost-benefit analysis methodology for phm applied to legacy commercial aircraft. In *2008 IEEE Aerospace Conference*, pages 1–13, 2008. doi: 10.1109/AERO.2008.4526599.
- M. S. Lebold, S. Pflumm, J. C. Banks, J. Bednár, K. M. Reichard, K. B. Fischer, and J. M. W. Stempnik. Detecting injector deactivation failure modes in diesel engines using time and order domain approaches. 2012.
- Y. LeCun, Y. Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- J. Lee, F. Wu, W. Zhao, M. Ghaffari, L. Liao, and D. Siegel. Prognostics and health management design for rotary machinery systems—reviews,

- methodology and applications. *Mechanical Systems and Signal Processing*, 42:314–334, 2014.
- L. Li, L. Qin, Z. Xu, Y. Yin, X. Wang, B. Kong, J. Bai, Y. Lu, Z. Fang, Q. Song, et al. Using artificial intelligence to detect covid-19 and community-acquired pneumonia based on pulmonary ct: evaluation of the diagnostic accuracy. *Radiology*, 296(2):E65–E71, 2020.
- Y. LI, T. KURFESS, and S. LIANG. Stochastic prognostics for rolling element bearings. *Mechanical Systems and Signal Processing*, 14(5):747–762, 2000. ISSN 0888-3270. doi: <https://doi.org/10.1006/mssp.2000.1301>. URL <https://www.sciencedirect.com/science/article/pii/S0888327000913013>.
- T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Doll'ar, and C. L. Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. URL <http://arxiv.org/abs/1405.0312>.
- H. Liu. *Feature Selection*, pages 402–406. Springer US, Boston, MA, 2010. ISBN 978-0-387-30164-8. doi: 10.1007/978-0-387-30164-8\_306. URL [https://doi.org/10.1007/978-0-387-30164-8\\_306](https://doi.org/10.1007/978-0-387-30164-8_306).
- J. Liu, W. Wang, and F. Golnaraghi. A multi-step predictor with a variable input pattern for system state forecasting. *Mechanical Systems and Signal Processing*, 23(5):1586–1599, 2009. ISSN 0888-3270. doi: <https://doi.org/10.1016/j.ymsp.2008.09.006>. URL <https://www.sciencedirect.com/science/article/pii/S0888327008002276>.
- J. Liu, W. Wang, F. Ma, Y. Yang, and C. Yang. A data-model-fusion prognostic framework for dynamic system state forecasting. *Engineering Applications of Artificial Intelligence*, 25(4):814–823, 2012. ISSN 0952-1976. doi: <https://doi.org/10.1016/j.engappai.2012.02.015>. URL <https://www.sciencedirect.com/science/article/pii/S0952197612000528>. Special Section: Dependable System Modelling and Analysis.
- J. S. Liu and R. Chen. Sequential monte carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93(443):1032–1044, 1998. doi: 10.1080/01621459.1998.10473765.
- L. Liu, K. Logan, and D. Cartes. Fault detection, diagnostics, and prognostics: software agent solutions. In *IEEE Electric Ship Technologies Symposium, 2005.*, pages 425–431, 2005. doi: 10.1109/ESTS.2005.1524710.

- L. Liu, D. Cartes, and J. Quiroga. Modeling and simulation for condition based maintenance: A case study in navy ship application. volume 1, pages 244–249, 01 2007. doi: 10.1145/1357910.1357950.
- S. Liu, A. Mallol-Ragolta, E. Parada-Cabaleiro, K. Qian, X. Jing, A. Kathan, B. Hu, and B. W. Schuller. Audio self-supervised learning: A survey. *Patterns*, 3(12):100616, 2022.
- Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- M. J. Loeffelholz and Y.-W. Tang. Laboratory diagnosis of emerging human coronavirus infections—the state of the art. *Emerging microbes & infections*, 9(1):747–756, 2020.
- C. Long, H. Xu, Q. Shen, X. Zhang, B. Fan, C. Wang, B. Zeng, Z. Li, X. Li, and H. Li. Diagnosis of the coronavirus disease (covid-19): rrt-pcr or ct? *European journal of radiology*, 126:108961, 2020.
