



TÉCNICO
LISBOA

Machine learning techniques for damage detection in wind turbine blades

Bernardo Henrique Santos Lopes

Thesis to obtain the Master of Science Degree in

Mechanical Engineering

Supervisors: Eng. André Guilherme Massano Tavares
Prof. Nuno Miguel Rosa Pereira Silvestre

Examination Committee

Chairperson: Prof. Paulo Rui Alves Fernandes
Supervisor: Prof. Nuno Miguel Rosa Pereira Silvestre
Member of the Committee: Prof. Elói João Faria Figueiredo

December 2021

Dedicado à minha família

Acknowledgments

First of all, I would like to thank Instituto Superior Técnico, particularly Professor Nuno Silvestre, for awakening my interest in structural mechanics and making this thesis possible.

Furthermore, I want to show my gratitude to Siemens Industry Software NV for the internship opportunity, which became, in every sense, an enriching life experience. With a special appreciation for the dedication and endless patience of Eng. André Tavares who closely followed the development of this thesis.

Resumo

A detecção de dano desempenha um papel vital no funcionamento da maquinaria e representa um dos maiores desafios nas indústrias mecânica, aeroespacial e aeronáutica. A indústria da energia eólica, onde as turbinas eólicas são sujeitas a enormes cargas mecânicas e aerodinâmicas e condições ambientais extremas durante o seu funcionamento, é uma das indústrias que mais pode beneficiar com este tipo de estudo. A instalação de turbinas eólicas requer investimentos significativos durante um longo período no que diz respeito a manutenção e reparações. De forma a reduzir estes custos e garantir a integridade e a longevidade de tais estruturas, a utilização de um sistema de monitorização de saúde estrutural fiável, associado a uma metodologia robusta de detecção de defeitos, é crucial para avaliar a viabilidade dos investimentos durante a sua vida útil. Para a posterior detecção de defeitos, inúmeras técnicas de machine learning e deep learning foram bem sucedidas numa vasta gama de aplicações.

Esta tese visa implementar algoritmos capazes de consistentemente detectar e avaliar falhas em duas pás de turbina eólica de material diferente. Técnicas como multivariate Gaussian anomaly detection e anomaly detection autoencoders foram incorporadas com sucesso através do reconhecimento de padrões de desvio do estado saudável para o seu estado danificado em dados adquiridos experimentalmente, tais como parâmetros modais e funções de resposta em frequência. Estes dados de vibração foram obtidos através de ensaios com um shaker modal, ensaios com um martelo modal e ensaios de pull-and-release aplicando análise modal experimental e análise modal operacional.

Palavras-chave: detecção de dano, pá de turbina eólica, análise modal, monitorização de saúde estrutural, machine learning, deep learning.

Abstract

Damage detection plays a vital role in the operation of machinery and represents one of the greatest challenges in the mechanical, aerospace, and aeronautical industries. The wind energy industry, where wind turbines are subjected to enormous mechanical and aerodynamic loads and extreme environmental conditions during operation, is one of the industries that can benefit most from this kind of study. The installation of wind turbines requires significant investments over a long period concerning maintenance and repairs. In order to reduce these costs and guarantee the integrity and longevity of such structures, the use of a reliable structural health monitoring system coupled with a robust anomaly detection methodology is crucial to assess the viability of the investments during their lifespan. Regarding the subsequent fault detection, numerous machine learning and deep learning techniques have succeeded in a broad range of applications.

This thesis aims to implement algorithms capable of consistently detecting and assessing faults in two wind turbine blades of different materials. Such techniques as multivariate Gaussian anomaly detection and anomaly detection autoencoders were successfully incorporated by recognising deviating patterns from the healthy state to its damaged state in experimentally acquired data such as modal parameters and frequency response functions. These vibration-based data were obtained through modal shaker testing, modal hammer testing and pull-and-release testing applying experimental modal analysis and operational modal analysis.

Keywords: damage detection, wind turbine blade, modal analysis, structural health monitoring, machine learning, deep learning.

Contents

- Acknowledgments v
- Resumo vii
- Abstract ix
- List of tables xiii
- List of figures xv
- Nomenclature xix

- 1 Introduction 1**
- 1.1 Motivation 1
 - 1.1.1 Wind energy 1
 - 1.1.2 Wind turbines 2
- 1.2 Topic overview 3
 - 1.2.1 Damage detection 3
 - 1.2.2 Machine learning 3
- 1.3 Objectives 5
- 1.4 Thesis outline 6

- 2 Theoretical background 7**
- 2.1 Modal analysis 7
 - 2.1.1 Single degree of freedom systems 7
 - 2.1.2 Multiple degrees of freedom systems 10
 - 2.1.3 Modal assurance criterion 12
 - 2.1.4 Experimental modal analysis 13
 - 2.1.5 Operational modal analysis 14
 - 2.1.6 Testing procedures 16
- 2.2 Machine learning 18
 - 2.2.1 Multivariate Gaussian anomaly detection 20
 - 2.2.2 Anomaly detection autoencoders 29

3	Experimental campaigns	32
3.1	Experimental setup	33
3.2	Modal hammer testing	36
3.3	Modal shaker testing	37
3.4	Pull-and-release testing	39
4	Modal parameter estimation	41
4.1	Titanium wind turbine blade	41
4.2	GFRP wind turbine blade	46
4.3	Modal tracking	51
4.3.1	Methodology	51
4.3.2	Results	54
5	Machine learning techniques	55
5.1	Multivariate Gaussian anomaly detection	55
5.1.1	Methodology	55
5.1.2	Results	61
5.2	Anomaly detection autoencoders	63
5.2.1	Methodology	63
5.2.2	Results	70
6	Conclusions	72
6.1	Achievements	73
6.2	Future work	74
	References	75

List of tables

2.1	Confusion matrix.	27
3.1	Testlab acquisition parameters.	32
3.2	Masses used in the tests of the titanium blade vs the GFRP blade.	35
5.1	MGAD accuracies for the different combinations of intervals of magnitude and position of the GFRP blade.	62
5.2	MGAD accuracies for the different combinations of intervals of magnitude and position of the Ti blade.	62
5.3	Accuracy comparison of the algorithm architectures.	70
5.4	Comparison of the preselection of algorithm architectures.	71

List of figures

1.1	European annual gross wind energy capacity installations per country according to WindEurope's scenario [3].	1
1.2	Horizontal axis wind turbine (HAWT) farm.	2
1.3	Wind turbine blade assembly (a) and maintenance (b).	2
1.4	Deep learning as a subset of machine learning in the context of artificial intelligence. . . .	4
1.5	Thesis overview schema.	5
2.1	Single degree of freedom system.	7
2.2	Real and imaginary part of a <i>receptance</i> function.	10
2.3	EMA schema.	14
2.4	OMA concept applied to a wind turbine.	14
2.5	Output matrix Y equal to the system matrix H, for white noise input.	15
2.6	Summary of EMA vs OMA processes [56].	16
2.7	Example of an experimental testing procedure for the acquisition of vibration-based data. . . .	17
2.8	Machine learning algorithm process.	18
2.9	Representation of a) supervised learning and b) unsupervised learning algorithms.	19
2.10	Illustration of the marginal and joint probability of a bivariate Gaussian distribution.	20
2.11	Multivariate Gaussian distribution for an identity covariance matrix and a null mean vector. . . .	21
2.12	Multivariate Gaussian distribution for a covariance matrix with the diagonal values decreased.	21
2.13	Multivariate Gaussian distribution for a covariance matrix with the diagonal values increased. . . .	22
2.14	Multivariate Gaussian distribution for a covariance matrix with first entry decreased and a unit value in the second entry of the mean vector.	22
2.15	Multivariate Gaussian distribution for a covariance matrix with first entry increased and a unit value in the second entry of the mean vector.	23
2.16	Multivariate Gaussian distribution for a covariance matrix with positive non-diagonal entries and the two entries of the mean vector changed.	23
2.17	Multivariate Gaussian distribution for a covariance matrix with negative non-diagonal entries and the two entries of the mean vector changed.	24

2.18 Univariate vs multivariate Gaussian distribution threshold contours.	25
2.19 Illustrative 2D dataset.	25
2.20 Gaussian distribution contours.	26
2.21 Anomaly classification.	27
2.22 Visualisation of an illustrative threshold for the marginal and joint probability of a bivariate Gaussian distribution.	28
2.23 Schema of a basic autoencoder.	29
2.24 Example of an autoencoder architecture for an image compression application [64].	30
3.1 Model 356A15 PCB accelerometer.	32
3.2 Simcenter SCADAS Recorder.	32
3.3 Titanium wind turbine blade.	33
3.4 Masses used to simulate damage on the Ti blade.	33
3.7 Mass positions used to perform the damaged experimental runs of the Ti blade.	33
3.5 Magnitude of the masses used in the Ti blade tests.	34
3.6 Mass of 127.0g on the position 7 of the Ti blade.	34
3.8 GFRP wind turbine blade.	34
3.9 Schema of the section of a WT blade [67].	34
3.10 Schematics of the manufacturing of a WT blade [68].	34
3.11 Masses used to simulate damage on the GFRP blade.	35
3.12 Mass positions used to perform the damaged experimental runs of the GFRP blade.	35
3.13 Modal hammer tips.	36
3.14 Model 86C03 PCB modal hammer.	36
3.15 Input spectrum of the impact testing performed with different hammer tip materials.	36
3.16 Coherence function of the impact testing performed with different hammer tip materials.	36
3.17 Random amplitude signal.	37
3.18 Chirp signal.	37
3.19 FRFs of the GFRP blade for a periodic chirp vs continuous random signal.	38
3.20 GFRP blade shaker setup.	38
3.21 Ti blade shaker setup.	39
3.22 FRFs of the Ti blade for a periodic chirp vs continuous random signal.	39
3.23 GFRP blade pull-and-release setup.	40
3.24 Ti blade pull-and-release setup.	40
4.1 FRFs of the three accelerometer directions and their sum.	41
4.2 Ti blade impact and shaker testing FRF sums.	42
4.3 Ti blade pull-and-release testing crosspower sum.	42

4.4	Ti blade FRFs from healthy vs damaged experimental runs.	42
4.5	Ti blade stabilisation diagram clustered with the representative poles.	43
4.6	Ti blade modal parameters.	44
4.7	Auto-MAC for a healthy experimental run of the Ti blade.	44
4.8	Ti blade mode shapes.	44
4.9	Natural frequencies of the 1 st mode of the Ti blade.	45
4.10	Natural frequencies of the 2 nd mode of the Ti blade.	45
4.11	Natural frequencies of the 4 th mode of the Ti blade.	46
4.12	Natural frequencies of the 7 th mode of the Ti blade.	46
4.13	GFRP blade impact and shaker testing FRF sums.	46
4.14	GFRP blade pull-and-release testing crosspower sum.	46
4.15	GFRP blade FRFs from healthy vs damaged experimental runs.	47
4.16	GFRP blade stabilisation diagram clustered with the representative poles.	47
4.17	GFRP blade modal parameters.	48
4.18	Auto-MAC for a healthy experimental run of the GFRP blade.	48
4.19	GFRP blade mode shapes.	49
4.20	Natural frequencies of the 1 st mode of the GFRP blade.	50
4.21	Natural frequencies of the 2 nd mode of the GFRP blade.	50
4.22	Natural frequencies of the 4 th mode of the GFRP blade.	51
4.23	Natural frequencies of the 12 th mode of the GFRP blade.	51
4.24	Modal tracking script methodology diagram.	52
4.25	Example of a <i>frequency score</i> assignment for the fourth Ti blade reference mode.	53
4.26	Example of a <i>MAC score</i> assignment for the fourth Ti blade reference mode.	53
5.1	MGAD methodology diagram.	55
5.2	MGAD subset processing diagram.	56
5.3	Multivariate Gaussian anomaly detection applied to the natural frequencies of the first two normal modes of the Ti blade.	57
5.4	Natural frequencies of the 14 th mode of the GFRP blade.	58
5.5	Natural frequencies of the 11 th mode of the GFRP blade.	58
5.6	Multivariate Gaussian distribution before the oversampling.	59
5.7	MGAD classification before the oversampling.	59
5.8	Transformation of a flat Gaussian distribution (blue) into a sharp distribution curve (red) due to oversampling.	60
5.9	Oversampled multivariate Gaussian distribution.	60
5.10	Oversampled MGAD classification.	61
5.11	ADAE methodology diagram.	63

5.12 ADAE reconstruction error principle.	64
5.13 ADAE FRF reconstructions for healthy inputs from the Ti blade.	65
5.14 ADAE FRF reconstructions for healthy inputs from the GFRP blade.	65
5.15 ADAE FRF reconstructions for damaged inputs from the Ti blade.	66
5.16 ADAE FRF reconstructions for damaged inputs from the GFRP blade.	66
5.17 Sum of the reconstruction RMSE from every healthy vs damaged experimental run for one of the accelerometer directions of the Ti blade.	66
5.18 Sum of the reconstruction RMSE from every healthy vs damaged experimental run for one of the accelerometer directions of the GFRP blade.	67
5.19 RMSE for the different experimental runs of the Ti blade.	67
5.20 RMSE for the different experimental runs of the GFRP blade.	68
5.21 ROC curve for the Ti and GFRP blades.	69
5.22 ROC curves for different classifier performances.	69
5.23 Threshold corresponding to the maximum IoU for one of the directions of an accelerometer.	69

Nomenclature

ADAE Anomaly detection autoencoder.

AE Autoencoder.

AMPS Automatic modal parameter selection.

AUC Area under the curve.

DL Deep learning.

DOF Degree of freedom.

EMA Experimental modal analysis.

FRF Frequency response function.

GFRP Glass-fiber reinforced plastic.

IoU Intersection over union.

MAC Modal assurance criterion.

MDOF Multiple degree of freedom.

MGAD Multivariate Gaussian anomaly detection.

ML Machine learning.

OMA Operational modal analysis.

RMSE Root-mean-square error.

ROC Receiver operating characteristic.

SDOF Single degree of freedom.

SHM Structural health monitoring.

Ti Titanium.

WT Wind turbine.

Chapter 1

Introduction

1.1 Motivation

1.1.1 Wind energy

The wind is one of the most appealing renewable energy sources for electricity production, representing a significant part of the power generation mix in some countries. For example, in 2017, wind energy in Portugal accounted for 21.6% of total electric energy production, covering 24.6% of the annual electricity demand [1]. In Europe, there is currently 205 GW of wind energy capacity. An additional 88 GW of net capacity will be added in Europe in the next five years, bringing the total capacity to 277 GW. Moreover, by 2030, the EU will have installed 323 GW of wind energy. Wind energy currently meets 15% of EU electricity demand; by 2030, it is expected to cover 30% of demand and employ 569,000 people [2]. European countries will begin implementing their 2030 National Energy and Climate Plans (NECPs) toward the 32% renewables target by the end of 2020. Wind energy will play a critical role over the coming years, allowing many European countries to continue decarbonising their power systems and transforming their energy systems in a cost-effective way [3].




Figure 1.1: European annual gross wind energy capacity installations per country according to WindEurope's scenario [3].

1.1.2 Wind turbines

Wind turbines are not only subject to mechanical and aerodynamic loads during operation but also extreme environmental conditions such as gusts, icing and lightning strikes. The main reasons for wind turbine (WT) damage are thermal cycling, sand, bird impacts, leading and trailing edge erosion, fatigue, moisture intrusion, and mechanical failure [2].

The installation of WTs requires significant investments during a minimum of a 20-year lifespan. To assess the reliability of the investments during this period, the use of structural health monitoring (SHM) systems to continuously monitor the condition of WTs is crucial. SHM can provide relevant information for rescheduling maintenance tasks and lowering its associated costs.

The maintenance of WT blades involves visual inspection, which can be dangerous and time-consuming. For horizontal axis wind turbines (HAWTs), standard scheduled maintenance can be performed twice a year and take up to 24 hours. In case of failure, unscheduled maintenances can be up to five times more expensive [4]. Regular WT *Operation and Maintenance* (O&M) costs are very high, accounting for 20–25% of the total



Figure 1.2: Horizontal axis wind turbine (HAWT) farm.

levelized cost of energy (LCOE) over the turbine's lifetime. High maintenance and repair costs raise the cost of energy, making wind energy less competitive [2, 5]. The average failure rate of an offshore WT is around eight times per year, and the blades are the fifth most common cause of failure, accounting for 6.2% of the failures (after the pitch and hydraulic system, auxiliary components, generator, and gearbox) [6]. In addition, around 3,800 WT blades are failing every year, representing 0.54% of the estimated 700,000 blades in operation worldwide [7]. The structural repair of a single wind blade can cost up to \$30,000, and the price of a new blade is, on average, around \$200,000 [2]. Therefore, developing a remote blade condition monitoring system is a topic that the industry is engaged in.



(a)



(b)

Figure 1.3: Wind turbine blade assembly (a) and maintenance (b).

1.2 Topic overview

1.2.1 Damage detection

Damage detected early on can prevent irreversible structural failures, hence the importance of a robust defect detection methodology for ensuring machinery's structural integrity and operational safety while lowering maintenance and repair costs. Reliable damage detection techniques are a significant challenge in the mechanical, aerospace, and aeronautical industries. In particular, for the WT industry, due to the high magnitude of the structures and respective investments [8–11].

Structural health monitoring

Structural health monitoring (SHM) can be defined as the acquisition, validation and analysis of technical data to facilitate life-cycle management of decisions. More broadly, it refers to a reliable system capable of detecting and evaluating adverse changes in a structure due to damage, or regular operation [7, 12, 13]. One way of carrying out this monitoring process is through the extraction and analysis of vibration-based data [1, 14–17].

Modal analysis

A general process of SHM based on modal analysis involves collecting relevant vibration response data from sensors attached to the surface undergoing an excitation from which is extracted modal information such as natural frequencies, damping ratios and mode shapes. Damaged structures will have these modal parameters affected due to local discontinuity created by the fault [18–22].

Experimental modal analysis (EMA) is a typical technique for estimating the modal properties of structures that have been used for many years. This technique assumes that the user knows both inputs and outputs; however, it is not always possible to measure both, particularly in operational conditions. For this reason, the operational modal analysis (OMA) technique was then introduced.

The OMA approach serves the same goals as the EMA, but it is more suitable for operating machines because it relies only on measured outputs and does not need to know about the input forces. The estimation of OMA parameters combined with a proper SHM strategy can be used to track unusual system variations and inform when unplanned maintenances are required. Considerable modal analysis research has already been done for WTs [4, 23–26].

1.2.2 Machine learning

In recent years, advanced computational artificial intelligence algorithms have achieved large success in detecting different types of faults [13, 27]. Fault detection using machine learning (ML) focuses on two main areas: anomaly detection and fault classification. While anomaly detection algorithms primarily focus on detecting whether or not an anomaly exists on the specimen, fault classification takes a step further in identifying both an anomaly and the type of anomaly that occurred [28]. The use of ML

techniques requires prior feature engineering. The choice of features and classifier has a significant impact on the performance of these damage detection algorithms. Furthermore, choosing and pre-processing fixed or hand-crafted features necessitates prior knowledge, expertise and may also require a substantial amount of computational power, which may limit their use in real-time analysis [29].

Deep learning

Deep learning (DL) is a subset of ML mainly focused on artificial neural network algorithms. DL goes a step further in terms of complexity and power, representing more complex architectures and algorithms such as deep neural networks, convolutional neural networks, and recurrent neural networks. As opposed to what was stated about ML, these techniques can learn a transformation or sequence of transformations from raw data, which does not imply better results than ML.




Figure 1.4: Deep learning as a subset of machine learning in the context of artificial intelligence.

Over the last few years, remarkable progress has been made in the field of DL with emerging applications for SHM and damage detection [30, 31]. The development of such DL techniques can be traced back to image recognition and signal analysis, where a large number of architecture designs for artificial neural networks have been created and successful results achieved [32–34]. In particular for damage detection algorithms using autoencoders [35–40]. Also, deep neural networks are being implemented in the mechanical and aerospace industries, for instance, in the damage detection of beam structures, [29, 41], damage assessment in rotating machinery and motor fault detection [28, 42–44], psychoacoustic analysis [45, 46], and other non-engineering applications [47]. Furthermore, feature extraction techniques have been employed to images generated from time-signals for SHM applications [46, 48, 49]. Studies have also been done for WT fault classification [50] and autonomous crack detection [51].

1.3 Objectives

- Understand the principles of modal analysis;
- Conduct experimental campaigns to collect vibration-based data from the blades;
- Perform modal analysis using state-of-the-art software;
- Explore different ML approaches to build a reliable and fully automatic damage detection model;
- Study which experimentally collected features are more relevant to the performance of fault detection;
- Compare the developed fault detection algorithms based on accuracy, sensibility and damage assessment capability;
- Assess the applicability of the developed models in a real-world scenario.

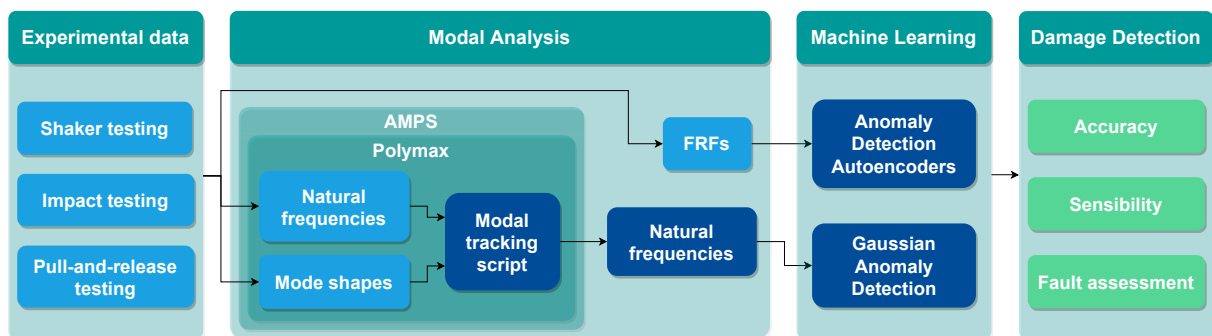


Figure 1.5: Thesis overview schema.

