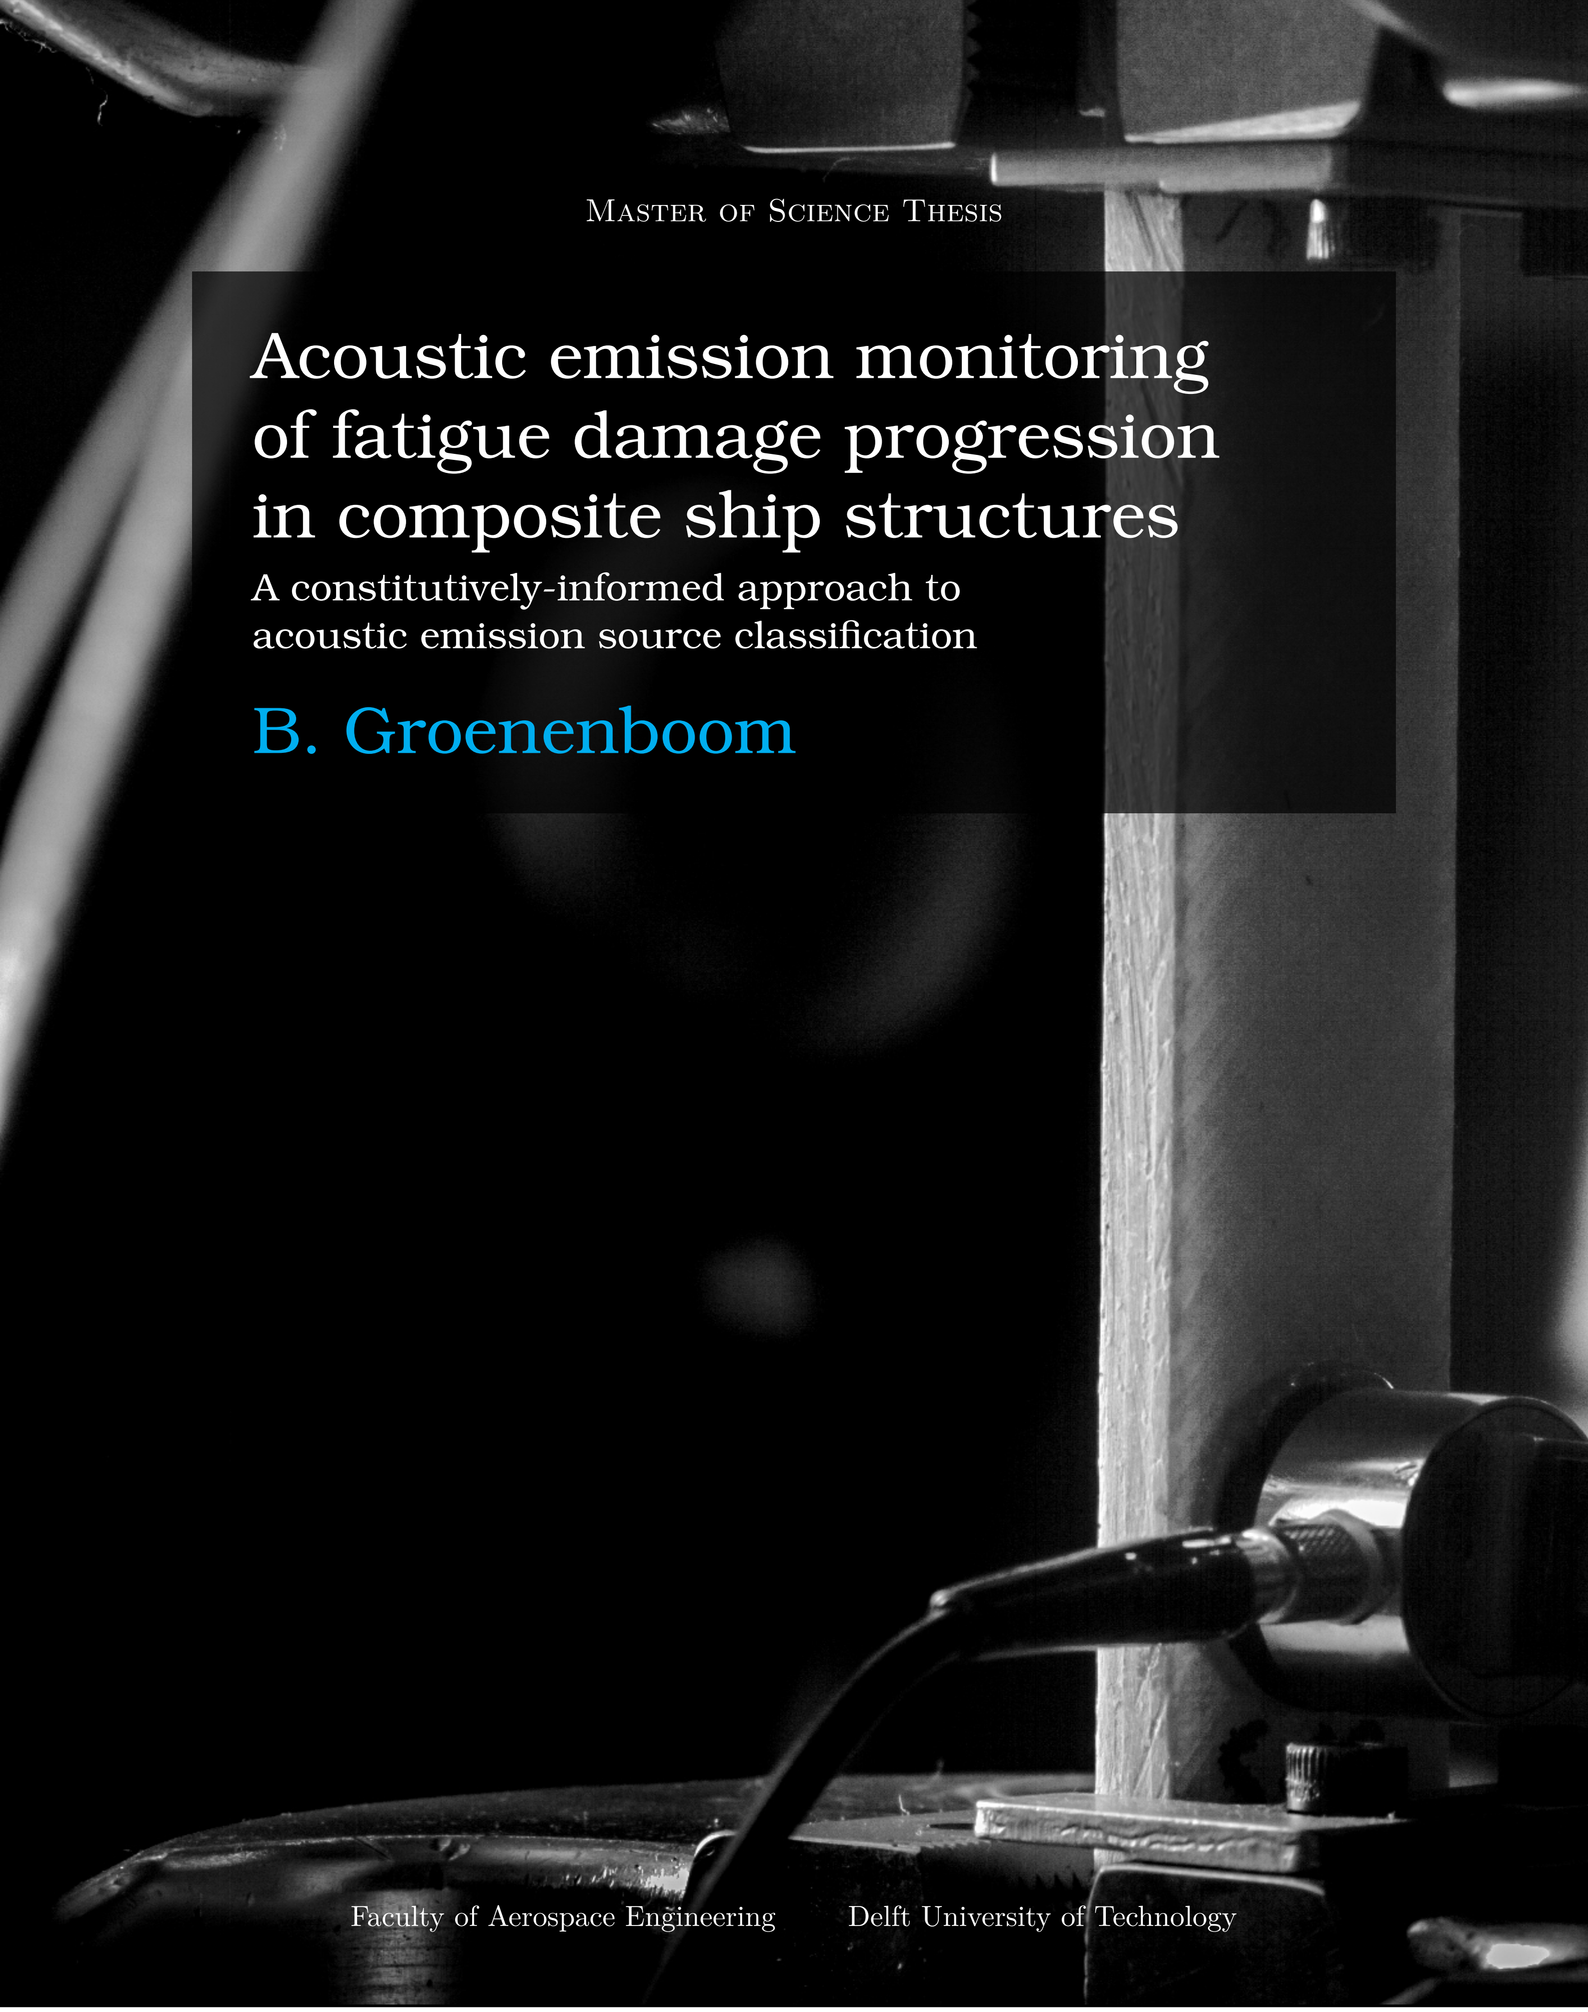

Master of Science Thesis

Acoustic emission monitoring of fatigue damage progression in composite ship structures

A constitutively-informed approach to acoustic emission source classification

B. Groenenboom

Faculty of Aerospace Engineering · Delft University of Technology

Acoustic emission monitoring of fatigue damage progression in composite ship structures

A constitutively-informed approach to acoustic emission source classification

Master of Science Thesis

For obtaining the degree of Master of Science in Aerospace Engineering
at Delft University of Technology

B. Groenenboom

16 September 2016

Faculty of Aerospace Engineering · Delft University of Technology

The work in this thesis was supported by Damen Schelde Naval Shipbuilding. Their cooperation is gratefully acknowledged.

Delft University of Technology
Faculty of Aerospace Engineering
Department of Aerospace Structures and Materials

GRADUATION COMMITTEE

Dated: 16 September 2016

Chair holder: Dr.ir. R.C. Alderliesten

Committee members: Dr. R.M. Groves

Dr. D. Zarouchas

Ir. L. Morel

Dr. O.A. Sharpanskykh

Abstract

Acoustic Emission (AE)-based structural integrity monitoring of composite structures potentially offers the means to gain knowledge on in-service damage accumulation, as well as mitigation of some of the risks associated with damage developing during operation. However, particularly for composite ship structures, this subject is sparsely represented in scientific literature. Moreover, one of the central challenges of AE monitoring—that of classification of AE events into source mechanisms—remains a largely unsolved problem. This can be attributed to the fact that validation of any identified classes of hits with regard to their underlying source mechanics is generally limited to data-centred metrics and post-mortem imaging of testing specimens, consequentially lacking the means to provide direct evidence for a constitutive link between material mechanics and AE hit classification. The current project aims to improve on the approach to the classification problem, by augmenting a data analysis approach to AE hit classification with a material mechanics perspective for the case of Tension-Tension (T-T) fatigue testing of Glass Fibre Reinforced Polymer (GFRP) coupons. This mechanics-informed perspective consists of using the instantaneous fatigue cycle phase as hit labels and validating this labelling with the coupon's stiffness degradation measured using Digital Image Correlation (DIC). From the fatigue cycle phase labelling of hits, at least two distinct classes are observed: first, the hits occurring mainly in the loading phase of the fatigue cycle, whose activity, when isolated, correlates positively with the rate of stiffness degradation in the test coupons; and second, the hits occurring in the unloading phase of the fatigue cycle, which consist of the vast majority of total AE activity in these fatigue tests. It is further found that this labelling strongly agrees with clustering of the hits in the 11-dimensional waveform feature-space of amplitude, energy, duration, rise time, counts, Root Mean Square (RMS) noise; and 5 partial power fractions in the frequency spectrum, when said feature-space is plotted by means of t-Distributed Stochastic Neighbour Embedding (t-SNE). It is concluded that the triple combination of validation of AE hit classification against stiffness degradation, the representation of the AE feature-space by t-SNE and the perspective offered by fatigue cycle phase labelling of hits shows promise as an effective and consistently verifiable approach to the AE classification problem. Therefore, the approach presented in the current project constitutes a novel methodology for AE source classification, providing potentially new insights that may not only be valuable to further development of AE-based structural integrity monitoring of composite structures, but for which its use may also lead to gains in the understanding of (fatigue) damage processes in composites in general.