- A. Mahmoud and A. Mohammed. A survey on deep learning for time-series forecasting. *Machine Learning and Big Data Analytics Paradigms: Analysis, Applications and Challenges*, pages 365–392, 2021.
- V. Majstorović and V. Milačić. Expert systems for maintenance in the cim concept. *Computers in Industry*, 15(1):83–93, 1990. ISSN 0166-3615. doi: [https://doi.org/10.1016/0166-3615\(90\)90086-5](https://doi.org/10.1016/0166-3615(90)90086-5). URL <https://www.sciencedirect.com/science/article/pii/0166361590900865>.
- M. Matesin, S. Loncaric, and D. Petravic. A rule-based approach to stroke lesion analysis from ct brain images. In *ISPA 2001. Proceedings of the 2nd International Symposium on Image and Signal Processing and Analysis. In conjunction with 23rd International Conference on Information Technology Interfaces (IEEE Cat.)*, pages 219–223. IEEE, 2001.
- Mayo clinic. Ultrasound, 2022. URL <https://www.mayoclinic.org/tests-procedures/ultrasound/about/pac-20395177>.
- R. L. McClelland, H. Chung, R. Detrano, W. Post, and R. A. Kronmal. Distribution of coronary artery calcium by race, gender, and age: results from the multi-ethnic study of atherosclerosis (mesa). *Circulation*, 113(1):30–37, 2006.

- Medical news today. How does a ct or cat scan work?, 2018. URL <https://www.medicalnewstoday.com/articles/153201>.
- O. Microbiology. The viral life cycle, 2023. URL <https://courses.lumenlearning.com/suny-microbiology/chapter/the-viral-life-cycle/>.
- C. Min and Z. Haijiang. Research on application of improved random forest in medical ultrasound image classification. In *Journal of Physics: Conference Series*, volume 1584, page 012007. IOP Publishing, 2020.
- Ministero della Salute. pcm-dpc/covid-19, 2023. URL <https://github.com/pcm-dpc/COVID-19>.
- R. Miron and M. E. Breaban. Revitalizing regression tasks through modern training procedures: Applications in medical image analysis for covid-19 infection percentage estimation. In *Image Analysis and Processing. ICIAP 2022 Workshops: ICIAP International Workshops, Lecce, Italy, May 23–27, 2022, Revised Selected Papers, Part II*, pages 473–482. Springer, 2022.
- Monib Co Ltd. Concrete corrosion repair, 2017. URL <https://medium.com/@monib/concrete-corrosion-repair-653582ab84b7>.
- J. J. Montero Jimenez, S. Schwartz, R. Vingerhoeds, B. Grabot, and M. Salaün. Towards multi-model approaches to predictive maintenance: A systematic literature survey on diagnostics and prognostics. *Journal of Manufacturing Systems*, 56:539–557, 2020. ISSN 0278-6125. doi: <https://doi.org/10.1016/j.jmsy.2020.07.008>. URL <https://www.sciencedirect.com/science/article/pii/S0278612520301187>.
- W. Moroder. Ultrasound image (sonogram) of a fetus in the womb, viewed at 12 weeks of pregnancy (bidimensional scan), 2012. URL [https://commons.wikimedia.org/wiki/File:CRL\\_Crown\\_rump\\_length\\_12\\_weeks\\_ecografia\\_Dr.\\_Wolfgang\\_Moroder.jpg](https://commons.wikimedia.org/wiki/File:CRL_Crown_rump_length_12_weeks_ecografia_Dr._Wolfgang_Moroder.jpg).
- myVMC. Mri (magnetic resonance imaging), 2018. URL <https://www.myvmc.com/investigations/mri-magnetic-resonance-imaging/>.
- P. Nair and I. Kashyap. Classification of medical image data using k nearest neighbor and finding the optimal k value. *International Journal of Scientific and Technology Research*, 9(4):221–226, 2020.

- S. Namuduri, B. N. Narayanan, V. S. P. Davuluru, L. Burton, and S. Bhansali. Deep learning methods for sensor based predictive maintenance and future perspectives for electrochemical sensors. *Journal of The Electrochemical Society*, 167(3):037552, 2020.