1.4 Thesis outline

The structure of this thesis is divided into five parts: a theoretical introductory part; a section about the experimental campaigns; a chapter regarding the post-processing of the experimental data; a section about the used ML techniques; and a final chapter with the relevant conclusions.

Firstly, section 2 introduces the theoretical concepts on which this thesis was based, namely modal analysis and the two ML algorithms used for damage detection.

Secondly, chapter Modal parameter estimation presents the two structures analysed: the titanium and glass-fiber reinforced plastic (GFRP) WT blades, as well as the experimental setups used to acquire the respective vibration-based data.

Thirdly, the post-processing and analysis of the modal parameters collected for the two blades are explained in section Modal parameter estimation.

In fourth place, chapter 5 shows in detail the ML algorithms used, their methodologies and results.

Finally, the general conclusions of the work and possible future improvements to be made are mentioned in the final chapter 6.

Chapter 2

Theoretical background

2.1 Modal analysis

When it comes to engineering design, structural vibration issues are one of the most critical concerns. This emphasises the importance of being able to predict and monitor the vibration levels during operation. As a result, a thorough understanding of each system's dynamic behaviour is essential. Modal analysis is a technique for examining a structure and estimating its dynamic properties such as natural frequencies, damping ratios, and mode shapes [52].

Almost all real-world situations can be modelled as systems with more than one degree of freedom (DOF). With that in mind, the elementary concept of single degree of freedom (SDOF) systems will be introduced, followed by the presentation of multiple degrees of freedom (MDOF) systems.

2.1.1 Single degree of freedom systems

The analysis of an oscillatory system's free response, or its behaviour when no external forces are applied, allows the modal parameters to be determined, which are, for an SDOF situation, the natural frequency ω_n and the damping ratio ξ .

Figure 2.7 illustrates a basic model for a viscously damped SDOF system composed of a mass m , a stiff spring k , and a damper with a damping coefficient c . The mass displacement is represented by $x(t)$, and the external force applied to the system is represented by $f(t)$, both as a function of time.




Figure 2.1: Single degree of freedom system.

The equilibrium of forces (balance of inertial, damping, elastic, and external forces) in this case is:

$$m\ddot{x}(t) + c\dot{x}(t) + kx(t) = f(t) \quad (2.1)$$

Where \dot{x} and \ddot{x} are, respectively, the velocity and acceleration of the system.

Considering that there is no displacement neither velocity in the initial moment, i. e. $x(0) = 0$ and $\dot{x}(0) = 0$, and applying the Laplace Transform, the equation 2.1 can be written as:

$$(ms^2 + cs + k)X(s) = F(s) \quad (2.2)$$

Where s represents the Laplace domain. Using Z as the dynamic stiffness of the system, it comes:

$$Z(s)X(s) = F(s) \quad (2.3)$$

Replacing $Z(s)$ by $H(s) = Z^{-1}(s)$ the result is:

$$X(s) = H(s)F(s) \quad (2.4)$$

$H(s)$ is the system's transfer function, and it relates the input force F to the output displacement X . In general, a system's transfer function is a complex-valued function that can be defined as the ratio between the Laplace Transform of the output (system response) and the input (system excitation), with zero initial conditions.

Going back to the equation 2.2 it is possible rewrite $H(s)$ in terms of the system dynamic properties:

$$H(s) = \frac{\frac{1}{m}}{s^2 + \frac{c}{m}s + \frac{k}{m}} \quad (2.5)$$

The roots of the *system characteristic equation* which is the denominator of the equation 2.5 represent the *system poles*. They are calculated using the following formula:

$$\lambda_{1,2} = -\frac{c}{2m} \pm \sqrt{\left(\frac{c}{2m}\right)^2 - \frac{k}{m}} \quad (2.6)$$

Given that the proposed system is conservative, i.e. it has no damping ($c = 0$), the positive solution of the equation 2.6 is:

$$\omega_n = \sqrt{\frac{k}{m}} \quad (2.7)$$

Where ω_n is the *undamped natural frequency* of the system.

The *critical damping*, c_0 , is obtained by equaling the term under the square root in the equation 2.6 to zero.

$$c_0 = 2m\sqrt{\frac{k}{m}} \quad (2.8)$$

The *damping ratio* ξ can also be used to define the system's damping in a dimensionless manner. This is the percentage of *critical damping* of the system:

$$\xi = \frac{c}{c_0} \quad (2.9)$$

The *damping ratio* value is used to classify systems because it determines their type of response. A system with $\xi = 0$ has no damping. The system is underdamped when the *damping ratio* is between 0 and 1. If the damping coefficient is equal to the *critical damping* ($\xi = 1$) it is classified as critically damped. The system is overdamped when the *damping ratio* is higher than one. Real-world systems with no active damping mechanism are normally underdamped, with a *damping ratio* of less than ten per cent ($\xi < 0.1$). As a result, all systems analysed from now on will be considered underdamped.

Going back to the equation of motion of the system (2.1) and solving it in the time domain:

$$x(t) = x_1 e^{\lambda_1 t} + x_2 e^{\lambda_2 t} \quad (2.10)$$

Where λ_1 and λ_2 are solutions of the homogeneous equation defined by the equation 2.6. That equation has two complex conjugate roots as a solution:

$$\lambda_1 = \sigma_1 + j\omega_1 \quad \text{and}; \quad \lambda_1^* = \sigma_1 - j\omega_1 \quad (2.11)$$

With σ_1 being the *damping factor* and ω_1 the *damped natural frequency*, also called *resonance*.

Knowing this, the transfer function defined by equation 2.5 can be written as:

$$H(s) = \frac{\frac{1}{m}}{(s - \lambda_1)(s - \lambda_1^*)} \quad (2.12)$$

According to partial fraction expansion theory:

$$H(s) = \frac{A_1}{(s - \lambda_1)} + \frac{A_1^*}{(s - \lambda_1^*)} \quad \text{with}; \quad A_1 = \frac{\frac{1}{m}}{j2\omega_1} \quad (2.13)$$

Where A_1 and A_1^* are the *residues*.

The previous concepts in the Laplace domain (s) may be rearranged in the frequency domain. The *Frequency Response Function (FRF)* is the transfer function analysed along the frequency axis $j\omega$ (when $\sigma = 0$):

$$H(p)|_{p=j\omega} = H(\omega) = \frac{A_1}{(j\omega - \lambda_1)} + \frac{A_1^*}{(j\omega - \lambda_1^*)} \quad (2.14)$$

It is possible to disregard the effect of the complex conjugate part and simplify the FRF expression

2.15 for frequencies close to the resonance frequency ($\omega \cong \omega_1$).

$$H(\omega) = \frac{A_1}{(j\omega - \lambda_1)} \quad (2.15)$$

According to the physic quantity used to measure the system response, there are three different types of FRFs [53]. The *receptance* is the function calculated in terms of the displacement relative to the input force. It can also called *dynamic compliance*. The FRF is referred to as *mobility* when measured in terms of velocity relative to force. When using the acceleration to measure the response, the FRF is called *accelerance* or *inertance*. Figure 2.2 depicts the typical behaviour of the FRF compliance near the resonance frequency for both the real and imaginary parts. A resonance peak is characterised in an FRF by a 0 in the real part, a peak in the imaginary part and a shift in phase of 180 degrees.




Figure 2.2: Real and imaginary part of a *receptance* function.

2.1.2 Multiple degrees of freedom systems

Problems with only one degree of freedom are extremely limited and nearly impossible to encounter in a real-world scenario. Hence, the concepts presented in the previous section 2.1.1 will be extended to systems with n degrees of freedom, where $n > 1$. In the time domain, the equation of motion for an MDOF system is the same as the equation of motion for an SDOF system, but it is a matrix equation:

$$[m]\{\ddot{x}\} + [c]\{\dot{x}\} + [k]\{x\} = \{f\} \quad (2.16)$$

Where $[m]$ is the mass matrix, $[c]$ the damping matrix and $[k]$ the stiffness matrix of the system. These matrices have $n \times n$ dimension. The vector $x(t)$ represents the displacement as a function of the time, and $f(t)$ is the excitation of the system also as a function of the time. Both have $n \times 1$ dimension.

The equation of motion can be transformed into the Laplace domain by assuming that all initial

displacements and velocities are zero for all degrees of freedom:

$$([m]s^2 + [c]s + [k])\{X(s)\} = \{F(s)\} \quad \text{or;} \quad [Z(s)]\{X(s)\} = \{F(s)\} \quad (2.17)$$

The dynamic stiffness matrix is represented by $[Z(s)]$. Inverting this equation results in:

$$\{X(s)\} = [H(s)]\{F(s)\} \quad (2.18)$$

Where $[H(s)]$ is the *transfer function matrix* because it represents the ratio of the system's response to the corresponding force input in the Laplace domain for each DOF. This matrix contains functions with complex values.

The *system characteristic equation* is the determinant of $[Z(s)]$, and its roots are the *system poles*, defining the values of the natural frequencies of the system. The solution for this problem is based on the eigenvalues of a matrix, so the first step is to define the equation 2.17 in the general eigenvalue problem formulation. Considering:

$$[A] = \begin{bmatrix} [0] & [m] \\ [m] & [c] \end{bmatrix} ; \quad [B] = \begin{bmatrix} -[m] & [0] \\ [0] & [k] \end{bmatrix} ; \quad \{Y\} = \begin{bmatrix} s\{X\} \\ \{X\} \end{bmatrix} \quad \text{and;} \quad \{F'\} = \begin{bmatrix} \{0\} \\ \{F\} \end{bmatrix}$$

While also taking into consideration that:

$$(s[m] - s[m])\{X\} = \{0\} \quad (2.19)$$

$$(s[A] + [B])\{Y\} = \{F'\} \quad (2.20)$$

Equation 2.20 is the general formulation for the eigenvalue problem when the *force vector* is null. The system's eigenvalues (λ_r , with $r = 1, \dots, n$) are the solutions of the equation 2.21, calculated for s .

$$\left| s[A] + [B] \right| = 0 \quad (2.21)$$

The number of solutions of this equation is determined by the system's degree of freedom, n . In complex conjugate pairs, there are $2n$ complex-valued eigenvalues. The *damping factor* is the real component of a pole, and the *damped natural frequency* is the imaginary part, as in the SDOF systems.

Each of the system's eigenvalues is associated with an eigenvector named *mode shape vector*, $\{\psi\}_r$. Like the system poles, they appear in complex conjugate pairs. The modal vectors contain complex-valued modal displacements, and the phase of these displacements might differ between elements. On the equation 2.17 these vectors make the *force vector* null by definition. The transfer function can also be

represented in terms of the system poles since they are the roots of the system characteristic equation:

$$[H(s)] = \frac{\text{adj}([Z(s)])}{\prod_{r=1}^n E(s - \lambda_r)(s - \lambda_r^*)} = \frac{\text{adj}([Z(s)])}{\prod_{r=1}^{2n} E(s - \lambda_r)} \quad (2.22)$$

Regarding SDOF systems, the partial fraction expansion theory can be applied to the expression 2.22 to obtain the *residues*, A_r . The following equation relates them to the *mode shape vectors*. For each residue:

$$[A]_r = Q_r \{\psi\}_r \{\psi\}_r^t \quad (2.23)$$

Since the *residues* $[A]_r$ are absolute quantities, it is possible to deduce from the equation 2.23 that the modal vectors are vectors with a scaling factor $[A]_r$. Each column j of the matrix $[A]$ has enough information to produce that matrix unless it coincides with a modal coefficient ψ_{jr} equal to zero. This occurs because all the columns are proportional to one another, but if the coefficient is 0, the corresponding row and column will be zero. As a result, if the excitation point is in a nodal point of one mode, that mode will not be recognised in experimental modal analysis. The FRF matrix should be obtained by evaluating the transfer function matrix along the frequency axis:

$$[H(j\omega)] = \sum_{r=1}^n \left(\frac{Q_r \{\psi\}_r \{\psi\}_r^t}{(j\omega - \lambda_r)} + \frac{Q_r^* \{\psi\}_r^* \{\psi\}_r^{*t}}{(j\omega - \lambda_r^*)} \right) \quad (2.24)$$

The matrix $[H(j\omega)]$ has complex values and depends on the input frequency, so its values are not constants. Taking into account that the FRF is defined as the output-to-input ratio. When an excitation F_k is applied on k DOF, $H_{jk}(\omega)$ is the response X_j for the j DOF. The input force and the system response (measured in displacement, velocity, or acceleration) are quantifiable. As a result, the equation 2.24 can be used to find the system's modal parameters. As previously stated, the FRF of a MDOF system is the sum of the FRFs of n SDOF systems.

2.1.3 Modal assurance criterion

The modal assurance criterion (MAC) is one of the most often used strategies for comparing modal vectors quantitatively [54]. The development of the MAC was modelled after the development of the conventional coherence calculation coupled with frequency response function computation. The MAC, like ordinary coherence, is a statistical indicator. This least squares-based form of linear regression analysis produces a more sensitive indicator to the largest difference between comparative values and a modal assurance criterion unaffected by small changes or small magnitudes. This criterion was first used in modal testing in conjunction with The Modal Scale Factor as an additional confidence factor when evaluating the modal vector from various excitation locations. The numerator of each term in an FRF matrix expressed in partial fraction expansion form corresponds to the matrix of residues or modal

constants. The MAC is calculated as the normalised scalar product of the vectors $\{\varphi_A\}$ and $\{\varphi_X\}$. The resulting scalars are arranged into the MAC matrix:

$$\text{MAC}(r, q) = \frac{\left| \{\varphi_A\}_r^T \{\varphi_X\}_q \right|^2}{\left(\{\varphi_A\}_r^T \{\varphi_A\}_r \right) \left(\{\varphi_X\}_q^T \{\varphi_X\}_q \right)} \quad (2.25)$$

Where the form of coherence function can be recognized, indicating the casual relationship between $\{\varphi_A\}$ and $\{\varphi_X\}$. The modulus in the numerator is taken after the vector multiplication, in this case, resulting in the squared absolute value of the sum of product elements. In addition, an equivalent formulation is:

$$\text{MAC}(A, X) = \frac{\left| \sum_{j=1}^n \{\varphi_A\}_j \{\varphi_X\}_j \right|}{\left(\sum_{j=1}^n \{\varphi_A\}_j^2 \right) \left(\sum_{j=1}^n \{\varphi_X\}_j^2 \right)} \quad (2.26)$$

The criterion has been used to quantify the accuracy of detected mode shapes as a Mode Shape Correlation Constant. For complex modes of vibration:

$$\text{MAC}(r, q) = \frac{\left| \{\psi_A\}_r^T \{\psi_X\}_q^* \right|^2}{\left(\{\psi_A\}_r^T \{\psi_A\}_r^* \right) \left(\{\psi_X\}_q^T \{\psi_X\}_q^* \right)} \quad (2.27)$$

The MAC assumes values between 0 when there is no consistent correspondence and 1 when representing a consistent correspondence. Small numbers imply poor resemblance between the two shapes, whereas values greater than 0.9 show consistent correspondence.

2.1.4 Experimental modal analysis

There are four possible approaches to analyse the vibrational behaviour of a system: analytically, using a finite element model, or through experimental modal analysis (EMA) and operational modal analysis (OMA).

The analytical method of conducting the modal analysis is to solve the equations that describe the system's dynamic because they can be written in terms of mass, stiffness, and damping matrices. One of the problems with this method is that very often (almost always for complex structures), it is not possible to build these matrices. Without them, it is impossible to calculate modal parameters analytically.

Another option is to use system models, which could be finite element models or modal models. Within the field of modal analysis, there is the OMA where the estimation of the modal parameters is done without knowing the input forces, and experimental modal analysis where both input and output forces are known. Regarding EMA, these measurements can be used to calculate the frequency domain ratio between the output and the input, resulting in the Frequency Response Function (FRF). EMA is based on the idea that these functions can be written in terms of modal parameters: natural frequency, damping ratio, and mode shapes [55]. The FRFs are complex frequency domain functions,

having both magnitude and phase information. The FRFs, can have the form of output acceleration, velocity or displacement over input force and describe how the structure moves at each measurement location per unit force at the input location using the input force as a reference. The resonance and damping information for all system modes is contained in the FRFs. Individual modes are chosen, the system's modal model is generated, and individual FRFs are synthesised from the modal model used for curve-fitting. The mode shapes for each of the identified modes can then be generated using the synthesised FRFs. One relevant feature of EMA is that it allows for the calculation of modal participation factors after the input force has been measured, and each mode shape can be properly mass-normalised or arbitrarily scaled [56].




Figure 2.3: EMA schema.

2.1.5 Operational modal analysis

Another modal analysis method is to measure the structure's response during its operation to determine the modal parameters. Operational modal analysis (OMA) is the term for this method, also known as output only modal analysis, because it is a method that does not require a measured external input to excite the structure. Thus, it can be used to monitor the condition of structures in operation for fatigue or fault detection and failure assessment, avoiding the interruption for measurements, relying solely on the structure's own ambient excitation. It is especially beneficial when the forces cannot be quantified or when exciting a structure is challenging [57]. This is the main advantage of this method. Another benefit of performing OMA is testing the structure in a real-world situation, challenging to replicate in a lab setting, making it a more realistic approach. Figure 2.4 illustrates the OMA concept, which depicts a WT excited by the wind. The purpose of OMA is, without knowing the input U , to be able to identify the structure H , from a dynamic perspective, by only studying the measured output signals of Y .




Figure 2.4: OMA concept applied to a wind turbine.

The principle of OMA, like EMA, is to create a mathematical model of the system's dynamic properties and use it to calculate its modal parameters. A curve-fitting technique is used to extract them. The model establishes the relation between the inputs $[F]$ (excitation applied to the system) and the outputs $[F']$ (system reactions to that excitation) in different points of the structure using frequency-domain functions. These functions are rearranged into the *system transfer matrix*, $[H]$, and are designated *auto-power spectra* and *cross-power spectra*. Since the input forces are not quantified in operational tests, the main difference between OMA and EMA is how this matrix is constructed. The input is known and measured in EMA; thus, the procedure is developed in a deterministic framework. However, the input is unknown and unmeasured in OMA; therefore, the procedures used are stochastic [58]. The concept underlying the OMA technique is to consider that the input forces are randomly dispersed in the space around the object and are in the form of white noise, knowing only the system's output during the test. This indicates that the input signal has the same magnitude across all frequencies in the bandwidth studied. With this


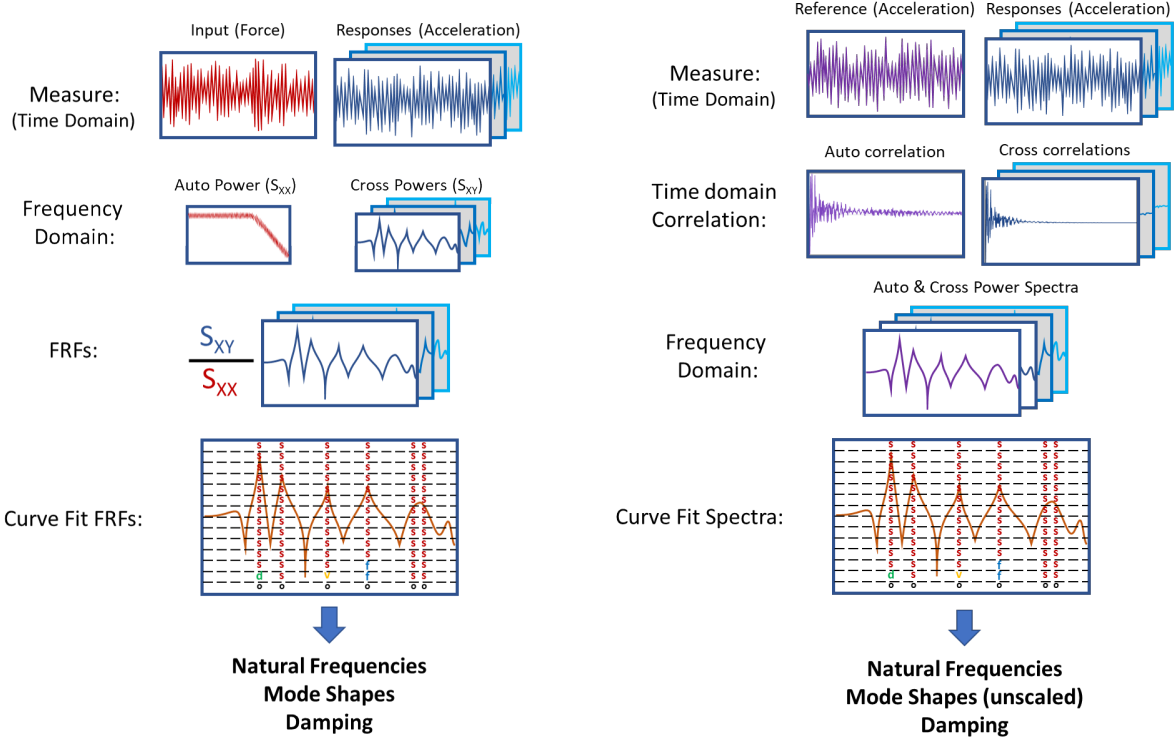


Figure 2.5: Output matrix Y equal to the system matrix H , for white noise input.

assumption, the system's responses contain all the information needed to characterise the system's dynamic, allowing the modal parameters to be extracted from a mathematical model generated solely by measuring the output. The FRFs cannot be calculated in this manner, as they are when using EMA, but the frequency information can still be extracted. The correlation functions are calculated using the structure's measured responses. These are time-domain functions that can be used as a statistical tool to identify recurring patterns in a signal that may be hidden inside a seemingly random signal. They can be used to compare a signal to itself in a delayed version with increasing time lags. Those functions are referred to as *autocorrelation function* (ACF). According to the level of correlation, this comparison returns a value between -1 and 1. When the signal contains periodic information, this number will be high at particular time lags. If the structure is excited with a random signal, increasing the delay between the signals should cause the correlation values to fall until they reach zero. The ACF will retrieve the common periodicities of the compared signals, which can then be translated to the frequency domain. The cross-correlation function works with the same principle as the ACF; the only difference is that it compares the signal to a reference signal rather than a delayed version of itself. This function extracts the common periodicities of responses obtained at two different points. One reference location is sufficient for OMA, but it is also possible to specify multiple measurement points as a reference. The system's measured response in different points of the structure is then used to calculate autocorrelation and cross-correlation functions and the common periods between the various measurements. These data is in the time domain; thus, it is converted to the frequency domain using the discrete Fourier transform

of the correlation functions. *Correlograms* are the functions that come from this procedure, and they display the dominant frequencies that are common across all measurement locations. These are *power spectra functions* (PSDs), which measure the signal's power content at each frequency, and the peaks may be calculated using curve-fitting technology, similarly to EMA.