Table of contents

List of acronyms

AE Acoustic Emission.

ASTM American Society for Testing and Materials.

CFRP Carbon Fibre Reinforced Polymer.

DIC Digital Image Correlation.

DSNS Damen Schelde Naval Shipbuilding.

FFT Fast Fourier Transform.

FLAP Fatigue Life Assessment Program.

GF Glass-Fibre.

GFRP Glass Fibre Reinforced Polymer.

IACS International Association of Classification Societies.

ICA Independent Component Analysis.

JIP Joint Industry Project.

k-NN k-Nearest-Neighbour.

KDE Kernel Density Estimation.

NDE Non-Destructive Evaluation.

NDT Non-Destructive Testing.

NN Neural Network.

PCA Principal Component Analysis.

PDF Probability Density Function.

RMS Root Mean Square.

SHM Structural Health Monitoring.

SOM Self-Organizing Map.

T-C Tension-Compression.

t-SNE t-Distributed Stochastic Neighbour Embedding.

T-T Tension-Tension.

TEU Twenty-foot Equivalent Unit.

UD Uni-Directional.

UTS Ultimate Tensile Strength.

List of symbols

A cross-sectional area (of a test specimen) $[m^2]$.

A_0 first asymmetric wave mode [-].

C_V coefficient of variation (σ/μ) of a distribution.

D Fatigue damage parameter: $1 - E_f/E_0$ [-].

E Young's modulus $[N/m^2]$.

E_0 initial ("zero" cycles) fatigue modulus $[N/m^2]$.

E_f fatigue modulus $[N/m^2]$.

F_N normal force $[N]$.

F_{UTS} ultimate tensile strength $[N]$.

F_a fatigue cycle amplitude $[N]$.

F_{max} fatigue cycle maximum load $[N]$.

F_{mean} fatigue cycle mean load $[N]$.

F_{min} fatigue cycle minimum load $[N]$.

L arbitrary length dimension $[m]$.

M applied moment $[Nm]$.

N number of fatigue cycles [-].

N_f number of fatigue cycles until failure [-].

R fatigue ratio: F_{min}/F_{max} [-].

S_0 first symmetric wave mode [-].

S_b micro buckling stress $[N/m^2]$.

S_t tensile strength $[N/m^2]$.

T time $[s]$.

V_f fibre volume fraction [-].

W weight $[N]$.

δ displacement $[m]$.

λ wavelength $[m]$.

ν Poisson's ratio [-].

ρ material density $[kg/m^3]$.

σ stress $[N/m^2]$.

θ arbitrary angle, e.g. of fibres w.r.t. loading direction [°].

ε strain [-].

ε_c strain in the composite [-].

$\varepsilon_{d,h}$ transverse fibre-matrix interface failure strain [-].

ε_{max} maximum fatigue strain [-].

ε_m strain in the matrix [-].

ε_{yy} (average) surface strain in the y-direction (along coupon axis) [-].

a crack dimension $[m]$.

c_1 symmetric wave velocity $[m/s]$.

c_2 asymmetric wave velocity $[m/s]$.

f frequency $[Hz]$.

h half plate thickness $[m]$.

k wave number [-].

t thickness $[m]$.

v (wave) velocity $[m/s]$.

List of figures

List of tables

Preface

Although my education has provided me with a decent background in aerospace engineering and composites, the subject of this report has been an entirely novel addition to the knowledge I have gained throughout both my BSc and MSc education. In tackling this project, it has become clear that the field of Acoustic Emission monitoring of composites consists of an overlapping area between the fields of data science and of material mechanics. This has made this project into an interesting exercise; but with my background it also makes sure that I am always lacking the mastery an expert would have of either of these two fields. Nonetheless, it is equally clear that an appreciation of (at least) both of these disciplines will be required to progress the field of Acoustic Emission monitoring of composite structures.

With any field of specialism, especially for an outsider it is easy to accept that if the established experts have not found a solution to a major challenge, the challenge must be impossible. In this project, it has become especially clear to me that there is in fact more truth to a converse statement: with only a moderate understanding of both composite damage mechanics and data science, it can be revealed that the bottleneck on progress is not defined by whether or not tasks are impossible; but rather by how much meaningful work can be put into the subject. There is an unfathomable number of problems available in science that warrant deeper investigation, even just within the project discussed in this document. All that is needed to advance not only structural integrity monitoring, but potentially the understanding of composite damage mechanics as well, is to dissect some of these problems with structured thinking, careful assertions (and possibly a fair amount of money). It is an inspiring and humbling thought that even though the amount of expertise in the modern world of science & technology may be overwhelming, the limiting factor in progress is still quite simply the number of people willing to spend enough time with the subject material. Or as some people may put it: we have a surprisingly good record at accomplishing the impossible. I am quite confident that any single person who would continue spending time on this subject from where I left off would be bound to make a difference sooner or later.

This document constitutes my Master of Science graduation thesis. This thesis project is supported by Damen Schelde Naval Shipbuilding and carried out in the *Structural Integrity and Composites* (SI&C) group and the *Aerospace Non-Destructive Testing Laboratory* (AeroNDT) of the *Aerospace Structures & Materials* department of *Delft University of Technology.* The structure of this report is conventional for reports of this type: First the project is introduced and framed, then background theory is presented, followed by testing- and analysis methodology, results; and finally conclusions and recommendations are presented. Readers who are already familiar with the fundamental concepts of Acoustic Emission monitoring and damage in composites may be interested to skip ahead to Chapter 4 (particularly Section 4.3), for details on the processing steps particular to the approach taken in this AE monitoring study; or to Chapter 5, detailing the analysis results attained through aforementioned approach. Readers interested in potential future prospects arising from the current approach to AE monitoring of composites may be particularly interested in the conclusions and recommendations chapters (resp. Ch. 6 and 7).

Various people who have taken part in supporting this project deserve to be thanked. First and foremost I would like to extend my gratitude towards my supervisors, Roger Groves and Dimitrios Zarouchas, who have not only supported me with critical feedback and insightful discussion, but also had to suffer my unyielding character throughout this project. I would like to thank Laurent Morel and Damen Schelde Naval Shipbuilding for supporting the project. The staff of the *Delft Aerospace Structures and Materials Laboratory* is thanked for their assistance in the project's production and testing campaign . Chirag Anand and Ping Liu deserve special thanks for trying to find the needles in the haystack. Finally, I would like to thank my dad, for supporting me through the many years of my academics and never showing a loss of confidence in my ability.

“Scientific progress goes ‘boink’?”

— *Calvin & Hobbes, by Bill Watterson*

Chapter 1

Introduction

An interest is voiced by industrial partner Damen Schelde Naval Shipbuilding (DSNS) to use Acoustic Emission (AE) monitoring as a means to evaluate and keep track of damage events over the lifetime of selected Glass Fibre Reinforced Polymer (GFRP) structures. While selected naval applications can stand to benefit from the use of composite structures for their high specific stiffness, -strength and fatigue resistance, the potentially attainable gains in structural efficiency are limited by the industry's limited means for full scale testing. Therefore, it is envisioned that AE monitoring of composite ship structures may be applied as a structural integrity monitoring tool through which knowledge can be gained on in-service damage accumulation, while also providing a potential means to mitigate structural damage risks in-service. [1] AE-based structural integrity monitoring, particularly for composite ship structures, is represented only to a limited degree in current literature. Therefore, in order to tackle this subject from the ground up in a structured manner, a building-block approach is taken for the investigation of the efficacy of AE as a tool for structural integrity monitoring of composite structures. This approach consists of building up damage classification confidence from coupon to structure and from relatively uni-modal damage mechanisms to more representative combinations of damage mechanisms. The current project aims to start at the base of this building-block approach by focusing on a coupon-sized structure which is subjected to stiffness degradation by a single dominant damage mechanism (matrix cracking) [2].

One of the major open challenges identified in the field of AE monitoring of composites is the classification of AE events into source mechanisms [3]. In many recent studies, machine learning techniques have been applied to be able to process the large and complex datasets produced by AE monitoring. However, with these techniques, validation of any identified classes of hits with regard to any underlying mechanics is generally limited to data-centred metrics and post-mortem imaging of testing specimens, consequentially lacking direct evidence for a constitutive link between material mechanics and AE hit classification. Furthermore, post-mortem observations are found to be not very suitable to detect damage for the specimens of relatively high ply thickness under consideration in this study. To resolve this shortcoming, the current project augments a data analysis approach to AE hit classification with a material mechanics perspective, by using the instantaneous fatigue cycle phase as hit labels and validating this labelling with the coupon's stiffness degradation measured using Digital Image Correlation (DIC). In addition, the classes in the supervised labelling are compared against the clustering of AE hits in 11-dimensional feature-space, using t-Distributed Stochastic Neighbour Embedding (t-SNE) as a means to visualise the high-dimensional data. The results of the current project therefore contribute to the development of AE-based structural integrity monitoring methodology by providing a more robust mechanics-driven approach to the analysis and classification of AE data. Furthermore, the methodology developed this project shows promise for use as a tool to further the understanding of fatigue damage processes in composites in general.

This document is structured as follows: In Chapter 2, the project framework is presented in detail. Thereafter, Chapter 3 provides a summary of the current understanding of damage mechanisms in composite plates, as well on acoustic emission monitoring. Subsequently, the project methodology is explained in Chapter 4. This consists of a description of coupon production and -testing processes, as well as a detailing of the processing and filtering steps applied to the empirical data acquired from testing. Analysis and results of this data is presented in Chapter 5, while conclusions to the project are presented in Chapter 6. Finally, this document is concluded with recommendations for future research in Chapter 7.

Chapter 2

Project framework

In this chapter, a contextual framework is presented for the project and its approach. The current project stems from a desire expressed by the project's industry partner, DSNS, to apply in-situ AE monitoring to composite ship structures. DSNS is part of the Netherlands-based Damen Schelde Group, which designs and builds ships of various types and purposes, including military and civilian ships, tugs, offshore support vessels, yachts and ferries. [4] The context of the project is rooted in the production methods used by DSNS; therefore, the production methods and materials of interest are based on those of DSNS and are presented in this chapter. Additionally, an introduction to typical design loads in a generic ship is given, in order to bring structural integrity monitoring to the context of relevant loading and damage types. But first, Section 2.1 describes the definitions and goals of the project in a little more detail, relating the desires of operators to research questions for the project. Following the current chapter, Chapter 3 provides in-depth descriptions of damage mechanisms in composite plates and AE monitoring, whereas the methodology of the current project is laid out in Chapter 4.

2.1 Research questions and structural integrity monitoring

In this section, the prospective goals of AE monitoring of composite structures are discussed in more detail, while a set of research questions is distilled. It has been established that this project stems from a desire by DSNS to evaluate and keep track of damage events in GFRP structures by use of AE monitoring. This may be described more formally as a desire to apply in-situ structural integrity monitoring to composites ship structures; and less formally as a desire to draw conclusions about a structure's damage state —whether that be in the form of stiffness, residual strength or otherwise— based on continuous in-service measurements on said structure. Envisioned gains from this objective can include a more thorough understanding of damage occurring during a (ship) structure's lifetime, diagnostic support to aid in a form of condition-based maintenance; and ultimately a potential for reductions in design safety factors by means of the former two gains. However, whether any particular goals as mentioned above are attainable with AE monitoring, depends foremost on the question of what information is contained in AE events.

From Section 3.2, reviewing current scientific literature on the topic of AE-based structural integrity monitoring of composites, follow some of the challenges and limitations of this technique: firstly and fundamentally, it may be important to note that AE monitoring is only able to document events as they are happening while the system is active; in other words, capturing transient behaviour. It does therefore not account for initial conditions of the structure and material. Secondly, although AE may on the outset be particularly suitable to monitor large sections of a structure, it does not follow unambiguously from existing scientific literature what information is exactly contained in the occurrence and the character of AE events originating from a composite laminate under loading. I.e. it is not entirely clear if the collection of AE events provides information that may be used to describe, for example, either what damage mechanisms are occurring, the amount of stiffness degradation; or remaining residual strength or -stiffness of the material. Therefore, a research question for the project may be posed in a general form:

Q1. How can in-situ AE measurements on a GFRP structure be used to meaningfully monitor damage development?