- M. A. Napoli Spatafora, A. Ortis, and S. Battiato. Mixup data augmentation for covid-19 infection percentage estimation. In *Image Analysis and Processing. ICIAP 2022 Workshops: ICIAP International Workshops, Lecce, Italy, May 23–27, 2022, Revised Selected Papers, Part II*, pages 508–519. Springer, 2022.
- S. J. Narayanan, R. Soundrapandiyan, B. Perumal, and C. J. Baby. Emphysema medical image classification using fuzzy decision tree with fuzzy particle swarm optimization clustering. In *Smart Intelligent Computing and Applications: Proceedings of the Second International Conference on SCI 2018, Volume 1*, pages 305–313. Springer, 2019.
- F. P. Nasution, S. Sævik, and J. K. Gjøsteen. Fatigue analysis of copper conductor for offshore wind turbines by experimental and fe method. *Energy Procedia*, 24:271–280, 2012.
- A. Neri, M. Camuglia, A. Tuninetti, E. Baralis, F. Giobergia, and D. Tricarico. Method and system for predicting system status, Aug. 23 2022. US Patent 11,423,321.
- NGK. Zfas-u2, 2023. URL [https://www.ngkntk.co.jp/english/product/sensors\\_plugs/wide\\_range\\_oxygen.html](https://www.ngkntk.co.jp/english/product/sensors_plugs/wide_range_oxygen.html).
- NIH. Ultrasound, 2016. URL <https://www.nibib.nih.gov/science-education/science-topics/ultrasound>.
- NIH. Computed tomography (ct), 2022. URL <https://www.nibib.nih.gov/science-education/science-topics/computed-tomography-ct>.
- NIH. Magnetic resonance imaging (mri), 2023. URL <https://www.nibib.nih.gov/science-education/science-topics/magnetic-resonance-imaging-mri>.
- S. Nixon, R. Weichel, K. Reichard, and J. Kozlowski. Machine learning approach to diesel engine health prognostics using engine controller data. *Annual Conference of the PHM Society*, 10, 09 2018. doi: 10.36001/phmconf.2018.v10i1.587.

- H. Nordal and I. El-Thalji. Lifetime benefit analysis of intelligent maintenance: Simulation modeling approach and industrial case study. *Applied Sciences*, 11(8), 2021. ISSN 2076-3417. URL <https://www.mdpi.com/2076-3417/11/8/3487>.
- F. S. Nowlan and F. Howard. Heap. reliability-centered maintenance. report number ad-a066579, 1978.
- P. Nunes, J. Santos, and E. Rocha. Challenges in predictive maintenance – a review. *CIRP Journal of Manufacturing Science and Technology*, 40:53–67, 2023. ISSN 1755-5817. doi: <https://doi.org/10.1016/j.cirpj.2022.11.004>. URL <https://www.sciencedirect.com/science/article/pii/S1755581722001742>.
- S. of Automotive Engineers. Sae ja1011: Evaluation criteria for reliability-centered maintenance (rcm) processes, 2009.
- M. F. B. Othman, N. B. Abdullah, and N. F. B. Kamal. Mri brain classification using support vector machine. In *2011 Fourth international conference on modeling, simulation and applied optimization*, pages 1–4. IEEE, 2011.
- D. W. Otter, J. R. Medina, and J. K. Kalita. A survey of the usages of deep learning for natural language processing. *IEEE transactions on neural networks and learning systems*, 32(2):604–624, 2020.
- A. Oussidi and A. Elhassouny. Deep generative models: Survey. In *2018 International conference on intelligent systems and computer vision (ISCV)*, pages 1–8. IEEE, 2018.
- T. Ozturk, M. Talo, E. A. Yildirim, U. B. Baloglu, O. Yildirim, and U. R. Acharya. Automated detection of covid-19 cases using deep neural networks with x-ray images. *Computers in biology and medicine*, 121:103792, 2020.
- P. Pal. Condition based maintenance of turbine and compressor of a codlag naval propulsion system using deep neural network. pages 01–12, 05 2019. doi: 10.5121/csit.2019.90601.
- I. Pan, H. H. Thodberg, S. S. Halabi, J. Kalpathy-Cramer, and D. B. Larson. Improving automated pediatric bone age estimation using ensembles of models from the 2017 rsna machine learning challenge. *Radiology: Artificial Intelligence*, 1(6):e190053, 2019.