(a) Experimental modal analysis.

(b) Operational modal analysis.

Figure 2.6: Summary of EMA vs OMA processes [56].

2.1.6 Testing procedures

Usually, while performing experimental modal analysis, it is impossible to measure the whole FRF matrix due to practical constraints. The number of modes observed during the testing is less than the number of outputs (number of points where the response is measured), which is also less than the number of degrees of freedom of the system [55].

Different types of inputs and outputs can be used to calculate the experimental FRFs. In most mechanical system tests, the inputs are in the form of force, while the outputs are acceleration, velocity, or displacement. A modal impact hammer or a modal shaker can apply force to the structure and quantify it. The method for measuring the response depends on the application. If the goal is to measure acceleration, the best approach is to utilise accelerometers. A laser Doppler vibrometer can be used for measuring the surface velocity [59], and string pots or Digital Image Correlation (DIC) [60] are two options for measuring displacements. The force spectrum of the excitation applied to the test object should be flat relative to frequency, meaning that the signal should excite all frequencies equally. This

way, the peaks in the response correspond to the system's natural frequencies. While using the hammer, the correct tip selection must be taken into account to stimulate the structure properly. The FRFs are normalised to the input force since they are all calculated with the same reference. Hence, the peaks of this function are the structure's resonant frequencies. It is common to calculate the FRFs numerous times and average them using the estimator method that best suits the test objectives to get more accurate modal parameter data. The coherence function is estimated to check the repeatability of each FRF since it indicates how much of the output is due to the input in the FRF and can be used as an indicator of the FRF's quality. The coherence is a number that ranges from 0 to 1. The measurements

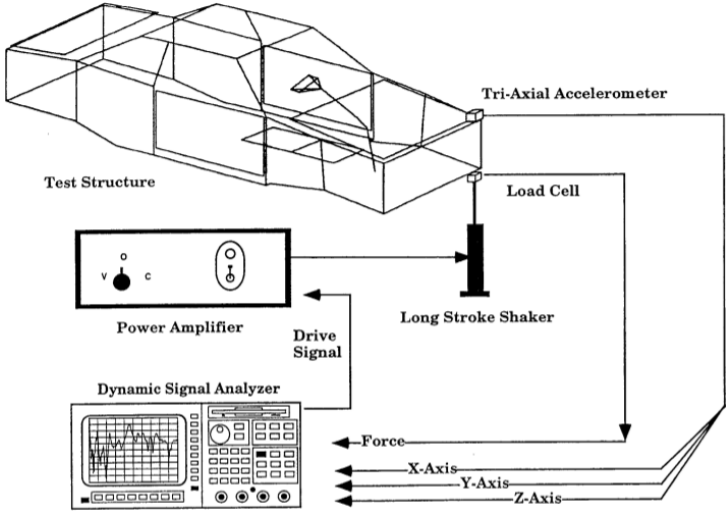


Figure 2.7: Example of an experimental testing procedure for the acquisition of vibration-based data.

are not repeatable when 0 or close to it, indicating that something is wrong with the setup. On the other hand, a coherence near 1 suggests that the FRF's amplitude and phase are consistent. The fact that the FRF can be replicated does not guarantee that the results will be identical. Structures are never entirely linear, and there are always extrinsic elements that influence experimental processes, such as instrumentation noise. Consequently, the results of repeated measurements will never be precisely the same. Therefore, estimators are required to calculate the average of the results and determine which is more accurate. The amplitude ratio (H) of the input to the output of FRFs is calculated using estimators; H1, H2, and HV are the most widely used. The H1 considers that all excitation measurements are accurate because there is no noise on the input and all of the noise is on the system's output. The estimator H2 assumes that there is no noise on the output, only in the input. Despite needing more processing effort, the HV estimator is better than the previous two since it takes into consideration noise on both the input and output. After obtaining the FRFs for the structure, the *natural frequencies* and *damping ratios* must be calculated. Different strategies can be utilised to estimate modal parameters in the frequency domain, but throughout this thesis, the algorithm used was the Polymax [61].

2.2 Machine learning

In this chapter, some introductory concepts of Machine Learning (ML) will be presented, as well as the algorithms used in this thesis project: multivariate Gaussian anomaly detection and anomaly detection autoencoders.

The practical use of an ML algorithm includes two phases. The algorithm is trained in the first phase, and then it is used to create predictions in the second phase. The training entails a set of data from which the algorithm will learn a specific task, while the testing involves a set of data to be analysed by the algorithm, which will then output a prediction that can be applied for a variety of purposes.

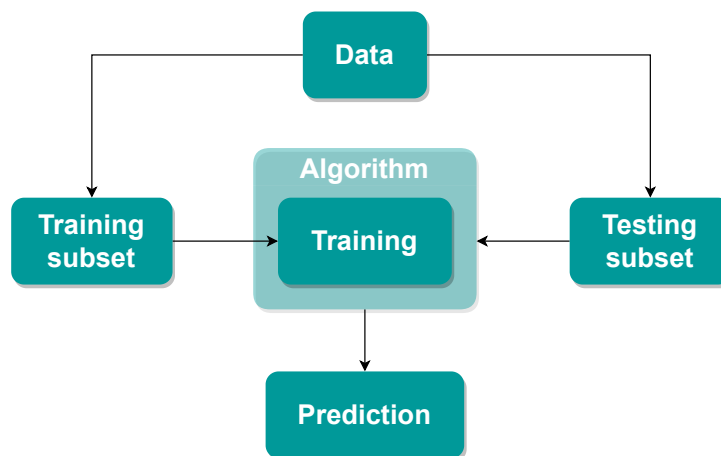


Figure 2.8: Machine learning algorithm process.

Supervised, unsupervised, semi-supervised, and reinforcement learning are the four main types of learning algorithms. Supervised and unsupervised learning algorithms were used in this thesis. The way the algorithm learns is the main difference between them. Below is a more detailed explanation of these types of learning [62, 63].

Supervised learning

The algorithm will be given labelled data, which means that the data from which it learns will have input features and labelled outputs for which the algorithm's predictions will be useful. The goal is for the algorithm to construct a function that correctly classifies the labelled outputs of the training set based on the input features while simultaneously being robust enough to classify unseen data appropriately. Furthermore, supervised learning can be divided into *classification* and *regression* problems based on the algorithm's prediction constraints.

Concerning *classification*, the training data is labelled for a certain number of finite classes. As a result, the algorithm's predictions are limited to these classes. An example of *classification* is the categorization of spam and non-spam emails.

Regarding *regression*, the prediction values are on a continuous interval; therefore, they can be different from the training set's labels. The price prediction of a house based on input characteristics like

its size or its number of rooms is a perfect illustration of a *regression* problem.

Linear regression, logistic regression, neural networks, and support vector machines are some examples of supervised learning techniques.

Unsupervised learning

The algorithm will be fed with unlabelled data containing only input features from which the algorithm should find structure in the data. This is accomplished through clustering algorithms, which put training examples with similar properties into the same cluster—for example, market segmentation algorithm training or astronomical data analysis. Unsupervised learning methods include algorithms like k-means, principal component analysis, hierarchical clustering, and self-organizing maps.

Semi-supervised learning

The algorithm will work with a small amount of labelled data and a large amount of unlabeled data. Semi-supervised learning falls between unsupervised learning (with no labelled training data) and supervised learning (with only labelled training data). It is a particular case of weak supervision. Speech analysis is a classic example of the value of semi-supervised learning models.

Reinforcement learning

Reinforcement learning is an ML training strategy that rewards desirable behaviours while penalising undesirable ones. A reinforcement learning agent can perceive and comprehend its environment, act, and learn through trial and error. Autonomous driving and personalised recommendation systems are examples of reinforcement learning.

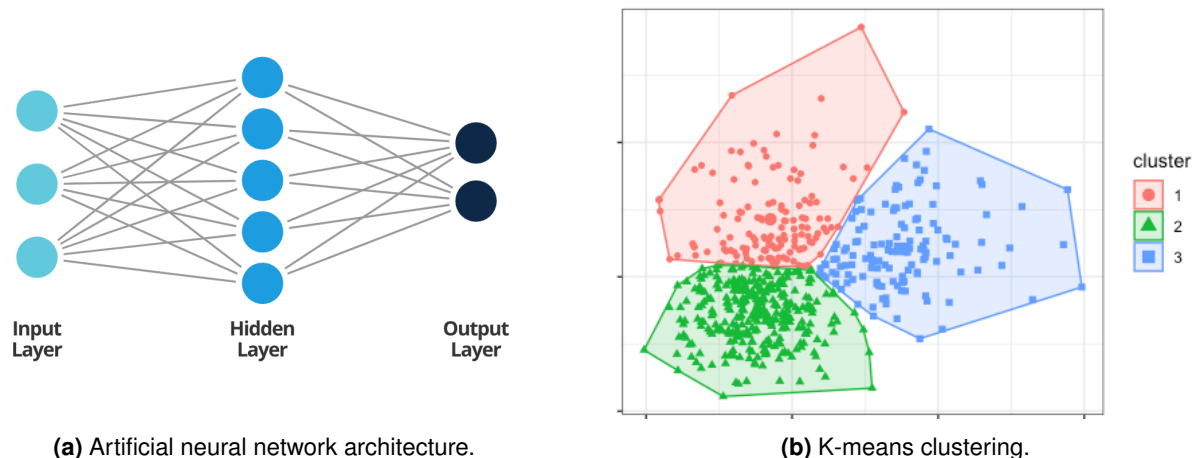


Figure 2.9: Representation of a) supervised learning and b) unsupervised learning algorithms.

2.2.1 Multivariate Gaussian anomaly detection

Gaussian anomaly detection is a type of ML algorithm frequently used when the datasets contain a relatively imbalanced number of examples between each class. It is primarily used to solve unsupervised learning problems, but it also includes some characteristics of supervised learning. The lack of a generally balanced number of instances per class is the primary reason for using such techniques as supervised learning. As a result, supervised learning algorithms like neural networks cannot be trained on an unbalanced set because they will be biased to classify the class with the most training samples. For example, the training of algorithm of a credit card fraud detection or a spam email classification. The last example is better suited to supervised learning because there are enough positive instances for an algorithm to classify them appropriately. One of the anomaly detection algorithms used in this thesis involves fitting the dataset to a multivariate Gaussian distribution. In order to differentiate training instances that correspond to anomalies from training examples that reflect common and acceptable values, a Multivariate Gaussian Distribution is fitted to the training set's characteristics, and a threshold selection process obtains the optimum threshold.

Multivariate Gaussian distribution

The density estimation $p(x)$ for a Multivariate Gaussian distribution is defined by Equation 2.28.

$$p(x, \mu, \Sigma) = \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1} (x - \mu)\right) \quad (2.28)$$

Where the mean vector of a feature's distribution ($\mu \in \mathbb{R}^n$) is:

$$\mu = \frac{1}{m} \sum_{i=1}^m x^{(i)} \quad (2.29)$$

And the covariance matrix ($\Sigma \in \mathbb{R}^{n \times n}$) is defined by:

$$\Sigma = \frac{1}{m} \sum_{i=1}^m (x^{(i)} - \mu)(x^{(i)} - \mu)^T \quad (2.30)$$




Figure 2.10: Illustration of the marginal and joint probability of a bivariate Gaussian distribution.

The shape, width, and orientation of the threshold boundary or contours will change as Σ varies. Nevertheless, Σ will always be a symmetric matrix. The graphics below show how the μ and Σ parameters affect the probability density of this distribution. An example of only two features was explored in order to plot this variation in a 3D plot. The distribution with a null mean vector and covariance matrix equal to the identity matrix, which will be used as a reference, is shown in figure 2.11.

$$\mu = \begin{bmatrix} 0 \\ 0 \end{bmatrix}; \Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

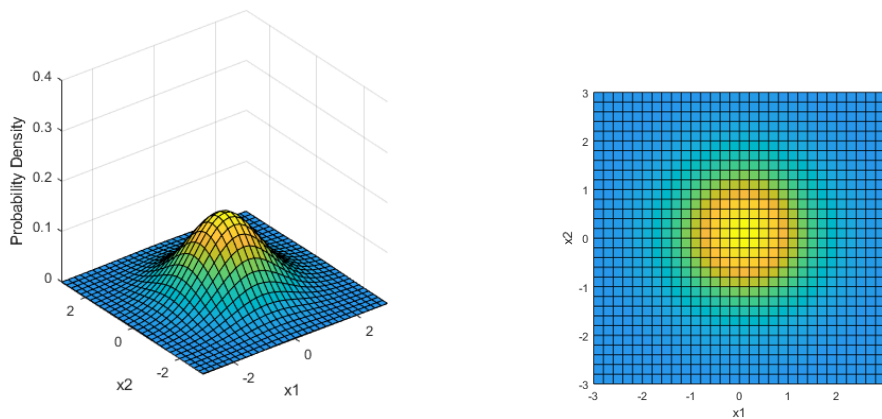


Figure 2.11: Multivariate Gaussian distribution for an identity covariance matrix and a null mean vector.

The variance of the features is measured by the Σ covariance matrix. Hence, reducing the diagonal values will diminish the diameter of the distribution. Furthermore, because the integral of the volume under the surface is 1, the height of the distribution will increase, as shown in figure 2.12.

$$\mu = \begin{bmatrix} 0 \\ 0 \end{bmatrix}; \Sigma = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.5 \end{bmatrix}$$

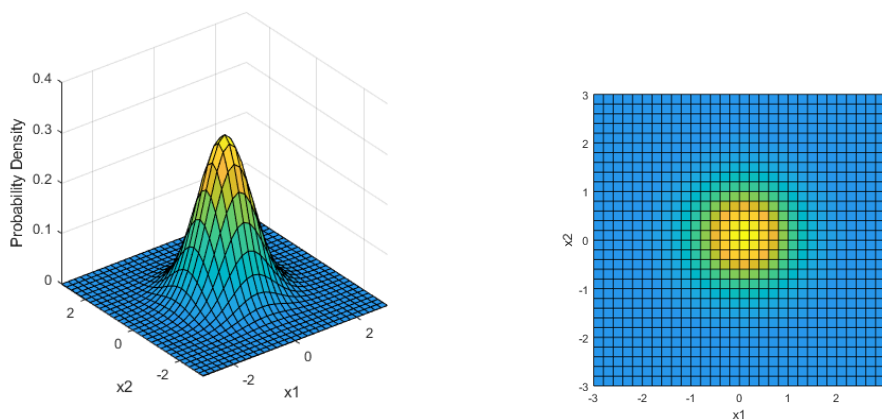


Figure 2.12: Multivariate Gaussian distribution for a covariance matrix with the diagonal values decreased.

Increasing the diagonal values of Σ , on the other hand, will increase the variance, and therefore the distribution will become flatter and wider, as shown in figure 2.13.

$$\mu = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad \Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$$




Figure 2.13: Multivariate Gaussian distribution for a covariance matrix with the diagonal values increased.

The following examples illustrate the effect of the mean vector μ and unequal values in Σ 's diagonal. In comparison to the reference, the second entry of μ was adjusted to 1. As a result, the distribution's centre, i.e. the yellowest part was shifted to the point (0, 1) as can be seen in the 2D plot of the figure 2.14.

$$\mu = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad \Sigma = \begin{bmatrix} 0.5 & 0 \\ 0 & 1 \end{bmatrix}$$




Figure 2.14: Multivariate Gaussian distribution for a covariance matrix with first entry decreased and a unit value in the second entry of the mean vector.

The first value of the covariance matrix was altered in both figures 2.14 and 2.15. The variance of the first feature was decreased in figure 2.14, while the variance of the first feature was increased in

the figure 2.15. As a result, the width of the distribution was decreased and increased in relation to the axis representing the first feature (x_1 axis). The 2D plots highlight this occurrence. Furthermore, as previously stated, when the width of the distribution narrows, the height rises because the integral of the volume under the surface must remain constant at 1. It can also be seen that these distributions are no longer circular when compared to the previous plots of figures 2.11, 2.12 and 2.13. As a result, the capability of a multivariate Gaussian distribution to form non-circular borders as opposed to a standard Gaussian distribution has been demonstrated.

$$\mu = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad \Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}$$

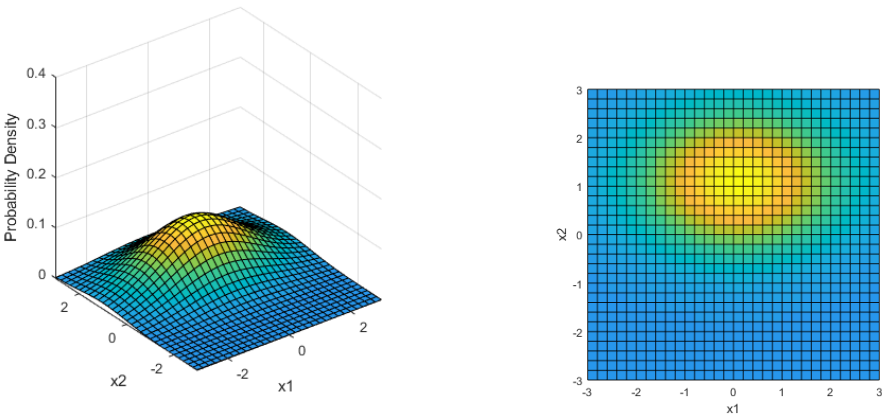


Figure 2.15: Multivariate Gaussian distribution for a covariance matrix with first entry increased and a unit value in the second entry of the mean vector.

$$\mu = \begin{bmatrix} -1 \\ 1 \end{bmatrix} \quad \Sigma = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix}$$

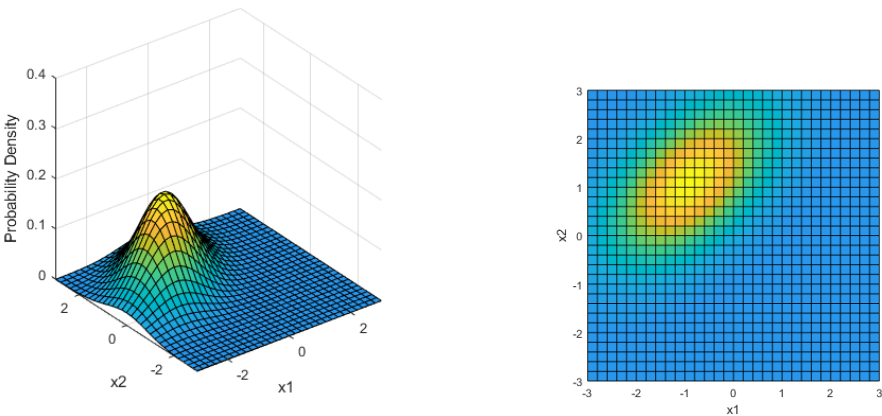


Figure 2.16: Multivariate Gaussian distribution for a covariance matrix with positive non-diagonal entries and the two entries of the mean vector changed.

The figures 2.16 and 2.17 depict the mean vector μ with two distinct values, which leads to the distributions' centres moving in both the x_1 and x_2 axes, as evidenced by the respective 2D plots. The non-diagonal values of the Σ matrix were also studied for the case of 0,5 and -0,8. Both non-diagonal values must be equal since the covariance matrix must be symmetric. The difference in the sign of these matrix entries varies the orientation of the distributions, as is evidenced in the 2D plots of 2.16 and 2.17.

$$\mu = \begin{bmatrix} -1 \\ 1 \end{bmatrix} \quad \Sigma = \begin{bmatrix} 1 & -0.8 \\ -0.8 & 1 \end{bmatrix}$$

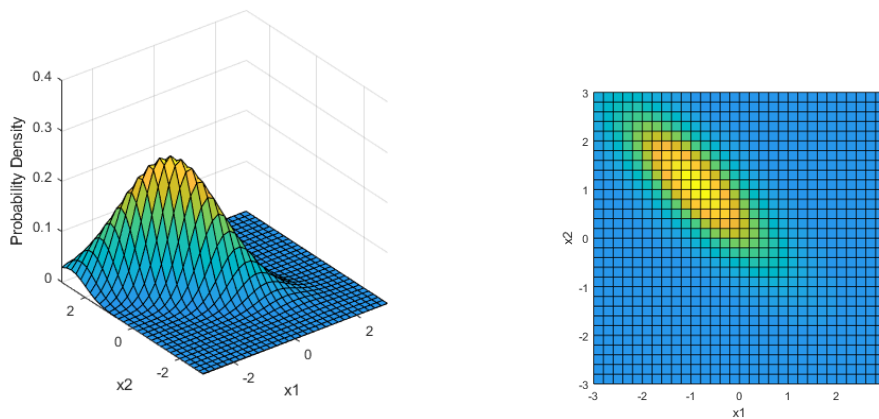


Figure 2.17: Multivariate Gaussian distribution for a covariance matrix with negative non-diagonal entries and the two entries of the mean vector changed.