An important keyword in Q1 is "meaningful". It is not unambiguous what the attainable meaningful information is. And whether information is meaningful, depends on the purpose of the operator. From the review of AE monitoring of composites in Section 3.2, it is furthermore concluded that the central major challenge of AE-based monitoring of composites is that of classification and validation of AE events into their respective source mechanisms. That is: it is possible to distinguish an arbitrary number of event classes in AE data, but it is non-trivial to validate what source mechanisms these event classes correspond to. In order to define a parameter against which AE data may be compared and validated, in this project DIC measurements are used to determine normalised stiffness degradation (see Section 4.3.1). Stiffness degradation is thereby chosen for this project as the "provisional" answer to the question of a meaningful measure of damage development implied in Q1. From the preceding, research question Q1 can then be expanded into two more precisely defined sub-questions:

Q1.1 Does the occurrence and/or the character of AE events in a composite structure contain information about damage development in the structure?

Q1.2 How can the classification of AE events contributing to damage be validated?

In this set of questions, note that Q1.1 pertains mainly to the question of what information is exactly contained in the AE hits, whereas Q1.2 concerns itself mainly with the problem of classification and validation; a challenge that is identified as a primary one for the field of AE-based monitoring of composite structures. The use of stiffness degradation as a measure for damage in this project then serves both questions, by being a measure of whether damage development is captured and a means to validate the classification of AE events. Furthermore, observing the somewhat fundamental nature of the questions to be answered, a simplified case for a structure and loading mode are considered, which will consist of low-amplitude fatigue loading on coupon specimens. These choices are substantiated in the chapters that follow.

2.2 Subject material and production parameters

The material type of interest is an E-glass fibre reinforced vinyl-ester (Dion VE 1057) composite in Uni-Directional (UD)-laminates. In structural applications by DSNS, the resin in this composite is used in a rubberised form. However, as producing the rubberised version of this resin constitutes a special production process which is not normally applied in small quantities, the experimental work described in this document is based on the un-rubberised base resin. In structural applications for DSNS, the typical laminate thickness varies between 4 and 120 mm, favouring quasi-isotropic layups. These laminates are produced in a vacuum infusion production process, using 2400 g/m^2 UD plies. Some structural details are finished in a hand lay-up process using 1200 g/m^2 plies. For the current project, the latter type of UD material is chosen for reasons rooted in economy and ease of handling. This material is depicted in Figure 2.1.

It may be noted that while no specific assertion of production quality from DSNS is available at this time, production quality in the marine and off-shore sector is considered relatively poor compared to other industrial sectors [5]. While the current project does not aim to investigate a dependency on production quality specifically, it is notable that it is likely to have an influence on damage initiation and/or progression. At the same time, the variability of production quality can also be understood as an additional argument supporting a demand for structural integrity monitoring in naval applications. Finally, while some of the structures built by DSNS are in fact not monolithic- but sandwich panels, the current project focuses on monolithic composites only, so as to simplify the number of different damage mechanisms under consideration. A detailed discussion of damage mechanisms in composite panels is presented in Section 3.1, whereas production of testing coupons out of the aforementioned materials is described in Section 4.1. [1]

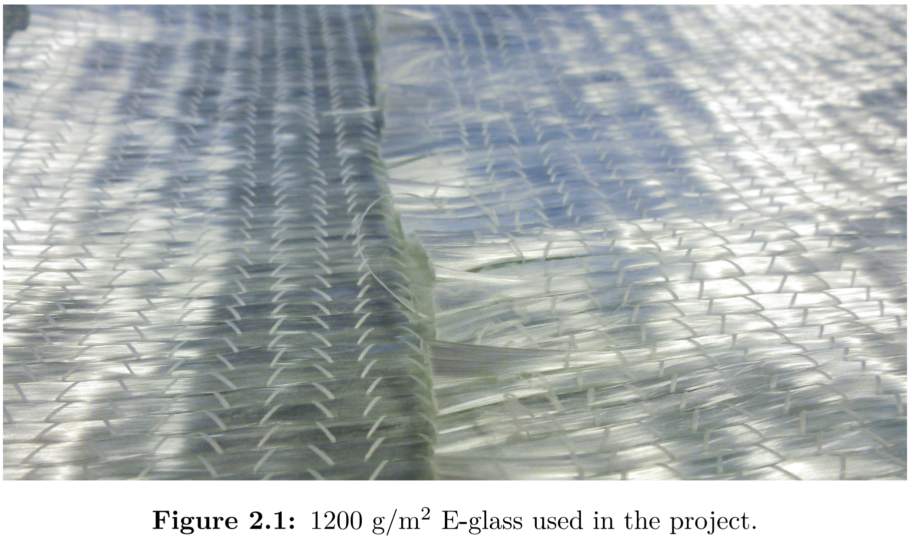

Figure 2.1: 1200 g/m^2 E-glass used in the project.

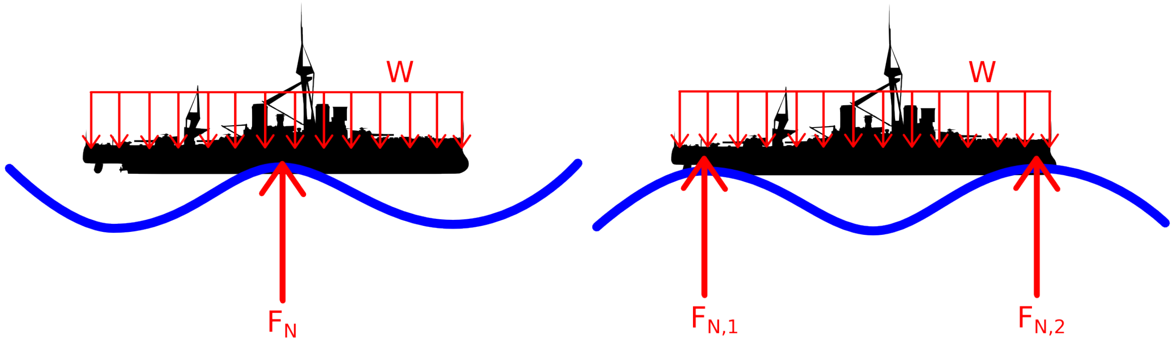

(a) Hogging **(b)** Sagging

Figure 2.2: Exaggerated schematic representation of a ship under vertical bending loads in hogging **(a)** and sagging **(b)** conditions. The ship weight W is balanced by normal buoyancy forces F_N, causing the ship to bend. Ship clip art from Wikimedia commons; CC BY-SA.

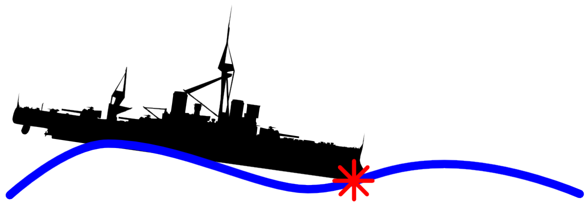

Figure 2.3: Schematic representation of the dynamic load case of a ship slamming into a wave as it is moving through the waves; the hull-wave impact indicated by the red star. Ship clip art from Wikimedia commons; CC BY-SA.

2.3 Ship structures and loads

In this section, a generic ship hull is considered, to gain an insight into the general set of load cases on a ship. From a structural perspective, a generic ship (hull) can be simplified to a slender, stiffened, thin-walled beam. The beam is bent both in the horizontal and vertical plane, for example by waves at an angle to the ship movement direction and by wave peaks and troughs respectively. Additionally, a torsional moment may be introduced, for example by asymmetric cargo loading or waves at an angle to the ship bearing. Evidently, when the hull of the ship bends in the vertical plane, the top and bottom of the “beam” are in tension or compression, while the sides are loaded in shear. [6]

Longitudinal loads on the beam model are introduced by hogging and sagging in passing waves, as depicted in Figure 2.2. When the ship is on top of a wave crest, its weight causes a bending moment to bend the ship downwards on both ends as shown in Figure 2.2a. In the opposite situation the ship is said to be sagging, as shown in Figure 2.2b. It may be noted that these are extrema; e.g. in a real ship in the hogging condition it is not necessarily supported by the wave in the middle; and the hydrostatic support of the ends is not necessarily completely zero. Even so, the actual maximum longitudinal bending loads do occur in these worst cases of a head wave of similar length as the ship. The design magnitude of these wave loads is standardized by the International Association of Classification Societies (IACS). In this example, the IACS specifies the bending moment distribution over the longitudinal direction as a function of ship length. [6]

Slamming of the ship in a wave is another typical load for a ship hull, as illustrated by Figure 2.3. This occurs particularly near the bottom front of the hull, and can result in dents in metal shell plating. The physics of slamming loads is governed by wave impact pressure, for which theoretical analysis is non-trivial. Therefore, these loads are often investigated experimentally. Indeed, the category of "unexpected loads", which may include rigid as well as hydrodynamic impact, is considered the most critical in current composite ship structures, as high safety factors on strength make failure from normal service conditions rare. [5, 6]

On a more detailed level, stiffened panels in the hull are balanced by the ship weight and the hydrostatic pressure. This causes bending of both panels and stiffeners. Similarly, in a slamming scenario these panels can be considered to be exposed to a hydrodynamic pressure impulse. Extending the static longitudinal loads as discussed before to a more detailed level, it can be seen again that panels must be loaded either in tension, compression or biaxial tension-compression (equivalently: shear).

This concludes discussion of the major loads on a ship; in the interest of conciseness, not all load cases are discussed here, but it should be brought to attention that a ship can be subject to many more. These include cargo handling loads (such as cargo-deck impact), sloshing of liquids, winching or other tool loads (such as for anchor handling or towing), moored conditions (such as mooring lines and rocking against a quay); and thermal loads - both steady state gradients and cycling. Of course, when military vessels should be considered, this list includes an extra set of more severe loads, as well as damage tolerance requirements. These conditions are not included in the current project. [6]

Finally it is worth noting that (composite) ship components are expected to be subjected to a significant amount of environmental exposure. Though environmental effects such as salt water ingress can significantly affect the strength and stiffness of GFRP structures, one study on composites finds no significantly change the character of AE events in E-glass/polyester samples. However, the frequency of occurrence of AE events is shown to be reduced [7]. While noting the potential importance of lifetime environmental effects in an in-situ AE monitoring system, the current project excludes environmental effects in order to limit its scope. [8–10]

In conclusion, typical ship structure loads can be summarized as shown in Table 2.1. Considering the ship in its simplified beam form, the major loading conditions can be said to be cases of beam bending, while out-of-plane panel loads are represented by loads as wave slamming or impact of tools or cargo. Particularly bending beam loads are potentially cyclic in character, as the typical use case of a ship is its movement through ocean waves. As a consequence, loading on ship panels can be identified as mainly cyclic modes of tension, compression, or biaxial tension-compression (equivalently: in shear). Out-of-plane loads occur as well and may be most critical in current ship design, but they are less predictable in magnitude and frequency. While the class of "unexpected" loads, most of which represent hard- or hydrodynamic impact, these types of loading may not be the most attractive candidates to reproduce in an AE event classification test. This is discussed in more detail in the following section. Nonetheless, also taking into account the discussion from Section 2.1, a higher level conclusion to take away from this section is the desirability of a structural integrity monitoring system to be able to detect transition from "less critical" to "more critical" damage in a structure, where possible, regardless of loading. Any assertion of whether this universal form would be possible based on AE monitoring is considered highly speculative at this time. Instead, the building-block

Table 2.1: Summary of typical (sub)structure loads on a generic beam-like ship.

(sub)structure		typical loading modes
Deck		Longitudinal tension or compression Cargo- or tool impact
Hull		Beam bending and -torsion Out-of-plane bending of panels under hydrostatic or -dynamic pressure
	Hull sides	Shear from beam bending or -torsion Quay-, cargo- or wave impact
	Hull bottom	Longitudinal tension or compression Out-of-plane slamming loads

approach is referred to here, by looking for the smallest 'block' on which empirical data can be collected. This should be a case where complexity of loading-, and consequentially of damage, is low, such that confidence on the internal damage mechanics is high. The dependency of damage mechanisms on loading is the subject of Section 3.1.

Chapter 3

Theory

In this chapter, the theoretical background of damage mechanisms in composite structures, and acoustic emission monitoring are presented. Section 3.1 presents an overview of damage mechanisms in composite structures, substantiated by recent studies on the subject, and organised in the loading modes of tension, compression, bending and impact. This section serves to place the current project in the wider context of current understanding of composite damage mechanisms, while also providing an argument for the layups and low level Tension-Tension (T-T) fatigue loading chosen for the current project.

Thereafter, Section 3.2 presents basic concepts of AE monitoring, as well as a summarized discussion of current and past data-centred approaches to AE monitoring in existing scientific publications. Again, this section serves to place the current project in a wider context. It also provides a succinct reference of some of the characteristics and limitations of AE, which evidently impact not only the results of this project, but the potential of AE-based structural integrity monitoring as a whole.