- A. Pandian and A. Ali. A review of recent trends in machine diagnosis and prognosis algorithms. In *2009 World Congress on Nature & Biologically Inspired Computing (NaBIC)*, pages 1731–1736, 2009. doi: 10.1109/NABIC.2009.5393625.
- H. Panwar, P. Gupta, M. K. Siddiqui, R. Morales-Menendez, P. Bhardwaj, and V. Singh. A deep learning and grad-cam based color visualization approach for fast detection of covid-19 cases using chest x-ray and ct-scan images. *Chaos, Solitons & Fractals*, 140:110190, 2020.
- P. Parikh, N. Shah, H. Ahmed, P. Schoenhagen, and M. Fares. Coronary artery calcium scoring: Its practicality and clinical utility in primary care. *Cleve Clin J Med*, 85(9):707–16, 2018.
- P. Paris and F. Erdogan. Closure to “Discussions of ‘A Critical Analysis of Crack Propagation Laws’” (1963, ASME J. Basic Eng., 85, pp. 533–534). *Journal of Basic Engineering*, 85(4):534–534, 12 1963. ISSN 0021-9223. doi: 10.1115/1.3656903. URL <https://doi.org/10.1115/1.3656903>.
- M. Pecht and R. Jaai. A prognostics and health management roadmap for information and electronics-rich systems. *Microelectronics Reliability*, 50(3):317–323, 2010. ISSN 0026-2714. doi: <https://doi.org/10.1016/j.microrel.2010.01.006>. URL <https://www.sciencedirect.com/science/article/pii/S0026271410000181>.
- S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J. B. Zimmerman, and K. Zuiderveld. Adaptive histogram equalization and its variations. *Computer vision, graphics, and image processing*, 39(3):355–368, 1987.
- J. Qiu, B. B. Seth, S. Y. Liang, and C. Z. Zhang. Damage mechanics approach for bearing lifetime prognostics. *Mechanical Systems and Signal Processing*, 16:817–829, 2002.
- J. Rabatel, S. Bringay, and P. Poncelet. Anomaly detection in monitoring sensor data for preventive maintenance. *Expert Systems with Applications*, 38(6):7003–7015, 2011. ISSN 0957-4174. doi: <https://doi.org/10.1016/j.eswa.2010.12.014>. URL <https://www.sciencedirect.com/science/article/pii/S0957417410013771>.
- L. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989. doi: 10.1109/5.18626.

- Radiological Society of North America. Body mri, 2022. URL <https://www.radiologyinfo.org/en/info/bodymr>.
- RF Wireless World 2012. Advantages of ultrasound | disadvantages of ultrasound, 2012. URL <https://www.rfwireless-world.com/Terminology/Advantages-and-Disadvantages-of-Ultrasound.html>.
- M. A. Ricci Lara, R. Echeveste, and E. Ferrante. Addressing fairness in artificial intelligence for medical imaging. *nature communications*, 13(1): 4581, 2022.
- A. Rimmer. Radiologist shortage leaves patient care at risk, warns royal college. *BMJ: British Medical Journal (Online)*, 359, 2017.
- B. D. Ripley. *Pattern recognition and neural networks*. Cambridge university press, 2007.
- M. Roberts, D. Driggs, M. Thorpe, J. Gilbey, M. Yeung, S. Ursprung, A. I. Aviles-Rivero, C. Etmann, C. McCague, L. Beer, et al. Common pitfalls and recommendations for using machine learning to detect and prognosticate for covid-19 using chest radiographs and ct scans. *Nature Machine Intelligence*, 3(3):199–217, 2021.
- V. Roger, J. Farinas, and J. Pinquier. Deep neural networks for automatic speech processing: a survey from large corpora to limited data. *EURASIP Journal on Audio, Speech, and Music Processing*, 2022(1):19, 2022.
- O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- S. Roy, W. Menapace, S. Oei, B. Luijten, E. Fini, C. Saltori, I. Huijben, N. Chennakeshava, F. Mento, A. Sentelli, et al. Deep learning for classification and localization of covid-19 markers in point-of-care lung ultrasound. *IEEE transactions on medical imaging*, 39(8):2676–2687, 2020.