When the magnitude of the non-diagonal elements is increased from 0,5 to 0,8, the distribution becomes more thinly peaked and narrow. Furthermore, the diagonal shape of the distribution lies on the contribution of the non-zero non-diagonal values.

The ability to link the features being evaluated for each training sample is provided by the fact that these probability distributions might change in shape when considering the covariance matrix calculation. This indicates that this algorithm can perform feature analysis, which is useful because it allows for more flexibility in feature selection and eliminates the need to correlate them manually. Compared to the simpler univariate Gaussian distribution, the multivariate Gaussian distribution has a significant advantage. The first makes a circular probability assumption, i.e. the threshold curves used to split cases into anomalies and non-anomalies can only be perfect circles. In figure 2.18 are illustrated the types of threshold contours of a univariate Gaussian distribution for a two-dimensional dataset, which, compared to the contours of the multivariate Gaussian distribution, would leave several non-anomalous points (in blue) outside the boundary and would not consider the two red points as anomalies. Therefore, the elliptical threshold curves of the multivariate distribution represent an optimal threshold, resulting in a more accurate anomaly detection algorithm.




Figure 2.18: Univariate vs multivariate Gaussian distribution threshold contours.

Anomaly detection

The steps of the Anomaly Detection algorithm can be summarized as follows, taking into account the information above mathematically demonstrated for fitting a Multivariate Gaussian distribution to the training set's features:

1. **Select x_j features that may be suggestive of unusual examples;**

An illustrative two-dimensional (2D) dataset is shown in figure 2.19 to visualize the logic behind the algorithm. It is possible to see a group of samples in the centre of the image and some outliers separated from this main group. These are examples with characteristics that differ significantly from the normal feature values.




Figure 2.19: Illustrative 2D dataset.

2. Calculate the mean vector μ and the covariance matrix Σ to fit a Multivariate Gaussian distribution;

A multivariate Gaussian distribution is fitted to the training set by calculating μ and Σ . The contours of the distribution are shown in Figure 2.20. Most examples are located inside the high probability elliptical region, whilst the anomalous examples are found in the lower probability regions.

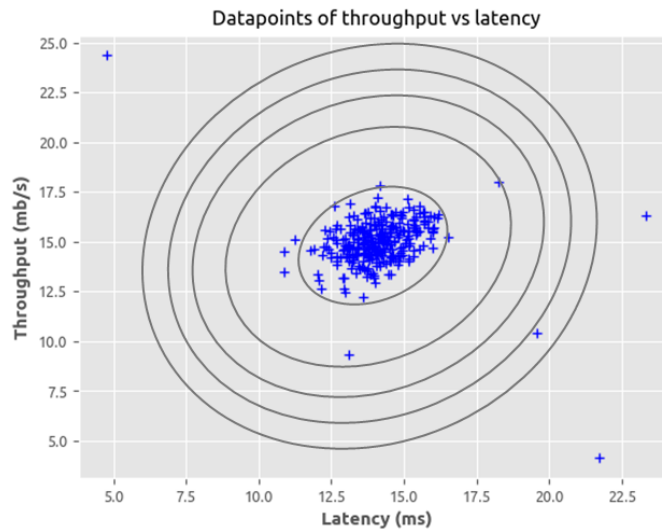


Figure 2.20: Gaussian distribution contours.

3. Calculate $p(\mathbf{x})$ for a new example \mathbf{x} ;

4. If $p(\mathbf{x}) < \epsilon$, identify an anomaly.

Where ϵ denotes a constant threshold value below an example's probability value is considered an anomaly.

$$p(x) < \epsilon \Rightarrow x \text{ is an anomaly}$$

$$p(x) \geq \epsilon \Rightarrow x \text{ is not an anomaly}$$

The next step of the algorithm is to find the optimal threshold by an iterative approach that uses a cross-validation subset with prior knowledge of the example's labels as anomalies or non-anomalies. This threshold value will be used to split the original dataset plotted in this section into anomalies and non-anomalies once that iterative procedure is completed. As shown in figure 2.21, the algorithm detected all anomalies outside of the high probability region, which are represented by red circles in the image.

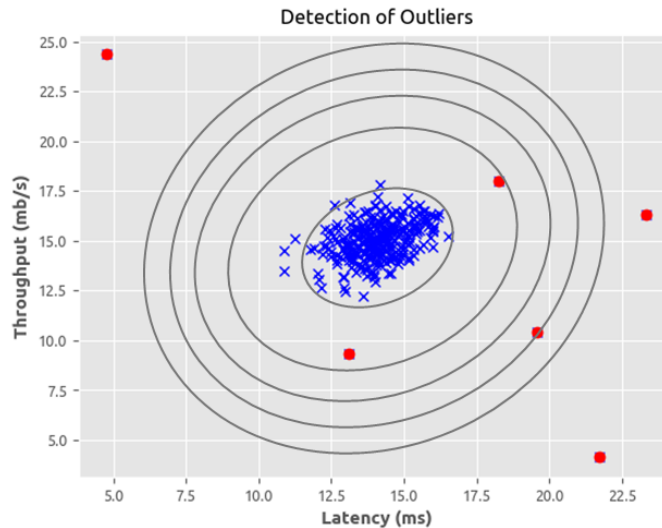


Figure 2.21: Anomaly classification.

After fitting the distribution, the threshold will be the parameter used to classify the testing subset of data. Hence the success of this algorithm’s classification is highly dependent on this value. If it is set too high, the algorithm will likely identify more samples as anomalies, leading to non-anomalies being misclassified as anomalies. If it is too low, the system will incorrectly label anomalies as non-anomalies.

A manual selection can be used, but this requires seeing the results step-by-step, which can be time-consuming and inconvenient. By analyzing a cross-validation subset of the total dataset, a better answer can be found by developing an automatic threshold procedure approach. Using this subset of data, the best threshold can be determined to provide the best anomalous classifications by applying evaluation metrics to various threshold results.

In order to explain the methodology of the automatic threshold estimation, some concepts need to be introduced first. Table 2.1 shows a confusion matrix, or error matrix. Based on the correspondence between the predicted values and the ground truth, the confusion matrix analyzes four classifications. To describe the problem of anomaly identification and future reference, the value 1 pertains to anomalies and 0 to non-anomalies.

Table 2.1: Confusion matrix.

		Ground truth	
		1	0
Predicted	1	True positive	False positive
	0	False negative	True negative

Using the presented classifications, the following metrics can be calculated:

$$Precision = \frac{True\ positives}{True\ positives + False\ positives} \tag{2.31}$$

$$Recall = \frac{True\ positives}{True\ positives + False\ negatives} \quad (2.32)$$

In regard to the common accuracy principle:

$$Accuracy = \frac{True\ positives + True\ negatives}{Total\ number\ of\ examples} \quad (2.33)$$

Precision is a metric that quantifies what fraction of all predicted anomalies was an anomaly in the context of this algorithm, whereas *recall* measures what fraction of all anomalies was accurately predicted. In order to classify an algorithm's success, it is also critical to be exact in the predictions provided without missing out on numerous anomaly occurrences. In order to simplify the threshold optimisation process, a new metric called *F₁ score* that combines *precision* and *recall* can be analysed.

$$F_1\ score = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall} \quad (2.34)$$

A methodology can be developed using the *F₁ score* to analyze different threshold (ϵ) values to make the threshold selection problem automatic:

$$for\ \epsilon = [max(\epsilon); min(\epsilon)]$$

$$Anomalies = p(x) < \epsilon$$

$$Calculate\ F_1\ score(\epsilon)$$

$$If\ F_1\ score(\epsilon) > F_1\ score(\epsilon')$$

$$Save\ new\ best\ \epsilon$$

Finally, the cross-validation subset is used to determine the optimal threshold value. Considering that previous knowledge of the cross-validation labels is required to calculate the *F₁ score*, this technique of anomaly detection is a clear demonstration of a supervised learning algorithm.

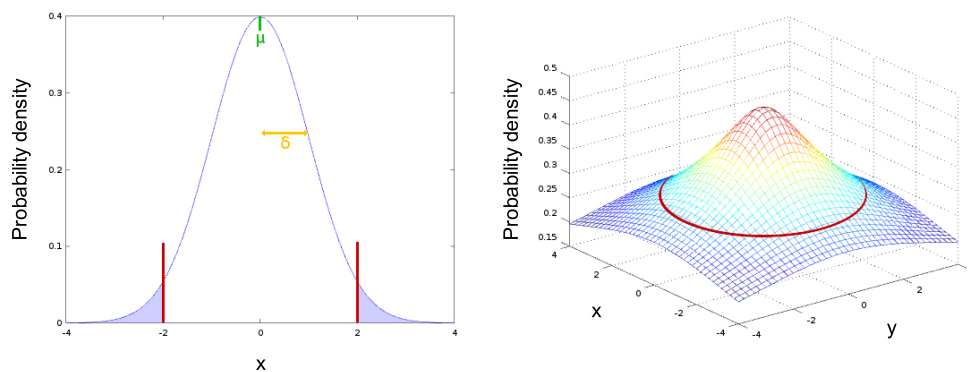


Figure 2.22: Visualisation of an illustrative threshold for the marginal and joint probability of a bivariate Gaussian distribution.

2.2.2 Anomaly detection autoencoders

Autoencoder

An autoencoder is a type of artificial neural network used to learn efficient codings of unlabeled data. By attempting to regenerate the input from the encoding, the encoding is validated and refined. By training the network to ignore insignificant data, the autoencoder learns a representation for a set of data, generally for dimensionality reduction. The algorithm is composed of two parts: an encoder that maps the input into code and a decoder that maps the code back into the input. Duplicating the signal is the simplest approach to do the copying work perfectly. Instead, autoencoders are frequently compelled to approximate the input, maintaining just the most significant aspects of the data in the copy.

The algorithm consists of a feedforward, non-recurrent neural network that employs an input layer and an output layer connected by one or more hidden layers, similar to single-layer perceptrons that participate in multilayer perceptrons. The number of nodes (neurons) in the output layer is the same as the input layer. Instead of forecasting a target value \mathbf{Y} given inputs \mathbf{X} , its goal is to reconstruct its inputs (minimizing the difference between the input and the output). Therefore, autoencoders learn unsupervised.




Figure 2.23: Schema of a basic autoencoder.

An autoencoder consists of two parts, the encoder and the decoder, which can be defined as transitions ϕ and ψ such that:

$$\begin{aligned}
 \phi &: \mathcal{X} \rightarrow \mathcal{F} \\
 \psi &: \mathcal{F} \rightarrow \mathcal{X} \\
 \phi, \psi &= \arg \min_{\phi, \psi} \|\mathcal{X} - (\psi \circ \phi)\mathcal{X}\|^2
 \end{aligned} \tag{2.35}$$

The encoder stage of an autoencoder takes the input $\mathbf{x} \in \mathbb{R}^d = \mathcal{X}$ and maps it to $\mathbf{h} \in \mathbb{R}^p = \mathcal{F}$ in the simplest case, given one hidden layer.

$$\mathbf{h} = \sigma(\mathbf{W}\mathbf{x} + \mathbf{b}) \tag{2.36}$$

This \mathbf{h} is usually referred to as code. σ is an element-wise activation function such as a sigmoid function or a rectified linear unit. \mathbf{b} is a bias vector, while \mathbf{W} is a weight matrix. Weights and biases are typically initialised at random and then modified iteratively by backpropagation during training. The autoencoder's decoder stage then maps \mathbf{h} to a reconstruction \mathbf{X}' of the same shape as \mathbf{X} :

$$\mathbf{x}' = \sigma'(\mathbf{W}'\mathbf{h} + \mathbf{b}') \quad (2.37)$$

Where σ' , \mathbf{W}' , and \mathbf{b}' for the decoder may be unrelated to the corresponding σ , \mathbf{W} , and \mathbf{b} for the encoder.

$$\mathcal{L}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|^2 = \|\mathbf{x} - \sigma'(\mathbf{W}'(\sigma(\mathbf{W}\mathbf{x} + \mathbf{b})) + \mathbf{b}')\|^2 \quad (2.38)$$

Where \mathbf{x} is typically averaged across the training set. Autoencoder training, like that of other feed-forward neural networks, is done through backpropagation of the error. The feature vector $\phi(x)$ can be thought of as a compressed version of the input x if the feature space \mathcal{F} has a lower dimensionality than the input space \mathcal{X} . Undercomplete autoencoders are an example of this. An autoencoder can potentially learn the identity function and become useless if the hidden layers are larger than or equal to the input layer or if the hidden units are given sufficient capacity. Overcomplete autoencoders, on the other hand, may nevertheless learn important features, according to experimental findings. The code dimension and model capacity should, in a perfect scenario, be determined by the complexity of the data distribution to be modelled. Regularized autoencoders are a model version that can be used to do this.

Autoencoders are used to solve various problems, including facial recognition, feature detection, anomaly detection, and word meaning interpretation. Moreover, an autoencoder can also be used as generative models: it can randomly generate new data similar to the input data, i.e. training data.

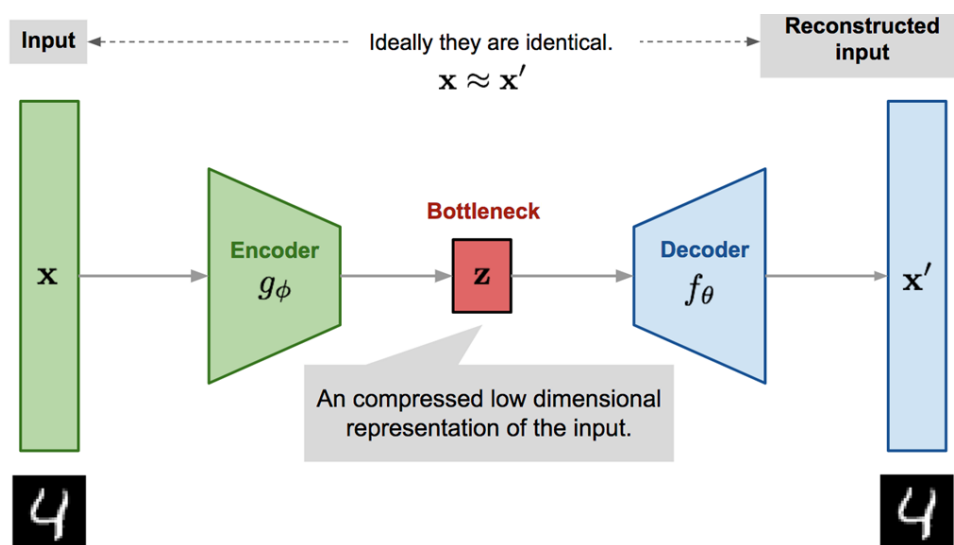


Figure 2.24: Example of an autoencoder architecture for an image compression application [64].

Anomaly Detection

Anomaly detection is one of several applications of autoencoders. The anomaly detection autoencoder (ADAE) is a semi-supervised learning algorithm. The model is encouraged to learn to exactly recreate the most frequently seen features by learning to replicate the most salient features in the training data under some of the limitations outlined previously. When confronted with anomalies, the model's reconstruction performance should be worse. In most circumstances, only data with normal examples is utilised for training the autoencoder; in others, the frequency of anomalies is low enough compared to the observation set that its contribution to the learnt representation can be ignored. The autoencoder will correctly reconstruct normal data after training, but unexpected anomalous input will fail to do so. To detect anomalies, the reconstruction error (the difference between the actual data and its low-dimensional reconstruction) is utilized as an anomaly score [36, 40].

Considering an implementation in the industry, these anomaly detection techniques are commonly confronted with two problems: the limited amount of labelled data and the lack of anomalous samples. Semi-supervised techniques try to solve these challenges. These techniques can work with access to only one type of class label, which is the normal, i.e. undamaged class [35]. The autoencoder training used in this thesis, presented in section 5.2.1, which is done with normal data only, is in line with this assumption.

Chapter 3

Experimental campaigns

The measurements done for this thesis are divided into two major experimental campaigns—one for the titanium WT blade and the other for the glass-fiber reinforced plastic (GFRP) WT blade. For each campaign, vibration-based data such as modal parameters and frequency response functions (FRFs) were obtained through modal hammer testing, modal shaker testing and pull-and-release testing. The tests were carried out for each campaign on the blades in their natural state and their damaged state. In order to simulate damage in the structures, masses of different magnitudes were placed at different blade positions. In order to better reproduce the operating conditions of the blade in which the blade is attached to the rotor, the tests were performed in clamped-free boundary conditions.

The system response was measured in terms of acceleration with model 356A15 PCB accelerometers as the one in figure 3.1. Since the objective of this thesis is to analyse the blade data in its altered state (with masses in the structure) in relation to its healthy state (without masses in the structure), the mass of the accelerometers can be neglected because it will be present in the measurements of the two strands. All the data acquisition was made using the Simcenter Testlab software with the parameters found in table 3.1 through a Simcenter SCADAS like the one in figure 3.2.



Figure 3.1: Model 356A15 PCB accelerometer.



Figure 3.2: Simcenter SCADAS Recorder.

Table 3.1: Testlab acquisition parameters.

Bandwidth [Hz]	800
Spectral lines	2048
Aquisition time [s]	2.56
Estimator	H1

3.1 Experimental setup

Titanium wind turbine blade

The first WT blade studied is made of a Ti6Al4V titanium alloy with a density of 4.429 Mg/m³, and Young's modulus of 110 GPa [65]; it is composed of four welded 3D printed parts manufactured by *3D Systems*, and it was designed and scaled-down to 1.265 m by *DTU Wind Energy*. The blade has also undergone a stress relief heat treatment and a polishing process.




Figure 3.3: Titanium wind turbine blade.

To perform the experimental tests of damage simulation on the Ti blade, the masses shown in figure 3.4 were used as illustrated in figure 3.6. The magnitude of each mass and the respective percentage of the total blade mass is specified in table 3.5.




Figure 3.4: Masses used to simulate damage on the Ti blade.

The masses presented were placed along the seven positions shown in figure 3.7, where it is also possible to see the ten triaxial accelerometers placed along its surface. Regarding the clamping, the blade was bolted through its flange to a concrete block.




Figure 3.7: Mass positions used to perform the damaged experimental runs of the Ti blade.

Figure 3.5: Magnitude of the masses used in the Ti blade tests.

Titanium blade mass [kg]	
2.48	

Mass [g]	Percentage of blade mass
15.5	0.6%
29.1	1.2%
38.5	1.6%
53.5	2.2%
68.5	2.8%
90.2	3.6%
127.0	5.1%




Figure 3.6: Mass of 127.0g on the position 7 of the Ti blade.

GFRP wind turbine blade

The second WT blade tested has the same geometry as the titanium blade and is made of composite material, specifically GFRP. The blade was designed, scaled-down and manufactured by *DTU Wind Energy*.



Figure 3.8: GFRP wind turbine blade.

Manufacturers typically produce two blade shells bonded together with at least one web (figure 3.9) to form the WT blade, as shown in figure 3.10 [66]. The shells used for this blade were made with uniaxial and biaxial GFRP, balsa core and glue. Aluminium was used at the root for the clamping to the rotor.




Figure 3.9: Schema of the section of a WT blade [67].




Figure 3.10: Schematics of the manufacturing of a WT blade [68].

The masses used for the damaged experimental runs of the GFRP blade are shown in figure 3.11, and the respective magnitude is detailed in table 3.2. Due to the favourable damage detection results obtained for the titanium blade presented further on in sections 5.1.2 and 5.2.2, for the GFRP blade experimental campaign that was carried out later, besides the equivalent mass magnitudes tested for the Ti blade, even smaller percentages of magnitude were tested in order to test the sensitivity of the algorithm further.




Figure 3.11: Masses used to simulate damage on the GFRP blade.

In figure 3.12, are represented the positions used for the experimental runs with mass as well as the arrangement of the ten accelerometers used. Concerning the clamping, the blade was screwed directly to an iron structure comparatively much heavier to ensure stability during testing, as can be seen in figure 3.20.




Figure 3.12: Mass positions used to perform the damaged experimental runs of the GFRP blade.

Table 3.2: Masses used in the tests of the titanium blade vs the GFRP blade.

Titanium blade mass [kg]		GFRP blade mass [kg]	
2.48		0.72	

Mass [g]	Percentage of blade mass	Mass [g]	Percentage of blade mass
-	-	0.4	0.06%
-	-	1.0	0.14%
-	-	2.1	0.29%
15.5	0.63%	4.3	0.60%
29.1	1.17%	8.1	1.12%
38.5	1.55%	10.4	1.44%
53.5	2.16%	14.8	2.06%
68.5	2.76%	19.5	2.71%
90.2	3.64%	26.5	3.68%
127.0	5.12%	35.4	4.92%

3.2 Modal hammer testing

The impact testing campaign was done using the modal hammer of the figure 3.14. Ten impacts were made for each run, so the FRFs obtained are an average ratio between the output and the input for each impact.



Figure 3.13: Modal hammer tips.



Figure 3.14: Model 86C03 PCB modal hammer.

GFRP wind turbine blade

Using the appropriate hammer tip (figure 3.13) is a key factor in obtaining a quality FRF measurement. For an ideal measurement, the input force must excite the frequency range of interest at high amplitude and have the amplitude evenly distributed along the frequency, as it does for the metal and plastic tip input spectrum, in figure 3.15.