In previous chapters the project has been introduced and framed in terms of its boundaries and anticipated long-term goals. The current chapter relates the approach of this project to that of other studies on AE monitoring of composites. In the next chapter, the approach of the current project is explained by describing production, testing- and analysis procedures in detail. Thereafter, Chapter 5 presents results of the analysis of the data.

3.1 Damage mechanisms in composite plates

In this section material damage mechanisms occurring in monolithic composite plate structures are reviewed. This section serves as a reference for the larger scope of damage mechanisms in a composite, illustrating the reasoning and boundary conditions behind the choice to concentrate on low load level fatigue for the first step of the building-block approach. Characterizing every single damage mechanism in a composite would constitute too large a scope for the current project, while introducing complexity to the damage process which could make validation difficult. It is therefore preferred to narrow down the current study to a relatively uni-modal damage mechanism. It may be noted this is still not a completely straightforward ambition: damage progression in a composite laminate is a complex process compared to fracture of homogeneous materials. This complexity is mainly manifested through interaction of different failure mechanisms. In terms of the stress field it may be described as follows: progressing local damages (discontinuities) alter the local stress field, which in turn gives rise to yet more new local damages. In contrast to (esp. fatigue) damage in metals, which is more localized, the damage process in composites can be described as an accumulation of diffuse damage, progressing into exponential growth of a dominating crack. [11, 12]

It is recalled that specifically quasi-isotropic layups are relevant to the project; though other layups are also discussed where they can contribute to the discussion. Given the particular importance of transition and interaction of different damage mechanisms in composite structures to the damage process, the emphasis of this section is on providing an insight in material damage mechanisms categorized by loading mode. From this, conclusions may be drawn as to what failure mechanisms are of interest to the project and particularly how they might be isolated. In the previous chapter it has been concluded that panels in a ship hull are potentially subject to the loading modes of tension, compression, bending and shear; and impact. Therefore, this section consists of a review of damage mechanisms in monolithic composite plate-like structures, organized by these loading modes. The findings and conclusions are summarised at the end of this section.

3.1.1 Damage mechanisms in tension

The fracture and damage process itself in a UD composite ply loaded in tension depends on the orientation of the load with respect to the fibres. Typically for a composite, the matrix has a lower failure strength and lower stiffness than the fibres, but a higher failure strain. Additionally, the fibre-matrix interface is weaker in the out-of plane direction than in shear mode. Therefore, failure of a UD fibre composite loaded parallel to the fibres may be expected to be governed by fibre failure, whereas in a transversely loaded composite, it can be expected to be governed by failure of matrix-fibre interface. These expected failure mechanisms are illustrated in Figures 3.1 and 3.2. In longitudinally loaded UD, fibre failure, matrix cracking and shear failure of the fibre-matrix interface govern the damage process. In off-axis orientations, fibre-matrix interface loading transitions from a shear mode (mode II) to an opening mode (mode I); though in particularly in the mixed region, whether this damage process is governed by matrix cracking or fibre-matrix interface failure depends strongly on their respective mechanical properties. [13, 14]

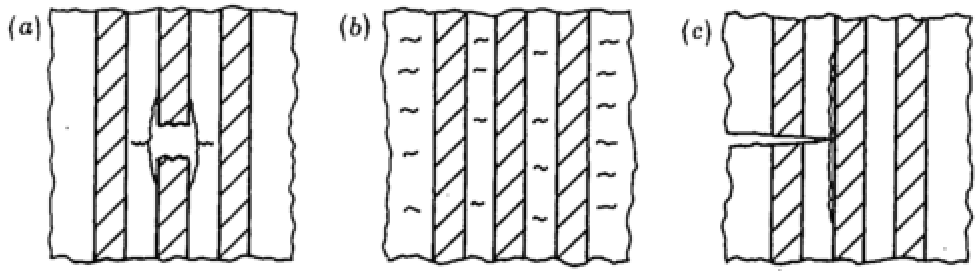

Figure 3.1: Fatigue damage mechanisms of UD composites loaded parallel to the fibre direction (vertically in the image): **(a)** fibre failure; **(b)** matrix cracking; **(c)** fibre-matrix interface shear failure. [13]

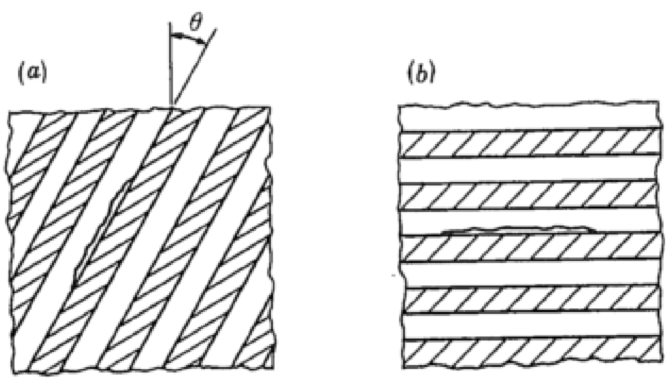

Figure 3.2: Matrix and interface cracking in UD composites loaded off-axis to the fibre direction (vertically in the image): **(a)** mixed opening and sliding mode crack growth under an off-axis angle $0° < \theta < 90°$; **(b)** opening mode (mode I) crack growth under transverse loading ($\theta = 90°$) [13]

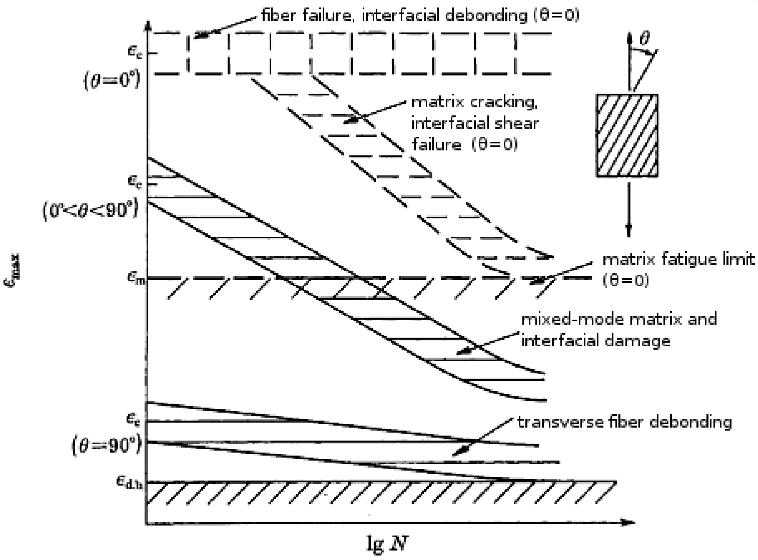

Figure 3.3: Generic Tension-Tension fatigue life diagram for UD composites, indicating regions of different mechanisms governing failure depending on fibre orientation and strain amplitude. In this figure ε_c is the failure strain of the composite, ε_m is the strain in the matrix; and $\varepsilon_{d,h}$ is the minimum strain at which fibre-matrix interface failure occurs. The dashed lines represent the damage mechanisms for the case where $\theta = 0$. The horizontal axis represents the log of the number of applied cycles, whereas ε_{max} on the vertical axis represents the maximum fatigue strain. [13]

The dependency of UD ply damage mechanisms on fibre orientation and nature of loading is investigated in more detail on the basis of Figure 3.3. It is important to note that the mechanisms discussed and depicted are understood as "governing" to the damage process: they are not exclusive and in reality most any type of loading leads to a mix of damage mechanisms. Even though in this way the validity of this figure may be approximate, it depicts in a concise way of some of those governing features of the UD composite damage process. Starting from the bottom of the figure: damage and failure are dominated by matrix and interface failure when the load is transverse to the fibres. As θ decreases, changing the orientation of the fibres towards parallel to the load, the stiff fibres start taking on more of the load, constraining the tensile stress in the matrix, transitioning to fibre-matrix interface shear failure and ultimately fibre failure for high amplitude strains. [13]

Though ultimate failure of a UD plies may be governed by the damage mechanisms depicted in Figures 3.1 and 3.2, damage accumulation in a laminate tends not to be limited to a single mechanism. In a more generalized laminate with more than one fibre direction, damage mechanisms occur in more or less distinct stages. The first stage is (transverse) matrix cracking. The sharpness and nature of a transition to other failure mechanisms depends on layup, nature of the constituents and loading conditions. In T-T fatigue, matrix cracks accumulate until a crack density saturation point is reached [2]. I.e. at a constant stress amplitude, the number of matrix cracks per volume of material (crack density) in transverse plies will not increase beyond the saturation point even if the number of load cycles (N) is still increasing. Notably in static testing of glass-polyester composites an initial "silent period" up to about 2% strain is found, where no cracking occurs. This behaviour may however be specific to this material configuration and can therefore not be fully generalized. Beyond the aforementioned saturation point, matrix cracks develop into other damage mechanisms, depending on the layup (see also

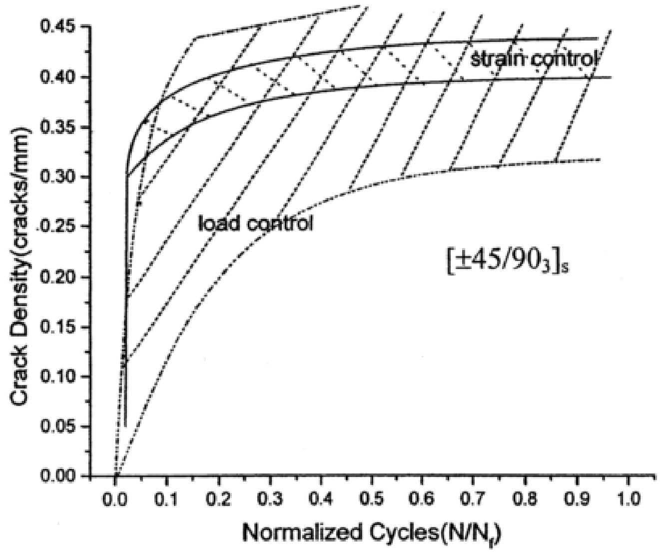

Figure 3.4: Matrix crack density development in E-glass/epoxy $[\pm45/90_3]_s$ laminates under either load- or strain controlled cyclic loading. A laminate including 0° plies behaves similar to the strain controlled case in terms of the spread of results. [2]

once more Figure 3.3). [2, 15–17]

After matrix crack saturation, failure of plies in a laminate at 0° to the loading direction tends to be dominated by fibre fracture, while failure of plies at 90° are governed by matrix or fibre-matrix interface failure. Consequently, 90° (transverse) plies are generally considered the most vulnerable in a laminate. In Carbon Fibre Reinforced Polymer (CFRP) specimens, transverse cracking is found to initiate in 90° layers and along the specimen edges, immediately propagating along the entire specimen width of thick layers. Also, surface preparation can have a significant influence on the location of crack initiation. A $[0_2/90_3]_s$ laminate tends to develop longitudinal cracking, while $[\pm45/90_3]_s$ layups may initiate local delamination even before the matrix crack saturation life is reached. After the first stage of matrix cracking, a second stage in the damage process is marked by a relatively unchanging stiffness and residual strength, which may start at $5-30\%$ of total cycles until failure (N_f), depending on the layup. It may be noted that in a $[\pm45/90_3]_s$ layup in tension, the ±45° plies contribute most to the stiffness in the loading direction and therefore carry most of the applied load. The difference in strain between ±45° and 90° ply orientations suggests an explanation for the propensity towards delamination initiation in these laminates. For the layup that does not include 0° plies and therefore is not strain-restricted, load-controlled fatigue testing produces much more scattered crack density results than a $[0_2/90_3]_s$ laminate would. However, when the test is strain-controlled from the load cell, this scatter is greatly reduced again, as illustrated by Figure 3.4. A straightforward explanation of this result is that the presence of 0° plies in itself limits the strain on the other plies. On the other hand it is also noted that their actual interaction in the material can be more complex by constraining the Poisson's contraction of the 0° plies. Also note the analogy of the behaviours described here with behaviour of off-axis UD specimens as shown in Figure 3.3. Finally, after this second stage, the third stage of the damage process consists of rapid damage growth until failure. [2, 16–19]

Gamstedt and Sjögren [20], as well as Asp et al. [21] have investigated the micro-level damage mechanisms in E-glass/epoxy and E-glass vinyl-ester specifically; the damage mechanism under loading transverse to the fibres has been found to be primarily governed by fibre-matrix interface failure rather than matrix failure. Studies in modelling attribute this to, quite straightforwardly, relative matrix- and interface strengths [14].