- G. D. Rubin, C. J. Ryerson, L. B. Haramati, N. Sverzellati, J. P. Kanne, S. Raoof, N. W. Schluger, A. Volpi, J.-J. Yim, I. B. Martin, et al. The role of chest imaging in patient management during the covid-19 pandemic: a multinational consensus statement from the fleischner society. *Radiology*, 296(1):172–180, 2020.

- C. Sammut and G. I. Webb. *Encyclopedia of machine learning*. Springer Science & Business Media, 2011.
- S. N. Saw and K. H. Ng. Current challenges of implementing artificial intelligence in medical imaging. *Physica Medica*, 100:12–17, 2022.
- J. Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- A. Shamayleh, M. Awad, and J. Farhat. Iot based predictive maintenance management of medical equipment. *Journal of Medical Systems*, 44:1–12, 2020.
- A. Signoroni, M. Savardi, S. Benini, N. Adami, R. Leonardi, P. Gibellini, F. Vaccher, M. Ravanelli, A. Borghesi, R. Maroldi, et al. Bs-net: Learning covid-19 pneumonia severity on a large chest x-ray dataset. *Medical Image Analysis*, 71:102046, 2021.
- K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- R. Smith-Bindman, D. L. Miglioretti, and E. B. Larson. Rising use of diagnostic medical imaging in a large integrated health system. *Health affairs*, 27(6):1491–1502, 2008.
- P. Soda, N. C. D’Amico, J. Tessadori, G. Valbusa, V. Guarrasi, C. Bortolotto, M. U. Akbar, R. Sicilia, E. Cordelli, D. Fazzini, et al. Aiforcovid: Predicting the clinical outcomes in patients with covid-19 applying ai to chest-x-rays. an italian multicentre study. *Medical image analysis*, 74: 102216, 2021.
- A. Sriram, M. Muckley, K. Sinha, F. Shamout, J. Pineau, K. J. Geras, L. Azour, Y. Aphinyanaphongs, N. Yakubova, and W. Moore. Covid-19 prognosis via self-supervised representation learning and multi-image prediction. *arXiv preprint arXiv:2101.04909*, 2021.
- S. A. Stansfield. Angy: A rule-based expert system for automatic segmentation of coronary vessels from digital subtracted angiograms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (2):188–199, 1986.
- S. Stephanie, T. Shum, H. Cleveland, S. R. Challa, A. Herring, F. L. Jacobson, H. Hatabu, S. C. Byrne, K. Shashi, T. Araki, et al. Determinants of

- chest radiography sensitivity for covid-19: a multi-institutional study in the united states. *Radiology: Cardiothoracic Imaging*, 2(5):e200337, 2020.
- S. Stephanie, T. Shum, H. Cleveland, S. R. Challa, A. Herring, F. L. Jacobson, H. Hatabu, S. C. Byrne, K. Shashi, T. Araki, J. A. Hernandez, C. S. White, R. Hossain, A. R. Hunsaker, and M. M. Hammer. Progressive pulmonary involvement, 2022. URL <https://doi.org/10.1148/ryct.2020200337>.
- R. Suarez-Bertoa, V. Valverde, M. Clairotte, J. Pavlovic, B. Giechaskiel, V. Franco, Z. Kregar, and C. Astorga. On-road emissions of passenger cars beyond the boundary conditions of the real-driving emissions test. *Environmental Research*, 176:108572, 2019. ISSN 0013-9351. doi: <https://doi.org/10.1016/j.envres.2019.108572>. URL <https://www.sciencedirect.com/science/article/pii/S001393511930369X>.
- G. A. Susto, A. Schirru, S. Pampuri, S. McLoone, and A. Beghi. Machine learning for predictive maintenance: A multiple classifier approach. *IEEE Transactions on Industrial Informatics*, 11(3):812–820, 2015. doi: 10.1109/TII.2014.2349359.
- C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- S. Tambe, A. Cao, R. McCaslin, and T. Edwards. An extensible cbm architecture for naval fleet maintenance using open standards. 2015.