Figure 3.15: Input spectrum of the impact testing performed with different hammer tip materials.




Figure 3.16: Coherence function of the impact testing performed with different hammer tip materials.

The coherence function can also be used as an indicator of the quality of the FRF. This function indicates how much of the output is predicted as a function of the input by checking the variation from measurement average to measurement average: a coherence of close to 1 indicates that the measurement is repeatable; on the other hand, a coherence of 0 indicates the measurement is not repeatable. The majority of the frequency range should ideally have a coherence value close to 1, with the only

exceptions occurring at anti-resonance frequencies where the response is low and affected by the measurement system noise, as is the case for the metal tip and the plastic tip, shown in figure 3.16.

Bearing these two factors in mind, both metal tip and plastic tip would be reasonable choices; thus, the tests on the GFRP blade were made using the metal tip. The best results were obtained for the impacts on the top surface of the blade, near position 3.

Titanium wind turbine blade

For the titanium WT blade, the hammering was made on the bottom surface of the blade between positions 2 and 3 using the plastic tip, which was the setup that presented the best results for the equivalent study done for the GFRP blade.

3.3 Modal shaker testing

GFRP wind turbine blade

Regarding the GFRP blade shaker testing, a signal was generated from the SCADAS and transmitted to the PCB K2007E01 electrodynamic shaker with an integrated power amplifier that allows adjusting the amplitude of the incoming signal. The shaker is attached to a load cell glued to the structure. The cell is also connected to the SCADAS and measures the force applied. This signal is used as the reference force for the FRFs calculation. Two types of excitation signals were tested: random and chirp, shown in figures 3.17 and 3.18.




Figure 3.17: Random amplitude signal.




Figure 3.18: Chirp signal.

The signal chosen for the GFRP blade was the periodic chirp. Both signals presented reasonable FRFs however the chirp presents slightly less noise, as evidenced in figure 3.19.




Figure 3.19: FRFs of the GFRP blade for a periodic chirp vs continuous random signal.

The load cell of the shaker was attached to the surface between positions 3 and 4 of the blade, as shown in figure 3.20.




Figure 3.20: GFRP blade shaker setup.

Titanium wind turbine blade

Concerning the titanium blade, figure 3.21 shows the setup where the PCB 2075E shaker actuated on the lower surface of the blade near position 3. Unlike the one used for the GFRP blade, this shaker required the use of an amplifier to process the signal coming from the SCADAS.




Figure 3.21: Ti blade shaker setup.

The continuous random signal was chosen for the titanium blade because it was the one that delivered the best results for this blade, as can be seen in the highlighted portion of the FRFs in figure 3.22, by the noise presented by the chirp signal excitation.




Figure 3.22: FRFs of the Ti blade for a periodic chirp vs continuous random signal.

3.4 Pull-and-release testing

As shown in figures 3.23 and 3.24, the pull-and-release testing was performed for both blades using an elastic and a cable tie to put the blade under tension; after that, the cable tie was cut, and the consequent acceleration was recorded.

GFRP wind turbine blade



Figure 3.23: GFRP blade pull-and-release setup.

Titanium wind turbine blade

For the titanium blade, in order to obtain an FRF capable of detecting the otherwise unrecorded modes, like the torsional and the in-plane bending, the elastic was attached in an inclined position to provide a diagonal excitation to the blade.




Figure 3.24: Ti blade pull-and-release setup.

Chapter 4

Modal parameter estimation

After the data acquisition, it is necessary to proceed with the data processing, which involves performing the modal analysis and extracting the relevant parameters for the implementation of ML. This chapter presents the estimated modal parameters and how they were obtained, both for the blade in their original state and with the different magnitude masses at different positions.

Each accelerometer of the setup records an FRF for each of its three directions, as shown in the example of figure 4.1. The FRFs represented in figure 4.2 correspond to the sum of the three directions of every accelerometer.




Figure 4.1: FRFs of the three accelerometer directions and their sum.

4.1 Titanium wind turbine blade

The bandwidth was analysed up to 800 Hz because, after several tests, it was the maximum frequency that guaranteed reliable results. Given the high stiffness of the titanium blade, the impact testing had the best performance because it was the only test capable of exciting all ten modes highlighted in figure 4.5 for the analysed bandwidth of 0-800 Hz. The modal shaker test could not decently excite

the ninth mode at 624.82 Hz since it was an in-plane bending mode which is a challenging mode to excite with the purely vertical excitation direction in which the shaker was set up, as seen in section 3.3. Furthermore, it incorrectly registered a small peak near 280 Hz and a double peak in the tenth mode at 756.65 Hz as shown in figure 4.2 due to coupling effects in the experimental setup. The pull-and-release test was unable to detect the torsional mode around the 500-600Hz bandwidth also because of the excitability limitations inherent to the type of test. Concerning the pull-and-release test, for the crosspower operational pre-processing, the vertical direction, i.e. the out-of-plane blade direction of the accelerometer closest to the tip, was used as a reference because it was the accelerometer direction that presented the best-defined autopower, so it was also the reference that produced the most prominent crosspower.

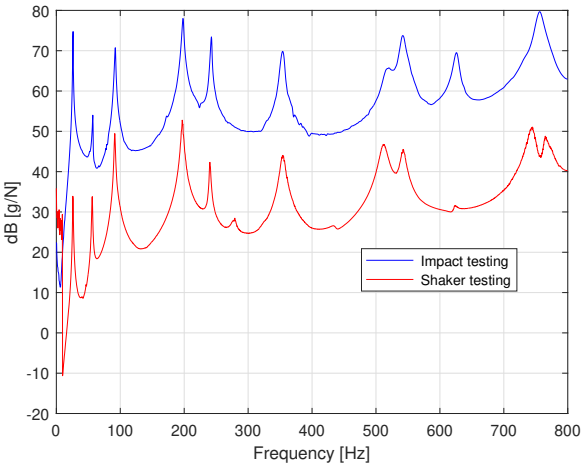


Figure 4.2: Ti blade impact and shaker testing FRF sums.




Figure 4.3: Ti blade pull-and-release testing crosspower sum.




Figure 4.4: Ti blade FRFs from healthy vs damaged experimental runs.

After analysing the FRFs in the healthy state, the experimental runs with masses, i.e. in the damaged state, will now be introduced. Figure 4.4 presents the FRFs obtained from all the healthy and damaged experimental runs for a single accelerometer direction, and the difference between the two is very clear.

For the following analysis of natural frequencies and damping ratios, the FRFs used were obtained from the impact testing and considering only the bandwidth of 15-800 Hz since it was the configuration that presented the best results for this blade, as explained above. Figure 4.5 shows the stabilisation diagram with the system poles obtained for a model size of 50 from the Polymax tool of the Simcenter Testlab. Ten reference normal modes represented in the figure were selected based on the stable poles, each corresponding to a peak of the sum of the 30 FRFs from the ten tri-axial accelerometers. The modal parameters presented in table 4.6 correspond to the average natural frequencies and damping ratios of all healthy experimental runs.




Figure 4.5: Ti blade stabilisation diagram clustered with the representative poles.

The Modal Assurance Criterion (MAC) analysis is used to determine the similarity between two mode shapes: If the mode shapes are identical, the MAC will have the value one represented in red in the plot of figure 4.7. On the other hand, if the mode shapes are very different, the MAC value will be close to zero, depicted in blue in the graph. A MAC analysis can flag potential issues with the modal analysis results. The MAC is used to analyse and identify the modes and areas that could benefit from acquiring more data points on the structure, which was not the case for the titanium blade experimental campaign. This criterion can also provide information about whether too many modes are being selected based on high correlation values between them, increasing the validity of the modal analysis [69]. The mode-set selected was compared to itself in the calculated Auto-MAC in figure 4.7. Ideally, each mode should have a unique mode shape. What can be observed in the figure, where all the non-diagonal values of the matrix have correlations below 30%. It can then be concluded that no extra modes have been selected.

Figure 4.6: Ti blade modal parameters.

Mode	Natural frequency [Hz]	Damping ratio
1	25.969	0.67%
2	56.187	0.89%
3	92.050	0.91%
4	197.803	1.00%
5	241.526	0.97%
6	353.559	0.99%
7	517.906	0.99%
8	541.696	0.86%
9	625.627	0.65%
10	756.772	0.89%




Figure 4.7: Auto-MAC for a healthy experimental run of the Ti blade.

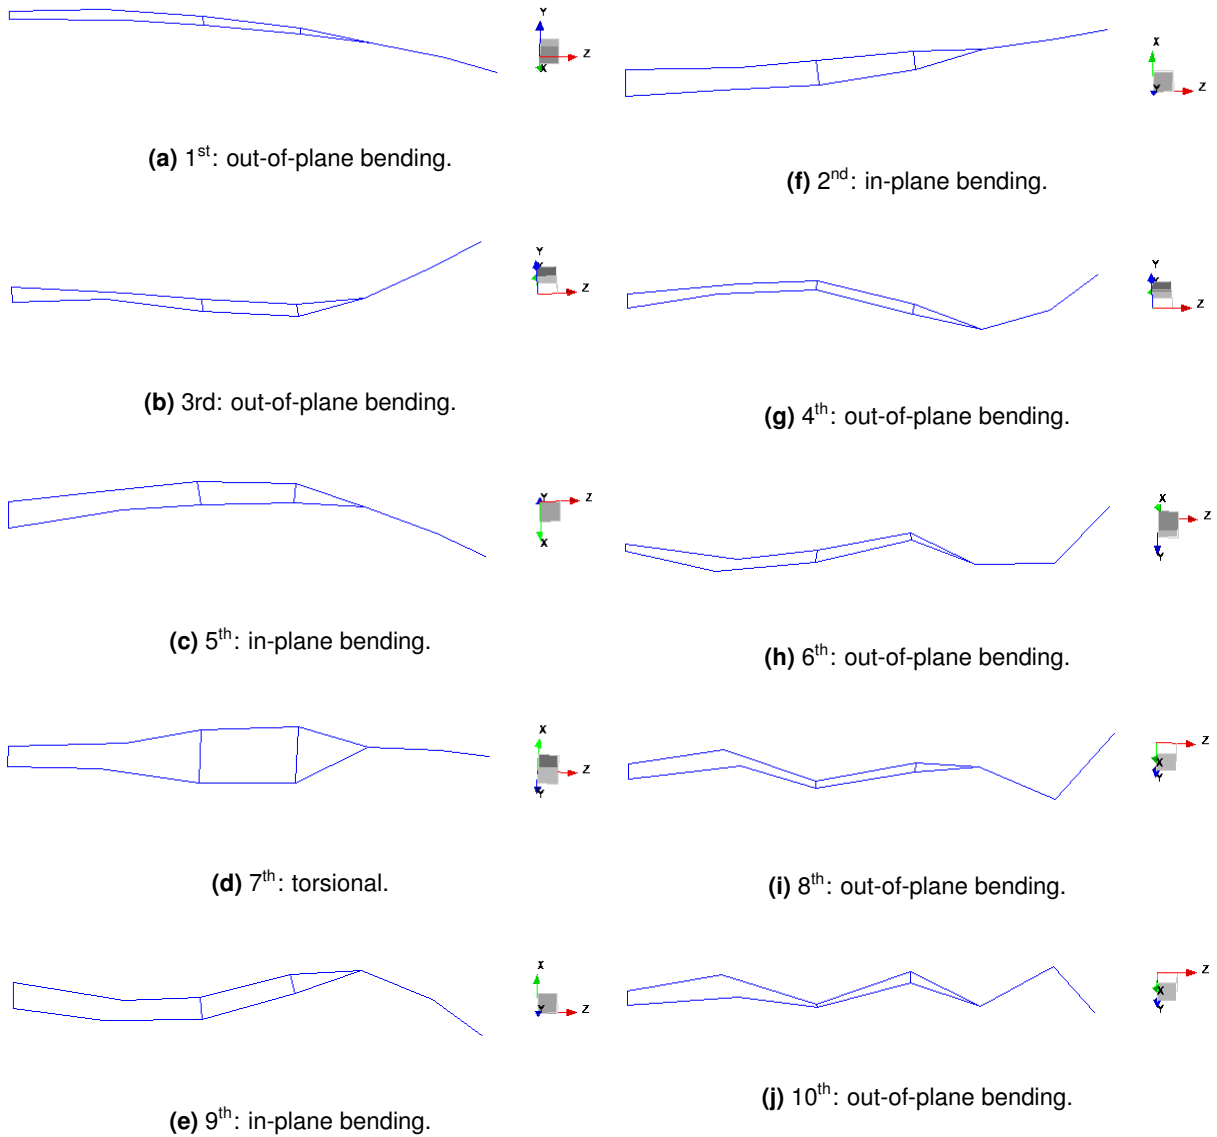


Figure 4.8: Ti blade mode shapes.

In the figure 4.8, are presented the modes shapes of the titanium WT blade obtained in Simcenter Testlab for the studied bandwidth. After analysing the FRFs and crosspowers of the three experimental test variants, the mode shapes and the respective MACs, it can be concluded, with a significant degree of certainty, that the selection of the reference modes is correct.

Having been defined the reference modes and the respective natural frequencies of the structure in the healthy state, the comparison with the damaged state will be made. In the following figures 4.9 and 4.10, it is represented in blue the frequency range of all the healthy experimental runs performed with the respective mean, minimum and maximum values represented by dashed lines. All points represented correspond to frequencies from damaged experimental runs, organised by position on the blade and magnitude of the mass used.




Figure 4.9: Natural frequencies of the 1st mode of the Ti blade.




Figure 4.10: Natural frequencies of the 2nd mode of the Ti blade.

The first evident conclusion is the deviation from the healthy reference frequencies for all masses tested for all the normal modes. The second interesting conclusion to be drawn from the first two modes of Ti represented in figure 4.9 and 4.10 is the almost linear increase of the frequency shift according to the increase of the mass's distance to the clamped side of the blade and the magnitude of the mass itself. This phenomenon only does not apply to the lowest magnitude mass of 15.5g, which, despite having sufficient magnitude to cause an observable deviation, does not have sufficient magnitude to cause a relevant deviation on par with the other masses so that any pattern can be observed.

The same tendency occurs until the fourth mode, from which the increase in the frequency shift is no longer related to the position of the mass in the structure and increases only with the increment of magnitude, as noticeable in figure 4.11. For high-order modes like the seventh mode in figure 4.12, the frequency distribution starts to be purely random.

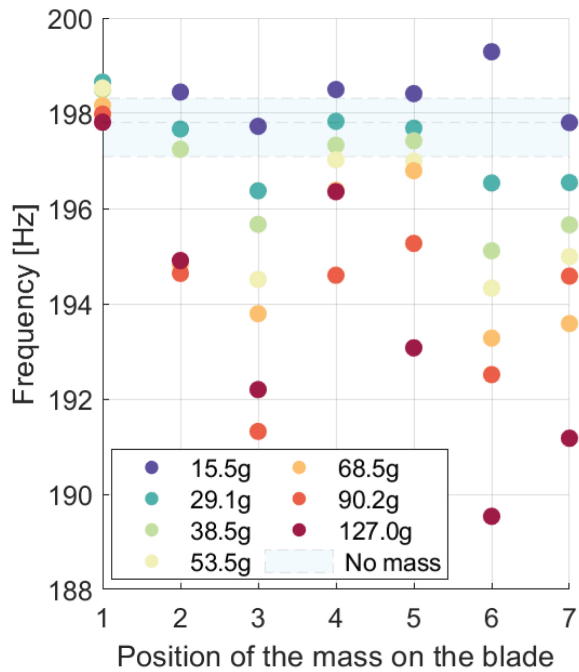


Figure 4.11: Natural frequencies of the 4th mode of the Ti blade.




Figure 4.12: Natural frequencies of the 7th mode of the Ti blade.

4.2 GFRP wind turbine blade

Due to its high malleability, the GFRP blade generated a crosspower from the pull-and-release test with much noise and was unable to evidence the fourteen normal modes depicted in figure 4.16. Both the impact and the shaker testing were able to detect those modes; however, the modal shaker was the one that did so most consistently throughout all the experimental runs. For the Crosspower operational pre-processing of the GFRP blade, the out-of-plane direction of the blade of the accelerometer closest to the tip was also used as a reference for the same reasons explained previously for the titanium blade.




Figure 4.13: GFRP blade impact and shaker testing FRF sums.




Figure 4.14: GFRP blade pull-and-release testing crosspower sum.

After analysing the FRFs of the blade in the healthy state, a comparison will now be made with the damaged state. Figure 4.15 depicts the clear distinction between the experimental runs carried out with the GFRP blade in both healthy and damaged states.




Figure 4.15: GFRP blade FRFs from healthy vs damaged experimental runs.

The GFRP blade was experimentally tested for the same frequency spectrum as the Ti Blade (0-800Hz) to allow for comparisons between the two under similar circumstances. However, given the entirely different nature of the two, the GFRP blade starts presenting inconsistencies for the higher frequencies of the bandwidth. For this reason, only a smaller, more reliable portion of the spectrum was considered henceforth.




Figure 4.16: GFRP blade stabilisation diagram clustered with the representative poles.

The FRFs used to find the eigenfrequencies and damping ratios of the GFRP blade were obtained from the modal shaker testing. Figure 4.16 shows the stabilisation diagram, for the 10-650 Hz bandwidth, with the system poles obtained for the model size of 50, where 14 reference modes were selected based on the stable poles, each corresponding to a peak of the sum of the FRFs.

Figure 4.17: GFRP blade modal parameters.

Mode	Natural frequency [Hz]	Damping ratio
1	15.346	0.83%
2	31.437	1.73%
3	47.890	1.51%
4	101.485	1.41%
5	114.879	2.41%
6	137.672	4.47%
7	178.737	1.53%
8	225.382	2.72%
9	247.519	1.50%
10	299.862	1.14%
11	388.350	2.01%
12	457.484	2.02%
13	560.414	2.38%
14	626.092	2.36%




Figure 4.18: Auto-MAC for a healthy experimental run of the GFRP blade.

The modal parameters presented in table 4.17 correspond to the average natural frequencies and damping ratios of all healthy experimental runs. The mode-set selected was compared to itself in the calculated Auto-MAC in figure 4.18, and some significant correlation values can be observed outside the diagonal, which was not verified for the Ti blade. These values are due to the spatial aliasing phenomenon, which causes the mode shapes to coincide in some of the sensors, not meaning that the mode shapes would be the same in the areas not covered by the sensors. What can be an indicator of an insufficient number of sensors for bandwidth analysed for the GFRP blade. Bearing this in mind, it is possible to assume that no extra modes have been selected.

The mode shapes for the analysed bandwidth are shown in the following figure 4.19. These animations were captured from Simcenter Testlab. Having analysed the FRFs and crosspowers of the three experimental test variants, the mode shapes and the respective MACs, it is possible to conclude, with a significant degree of certainty, that the selection of the reference modes for the GFRP blade is correct.