Some empirical studies make use of an inherent bias towards particular damage mechanisms that some combinations of layup and loading have. For example, fracture of a UD specimen loaded transversely to the fibres is governed by matrix cracking or matrix-fibre interface failure, depending on the relative mechanical properties of the constituents. Godin and Huguet use the expected relative dominance of matrix cracking and fibre-matrix interface failure under tensile loading of 90° and 45° specimens as a means to characterize AE response [15, 22]. However, it may be noted that this method in itself provides limited reliability in damage mechanism validation because in reality the damage mechanisms in these two different laminates are not distinct per se; e.g. in both laminate types, matrix cracking can certainly be expected to occur. External validation, such as ultrasonic C-scan or section microscopy is still preferred to validate the damage mechanisms actually occurring in the tested material. An interesting recent development in Non-Destructive Testing (NDT) is the use of optical coherence techniques. Optical coherence tomography and -elastography show promise as a safe alternative to radiographic techniques in that they can be used to produce a 3D image of the fracture surface, even while the material is being loaded [23, 24].

Tensile loading traditionally tends to receive most attention in literature compared to other loading modes. As a result, damage mechanisms in this mode are relatively well documented. More importantly for the current project: the dependency of damage mechanisms on loading and fibre orientation, such as presented in the discussion on Figure 3.3, presents the potential to generate relatively isolated damage mechanisms (matrix cracking) by testing in T-T fatigue at low load levels. Moreover, a transition from matrix cracking to other damage mechanisms for this case presents a relatively straightforward and unambiguous answer to the question of how to distinguish critical from less-critical damage, as discussed in the introduction to this section.

3.1.2 Damage mechanisms in compression

While tensile loading might be the most straightforward and common material test, compressive loading introduces its own damage mechanisms which should be considered for completeness. This includes fibre or matrix splitting and fibre micro-buckling, which is governed by fibre bending stiffness and leads to kink bands in the direction parallel to the fibres. For composites transversely loaded to the fibre direction, the matrix is often the weakest link. Consequently damage occurs mainly in the matrix and the fibre-matrix interface. Micrographs of kink band formation, transverse compressive failure and fibre splitting are shown in Figure 3.5.

An experimental study by Oguni et al. [27] examines the failure mechanisms of compressively loaded UD E-glass/vinyl-ester specimens. Under uniaxial compression the main failure mechanism is axial splitting. But when the specimens are subjected to proportional multiaxial compression, the failure mechanism changes to kink band formation without splitting. A different study by Kosmann et al. [28] subjects E-glass/epoxy laminates to uniaxial compression. It is concluded that the presence of voids is a major contributor to local fibre kinking and early matrix cracks and that only stiffness, not strength, is significantly affected by the void content. These studies highlight the documented increase in complexity of composite compression failure compared to tensile failure. These additional and potentially distinct failure mechanisms should be considered in order to apply AE monitoring to in-situ applications. [17]

The aforementioned study by Gamstedt and Sjögren [20] specifically investigates the micro-scale damage mechanisms of Tension-Compression (T-C) fatigue compared to T-T fatigue in

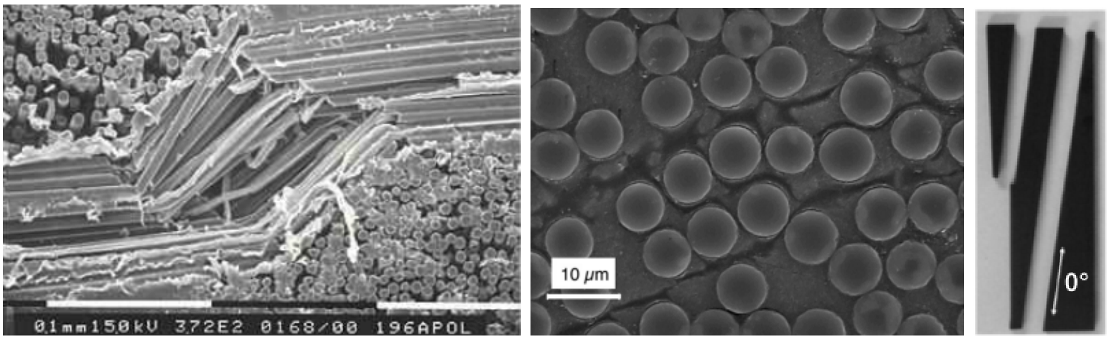

(a) [17] **(b)** [25] **(c)** [26]

Figure 3.5: **(a)** Fibre micro-buckling clustered in kink bands in self-reinforced polypropylene, due to compressive loading (in the horizontal direction in the image).; **(b)** Fibre-matrix interface failure and the formation of shear bands in graphite/epoxy composite due to compressive loading transverse to the fibres (in the horizontal direction in the image); **(c)** Fibre splitting.

E-glass/vinyl-ester. It has been found that T-C fatigue gives rise to more rapid cyclic debond growth around transverse fibres, causing composite laminates containing transverse plies to degrade faster than in a T-T loading case. Transverse cracks have been found to be initiated by interface debonds, consequently accelerating transverse cracking in T-C fatigue loaded specimens.

In conclusion, damage and failure of composites under compression is governed by "extra" failure mechanisms that are quite distinct from those in tension. Some of these might be viewed as another form of matrix cracking, fibre failure or interface failure; for example, transverse loading leading to interface failure and matrix cracking. However, mechanisms such as fibre splitting and kink band formation may be argued not to be representable by these categories. On the most straightforward level, it may be said that damage mechanisms in compression present a more complex interaction, simply because there are more mechanisms to be considered.

3.1.3 Damage mechanisms in bending and out-of-plane shear

In Section 2.3 bending in ship hulls has been mainly attributed to hydrostatic loads on hull panels. Shear is taken here together with bending, because out-of-plane shear is mainly associated with the bending of thick laminates and ship hull laminates of up to 120 [m] in length are definitely in this category. It may be noted that in composite panels, shear is also represented in the in-plane direction. In a laminate, in-plane shear is governed on the micro-/meso level by the inter-ply and intra-ply interfaces, resulting in mainly mode II and mixed-mode crack opening respectively. This type of damage has been discussed in the previous subsection. This can also be seen when recalling the alternate representation of in-plane shear: bi-axial tension-compression. [29]

Bending of a plate causes a linear variation of stress between compression on one free surface and tension on the other. Consequently, the damage mechanisms discussed in sections 3.1.1-3.1.2 potentially also occur here. For example, micro-buckling of longitudinal fibre can occur on the compressive side; this is illustrated by Figure 3.6. Since micro-buckling is also a stability-governed phenomenon, the effective stiffness of the micro-buckled section is also reduced. This shifts the neutral axis of the plate, but also potentially makes bending failure through micro-buckling an unstable process. [17, 30]

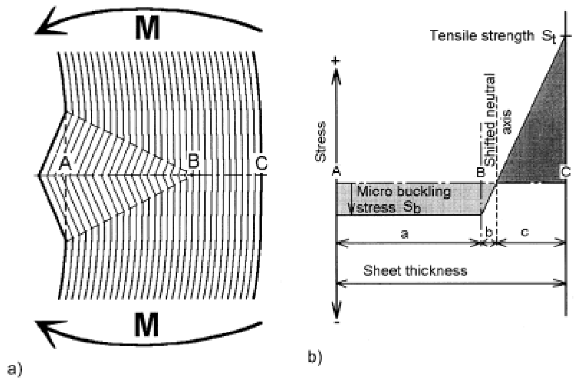

Figure 3.6: 0° fibres in the compressive side are vulnerable to micro buckling under bending moment M at the micro buckling stress S_b even before the tensile strength S_t is reached. **(a)** Schematic depiction of a UD composite bending causing micro buckling of fibres on the compressive side; **(b)** through-the-thickness axial stress profile shows a shifting in the neutral line caused by a reduction in compressive stiffness on the micro-buckled side. [30]

Bending also introduces new damage mechanisms beyond those represented through compressive and tensile loading. Linear variation of the stress through the thickness needs to be balanced by out-of-plane shear. For thin laminates, this stress is usually negligible. Though particularly for thick laminates this stress can significantly load the laminate in inter-ply shear, leading to a delamination damage mechanism. Inter-ply delamination is categorized by crack opening in modes I, II and III. Although in general all three fracture modes can occur in a composite structure; delamination propagation occurs mainly in modes I, II and mixed (I+II) [31]. Fracture under mode I loading is the most critical mode in a real structure because the energy required to initiate delamination is lower in this case than in mode II [32]. Therefore most publications studying delamination focus on mode I and II, commonly represented by a double-cantilever beam or three-point bending tests respectively - these tests are depicted in Figure 3.7. Particularly with three-point bending tests, it is noted that fracture can be a complex process because of the interaction of out-of-plane shear and bending stresses. On the other hand, this feature is taken advantage of in testing by varying the ratio of support span (L) over thickness (t) of the three-point bending specimen. For low L/t (about < 5 for GFRP), fracture is dominated by inter-laminar shear failure (mode I+II); whereas for high L/t (around > 10), fracture is dominated by intra-laminar tensile and compressive failure (mode II). Although fracture in these specimen geometries may be dominated by said failure mechanisms, it is not conclusive that these are the only occurring mechanisms. [17, 29]

Examining these bending tests in a little more detail reveals three consecutive damage stages for quasi-static three point bending tests of ±45° CFRP laminates [34]: In the first stage, which starts immediately from 0% strain, matrix cracks initiate within the plies. In this stage, the specimen-stress strain relation behaves non-linearly. Subsequently in the second stage, transverse matrix cracks propagate and multiply rapidly. In locations where shear stresses are high, these transverse cracks develop into the inter-ply interface (delaminations). An important observation in the second stage is that accumulation of damage is associated with coupling between damage mechanisms in this stage only. Finally, stage three is associated with the opening of delaminations and slipping between plies.

Figure 3.8 shows the increase in electrical resistance of a CFRP laminate subject to flexural loading. This illustrates the damage stages as discussed above (in the figure designated as “zones”),

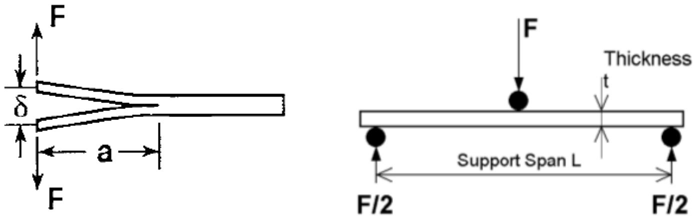

(a) [33] (b) [17]

Figure 3.7: Schematic representation of **(a)** Double cantilever beam test. Inter-laminar fracture is forced in mode I, resulting in a crack of dimension a and a displacement δ. The beam is constrained on the right side. **(b)** Three point bending test. Fracture mode depends on the ratio of span over thickness (L/t).

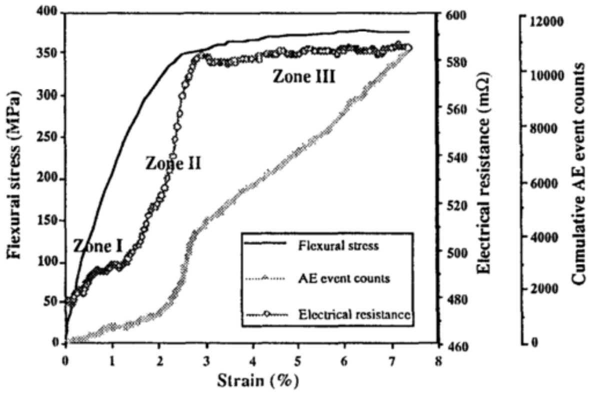

Figure 3.8: Results of AE and electrical resistance measurements on a ±45° CFRP laminate subjected to a flexural load. The stress-strain curve is shown, as well as the dependent electrical resistance and cumulative AE event counts. [34]

as well as a typical progression of damage events. The three stages of damage progression can be clearly distinguished from the figure by their different degrees of damage activity, which is measured through both AE activity and electrical resistance.

Loading rate dependency of mode I delamination in glass-epoxy laminates has been investigated by Benmedakhene et al. [32]: due to matrix brittleness increasing with loading velocity, fracture is found to move from mainly matrix failure at low velocity to fibre-matrix and inter-ply interfaces for higher velocities. However, this study makes note of other works sometimes finding contradictory results with regard to the velocity-dependency of the mode I critical energy release rate. These differences are attributed to variations in subject material or experimental conditions.

It may be concluded that in the context of distinguishing damage mechanisms, bending modes are of interest mainly for distinguishing between mode I and II delamination. But because delamination itself is a relatively critical damage mechanism and bending of thick laminates is associated with delamination, bending tests may be of interest in generating relatively pure delamination modes in future studies. Since bending is a combination of tension and compression, the same conclusion from Section 3.1.2 applies: the extra complexity introduced by this loading mode does not constitute a suitable first block in the building-block approach.