- X. Tang, M. Xiao, Y. Liang, H. Zhu, and J. Li. Online updating belief-rule-base using bayesian estimation. *Knowledge-Based Systems*, 171:93–105, 2019. ISSN 0950-7051. doi: <https://doi.org/10.1016/j.knosys.2019.02.007>. URL <https://www.sciencedirect.com/science/article/pii/S0950705119300528>.
- J. M. Tanner. Assessment of skeletal maturity and predicting of adult height (tw2 method). *Prediction of adult height*, pages 22–37, 1983.
- The Johns Hopkins University. Positron emission tomography (pet), 2023. URL <https://www.hopkinsmedicine.org/health/treatment-tests-and-therapies/positron-emission-tomography-pet>.

- T. Tinga. Application of physical failure models to enable usage and load based maintenance. *Reliability engineering & system safety*, 95(10):1061–1075, 2010.
- T. Tinga. Predictive maintenance of military systems based on physical failure models. *CHEMICAL ENGINEERING*, 33, 2013.
- T. Tinga, A. Homborg, M. Woldman, N. Heerink, and M. Smeding. Advanced predictive maintenance concepts based on the physics of failure. *Optimal deployment of military systems. Technologies for military missions in the next decade*, pages 291–314, 2014.
- E. A. Tjoa, I. P. Y. N. Suparta, R. Magdalena, and N. K. CP. The use of clahe for improving an accuracy of cnn architecture for detecting pneumonia. In *SHS Web of Conferences*, volume 139, page 03026. EDP Sciences, 2022.
- J. F. Torres, D. Hadjout, A. Sebaa, F. Martínez-Álvarez, and A. Troncoso. Deep learning for time series forecasting: a survey. *Big Data*, 9(1):3–21, 2021.
- D. Tricarico, S. Longo, M. Giroto, and M. Melis. Tool and method to analyse medical images of patients, Oct. 21 2021. EP3901962A1.
- D. Tricarico, M. Calandri, M. Barba, C. Piatti, C. Geninatti, D. Basile, M. Gatti, M. Melis, and A. Veltri. Convolutional neural network-based automatic analysis of chest radiographs for the detection of covid-19 pneumonia: A prioritizing tool in the emergency department, phase i study and preliminary “real life” results. *Diagnostics*, 12(3):570, 2022a.
- D. Tricarico, H. A. H. Chaudhry, A. Fiandrotti, and M. Grangetto. Deep regression by feature regularization for covid-19 severity prediction. In *Image Analysis and Processing. ICIAP 2022 Workshops: ICIAP International Workshops, Lecce, Italy, May 23–27, 2022, Revised Selected Papers, Part II*, pages 496–507. Springer, 2022b.
- D. Tricarico, A. Neri, G. Tomasino, D. P. Cavallo, and D. Gionta. Image identification and retrieval for component fault analysis, June 7 2022c. US Patent 11,354,796.
- M. Tutuianu, P. Bonnel, B. Ciuffo, T. Haniu, N. Ichikawa, A. Marotta, J. Pavlovic, and H. Steven. Development of the world-wide harmonized light duty test cycle (wltc) and a possible pathway for its introduction

- in the european legislation. *Transportation Research Part D: Transport and Environment*, 40:61–75, 2015. ISSN 1361-9209. doi: <https://doi.org/10.1016/j.trd.2015.07.011>. URL <https://www.sciencedirect.com/science/article/pii/S1361920915001030>.
- University of Freiburg. U-net: Convolutional networks for biomedical image segmentation, 2011. URL <https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/>.
- A. S. C. d. B. Università degli Studi di Brescia. The brixia covid-19 project - dataset, 2022. URL <https://brixia.github.io/#dataset>.
- Università degli studi di Torino. Intelligenza artificiale a fianco del radiologo per "smascherare" rapidamente il paziente covid con malattia polmonare: parte la sperimentazione al san luigi, 2020. URL [https://www.unito.it/comunicati\\_stampa/intelligenza-artificiale-fianco-del-radiologo-smascherare-rapidamente-il-](https://www.unito.it/comunicati_stampa/intelligenza-artificiale-fianco-del-radiologo-smascherare-rapidamente-il-)
- G. Vachtsevanos, F. Lewis, M. Roemer, A. Hess, and B. wu. Intelligent fault diagnosis and prognosis for engineering systems: Methods and case studies. 13, 01 2006. doi: 10.1002/9780470117842.