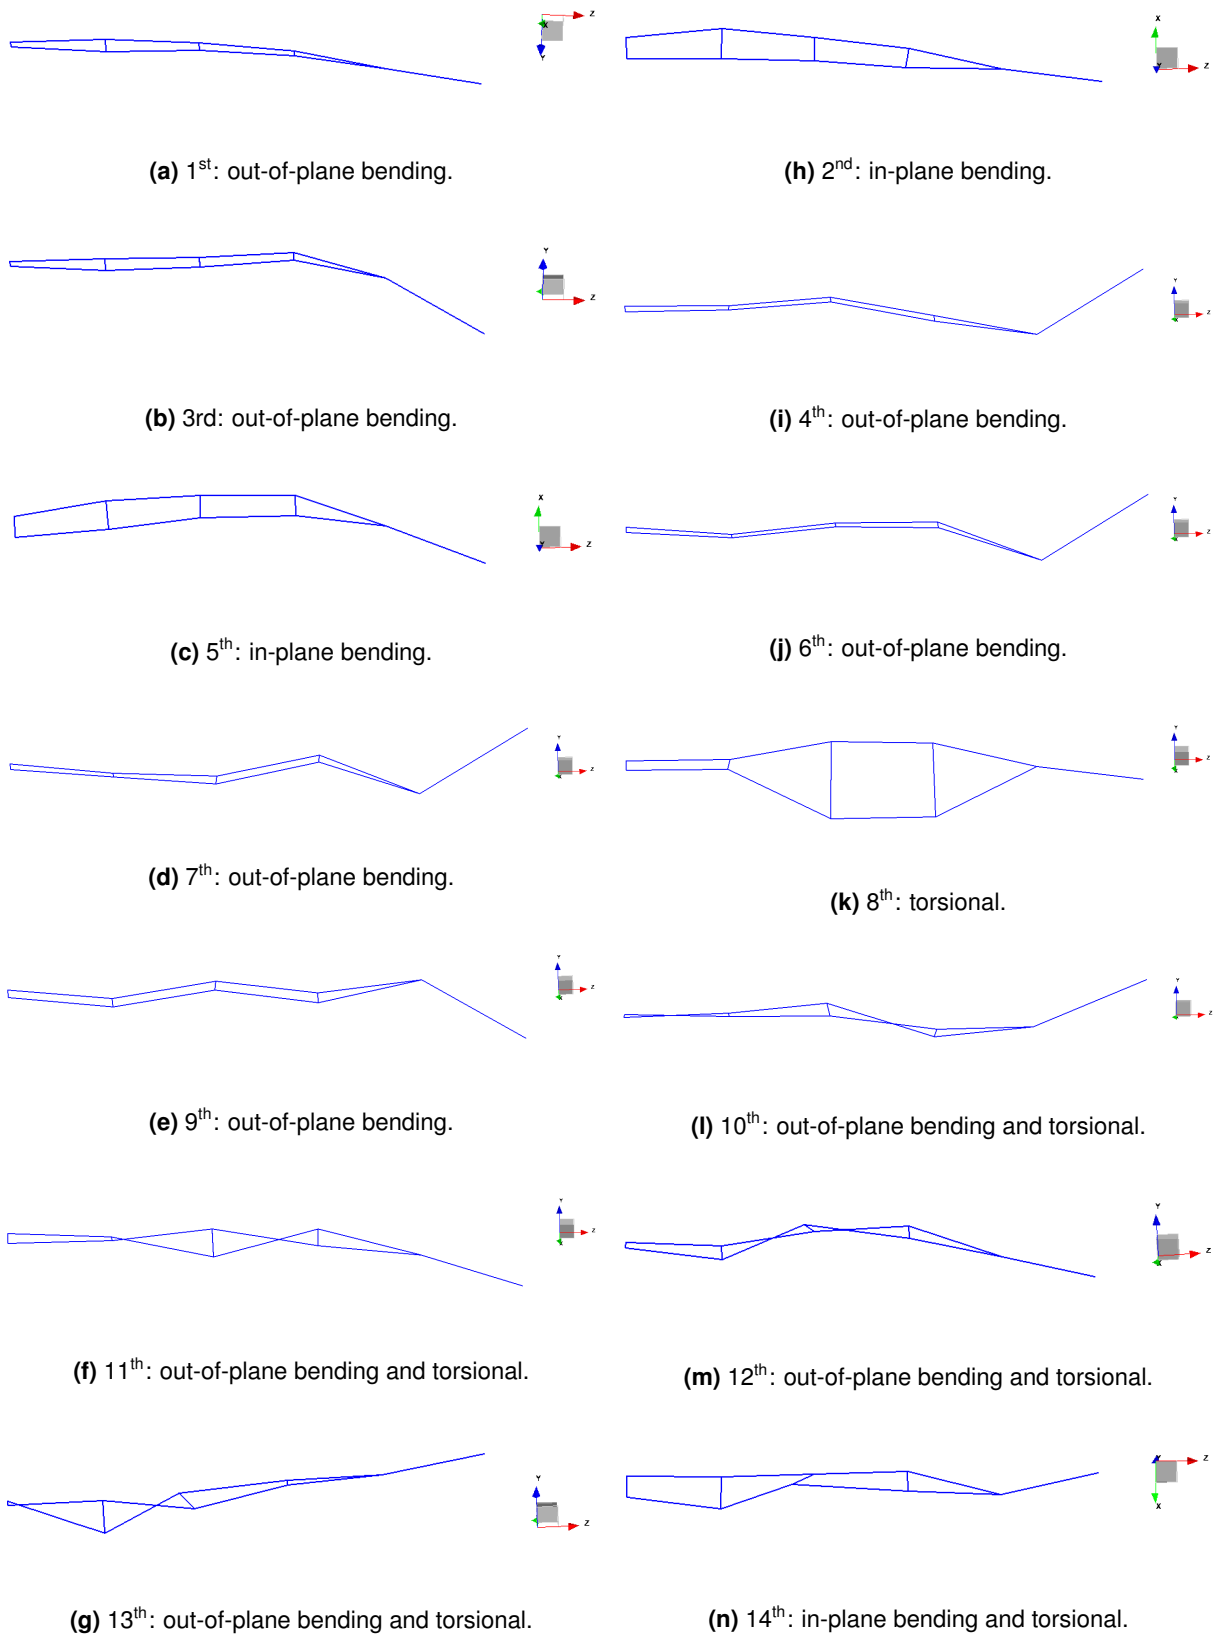


Figure 4.19: GFRP blade mode shapes.

After selecting the reference modes of the healthy state, the comparison to the damaged frequencies can now be made.

As seen for the titanium blade, for the GFRP blade, it is also possible to confirm the deviation of the damaged frequencies from the blue frequency band corresponding to the healthy experimental runs. This deviation is evident for all the masses, including the smaller percentage ones not analysed for the titanium blade. However, this deviation is only evident for most of the masses in positions five onwards, which correspond to the last third of the blade, near the tip.




Figure 4.20: Natural frequencies of the 1st mode of the GFRP blade.




Figure 4.21: Natural frequencies of the 2nd mode of the GFRP blade.

The second conclusion about the increase of this deviation with the approach of the mass to the tip and the increase of the mass magnitude, which was verified for the Ti blade, is also verifiable for the GFRP blade. The fact that more mass magnitudes were analysed for this blade highlights, even more, this linear correlation of the frequency shift with the mass magnitude and with the mass position, as it is possible to observe in figures 4.20 and 4.21.

This tendency, similar to the titanium blade, is verified until the fourth mode, where the correlation happens only with the increase in magnitude, as shown in figure 4.22. For high-order modes like the twelfth mode in figure 4.23, there is no longer any correlation.




Figure 4.22: Natural frequencies of the 4th mode of the GFRP blade.




Figure 4.23: Natural frequencies of the 12th mode of the GFRP blade.

4.3 Modal tracking

4.3.1 Methodology

Monitoring structures health concerning modal analysis is a process that can be divided into three parts: modal parameter estimation, automated modal parameter selection, modal tracking.

Concerning the modal parameter estimation, in this thesis, it is done through the Polymax algorithm of the Simcenter Testlab software [61], where the poles of the system, such as those seen in the stabilisation diagram of figure 4.16, are calculated.

Regarding the automated modal parameter selection, after calculating the poles, the selection is made through the automatic modal parameter selection (AMPS) tool, also in Simcenter Testlab. The AMPS is an algorithm with a clustering approach that analyses the poles generated by Polymax, associating and agglomerating them into the respective modes [70].

One of the methodologies that this thesis intends to implement involves the analysis of the healthy and damaged natural frequencies seen throughout this chapter 4. In order to proceed with this modal tracking part for the large number of experimental runs performed for each of the states and each of the blades, the two previous processes of modal parameter estimation and selection must be fulfilled efficiently. Proceeding with the individual analysis of every run through the software - making the selection of the normal modes using the AMPS along with the Polymax algorithm and the subsequent export of each file - would be inefficient and time-consuming.

Hence, the need to create a script to make this process automatic and time-efficient arose. The methodology behind this script is schematised in figure 4.24.

The first step of the script is to process the files used as a reference to perform the modal tracking,

which correspond to the natural frequencies and mode shapes of the structure in its healthy state. At the same time, the FRF files from which the modal parameters will be extracted are processed. The Polymax algorithm, together with the AMPS tool are applied to these files with a purposely exaggerated modal order to obtain an extensive pre-selection of detected normal modes. These modes are then compared with the reference parameters mentioned above; a filtering process is done based on a score that considers the correlation of the reference modal parameters with the extensive list of pre-filtered ones.

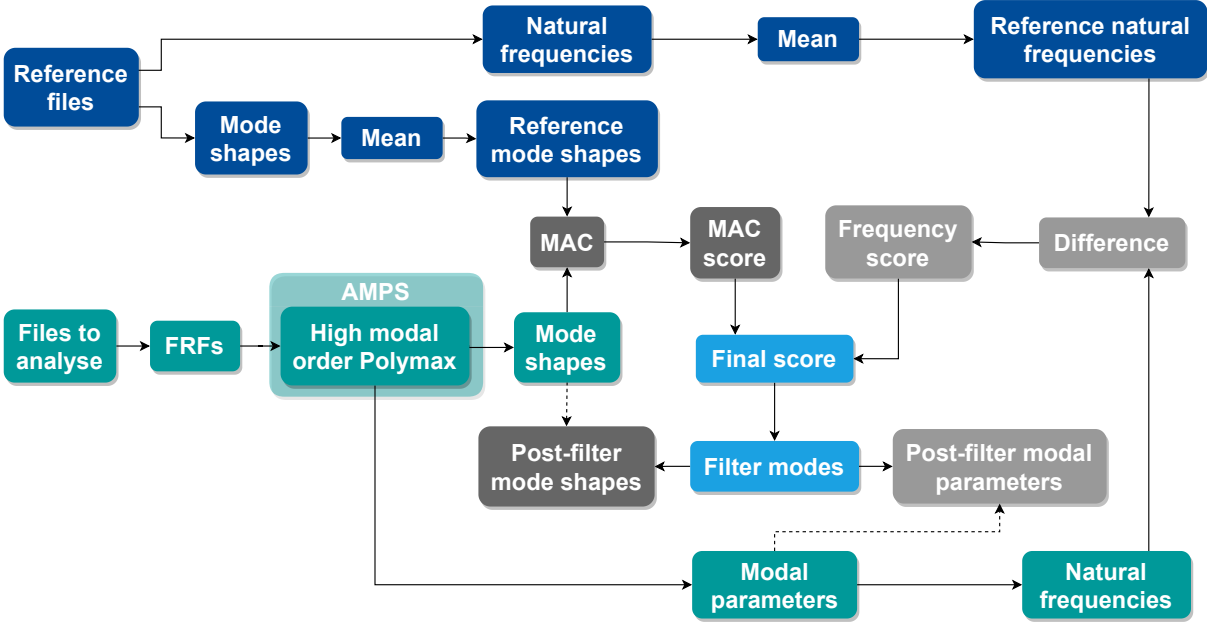


Figure 4.24: Modal tracking script methodology diagram.

The final score based on which the modes are filtered is the combination of a *frequency score* and a *MAC score*. In order to explain the methodology behind each one, the titanium blade reference modes will be used as an example. The *frequency score* is simply an assignment of values from 1 up to the total number of detected modes pre-filtering based on their difference from the reference frequency. As shown in figure 4.25 for only a portion of the pre-filter modes for the sake of simplification, regarding the fourth reference mode, the difference between all pre-filter modes and reference frequency is calculated. The highest score value is assigned to the mode whose frequency is closer to the reference, which, in this case, is the seventh pre-filter mode that differs only 0.5 Hz.

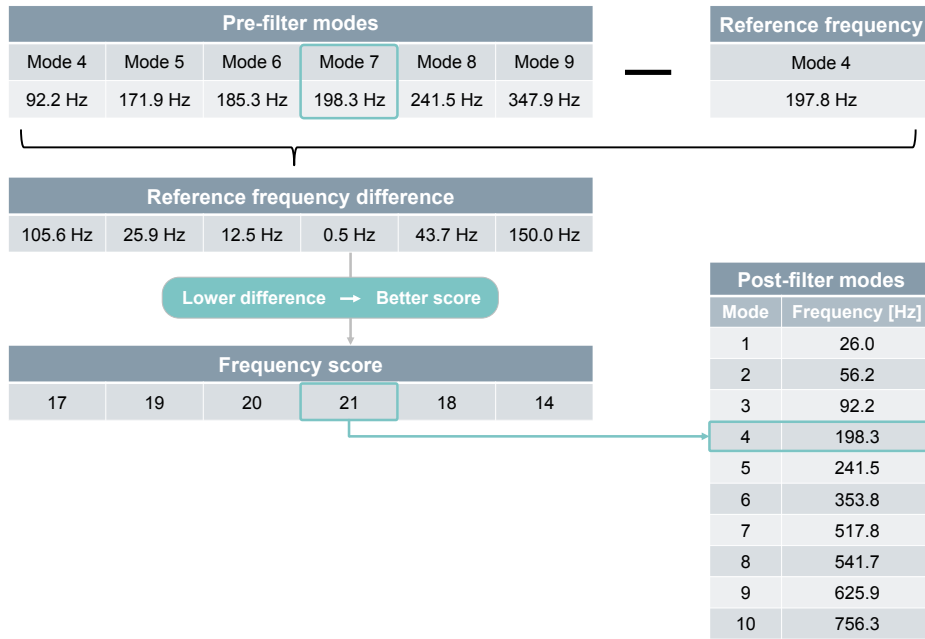



Figure 4.25: Example of a *frequency score* assignment for the fourth Ti blade reference mode.

As for the *MAC score*, it is also an assignment of values from 1 to the total number of pre-filter modes where the highest value is assigned to the highest MAC correlation between the pre-filter modes and the reference modes. For the Ti blade case shown in figure 4.26, likewise the frequency score, the highest correlation value, i.e. the reddest element of the fourth row in the MAC plot, corresponds to the seventh pre-filter mode.



m

Figure 4.26: Example of a *MAC score* assignment for the fourth Ti blade reference mode.

4.3.2 Results

The final score is an equation that takes into account the two scores presented above. In this equation, the *frequency score* and the *MAC score* can assume different weights. The configuration that presented the best results was attributing a factor of 2 to the *MAC score* to give more preponderance to the mode shapes when filtering the modes. This modal tracking script worked perfectly for all healthy experimental runs; however, for the damaged runs where some of the FRFs do not present all the intended reference modes, it was necessary constant monitoring and adjustment of the scores weight parameters mentioned above.

Chapter 5

Machine learning techniques

5.1 Multivariate Gaussian anomaly detection

5.1.1 Methodology

The idea behind this methodology is based on the natural frequency shift of a damaged structure in relation to its healthy state, which was evidenced in the chapter Modal parameter estimation. Therefore the premise is: if the algorithm can learn the natural frequencies of a structure in its undamaged state, it should be able to detect divergent frequencies of the same structure and consider them as anomalies.

As portrayed in the diagram of figure 5.1, the algorithm starts by reading the files from each experimental run, in the damaged and undamaged state separately, containing the natural frequencies of the blade. Each file corresponds to an experimental run containing the eigenfrequencies for every mode.

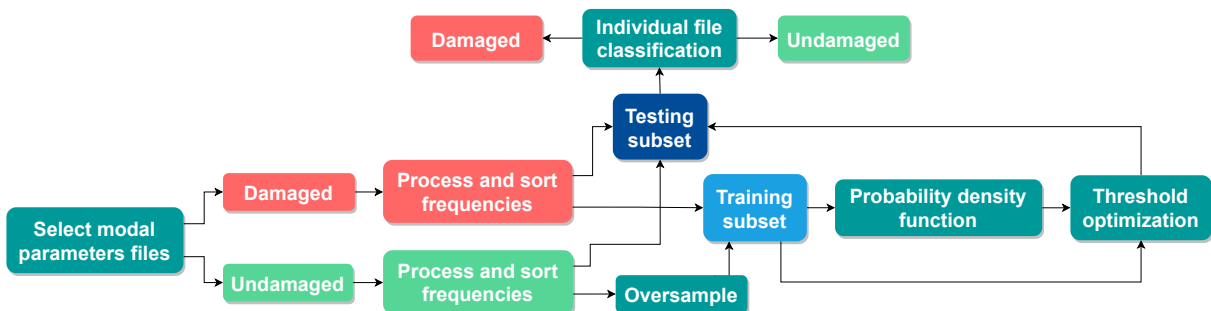


Figure 5.1: MGAD methodology diagram.

The data from the selected files are then divided equally into two subsets, both containing damaged and undamaged data: one used for the training of the algorithm - finding the optimal threshold; and the other for testing - classifying the files based on the optimal threshold. The following diagram in figure 5.2 illustrates how the dataset is processed.

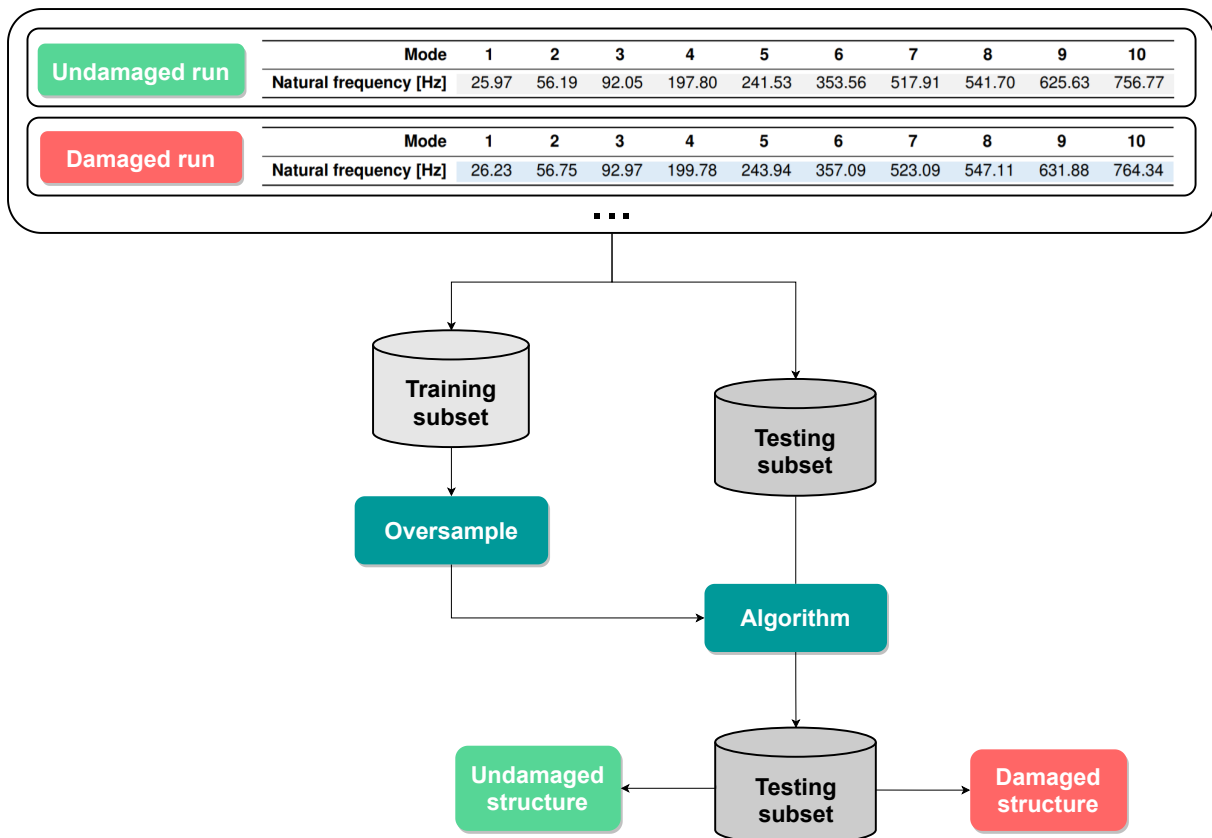




Figure 5.2: MGAD subset processing diagram.


The process of quantifying how divergent a natural frequency should be to be considered an outlier was done using a multivariate Gaussian distribution. This threshold optimisation process is described in the theoretical chapter 2.2.1. It was explained that the Gaussian distribution could be used considering more than one feature resulting in a multivariate Gaussian distribution that will provide a more optimised threshold; therefore, a more accurate algorithm. For the implementation within the scope of this thesis, the features used to perform the multivariate Gaussian distribution were the natural frequencies of all the modes for each blade, estimated in section 4. A visual representation of this approach can be seen in figure 5.3. Only two modes are represented to help visualise and understand the concept since it is impossible to visualise a multivariate Gaussian distribution with more than two variables.



(a) 3D visualisation.



(b) 2D visualisation.



(c) Anomaly classification.

Figure 5.3: Multivariate Gaussian anomaly detection applied to the natural frequencies of the first two normal modes of the Ti blade.

Regarding the experimental runs of the blade in its healthy state, some normal modes may present considerable frequency intervals, as is the case of the fourteenth mode of the GFRP blade shown in figure 5.4, which jeopardise the performance of the algorithm.




Figure 5.4: Natural frequencies of the 14th mode of the GFRP blade.





Figure 5.5: Natural frequencies of the 11th mode of the GFRP blade.

That is why it is more beneficial to use a multivariate Gaussian distribution that can consider all the modes instead of the simpler univariate Gaussian distribution approach, which considers only one of the modes. This way, the algorithm, based on the majority of analysed modes, will have a more reasoned and balanced approach taking weight off these residual and unideal types of modes. In figure 5.5 is an example of a more favourable type of mode for the objective of the methodology having an almost null healthy frequency confidence interval.

Oversampling


The first attempt to implement this methodology obtained the Gaussian distribution represented by the colour bar of figure 5.6, resulting in the classification represented in figure 5.7. It is possible to observe that the probability distribution function of the healthy frequencies assumes values very similar to the damaged frequencies, which means it has a flat Gaussian curve, making it challenging to select an adequate threshold. This happens because the number of healthy experimental runs is much smaller than the number of damaged runs, so the algorithm cannot consider the healthy state as the normal behaviour of the structure. Another reason for this Gaussian distribution problem is that, for some masses, the frequency shift of the damaged frequencies referred to in chapter 4 is insufficient to distinguish the damaged frequencies from the healthy ones.



(a) Healthy frequencies.

(b) Damaged frequencies.

Figure 5.6: Multivariate Gaussian distribution before the oversampling.



(a) Healthy frequencies.

(b) Damaged frequencies.

Figure 5.7: MGAD classification before the oversampling.

The solution to this problem was to make an oversampling of the healthy subset of the data. Oversampling consists of supplementing the training data with multiple copies of some of the minority classes, creating a transformation in the Gaussian distribution as illustrated in figure 5.8.




Figure 5.8: Transformation of a flat Gaussian distribution (blue) into a sharp distribution curve (red) due to oversampling.

The oversampling factors, i.e. the minimum number of times the healthy data were multiplied to obtain the best accuracy possible, were around 500 for the Ti blade and 2000 for the GFRP blade. The following figures 5.9 and 5.10 illustrate the significant difference in the distribution after the oversampling and its impact on the subsequent classification. To give a numerical sense of the improvement, considering all of the damaged experimental runs, the accuracy of the Ti blade algorithm went from 55.6% to 100% and GFRP blade one from 66.7% to 94.1%, representing an increase in the orders of 80% and 40% respectively.