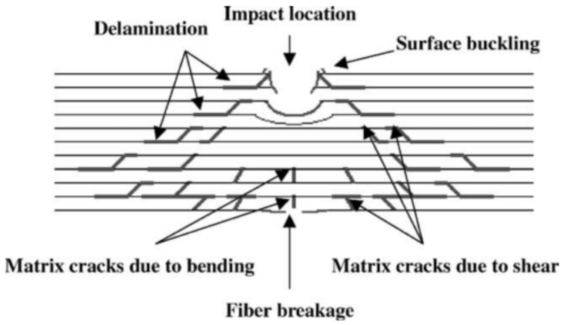

Figure 3.9: Schematic representation of a typical composite laminate impact damage mode. [37]

3.1.4 Damage mechanisms in impact

As discussed in Section 2.3, impact is relevant to naval structures in the form of "unexpected loads", as well as hull-wave collisions. For naval applications, a distinction is made between hard and soft impactors, respectively representing collisions with rigid bodies or wave loading. In addition, the field of impact studies make a distinction between high-velocity and low-velocity impact. High velocity impact is often (but not always; the distinction is debated) characterized as being dominated by stress wave propagation; the structure does not have enough time to respond, leading to localized energy absorption and damage. By contrast, a low velocity event can be treated as quasi-static and is governed by the elastic response of the entire structure. In composites, low-velocity impact is also associated with mainly matrix damage, whereas for high velocities, fracture progressively occurs more in the form of fibre-matrix failure (both intra and inter-ply). Impact studies for marine applications reveal that even for soft impactors, damage tends to be localized. This suggests that impacts in marine applications tend to be in the high-velocity camp. However, as has been noted before, soft impacts for naval applications are not very well understood and even the availability of relevant empirical methodology is limited. [5, 35, 36]

Impact on a composite panel causes damage in four modes, identified as: matrix-, delamination-, fibre- and penetration mode. Generally inter-laminar failure will be the first mechanism, as the weakest link for impact in composite panels is often their low in-plane strength. One solution that is used to combat this is through the use of "z-pins" in composite panels. These z-pins are also used in stiffener attachment to postpone skin-stiffener separation. For low velocity impacts, damage is limited to matrix cracking. Fibre-matrix interface failures and fibre fractures require respectively higher impact energy. A typical through-the-thickness damage profile after impact on a composite is depicted in Figure 3.9. What is specifically typical is that the damaged area conically extends from the impact location to the opposite face of the laminate. There is a (nearly) undamaged zone (centre in the figure) and farther down, the bending of the impact has been carried by membrane stresses, causing fibre fracture at the lower surface. [1, 5, 35, 36].

In light of the project being concerned with monitoring of damage as it occurs, it may be noted that particularly impact tests are a poor initial candidate for an AE classification study because of its characteristic instantaneous damage evolution, while the impact itself is also expected to add acoustic energy to the measurements. Alternatively, loading after impact may be considered. In damage tolerance of a ship, this is may be of significance as well. Though the process of damage progression in loading after impact is considered to more complex even than impact itself, delamination growth is characteristic in cyclic compression loading after impact. Delamination growth may be envisioned to fit a project focus of tracking criticality of damage;

however this is definitely not the only damage mechanism at work in this load case. As a consequence, compression after impact may be seen as of limited use as a classification test by itself. [38, 39]

3.1.5 Summary and conclusions

In most studies reviewed in this section, as well as for those in Section 3.2.5 on the subject of AE source classification, three damage mechanisms tend to be significantly more represented in literature than any others. These mechanisms are matrix cracking, fibre failure and fibre-matrix interface failure. Many studies distinguish other mechanisms, such as: matrix crack initiation and propagation, longitudinal matrix splitting, fibre pull-out and mode I or II delamination. This skewed attention may be interpreted in that these three are most fundamental to the damage process, while the latter list of mechanisms can be argued to be a combination of the "fundamental" mechanisms. Alternatively, the predominance of these three modes may not be more than that observation in itself; i.e. these three are most well-described in literature, whereas others are investigated in less detail as of this time. Regardless, it may be concluded that consensus on which categorization of mechanisms is most meaningful for the purposes of structural integrity monitoring is yet to be reached.

Table 3.1: Summary of loading modes on composite panel-structure and associated damage mechanisms as discussed in this section. Note any categorization like this is a simplification. Some parameters that strongly influence the damage mechanisms occurring that are not explicitly mentioned in this table include: layup, production parameters and -defects, specimen/structure boundary conditions and composite constituents.

Loading	**Load orientation**	**Governing damage mechanisms**
Tension	parallel to fibres	Fibre failure, matrix cracking, fibre-matrix interface shear failure
	transverse to fibres	Mode I transverse fibre-matrix interface failure
Compression	parallel to fibres	Dominated by fibres: fibre or matrix splitting, fibre micro-buckling into kink bands
	transverse to fibres	Dominated by matrix failure, fibre-matrix interface shear failure
In-plane shear		Inter-ply delamination, fibre-matrix interface shear failure
Bending & out-of-plane shear		Tension & compression failure mechanisms on resp. tension- and compression side of bending moment. Inter-ply in-plane shear failure (for thick laminates).
Impact		Matrix cracking at low energy; more severe damage s.a. delamination and penetration at higher energies.

In conclusion, the damage mechanisms as they occur in different loading modes can be summarized by Table 3.1. Relevant conclusions to the project are to select loading modes and damage mechanisms of interest, both in the context of the project framework, as well as in the interest of generating damage mechanisms suitable for AE detection. As stated before, although impact has definite relevance to ship loading, the instantaneous nature of damage occurring in this case is considered an unattractive feature if damage should be characterized by AE data. Contrarily, a specimen loaded in fatigue might present a more gradual introduction of damage, while loading a specimen only in tension limits the number of damage mechanisms in the material. This presents the advantage of isolation of damage mechanisms, as much as possible. Even though uniaxial fatigue loaded specimens are not necessarily representative for critical ship loads (see Section 2.3), they are nonetheless chosen as AE characterization specimens in an effort to limit complexity of damage accumulation. In addition, the relatively gradual sequential nature of damage can be used to isolate predominantly matrix cracking different UD laminates, leaving an opportunity to potentially detect transitions to other damage mechanisms. For the purposes of test specimen layup selection, it may be important to realize here that there is a trade-off at work: it may be more representative for the recorded AE response to use quasi-isotropic laminates in AE characterization tests. However, in these layups a more complex damage accumulation may also be expected, complicating characterization. Therefore selected simpler laminates are preferred, built up out of only few different fibre directions.

3.2 Acoustic Emission monitoring

The purpose of this section is to provide a basic understanding of the relevant mechanics and methodology for monitoring of composite structures. First, AE is introduced as a technique for NDT and monitoring. Though the interest of the current project in AE is for the purpose of structural integrity monitoring, AE is introduced here from the wider perspective of NDT. AE fundamental mechanics, as well as a typical monitoring set-up and equipment are presented. Finally, current state-of-the-art AE source classification approaches for AE monitoring of composites are reviewd.

NDT is described by the American Society for Testing and Materials (ASTM) as: "the development and application of technical methods to examine materials or components in ways that do not impair future usefulness and serviceability in order to detect, locate, measure and evaluate flaws; to assess integrity, properties and composition; and to measure geometrical characteristics." [40] It may be noted that this definition includes other applications besides damage detection. AE techniques are no exception to this, having been used in a variety of applications like process monitoring, leak-testing, weld quality monitoring and condition monitoring. [41] The use of AE in structural integrity monitoring then represents the use of in-situ continuous AE monitoring of a structure and processing of this data to make statements about its structural integrity; particularly to support the structure's maintenance.

A distinct feature that sets AE apart from many other NDT or Non-Destructive Evaluation (NDE) methods is that it is passive. I.e. rather than emitting energy which is reflected back to a sensor, like for example in ultrasonic or radiographic techniques, AE monitoring depends on the release of elastic energy from events within the material. Elastic waves then propagate through the material to be registered by a transducer mounted to the structure surface or embedded in the material. A typical AE event as captured by a transducer is shown in Figure 3.10. In general, an energy release can represent the initiation or progression of damage, but can also be due to, for example, friction, turbulent flow or impact. An important consequence of the dependency of AE techniques on release of energy is that AE will only detect changing (transient) activity in a subject. This has an important consequence for the application of AE to structural monitoring: it is inherently impossible to know the initial state of the material (i.e. damages, inclusions, etc.) merely by AE monitoring. [41, 42]

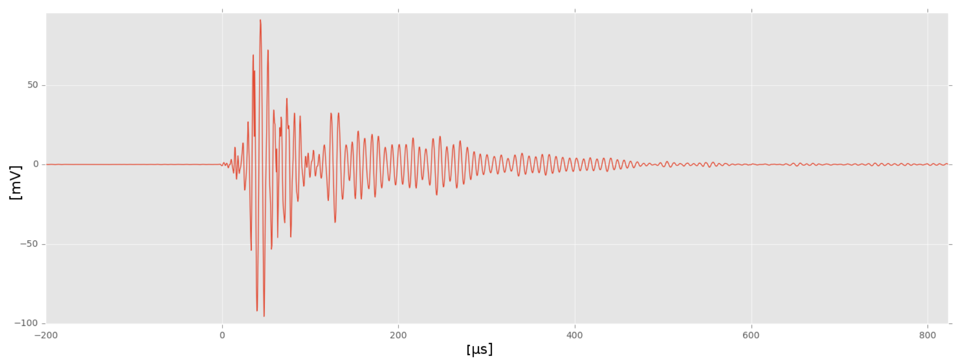

Figure 3.10: Typical response of a transducer to an AE event in a composite.

In the current project, AE monitoring is used as a means to monitor damage progression. I.e. the interest is in recording the acoustic energy release from crack initiation and progression in the material. Of particular interest is an approach that correlates specific types of damage mechanisms with the characteristics of an AE event. Many past studies have shown that the basic ingredients to enable this type of strategy exist: AE activity positively correlates with the level of damage in a material (including its initial state) and waveforms do carry information about their source events. However, the translation of an AE event as captured by a sensor into a description of a source type remains a central challenge to AE monitoring. [18, 42–44]

The sections below are structured as follows: first, some of the fundamental physics of waves and their interaction with material is presented. Then, Section 3.2.2 presents a discussion of the Kaiser effect, which may be considered one of the first and fundamental concepts of AE monitoring. Subsequently, Section 3.2.3 gives an overview of typical characteristics of AE sensors and an AE monitoring set-up. In Section 3.2.4, an overview of some of the applications of AE monitoring for structural integrity monitoring in ships is given. Finally, Section 3.2.5 presents a concise review of current state-of-the-art AE source classification approaches for AE monitoring of composites.

3.2.1 Mechanics of acoustic waves in a plate

In this section, some of the fundamental mechanics of wave propagation generally relevant to AE are laid out. For the current project, the focus is on plate structures, in which waves predominantly propagate in 2D. First, some of the formal physical fundamentals are given. Thereafter, the internal interaction of waves with the material and its boundaries is discussed, which includes the effects of reflection, attenuation, dispersion, scattering and diffraction. In the next section, the Kaiser effect is discussed.

Wave velocity and wave governing mechanics

In a solid, acoustic wave propagation is organized into four principal modes. These are longitudinal waves, shear waves, surface waves and plate waves in thin materials. Since the current project is mainly concerned with plate-like geometries, the only waves considered here will be plane or plate waves. Plane waves are a simplified propagation case which is valid for plate thicknesses of only a few wavelengths. In NDT and AE in particular, Lamb waves are used most extensively. Lamb waves are in fact a superposition of longitudinal- and shear waves. They attract attention for ultrasonic-based studies because Lamb waves can propagate over long distances with low attenuation. As a consequence, Lamb waves can be detected over long distances, even in relatively dispersive and attenuating mediums such as (the matrix component of) composites. [42, 45]

A rough check can provide an indication of validity of plane waves for the case of the project: the frequency domain (f) of acoustic events is of $O(10^5)$ s^{-1} [46], while the velocity (v) is of $O(10^3)$ [m/s] [22]. Therefore, per $\lambda = v/f$, the wavelength (λ) must be of $O(10^-2)$ [m]. This puts the lower end of the laminate thickness ranged used by DSNS (4 to 120 [mm]) still in the right range for "one or a few wavelengths". In the higher end of this thickness range, the validity of plane waves may be limited.