- G. J. Vachtsevanos and K. P. Valavanis. A novel approach to integrated vehicle health management. In *AVT-305 Research Specialists Meeting on Sensing Systems for Integrated Vehicle Health Management for Military Vehicles*, Belgium, 2018. NATO Science & Technology Organization.
- N. Vafaei, R. A. Ribeiro, and L. M. Camarinha-Matos. Fuzzy early warning systems for condition based maintenance. *Computers & Industrial Engineering*, 128:736–746, 2019. ISSN 0360-8352. doi: <https://doi.org/10.1016/j.cie.2018.12.056>. URL <https://www.sciencedirect.com/science/article/pii/S0360835218306594>.
- R. A. Vingerhoeds, P. Janssens, B. D. Netten, and M. A. Fernández-Montesinos. Enhancing off-line and on-line condition monitoring and fault diagnosis. *Control Engineering Practice*, 3:1515–1528, 1995.
- S. VJ et al. Deep learning algorithm for covid-19 classification using chest x-ray images. *Computational and Mathematical Methods in Medicine*, 2021, 2021.
- J. Waller, A. O'Connor, E. Raafat, A. Amireh, J. Dempsey, C. Martin, and M. Umair. Applications and challenges of artificial intelligence in

- diagnostic and interventional radiology. *Polish Journal of Radiology*, 87 (1):113–117, 2022.
- G. Wang, X. Liu, J. Shen, C. Wang, Z. Li, L. Ye, X. Wu, T. Chen, K. Wang, X. Zhang, et al. A deep-learning pipeline for the diagnosis and discrimination of viral, non-viral and covid-19 pneumonia from chest x-ray images. *Nature biomedical engineering*, 5(6):509–521, 2021a.
- S. Wang, B. Kang, J. Ma, X. Zeng, M. Xiao, J. Guo, M. Cai, J. Yang, Y. Li, X. Meng, et al. A deep learning algorithm using ct images to screen for corona virus disease (covid-19). *European radiology*, pages 1–9, 2021b.
- M. N. Wernick, Y. Yang, J. G. Brankov, G. Yourganov, and S. C. Strother. Machine learning in medical imaging. *IEEE signal processing magazine*, 27(4):25–38, 2010.
- WHO. Who director-general’s opening remarks at the media briefing on covid-19 - 11 march 2020, 2020. URL [https://www.who.int/director-general/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19-](https://www.who.int/director-general/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19)
- WHO. Coronavirus, 2023a. URL [www.who.int/health-topics/coronavirus](http://www.who.int/health-topics/coronavirus).
- WHO. Who coronavirus (covid-19) dashboard, 2023b. URL <https://covid19.who.int/>.
- WHO. Strengthening medical imaging, 2023c. URL <https://www.who.int/activities/strengthening-medical-imaging>.
- Wikipedia. Dna damage leads to aging, cancer or apoptosis - dna damage (naturally occurring) - wikipedia, 2018. URL [https://commons.wikimedia.org/wiki/File:DNA\\_damage\\_leads\\_to\\_Aging,\\_Cancer\\_or\\_Apoptosis.jpg](https://commons.wikimedia.org/wiki/File:DNA_damage_leads_to_Aging,_Cancer_or_Apoptosis.jpg).
- Wikipedia. Common rail, 2023. URL [https://en.wikipedia.org/wiki/Common\\_rail](https://en.wikipedia.org/wiki/Common_rail).
- Wikipedia contributors. Oversampling and undersampling in data analysis — Wikipedia, the free encyclopedia, 2023a. URL [https://en.wikipedia.org/w/index.php?title=Oversampling\\_and\\_undersampling\\_in\\_data\\_analysis&oldid=1172771210](https://en.wikipedia.org/w/index.php?title=Oversampling_and_undersampling_in_data_analysis&oldid=1172771210). [Online; accessed 10-September-2023].

- Wikipedia contributors. Evaluation measures (information retrieval) — Wikipedia, the free encyclopedia, 2023b. URL [https://en.wikipedia.org/w/index.php?title=Evaluation\\_measures\\_\(information\\_retrieval\)&oldid=1146187267](https://en.wikipedia.org/w/index.php?title=Evaluation_measures_(information_retrieval)&oldid=1146187267). [Online; accessed 2-June-2023].