Figure 5.9: Oversampled multivariate Gaussian distribution.




Figure 5.10: Oversampled MGAD classification.

This technique is especially relevant given the context in which this thesis is set. This oversampling process is essentially what happens naturally in a real-world SHM application where most of the extracted data is done for structures in their healthy state.

5.1.2 Results

In the end, the testing subset of natural frequencies is classified based on the selected threshold, i.e. separated into damaged and non-damaged. This classification, compared with the ground truth (the actual state of health of the structure for the corresponding frequencies), leads to the accuracy percentage of the algorithm. This accuracy can vary drastically depending on the damaged experimental runs analysed, hinging on the positions and magnitude of the masses considered.

GFRP turbine blade

In table 5.1, it is presented the sensitivity study of the algorithm's damage detection as a function of the mentioned variables. The results presented are from the shaker testing since it was the one that presented the most reliable data, as seen in section 4.2, and consequently the best results. On the left are the intervals of masses, and in opposition, are the intervals of mass positions (figures 3.7 and 3.12) considered to reach the corresponding accuracy. For instance, the first column of the first row of the table corresponds to the algorithm's accuracy when running, considering all the damaged experimental runs of the, in other words, every combination of magnitude and position of the mass. Since the tests were performed with ten different mass magnitudes in seven different positions, the algorithm ran with

seven times ten, i.e. 70 different files, representing 70 experimental runs of the blade in the damaged state. On the other hand, the entry of the last column of the last row corresponds to the algorithm's accuracy, taking only into consideration the experimental run with the mass of 35.4g at position 7.

Table 5.1: MGAD accuracies for the different combinations of intervals of magnitude and position of the GFRP blade.

Mass [g]	Position of the mass on the blade						
	1 to 7	2 to 7	3 to 7	4 to 7	5 to 7	6 to 7	7
0.4 to 35.4	94.1%	97.8%	95.1%	94.4%	93.8%	92.6%	100.0%
1.0 to 35.4	95.7%	97.7%	94.7%	94.1%	93.3%	92.3%	100.0%
2.1 to 35.4	95.5%	97.5%	94.4%	93.8%	93.1%	92.0%	100.0%
4.3 to 35.4	95.0%	100.0%	100.0%	100.0%	100.0%	91.7%	100.0%
8.1 to 35.4	94.6%	100.0%	100.0%	100.0%	100.0%	91.3%	100.0%
10.4 to 35.4	93.9%	100.0%	100.0%	100.0%	100.0%	90.9%	100.0%
14.8 to 35.4	93.3%	100.0%	100.0%	100.0%	100.0%	90.5%	100.0%
19.5 to 35.4	92.6%	100.0%	100.0%	100.0%	100.0%	90.0%	100.0%
26.5 to 35.4	91.3%	100.0%	100.0%	100.0%	100.0%	89.5%	100.0%
35.4	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

As seen in chapter 4, defects in positions closer to the free extremity of the blade and higher magnitude lead to higher frequency shifts. The opposite is true for natural frequencies from experimental runs with mass positions closer to the clamped side (position 1 of figure 3.12) and lower magnitude, with insufficient variation in comparison with its healthy state. This damaged frequency shift will strongly influence the thresholding process. A high-frequency shift will result in a more optimal threshold, whereas a low-frequency shift will result in a poor threshold, which will be reflected in the algorithm's accuracy. The results obtained shown in table 5.1 validate this reasoning.

Titanium turbine blade

Regarding the titanium blade, the presented results are from the hammer testing, which presented the most reliable data, shown in section 4.1, and, therefore, the best results. As seen for the GFRP WT blade, the worst performances are noticed when considering masses of 0.4, 1.0 and 2.1 grams, which correspond to a maximum of 0.3% of the GFRP blade's mass.

Table 5.2: MGAD accuracies for the different combinations of intervals of magnitude and position of the Ti blade.

Mass [g]	Position of the mass on the blade						
	1 to 7	2 to 7	3 to 7	4 to 7	5 to 7	6 to 7	7
15.5 to 127.0	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
29.1 to 127.0	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
38.5 to 127.0	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
53.5 to 127.0	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
68.5 to 127.0	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
90.2 to 127.0	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
127.0	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

As explained in the chapter on the Experimental campaigns the titanium WT blade was only tested for masses as light as 0.6% of its total weight. For this reason, given the results for the GFRP blade, it makes sense that the results for the Ti blade would be close to the perfect accuracy, as was the case for the GFRP blade for most combinations of runs above 0.3% of the total mass. Which was verified as shown in table 5.2.

5.2 Anomaly detection autoencoders

5.2.1 Methodology

Another technique developed was the Anomaly Detection Autoencoder (ADAE). The ADAE is a semi-supervised learning approach that involves training the algorithm with normal features, i.e. data from the structure in its healthy state and then classifying the rest of the testing dataset based on its reconstruction error. If the error of a certain point exceeds an established threshold, it is classified as an anomaly.

As schematically shown in the diagram of figure 5.11, the algorithm starts by reading the files from each experimental run, in the damaged and undamaged state separately, containing the FRFs from the blade. Each experimental run file has data from the three directions (x , y and z) of each accelerometer. Considering that ten accelerometers were used, there are three times ten FRFs, i.e. 30 per experimental run. Then, the data are normalised along the accelerometer directions to reduce the experimentally inevitable differences in data acquisition from run to run.

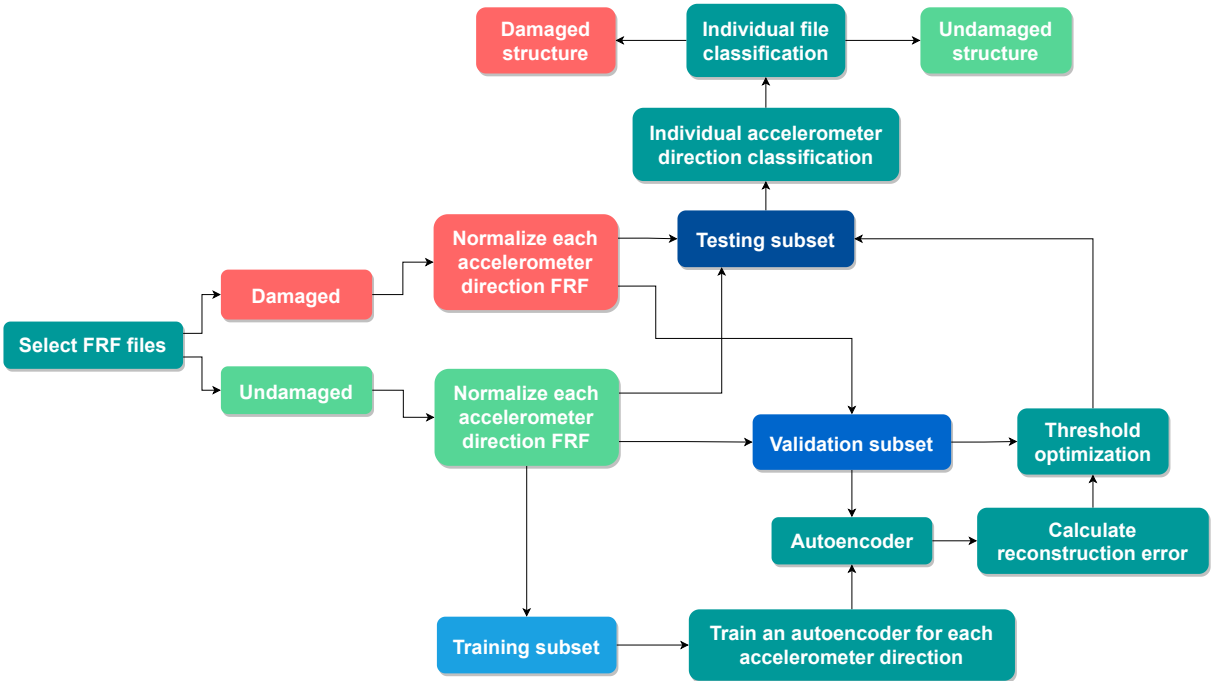


Figure 5.11: ADAE methodology diagram.

After that, the data from the selected files are divided into three subsets. The first one, consisting of one-third of the healthy data, is used for the training of the AEs. An AE is trained for each direction of each accelerometer. The second subset containing one-third of healthy data and half of the damaged data is used for the validation process - finding the optimal threshold. Finally, the last subset, also consisting of half of the damaged data and one-third of the undamaged data is used for testing - classifying the files based on the threshold found.

In figure 5.12 is illustrated the fundamental principle of an AE applied to our use case. At the top of the diagram, in green, it is represented the input signal corresponding to an FRF in the healthy state and the respective AE reconstruction, which exhibits an almost perfect reconstruction, originating only residual error. While at the bottom, in red, the input signal of the structure in its damaged state manifests a significant difference in the AE reconstruction and corresponding error.

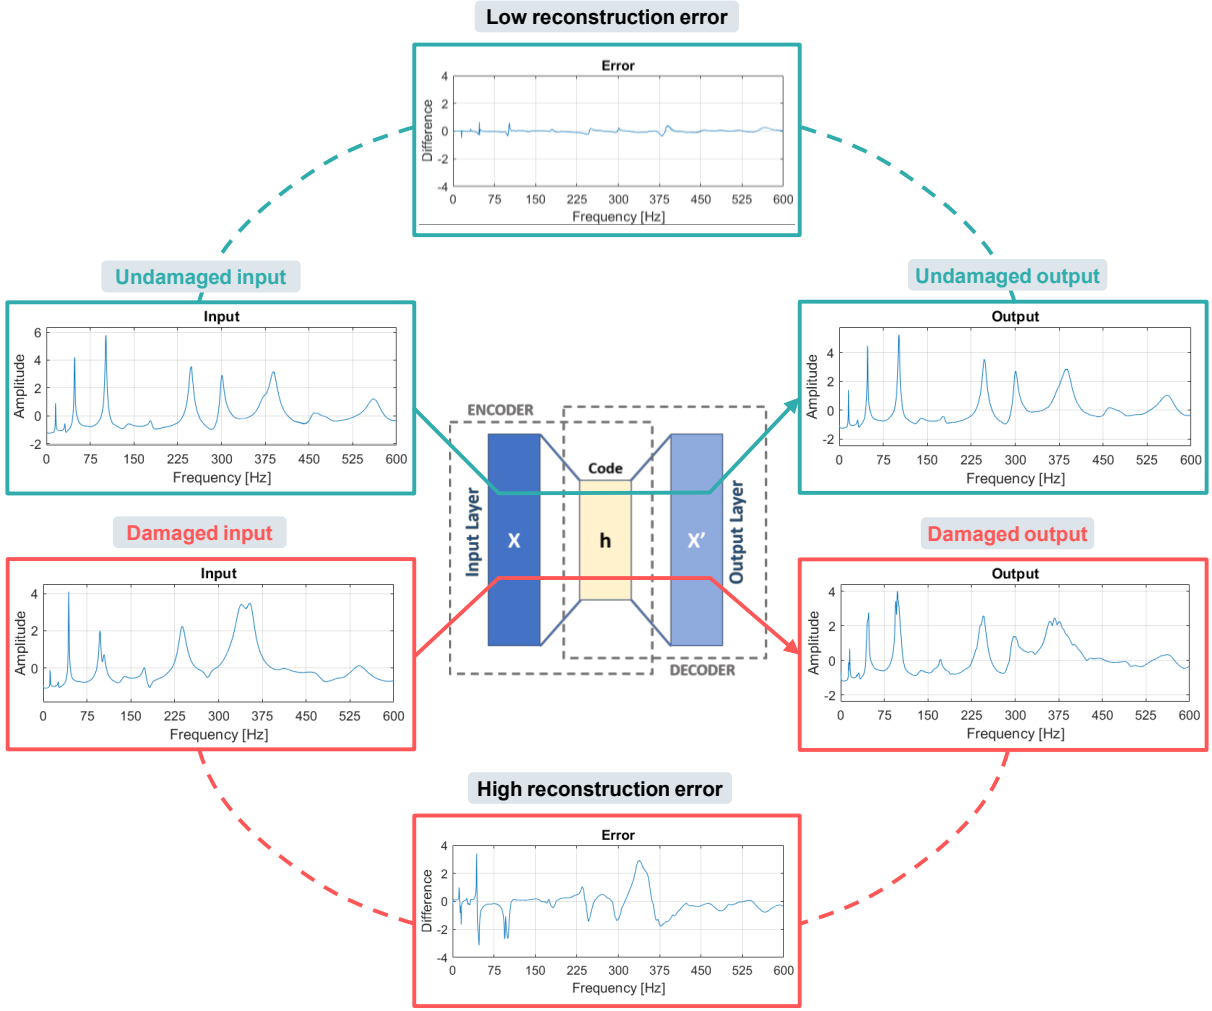


Figure 5.12: ADAE reconstruction error principle.

The error shown in figure 5.12 is simply the difference between the input and the output. However, for this thesis algorithm implementation, a more reliable form and better-suited method of estimating the

reconstruction error was used: the root-mean-square error (RMSE). The way to calculate it is described in equation 5.1 where \hat{y} represents the predicted values, y the observed values and n the number of observations.

$$RMSE = \sqrt{\frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{n}} \tag{5.1}$$

The following figures 5.14 and 5.13 show the reconstructions of the FRFs of the structure in its healthy state, and it is visible that the reconstruction is near perfect; i.e., the output FRF is roughly the same as the input FRF.




Figure 5.13: ADAE FRF reconstructions for healthy inputs from the Ti blade.




Figure 5.14: ADAE FRF reconstructions for healthy inputs from the GFRP blade.

While in figures 5.15 and 5.16, it is represented the case for the input FRFs of the structure in its damaged state, and the difference in their reconstruction is evident and significant.




Figure 5.15: ADAE FRF reconstructions for damaged inputs from the Ti blade.




Figure 5.16: ADAE FRF reconstructions for damaged inputs from the GFRP blade.

In order to give an insight into the distribution of the reconstruction error along the spectral lines of the FRF, i.e. along its bandwidth, figures 5.17 and 5.18 show the RMSE between all the collected experimental runs and their respective output reconstructions for the damaged and healthy state separately.




Figure 5.17: Sum of the reconstruction RMSE from every healthy vs damaged experimental run for one of the accelerometer directions of the Ti blade.




Figure 5.18: Sum of the reconstruction RMSE from every healthy vs damaged experimental run for one of the accelerometer directions of the GFRP blade.

As expected, it is possible to notice the RMSE corresponding to undamaged data remaining approximately null the whole spectrum. In contrast, the damaged data part remains significant over the whole bandwidth, with the largest spikes in the zones near the natural frequencies of the blade in its original state.




Figure 5.19: RMSE for the different experimental runs of the Ti blade.

The graphs in figures 5.19 and 5.20 shows the variation of RMSE over the different experimental runs for both healthy and damaged data. Analogously to what was observed in chapter 4, where the natural frequencies of the damaged structure presented a larger deviation for higher magnitude and masses closer to the tip of the blade, the same applies to the reconstruction error of the AEs.




Figure 5.20: RMSE for the different experimental runs of the GFRP blade.

It can be seen that for the titanium turbine blade, it is quite simple to define a threshold value between the errors of the healthy and damaged structure due to the large difference between the two. Regarding the GFRP turbine blade, the same is true but only for masses above 4.3g, corresponding to 0.6% of the blade mass. For the algorithm to be able to identify damage below those magnitudes, the thresholding process addressed in the following subchapter has to be taken into account.

Thresholding

The selection of the reconstruction error threshold, which defines what points are considered faults, is one of the most significant challenges when using ADAEs; this value must be determined before any classification results can be obtained. When labels, i.e. the ground truth data, are known and exploitable, as with the training and validation sets, the threshold can be set by maximising any given metric, a process known as *supervised thresholding*. For each possible threshold, a corresponding value of true positives and false positives exists; therefore, a receiver operating characteristic (ROC) curve like the ones shown in figure 5.21 can be created. As illustrated in figure 5.22, the best classifier possible corresponds to a point in the coordinate (0,1), representing 100% sensitivity (no false negatives), 100% specificity (no false positives) and an area under the curve (AUC) of 1.

Analysing the number of true positives and false positives can be more informative than focusing only on the *accuracy* (equation 2.33) when dealing with datasets with a large percentage of non-anomalous data. To facilitate the study of the threshold selection considering both variables at the same time, a score based on the Intersection over Union (IoU) method defined in equation 5.2 was used.

$$IoU = \frac{True\ positives}{True\ positives + False\ negatives + False\ positives} \quad (5.2)$$

The IoU value can be plotted as a function of the multiple thresholds analysed by the algorithm, as shown in figure 5.23. The threshold selected for each AE corresponds to the maximum of the IoU function.




Figure 5.21: ROC curve for the Ti and GFRP blades.




Figure 5.22: ROC curves for different classifier performances.




Figure 5.23: Threshold corresponding to the maximum IoU for one of the directions of an accelerometer.

A threshold for each of the accelerometer directions is obtained with this process, which means there will also be a classification for each of them. Knowing that there are ten accelerometers, each file, i.e. experimental run, will have 30 classifications (3 directions per accelerometer). Now, to decide if a file should be considered healthy or damaged based on these 30 classifications, a methodology that consists in finding the ratio that estimates how much an experimental run is likely to be considered or not damaged was developed. Given that each one of the 30 classifications of an experimental run is labelled as 0 if it is considered healthy and as 1 when it is considered damaged; the ratio is calculated simply by summing all the 30 classification labels of the run and dividing it by the total number of accelerometer direction classifications (i.e. 30). In other words, this ratio will range from 0 to 1, whereby an experimental run with a ratio of 0 means that it had all its accelerometers classified as healthy, while a run with a ratio of 1 means that it had all its accelerometers rated as damaged. With this in mind, experimental runs with ratios below 0.5 are considered healthy, and above 0.5 are deemed damaged.

An alternative considered for the thresholding process was to evaluate the distribution of RMSE values across the input data and find the midpoint between two clusters. However, this method presupposes anomalies to be found in the first place, whereas in a real-world application, most of the tested structures will be defect-free, so it was not implemented.

5.2.2 Results

Finally, the testing subset consisting of FRFs is classified based on the threshold found, i.e. separated into damaged and non-damaged. This classification, compared with the ground truth labels, leads to the algorithm's accuracy.

Contrary to what was seen regarding MGAD in the chapter 5.1.2, this algorithm was able to achieve perfect accuracy for certain configurations when analysing the entire dataset of masses tested. Thus, those configurations, i.e. AE architectures, will be presented, and the most optimal for each of the WT blades will be ascertained.

The algorithm architecture and the tuning of its hyperparameters is one aspect that significantly impacts its accuracy. The two hyperparameters that most affected the final output of the algorithm were the *code size* and the maximum number of *epochs*. An epoch is a term used in ML that indicates the number of times that the learning algorithm will work through the entire training dataset, and the *code size* represents the number of neurons in the hidden layer, i.e. the number of nodes in the middle layer.

GFRP turbine blade

In table 5.3, it is shown the study performed for these two parameters and how they influenced the accuracy of the algorithm concerning the GFRP blade data. The results presented are from the shaker testing since it presented the the best results, as presented in section 4.2.

Taking into account that the algorithm execution time and its computational requirements increase with the increasing of the hyperparameters, i.e., along the lines and columns of table 5.3, a preselection of the four configurations, which presented perfect accuracy for the less computational impact possible, was made.

Table 5.3: Accuracy comparison of the algorithm architectures.

Epochs	Code size							
	1	5	10	20	30	40	50	60
1	97.8%	91.3%	95.7%	89.1%	93.5%	95.7%	93.5%	93.5%
2	95.7%	95.7%	87.0%	87.0%	89.1%	91.3%	89.1%	91.3%
3	97.8%	91.3%	87.0%	87.0%	87.0%	87.0%	91.3%	87.0%
4	91.3%	91.3%	89.1%	89.1%	87.0%	91.3%	87.0%	87.0%
5	89.1%	89.1%	91.3%	91.3%	87.0%	93.5%	87.0%	93.5%
10	93.5%	91.3%	91.3%	91.3%	97.8%	97.8%	97.8%	91.3%
15	100.0%	97.8%	100.0%	97.8%	97.8%	97.8%	97.8%	97.8%
20	95.7%	100.0%	100.0%	97.8%	100.0%	100.0%	95.7%	97.8%
30	100.0%	100.0%	100.0%	97.8%	100.0%	100.0%	100.0%	97.8%
50	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
100	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
200	100.0%	100.0%	97.8%	100.0%	100.0%	100.0%	100.0%	100.0%
300	97.8%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
500	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
800	100.0%	97.8%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
1000	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

In order to differentiate the 100% accuracy algorithms, their execution time and classification errors regarding each accelerometer and experimental run were estimated.

The *run classification error* corresponds to the sum of the difference between the previously explained ratio and the respective ideal classification of every experimental run. For instance, an experimental run with nine of its ten accelerometers classified as damaged (ratio of 0.9), in relation to its respective ideal classification (ten of ten accelerometers classified as anomalies, i.e. a ratio of 1) represents an error of 0.1. The symmetrical reasoning applies to the healthy state. The error values in table 5.4 are the sum of all experimental run errors for the respective configuration. A similar approach was made for the *accelerometer classification error*, where the values presented in the table are the sum of the differences between the classification ratios of each accelerometer and its respective ideal classification.

Table 5.4: Comparison of the preselection of algorithm architectures.

Code size	Epochs	Runtime [s]	\sum Run classification error	\sum Accelerometer classification error
1	15	14.2	4.54	1.98
1	30	16.7	3.96	1.80
5	20	16.3	4.25	1.63
10	15	15.7	4.54	1.63

Based on table 5.4, the architecture that presented the most balanced amount of error and runtime was the one with 30 *epochs* and a *code size* of 1, so it was the implemented configuration. The chosen architecture is optimal for the data collected in the experimental environment shown in chapter 3. For less controlled environment data, as is the case of a real-world scenario, where there is more variance like changing atmospheric conditions and longer acquisition time intervals, more demanding and powerful architectures may be required to distinguish the healthy from the damaged autoencoder reconstructions.