Lamb waves consist of a number of vibrational modes, but two modes are most commonly used. These are the symmetric (extensional) mode and the asymmetric (flexural) mode; their fundamental modes, respectively S_0 and A_0, as depicted in Figure 3.11, are of particular practical

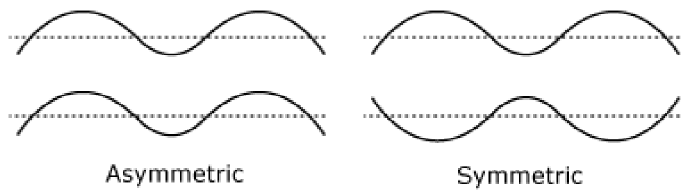

Figure 3.11: Schematic representation of flexural (asymmetric, A_0) and extensional (symmetric, S_0) Lamb wave modes in a plate. [42]

use because they are triggered regardless of the frequency of the plate wave. Both modes also travel at different velocities, analytically described by Eq. 3.1 and 3.2 respectively. [41, 47]

$$c_1 = \sqrt{\frac{\lambda + 2\mu}{\rho}} \tag{3.1}$$

$$c_2 = \sqrt{\frac{\mu}{\rho}} \tag{3.2}$$

Where c_1 is the velocity of the symmetric wave and c_2 is the velocity of the asymmetric wave. These relations are derived from the equations of motion for displacements in a linear elastic homogeneous isotropic body [41], where ρ is the material density. λ, μ are Lamé's constants as defined by Eq. 3.3 and 3.4:

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)} \tag{3.3}$$

$$\mu = \frac{E}{2(1+\nu)} \tag{3.4}$$

Here E is the material's Young's modulus and ν its Poisson's ratio. The symmetric and asymmetric Lamb waves can be described by their set of governing equations, which is given below in Eq. 3.5 for the symmetric mode and Eq. 3.6 for the asymmetric mode [45]. In this set of equations, $p^2 = f^2/c_1^2 - k^2$, $q^2 = f^2/c_2^2 - k^2$ with k the wave number and h the half plate thickness.

$$\frac{\tan qh}{\tan ph} = \frac{-4k^2pq}{(q^2-k^2)^2} \tag{3.5}$$

$$\frac{\tan qh}{\tan ph} = \frac{-(q^2-k^2)^2}{4k^2pq} \tag{3.6}$$

The wave speed dependency on stiffness is also manifested on a local level. Fibres tend to be much stiffer than the matrix, adding to the complexity of wave propagation in a composite. In a single ply, fibres propagate much faster in the fibre direction than transversely to the fibres, while transmission between plies also occurs. If the laminate is quasi-isotropic, propagation might be approximated by isotropic behaviour. But in any other case, it is important to account for the directional dependency of the propagation velocity. Analytical solutions for Lamb wave propagation exist, even for anisotropic materials [48, 49]. Because propagation speed is a function of material stiffness, it is important to note that propagation behaviour in a composite depends strongly on the layup. As an indicative example: in an E-glass/epoxy

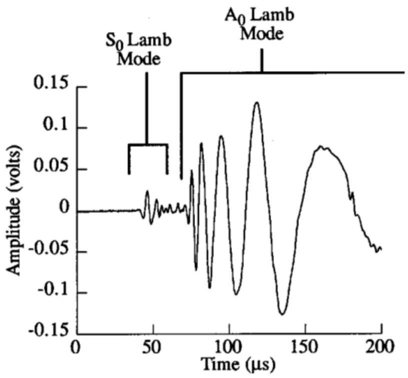

Figure 3.12: Typical AE waveform generated by a pencil lead break source in a graphite/epoxy composite panel. The arrivals of extensional (S_0) and flexural (A_0) wave modes are indicated. [54]

UD laminate the wave propagation speed may vary between 2750 [m/s] and 3400 [m/s] in the longitudinal direction of a 90° and a ±35° layup respectively [22].

The relative dominance of wave modes can be used to distil information from an AE signal. This method of analysis is called modal AE and has been introduced to AE monitoring by Gorman [47, 50] and Pollock [51]. AE analysis based on the relative dominance of the extensional and flexural wave has been shown to be effective at discriminating source mechanisms [52,53]. Figure 3.12 illustrates this principle with a pencil lead break event as captured in a graphite/epoxy plate: it can be seen from this figure that the extensional mode indeed arrives before the flexural one, while its amplitude is lower and its frequency is higher. The wave modes are typically separated based on their relative arrival times or by splitting the signal into two frequency bands. With knowledge about the relative dominance of flexural and extensional wave components for different damage mechanisms, this can be used to aid in signal classification. [19, 50, 52, 53]

However, some limitations should be brought to attention. Firstly, because dispersion and attenuation properties of both wave modes differ, care must be taken when applying modal AE over longer propagation distances. This also applies when dealing with structures that include complexities such as joints or stiffeners. Second, matrix cracks can still generate flexural mode waves when they are initiated at the plate edges; particularly thick plates in bending are sensitive to this. It can also be noted that, in graphite/epoxy tensile test samples, matrix cracking is shown to initiate along plate surfaces anyway. Though using source location information can help alleviate these sources of misclassification. [52, 53]

Reflection, attenuation, dispersion, scattering and diffraction

In general for waves in a medium, attenuation, reflection at interfaces, and transmission are governing mechanics of propagation. These can be especially important for AE monitoring, as these mechanics influence the characteristics of the waveform measured by an AE transducer. That is, a combination of mode conversion during reflection and frequency dispersion may cause an initially coherent AE waveform to be “stretched out”. Meanwhile attenuation may change the relative amplitude and character of the waveform components, such that the waveform as recorded by the transducer is significantly different than the emitted one. Reflection, attenuation, dispersion and scattering & diffraction are briefly discussed below.

Reflection of waves at a material interface is governed by Snell’s law. Even though the behaviour of the reflection itself is identical for both wave modes, the symmetric and asymmetric components of the wave will reflect differently because of their respective difference in propagation speed (by Eq. 3.1 and 3.2). Furthermore, waves may exhibit mode conversion in the reflection process. This is illustrated in Figure 3.13 and represented by Snell's law in Eq. 3.7 for an incident asymmetric wave; in this equation, a subscript r denotes a reflected wave, i an incident wave; and a and s denote asymmetric and symmetric components respectively. Figure 3.14 illustrates the dependency of mode conversion behaviour for an incident asymmetric wave in an isotropic material with Poisson's ratio 0.3. In the example of this figure, a critical reflection angle can be observed at an incident angle of about 30°, above which a asymmetric wave is reflected without attenuation. This critical angle is material-specific however, so this figure should not be generalized. Moreover, from Figure 3.14b it is apparent that for practically for all incidence angles there will be some reflected asymmetric wave component. The case for an incident symmetric wave behaves similarly, though the curve has a different shape. [41]

$$\frac{\sin\theta_{i,a}}{c_2} = \frac{\sin\theta_{r,s}}{c_1} = \frac{\sin\theta_{r,a}}{c_2} \tag{3.7}$$

Energy loss of elastic waves is categorized in a set of mechanics. Some literature sources make slightly different distinctions [40, 52], but here four mechanisms will be distinguished [41]: geometric spreading, dispersion, scattering and diffraction; and attenuation. Attenuation in general can be globally summarized by the non-conservative effects of wave propagation. This energy is lost as heat and is said to be absorbed by the material.

In geometric attenuation, a disturbance in an infinite medium propagates outward, distributing its energy such that it reduces proportionally the propagation distance squared. However, for surface or plate waves, the amplitude only decreases by $\frac{1}{r}$ if the wave is sufficiently far from the source. Waves propagating in one-dimensional media are unaffected by geometric attenuation. The latter statement may seem trivial and merely theoretical, but in certain simplified conditions it may still apply, such as in a vibrating string or a wave from a line source in the plane of a plate. [41]

The velocity of acoustic waves is frequency-dependent, which causes dispersion. An initially short-wavelength broadband acoustic burst is "stretched" as it propagates further, due to the difference in velocity between frequency components. Moreover, this behaviour is also different for flexural and extensional modes. Figure 3.15 illustrates this for a theoretical example. It can be seen in this figure that while the fundamental flexural mode (A_0) is mostly sensitive to dispersion in lower frequencies and dispersion-free in higher ones, the fundamental extensional mode (S_0) has a pronounced dependency of velocity on frequency - though only at high

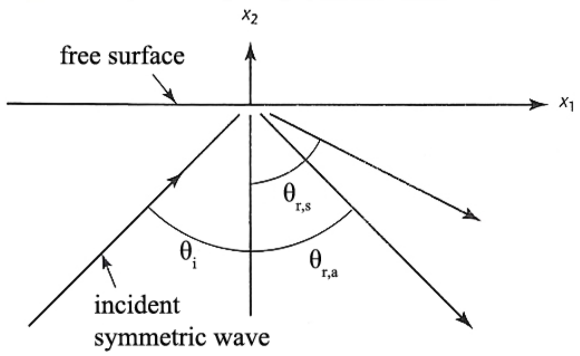

Figure 3.13: Reflection of an asymmetric wave at a free free surface of a plate into a symmetric and an asymmetric component. Adapted from [41].

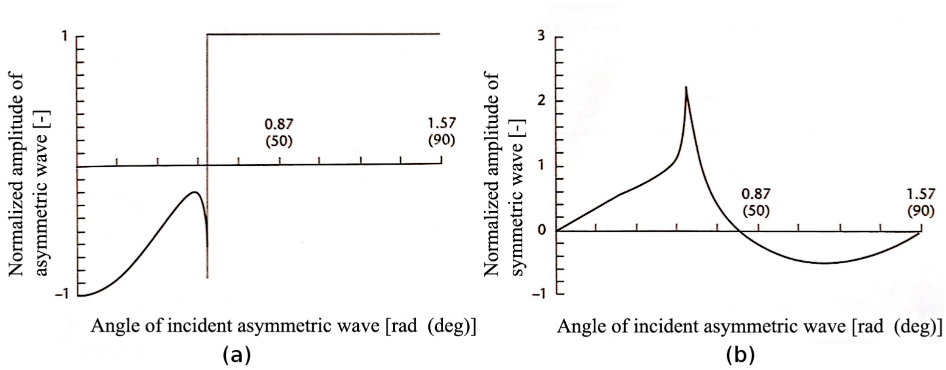

Figure 3.14: Reflection coefficients for an incident flexural wave in a material with Poisson's ratio 0.3 as a function of incidence angle; **(a)** Relative amplitude of the reflected flexural wave; **(b)** Relative amplitude of the reflected extensional wave. [41]

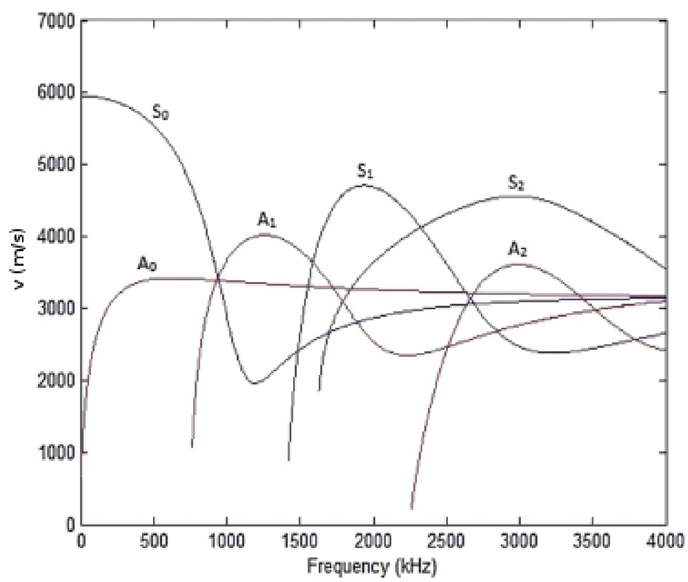

Figure 3.15: Theoretical group velocity dispersion curves for the first three symmetric and asymmetric Lamb wave modes in an isotropic material with a symmetric wave velocity of 7200 [m/s], an asymmetric wave velocity of 3368 [m/s] and a thickness of 2.24 [mm]. [55]

frequencies. Note again that though higher-order modes may be of interest, most AE studies that consider modal analyses are concerned mainly with the fundamental modes, as these are triggered for any plate wave frequency. Velocity dispersion is usually represented in dispersion curves like these, which can be determined either theoretically [56] or experimentally. [50, 52]

Scattering and diffraction are mechanisms associated with boundaries and discontinuities. These may include cavities, cracks and inclusions. But in composites, this also includes the many matrix-fibre interfaces. An acoustic wave encountering any discontinuities will be (partially) reflected. A plane wave encountering a void or inclusion will actually dissipate energy in all directions; this is called scattering. Waves encountering a sharp discontinuity, such as a fractured fibre edge, may experience partial diffraction of its energy around the edge. Both scattering and diffraction cost energy, thereby attenuating the acoustic wave. As fibre composites are heterogeneous materials level, this type of attenuation can be expected to significantly impact wave propagation - and indeed, it has been found that AE event amplitude is strongly influenced by wave propagation in composites [44]. [41]

3.2.2 The Kaiser effect and Felicity ratio

One of the early fundamental mechanics found in AE monitoring is the Kaiser effect. Josef Kaiser found in the 1950s that when loading any material type [57]: "low level [acoustic] emissions begin even at the lowest stress levels (< 1 [MPa]). They are detectable all the way through to the failure load, but only if the material has experienced no previous loading." In other words: in general a material will emit acoustic signals when it is loaded at a level to which it has not been loaded before. The inverse statement is also true: once a load is applied and the acoustic emissions from the resulting stress has ceased, the material will not emit new acoustic events unless the load is increased - even when the load is completely removed and reapplied. Consequently, the Kaiser effect is a potentially very powerful observation. However, there are additional mechanics involved, leading to, among others, some exceptions for composite materials.