- Wikipedia contributors. Medical imaging — Wikipedia, the free encyclopedia, 2023c. URL [https://en.wikipedia.org/w/index.php?title=Medical\\_imaging&oldid=1148816756](https://en.wikipedia.org/w/index.php?title=Medical_imaging&oldid=1148816756). [Online; accessed 14-April-2023].
- Wikipedia contributors. Self-supervised learning — Wikipedia, the free encyclopedia, 2023d. URL [https://en.wikipedia.org/w/index.php?title=Self-supervised\\_learning&oldid=1149327474](https://en.wikipedia.org/w/index.php?title=Self-supervised_learning&oldid=1149327474). [Online; accessed 17-April-2023].
- Wikipedia contributors. Support vector machine — Wikipedia, the free encyclopedia, 2023e. URL [https://en.wikipedia.org/w/index.php?title=Support\\_vector\\_machine&oldid=1144271534](https://en.wikipedia.org/w/index.php?title=Support_vector_machine&oldid=1144271534). [Online; accessed 17-April-2023].
- Wikipedia contributors. Transfer learning — Wikipedia, the free encyclopedia, 2023f. URL [https://en.wikipedia.org/w/index.php?title=Transfer\\_learning&oldid=1148132688](https://en.wikipedia.org/w/index.php?title=Transfer_learning&oldid=1148132688). [Online; accessed 17-April-2023].
- M. Woo, M. Heo, A. M. Devane, S. C. Lowe, and R. W. Gimbel. Retrospective comparison of approaches to evaluating inter-observer variability in ct tumour measurements in an academic health centre. *BMJ open*, 10(11):e040096, 2020.
- W.-J. Wu, S.-W. Lin, and W. K. Moon. Combining support vector machine with genetic algorithm to classify ultrasound breast tumor images. *Computerized Medical Imaging and Graphics*, 36(8):627–633, 2012.
- W. Yiwei, G. Christian, N. Binaud, B. Christian, and F. Jian. A model-based prognostics method for fatigue crack growth in fuselage panels. *Chinese Journal of Aeronautics*, 32(2):396–408, 2019.
- M.-Y. You, F. Liu, W. Wang, and G. Meng. Statistically planned and individually improved predictive maintenance management for continuously monitored degrading systems. *IEEE Transactions on Reliability*, 59(4):744–753, 2010.



- W. Yu and T. Harris. A new stress-based fatigue life model for ball bearings. *Tribology Transactions - TRIBOL TRANS*, 44:11–18, 01 2001. doi: 10.1080/10402000108982420.
- M. F. Zarandi, M. Zarinbal, and M. Izadi. Systematic image processing for diagnosing brain tumors: A type-ii fuzzy expert system approach. *Applied soft computing*, 11(1):285–294, 2011.
- J. M. Zerlin and R. J. Hernandez. Approach to skeletal maturation. *Hand clinics*, 7(1):53–62, 1991.
- H. Zhang, C. Wu, Z. Zhang, Y. Zhu, H. Lin, Z. Zhang, Y. Sun, T. He, J. Mueller, R. Manmatha, et al. Resnest: Split-attention networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2736–2746, 2022.
- J. Zhang, K.-K. Ma, M.-H. Er, and V. Chong. Tumor segmentation from magnetic resonance imaging by learning via one-class support vector machine. In *International Workshop on Advanced Image Technology (IWAIT'04)*, pages 207–211, 2004.
- Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, and J. Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 3–11. Springer, 2018.
- H. Zhu, M.-D. Luo, R. Wang, A.-H. Zheng, and R. He. Deep audio-visual learning: A survey. *International Journal of Automation and Computing*, 18:351–376, 2021.
- M. Zhu. Recall, precision and average precision. *Department of Statistics and Actuarial Science, University of Waterloo, Waterloo*, 2(30):6, 2004.
- H. Zunair and A. B. Hamza. Synthesis of covid-19 chest x-rays using unpaired image-to-image translation. *Social network analysis and mining*, 11:1–12, 2021.