Titanium turbine blade

Considering that the ADAEs results were favourable for the GFRP WT blade when analysing all the tested masses, the same results were expected for the titanium WT blade since the masses used in this blade tests were heavier (table 3.2), meaning they would be easier to detect. This was confirmed with the perfect accuracy results obtained considering all the masses and using the minimal configuration of 1 *epoch* and *modal size* 1; the fastest and less computationally demanding type of AE possible. Therefore, optimisation analysis as the one done for GFRP blade architecture is not necessary for the Ti blade. The data analysed to draw these conclusions was from the impact hammer testing since it presented the best results.

Chapter 6

Conclusions

After all the research work developed throughout this thesis, it was concluded that implementing ML techniques concerning damage detection on vibration-based data was successful.

The fact that this study benefited from experimental tests on two different material WT blades with the same geometry was particularly useful to validate some of the reasonings developed throughout the project; coupled with the fact that these algorithms run automatically, without the need for manual interaction, allows them to be applicable to other WT blades.

The conclusions about the results obtained can be structured according to the two algorithms developed: MGAD and ADAEs.

Regarding the MGAD, its initial premise centred on the variation of natural frequencies of a damaged structure was soon confirmed in the chapter Modal parameter estimation. However, with regard to the GFRP blade, this frequency variation was insufficient for mass positions closer to the clamped side of the blade and lower magnitude masses, driving the algorithm to worse performance for masses up to 0.29% of the blade's weight. Since the Ti blade was only tested with masses of 0.6% of its total weight, the algorithm detected all the masses analysed in every position of the blade. When considering the application of this algorithm in a real-world scenario, in which the amount of undamaged data is much greater since the data collection is done most of the time in its healthy state, the confidence interval of frequencies considered healthy may become more accurate and sensitive to the point of detecting frequency deviations as small as those from the masses not detected by the MGAD algorithm. However, for a realistic application of this algorithm, it would need to be implemented in conjunction with a modal tracking algorithm like the one presented in section 4.3 so that the whole process from data acquisition until the detection of the anomalous natural frequencies could run in a fluid and automatic manner. This process should ideally be done in real-time so that the fault detection can be done in time to avoid significant damages. The heavy reliance on precise modal parameter estimation is a drawback of MGAD.

Concerning the ADAEs results, it can be concluded that its premise, based on the variation of the FRF of a damaged structure, verified in chapter 4, was proven to be quite reliable, enabling the algorithm to detect all the masses tested in every position of the two blades. The only aspect to consider that affects the algorithm's performance is the architecture and the optimisation of the hyperparameters. For

the Ti blade, the algorithm presented perfect accuracy for its simplest architecture. On the other hand, for the GFRP blade, it was necessary to adjust the number of epochs and the code size to obtain the same results. This adjustment was necessary because the masses tested on the GFRP blade were proportionally lighter, making them harder to detect.

It should be taken into account that the best damage detection results were obtained with experimental data from the modal hammer and modal shaker testings, which require EMA that is the most challenging form of modal analysis to implement in a real scenario in opposition to OMA, as explained in the chapter Modal analysis.

Overall, the ML implementation in the damage detection methodologies resulted in excellent classification accuracies. Although the ADAEs algorithm presented better damaged detection results, both techniques have the potential for achieving even better performances.

6.1 Achievements

This thesis's objectives, specifically the conducting of experimental campaigns to collect the vibration-based from the blades and the afterwards performance and understanding of the modal analysis principles, were fulfilled.

Regarding the implementation of ML algorithms, these were able to successfully accomplish the mission of detecting damage and interpreting its magnitude; however, no relevant progress was achieved concerning the part of localising the fault. The algorithms were then compared in terms of accuracy, sensitivity, damage assessment and applicability in a real-world scenario.

The main lesson learned with this project was how to research, understand and implement ML techniques in a short period, have a critical analysis of the results and come up with different approaches or new directions to take. It was also interesting working with large amounts of data of so many different kinds, namely data from different types of testing (impact; shaker; and pull-and-release), setups (without masses; with different magnitude masses; and different mass positions), structures (Ti WT blade; and GFRP WT blade), modal parameters (frequency; damping; and mode shapes) which lead to an almost infinite number of possible approaches for implementing ML. Organising and visualising such diverse data types was quite challenging and required a more profound knowledge of the theoretical background to interpret its behaviour and conceptualise possible improvements to the models.

6.2 Future work

The methodologies developed in this thesis still have a long path to take before they can be reliably implemented as a form of damage detection in a real-world application. Thus, several developments and future steps to improve their performance and reliability will be suggested:

- Introduce the modal parameter of damping as another variable to be taken into account in the MGAD algorithm. The multivariate Gaussian distribution was done considering the natural frequencies of all the normal modes of the structure. This ML technique can also be done considering both the frequency and damping values of each mode;
- In order to improve the accuracy and optimise the architecture of the algorithms, implement Monte Carlo methods to assess the best data parameters to analyse. Namely, which normal modes, frequency bandwidth, experimental runs, i.e. masses and positions, should be considered or have more or less weight in the thresholding and classification process of the algorithm. In addition, experiment with other ML techniques;
- Complement the study done for Ti blade, testing it with lower magnitude masses to match the blade mass percentages studied for the GFRP blade tests;
- Perform experimental tests on a blade with permanent damage besides the simulated damage from masses analysed in this thesis. It would also be interesting to do experimental tests on a blade bolted to a small-scale operating WT;
- For this thesis, only accelerometer data was used. There are other ways of acquiring data that may be relevant to analyse—for instance, surface velocity from a laser Doppler vibrometer or displacements from digital image correlation;
- Further elaborate on the methodologies developed aggregating more than one ML technique and analysing more than one type of modal parameter to achieve an algorithm capable of identifying damaged structures with a higher degree of certainty. Furthermore, apply these methodologies to a structure other than a WT blade.

References

- [1] G. Oliveira, F. Magalhães, Cunha, and E. Caetano. Vibration-based damage detection in a wind turbine using 1 year of data. *Structural Control and Health Monitoring*, 25(11), 11 2018. ISSN 15452263. doi: 10.1002/stc.2238.
- [2] M. McGugan and L. Mishnaevsky. Damage mechanism based approach to the structural health monitoring of wind turbine blades. *Coatings*, 10(12):1–23, 12 2020. ISSN 20796412. doi: 10.3390/coatings10121223.
- [3] Wind Energy in Europe: Outlook to 2023. Technical report, 2019.
- [4] E. Di Lorenzo, S. Manzato, B. Peeters, F. Marulo, and W. Desmet. Operational modal analysis for rotating machines: Challenges and solutions. *IOMAC 2017 - 7th International Operational Modal Analysis Conference*, (April):106–109, 2017.
- [5] Y. Feng, P. J. Tavner, and H. Long. Early experiences with uk round 1 offshore wind farms. *Proceedings of the Institution of Civil Engineers-energy*, 163(4):167–181, 2010.
- [6] J. Carroll, A. McDonald, and D. McMillan. Failure rate, repair time and unscheduled O&M cost analysis of offshore wind turbines. *Wind Energy*, (August 2015):1–20, 2013. doi: 10.1002/we.
- [7] W. Yang, Z. Peng, K. Wei, and W. Tian. Structural health monitoring of composite wind turbine blades: Challenges, issues and potential solutions. *IET Renewable Power Generation*, 11(4):411–416, 2017. ISSN 17521424. doi: 10.1049/iet-rpg.2016.0087.
- [8] F. L. dos Santos, B. Peeters, H. Van der Auweraer, L. C. Góes, and W. Desmet. Vibration-based damage detection for a composite helicopter main rotor blade. *Case Studies in Mechanical Systems and Signal Processing*, 3:22–27, 6 2016. ISSN 23519886. doi: 10.1016/j.csmssp.2016.01.001.
- [9] X. Chen, S. Semenov, M. McGugan, S. Hjelm Madsen, S. Cem Yeniceli, P. Berring, and K. Branner. Fatigue testing of a 14.3 m composite blade embedded with artificial defects – Damage growth and structural health monitoring. *Composites Part A: Applied Science and Manufacturing*, 140, 1 2021. ISSN 1359835X. doi: 10.1016/j.compositesa.2020.106189.
- [10] Y. Du, S. Zhou, X. Jing, Y. Peng, H. Wu, and N. Kwok. Damage detection techniques for wind turbine blades: A review, 7 2020. ISSN 10961216.

- [11] L. Doliński and M. Krawczuk. Damage detection in turbine wind blades by vibration based methods. In *Journal of Physics: Conference Series*, volume 181. Institute of Physics Publishing, 2009. doi: 10.1088/1742-6596/181/1/012086.
- [12] M. Martinez-Luengo, A. Kolios, and L. Wang. Structural health monitoring of offshore wind turbines: A review through the Statistical Pattern Recognition Paradigm, 10 2016. ISSN 18790690.
- [13] C. R. Farrar and K. Worden. *Structural health monitoring: a machine learning perspective*. John Wiley & Sons, 2012.
- [14] D. García and D. Tcherniak. An experimental study on the data-driven structural health monitoring of large wind turbine blades using a single accelerometer and actuator. *Mechanical Systems and Signal Processing*, 127:102–119, 7 2019. ISSN 10961216. doi: 10.1016/j.ymssp.2019.02.062.
- [15] D. Tcherniak and L. L. Mølgaard. Active vibration-based structural health monitoring system for wind turbine blade: Demonstration on an operating Vestas V27 wind turbine. *Structural Health Monitoring*, 16(5):536–550, 9 2017. ISSN 17413168. doi: 10.1177/1475921717722725.
- [16] G. Oliveira, F. Magalhães, Cunha, and E. Caetano. Continuous dynamic monitoring of an onshore wind turbine. *Engineering Structures*, 164:22–39, 6 2018. ISSN 18737323. doi: 10.1016/j.engstruct.2018.02.030.
- [17] S. Hall. The effective management and use of structural health data. In *Proceedings of the 2nd International Workshop on Structural Health Monitoring*, pages 265–275, 1999.
- [18] Z. Liu, M. Ardabilian, A. Zine, and M. Ichchou. Crack-type damage detection and localization of a thick composite sandwich structure based on Convolutional Neural Networks. Technical report.
- [19] G. R. Gillich, Z. I. Praisach, M. Abdel Wahab, and O. Vasile. Localization of transversal cracks in sandwich beams and evaluation of their severity. *Shock and Vibration*, 2014, 2014. ISSN 10709622. doi: 10.1155/2014/607125.
- [20] G. C. Larsen, P. Berring, D. Tcherniak, H. Nielsen, K. Branner, and K. B. Effect. Effect of a Damage to Modal Parameters of a Wind Turbine Blade. Technical report, 2014. URL <https://hal.inria.fr/hal-01020347>.
- [21] Y. Wang, M. Liang, and J. Xiang. Damage detection method for wind turbine blades based on dynamics analysis and mode shape difference curvature information. *Mechanical Systems and Signal Processing*, 48(1-2):351–367, 10 2014. ISSN 10961216. doi: 10.1016/j.ymssp.2014.03.006.
- [22] S. S. Kessler, S. M. Spearing, M. J. Atalla, C. E. Cesnik, and C. Soutis. Damage detection in composite materials using frequency response methods. *Composites Part B: Engineering*, 33(1): 87–95, 2002.
- [23] E. D. Lorenzo, G. Petrone, S. Manzato, B. Peeters, W. Desmet, and F. Marulo. Damage detection in wind turbine blades by using operational modal analysis. *Structural Health Monitoring*, 15(3): 289–301, 5 2016. ISSN 17413168. doi: 10.1177/1475921716642748.

- [24] F. Marulo, G. Petrone, E. Di Lorenzo, and A. Cutolo. Operational modal analysis for SHM of a wind turbine blade. *SHMII 2015 - 7th International Conference on Structural Health Monitoring of Intelligent Infrastructure*, 2015.
- [25] F. Marulo, G. Petrone, V. D'alessandro, and E. D. Lorenzo. Operational modal analysis on a wind turbine blade. Technical report.
- [26] M. D. Ulriksen, D. Tcherniak, P. H. Kirkegaard, and L. Damkilde. Operational modal analysis and wavelet transformation for damage identification in wind turbine blades. *Structural Health Monitoring*, 15(4):381–388, 7 2016. ISSN 17413168. doi: 10.1177/1475921715586623.
- [27] J. Chatterjee and N. Dethlefs. Scientometric review of artificial intelligence for operations & maintenance of wind turbines: The past, present and future, 7 2021. ISSN 18790690.
- [28] O. Janssens, V. Slavkovikj, B. Vervisch, K. Stockman, M. Loccufier, S. Verstockt, R. Van de Walle, and S. Van Hoecke. Convolutional neural network based fault detection for rotating machinery. *Journal of Sound and Vibration*, 377:331–345, 2016.
- [29] O. Abdeljaber, O. Avci, S. Kiranyaz, M. Gabbouj, and D. J. Inman. Real-time vibration-based structural damage detection using one-dimensional convolutional neural networks. *Journal of Sound and Vibration*, 388:154–170, 2017.
- [30] R. Zhao, R. Yan, Z. Chen, K. Mao, P. Wang, and R. X. Gao. Deep learning and its applications to machine health monitoring, 1 2019. ISSN 10961216.
- [31] Y. Lei, F. Jia, J. Lin, S. Xing, and S. X. Ding. An Intelligent Fault Diagnosis Method Using Unsupervised Feature Learning Towards Mechanical Big Data. *IEEE Transactions on Industrial Electronics*, 63(5):3137–3147, 5 2016. ISSN 02780046. doi: 10.1109/TIE.2016.2519325.
- [32] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [33] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [34] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [35] M. S. Minhas and J. Zelek. Semi-supervised Anomaly Detection using AutoEncoders. 2020. URL <http://arxiv.org/abs/2001.03674>.
- [36] M. Sakurada and T. Yairi. Anomaly detection using autoencoders with nonlinear dimensionality reduction. *ACM International Conference Proceeding Series*, 02-December:4–11, 2014. doi: 10.1145/2689746.2689747.

- [37] W. Sun, S. Shao, R. Zhao, R. Yan, X. Zhang, and X. Chen. A sparse auto-encoder-based deep neural network approach for induction motor faults classification. *Measurement: Journal of the International Measurement Confederation*, 89:171–178, 7 2016. ISSN 02632241. doi: 10.1016/j.measurement.2016.04.007.
- [38] C. Lu, Z. Y. Wang, W. L. Qin, and J. Ma. Fault diagnosis of rotary machinery components using a stacked denoising autoencoder-based health state identification. *Signal Processing*, 130:377–388, 1 2017. ISSN 01651684. doi: 10.1016/j.sigpro.2016.07.028.
- [39] L. Wen, L. Gao, and X. Li. A new deep transfer learning based on sparse auto-encoder for fault diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(1):136–144, 1 2019. ISSN 21682232. doi: 10.1109/TSMC.2017.2754287.
- [40] J. An and S. Cho. Variational autoencoder based anomaly detection using reconstruction probability. *Special Lecture on IE*, 2(1):1–18, 2015.
- [41] O. Avci, O. Abdeljaber, S. Kiranyaz, and D. Inman. Structural damage detection in real time: implementation of 1d convolutional neural networks for shm applications. In *Structural Health Monitoring & Damage Detection*, volume 7, pages 49–54. Springer, 2017.
- [42] T. Ince, S. Kiranyaz, L. Eren, M. Askar, and M. Gabbouj. Real-time motor fault detection by 1-d convolutional neural networks. *IEEE Transactions on Industrial Electronics*, 63(11):7067–7075, 2016.
- [43] W. Zhang, C. Li, G. Peng, Y. Chen, and Z. Zhang. A deep convolutional neural network with new training methods for bearing fault diagnosis under noisy environment and different working load. *Mechanical Systems and Signal Processing*, 100:439–453, 2018.
- [44] F. Jia, Y. Lei, J. Lin, X. Zhou, and N. Lu. Deep neural networks: A promising tool for fault characteristic mining and intelligent diagnosis of rotating machinery with massive data. *Mechanical Systems and Signal Processing*, 72:303–315, 2016.
- [45] B. Lopes, C. Colangeli, K. Janssens, A. Mroz, and H. Van der Auweraer. Neural network models for the subjective and objective assessment of a propeller aircraft interior sound quality. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, volume 259, pages 4124–4135. Institute of Noise Control Engineering, 2019.
- [46] A. Tavares, E. D. Lorenzo, B. Peeters, G. Coppotelli, and N. Silvestre. Damage detection in lightweight structures using Artificial Intelligence techniques. Technical report.
- [47] S. Kiranyaz, T. Ince, and M. Gabbouj. Real-time patient-specific ecg classification by 1-d convolutional neural networks. *IEEE Transactions on Biomedical Engineering*, 63(3):664–675, 2015.
- [48] M. de Oliveira, A. Monteiro, and J. Vieira Filho. A new structural health monitoring strategy based on pzt sensors and convolutional neural network. *Sensors*, 18(9):2955, 2018.

- [49] J. C. Aldrin and D. S. Forsyth. Demonstration of using signal feature extraction and deep learning neural networks with ultrasonic data for detecting challenging discontinuities in composite panels. In *AIP Conference Proceedings*, volume 2102, page 020012. AIP Publishing, 2019.
- [50] A. Shihavuddin, X. Chen, V. Fedorov, A. Nymark Christensen, N. Andre Brogaard Riis, K. Branner, A. BJORHOLM DAHL, and R. Reinhold Paulsen. Wind turbine surface damage detection by deep learning aided drone inspection analysis. *Energies*, 12(4):676, 2019.
- [51] C. V. Dung et al. Autonomous concrete crack detection using deep fully convolutional neural network. *Automation in Construction*, 99:52–58, 2019.
- [52] D. J. Ewins. *Modal Testing: Theory, Practice and Application*. John Wiley Sons, 2nd edition, 2009.
- [53] P. Avitabile. *Modal Testing: A Practitioner's Guide*. 2017. ISBN 9781119222897. doi: 10.1002/9781119222989.
- [54] M. Pastor, M. Binda, and T. Harčarik. Modal assurance criterion. *Procedia Engineering*, 48: 543–548, 2012. ISSN 1877-7058. doi: <https://doi.org/10.1016/j.proeng.2012.09.551>. URL <https://www.sciencedirect.com/science/article/pii/S1877705812046140>. Modelling of Mechanical and Mechatronics Systems.
- [55] W. Heylen, S. Lammens, and P. Sas. *Modal Analysis Theory and Testing*. Katholieke Universiteit Leuven, Faculty of Engineering, Department of Mechanical Engineering, Division of Production Engineering, Machine Design and Automation, Leuven, 1998.
- [56] What is operating modal analysis. URL <https://community.sw.siemens.com/s/article/OMG-What-is-OMA-Operating-Modal-Analysis>. 2021-09-30.
- [57] R. Brincker and C. Ventura. *Introduction to Operational Modal Analysis*. John Wiley Sons, 2015.
- [58] C. Rainieri and G. Fabbrocino. *Operational Modal Analysis of Civil Engineering Structures: An Introduction and Guide for Applications*. Springer-Verlag, New York, 1st edition.
- [59] A. Tavares. Artificial Intelligence Approach to Damage Detection in Lightweight Structures. (July), 2020.
- [60] C. Ferreira. Digital Image Correlation for Vibration Analysis. 2021.
- [61] B. Peeters, H. Van Der Auweraer, P. Guillaume, and J. Leuridan. The PolyMAX frequency-domain method: A new standard for modal parameter estimation? *Shock and Vibration*, 11(3-4):395–409, 2004. ISSN 10709622. doi: 10.1155/2004/523692.
- [62] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [63] C. Bishop. *Pattern Recognition and Machine Learning: All "just the Facts 101" Material*. Information science and statistics. Springer (India) Private Limited, 2013. ISBN 9788132209065. URL <https://books.google.co.in/books?id=HL4HrgEACAAJ>.

- [64] From Autoencoder to Beta-VAE. URL <https://lilianweng.github.io/lil-log/2018/08/12/from-autoencoder-to-beta-vae.html>.
- [65] Properties: Titanium Alloys - Ti6Al4V Grade 5. URL <https://www.azom.com/properties.aspx?ArticleID=1547>.
- [66] F. Sayer, A. Antoniou, S. Goutianos, I. Gebauer, K. Branner, and C. Balzani. ReliaBlade Project: A Materials Perspective towards the Digitalization of Wind Turbine Rotor Blades. *IOP Conference Series: Materials Science and Engineering*, 942(1), 2020. ISSN 1757899X. doi: 10.1088/1757-899X/942/1/012006.
- [67] B. F. Sørensen. Materials and structures for wind turbine rotor blades - An overview. *ICCM International Conferences on Composite Materials*, 2009.
- [68] L. Mishnaevsky, K. Branner, H. N. Petersen, J. Beauson, M. McGugan, and B. F. Sørensen. Materials for wind turbine blades: An overview. *Materials*, 10(11):1–24, 2017. ISSN 19961944. doi: 10.3390/ma10111285.
- [69] Modal Assurance Criterion (MAC). URL <https://community.sw.siemens.com/s/article/modal-assurance-criterion-mac>.
- [70] A. Tavares, E. D. Lorenzo, B. Cornelis, B. Peeters, K. Gryllias, and W. Desmet. Damage localization on lightweight structures with non-destructive testing and machine learning techniques. In *Data Science in Engineering, Volume 9*, pages 71–74. Springer, 2022.