The Kaiser principle refers to the stress state of a material, while the stress state is also a function of the condition of the material. It is helpful to note that this is potentially a local vector property in the material. An early application of the Kaiser effect is the Dunegan corollary and it illustrates this local stress aspect. The Dunagan corollary states that if acoustic emissions are observed before a previous maximum load is reached, new damage must have occurred [42]. This

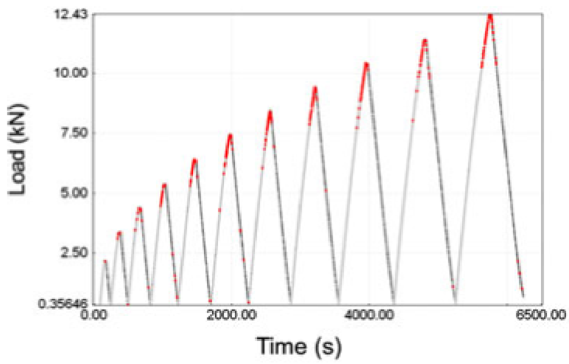

Figure 3.16: Acoustic bursts (red markings) during the depicted loading-unloading-reloading spectrum of E-glass/epoxy GFRP specimens reveals the Kaiser effect and the Felicity effect. [58]

is used for example in proofing of pressure vessels: in use they are subjected to a working load, while periodically the pressure vessels are tested by loading them to a specified higher proof load. The very first proof load will of course generate acoustic events, while any subsequent proof load will not. But when discontinuity growth has occurred during service (working loads), this will change the stress path and thereby the stress level when the vessel is subjected to the next proof load. In this application, the occurrence of AE events during the proof load is a measure of damage experienced in service.

For many materials and conditions the Kaiser effect is not "perfect": a measure of the Kaiser effect breaking down is the Felicity effect. Depending on the material, AE events can occur at levels below the last applied maximum load or even not occur at all. Internal friction between cracked surfaces can contribute to this; materials are especially prone to the Felicity effect when they are close to failure or in poor condition. Moreover, the Kaiser effect notably also fails when time-dependent mechanisms such as rheological flow or matrix relaxation are involved. The Felicity ratio is used to quantify the breakdown of the Kaiser effect, where Felicity ratio of 1 signifies the special case of the perfect Kaiser effect. Or in mathematical terms, it is the ratio of the stress at which AE activity starts over the previous maximum stress. [41, 59]

A manifestation of the Kaiser effect and the Felicity effect in a GFRP specimen can be clearly observed from Figure 3.16. In this study, an E-glass/epoxy test specimen is loaded in tension in a loading-unloading-reloading spectrum as can be observed from the figure. The red markings are the AE event picked up by a transducer mounted to the specimen. AE events can be seen to mainly occur when the specimen is loaded to a higher level than it has been loaded before. This is the Kaiser effect. But it is not "perfect": some of the AE events occur before the new load has exceeded the previous maximum load. This means the Felicity ratio is less than unity. Some AE events can also be seen to occur during unloading, suggesting there might be acoustic activity from internal friction; though it could also mean simply that damage occurs below the last maximum load. A similar AE response can be observed when the same type of loading spectrum is applied in compression. [58]

The Kaiser effect seems an attractive candidate to take into account in structural integrity monitoring because of its simplicity: it is merely based on counting AE events, requires no very detailed signal analysis (discounting noise filtering). On the other hand, (saline) water exposure and natural ageing has been known to affect AE activity in composites [7, 10]. Therefore it may not be straightforward to depend on this type of approach in structural integrity monitoring. The Felicity effect itself may on the other hand be a valuable metric for composite ship structures in particular. Because these structures tend to be of poor - and therefore varying - quality, the relation between Felicity and material condition could contribute an interesting metric to take into account.

3.2.3 Acoustic Emission set-up and equipment

In this section, a typical set-up for AE monitoring in experimental conditions is presented. Though a potential ultimate goal of in-situ monitoring might have need for a slightly different set-up, currently only laboratory conditions will be considered. Additionally, it may be noted that only a few basic elements of AE testing are discussed here. Many resources exist that provide a comprehensive and more detailed guide to AE testing than can be presented in this document [41, 60–62]. A detailed set-up for the testing in this project is given in Section 4.2. This section first describes sensor hardware before going into typical AE hardware set-up.

Sensors

Sensors used to capture AE events are piezoelectric transducers, converting dynamic motion into a current by the piezoelectric effect, caused by deformation of a piezoelectric material. Other types of transducers have also been used, such as laser interferometry and fibre-optic sensors based on microbending of the fibre. However, piezoelectric transducers are currently most commonly used by far. [41]

A schematic representation of a typical AE transducer is shown in Figure 3.17. The active element is usually a piezoelectric crystal made of lead zirconate titanate (PZT) and is marked in the figure. It is sandwiched between two electrodes. A coupling agent helps transmission of acoustic waves from the test subject to the active element. A wear plate protects the active elements, and the damping material is added to reduce acoustic reflections back to the active element. The wear plate can also serve to load the active element such that it is less resonant (more broadband): piezoelectric crystals normally have a resonant frequency to which the electronic response is much more sensitive than other frequencies. In order to achieve a broadband response, a crystal is heavily damped, while designing its resonant frequency such that it lies outside the operating frequency range for the sensor. In AE monitoring, resonant sensors are typically only used when the frequency content of recorded events is not of interest. Broadband sensors are most common and they are typically designed to capture motion in the 30 [kHz] − 1 [MHz] range, though for materials with high attenuation such as fibre reinforced plastics, sometimes lower frequencies are used because higher frequencies tend to suffer more dissipation than lower ones. On the other hand, higher frequencies tend to contain less background noise. [42, 59, 62]

An acoustic wave propagating through a material can be decomposed into two vibrations in the plane of the surface and one normal to it. In principle, AE sensors can be designed to respond to each of these modes; yet virtually all commercially available AE sensors are designed to be sensitive to the normal component only [62]. An AE sensor can rely on registering only the normal component because longitudinal, Rayleigh and shear waves typically all have a component of excitation normal to the surface [41]. Even so, it is important to realize that when processing AE monitoring data, in fact only plate-normal components of acoustic waves tend to be measured.

Testing set-up

A typical AE set-up consists of a set of piezoelectric transducers, preamplifiers, main amplifier(s), measurement circuitry feeding into data buffers and fed into a PC to store and present data. A representation of the AE measurement chain is depicted in Figure 3.18. AE transducers are discussed in the previous subsection. The preamplifier must be close to the transducer to

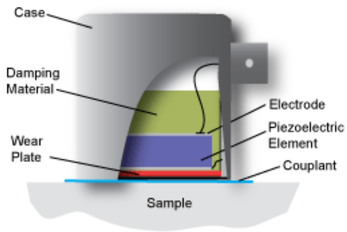

Figure 3.17: Diagram of the design and mounting of a typical AE transducer. [42]

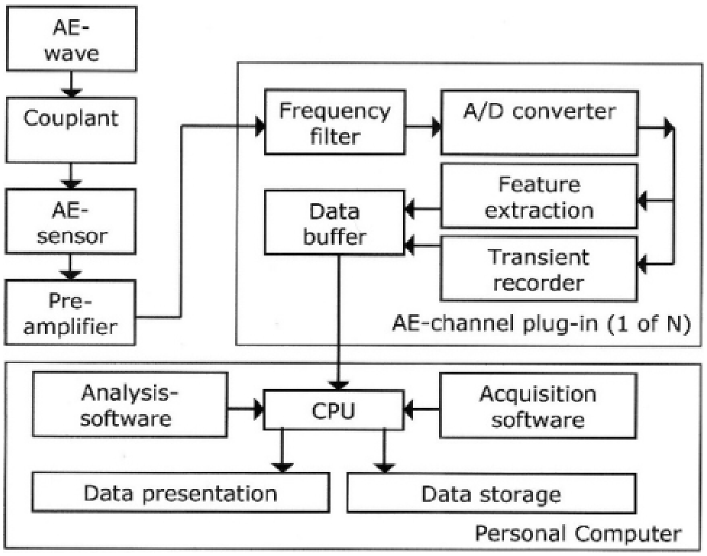

Figure 3.18: Schematic measurement chain for AE testing. Note that an AE monitoring set-up can contain multiple AE sensors; their transducer signals are processed in parallel (N channels). [63]

minimize the loss in signal-to-noise ratio. Many current AE transducers have a preamplifier integrated in the transducer casing. The preamplifier also typically provides some filtering as well as gain, which is commonly 40 [dB]. [59, 63]

AE signals from materials are measured in the 10 [kHz] -1.5 [MHz] range. Although generally 90% of AE activity occurs in the $10 - 550$ [kHz] [64]. Bandpass filtering is usually applied as desired for the application, i.e. to filter known sources of high- and low-frequency noise. Noise can also be filtered by location (e.g. to filter specimen grip fretting noise) or using signal classification techniques, which are discussed in Section 3.2.5. Finally, note that most post-processing of signal data can still be done using a software; though care must be taken when relying on this because of discretisation in the A-D conversion step. Data processing is discussed more elaborately in relation to AE event classification in Section 3.2.5, and concretely for the current project in Section 4.3. [41, 42, 60]

For a couple of different purposes, a pencil lead break test (also known as a Hsu-Nielsen source) is often performed prior to AE testing. This is a standardized procedure covered by ASTME976 [61], and involves breaking the end of a $0.3 - 0.5$ [mm] 2H pencil lead on the surface of the subject material. The mechanical pencil's head is supported by a Teflon cone which keeps the pencil at a 30° angle with the specimen surface. This procedure creates reproducible acoustic waves in the material, which can be used to check if an AE measurement set-up is working as desired, or to measure signal attenuation and wave propagation speed.

For AE event classification tests in the current project, a straightforward conventional AE set-up is chosen. In terms of hardware, the *Vallen AMSY-6* AE system available to the *Structural Integrity & Composites* group at TU Delft is popularly used in AE monitoring studies (see many of the AE studies in the references to this document) and compatible with the project's needs. For sensors, broadband (VS900-M) sensors are used, as the frequency content of damage events is not ruled out to be a feature contributing to AE event classification. For a full description of the test set-up for the experimental work in this project, refer to Section 4.2. In the next subsection, current documented work on structural integrity by AE monitoring on ships is reviewed.

3.2.4 Acoustic Emission monitoring in ships

Many current studies on AE as a means for structural integrity monitoring are focused on generic composite structures, or are aimed at aerospace applications. But AE monitoring also receives some interest from the marine and off-shore industry. Because composite materials have not been very common in ship structures, AE monitoring of composite ship structures has not been documented extensively. There is some documented work on AE monitoring of vinyl-ester coupon types, finding water absorption to influence AE activity and correlating (but not so much detecting) the onset of delamination in environmentally exposed flexure tests with an increase in AE activity [9, 10]. [5, 65]

While data on composite ship structures may be limited, metal ship structures are subject to both lab studies and in-situ field trials. A trial on a 4800 Twenty-foot Equivalent Unit (TEU) container ship investigates the viability of an AE monitoring system as a screening tool for surveys and inspection planning [65]. The study concludes generally positively, in that crack and corrosion growth detection is verified to be successful. No mention is made of any false negatives, however. Furthermore, it is remarked that the marine environment gives rise to significant amounts of background noise, stressing successful signal discrimination and noise reduction to be essential for the success of AE structural integrity monitoring in maritime industry. Particularly noise from hydrodynamic evens such as waves, sloshing water and rain are identified as difficult to discriminate from crack growth. As a final note, this study finds the effective maximum distance between an AE event and a transducer to be about 4.0 m. In a similar trial on an 182 [m] oil tanker in fact a sort of building block approach is used to scale from mode I/II three-point fatigue bending tests to an in-situ field trial [66]. AE monitoring is performed during cruise at different speeds, anchoring, different weather conditions and during operation of devices on the ship. A classification tool based on pattern recognition (see Section 3.2.5) has been used to successfully detect AE sources of corrosion damage.

Finally, it is worth noting that the US coast guard is leading similar trials in order to assess structural fatigue life and evaluate the potential of Structural Health Monitoring (SHM) via AE monitoring on a US coast guard test bed. These are the Valid Joint Industry Project (JIP) and the US Coast Guard's Fatigue Life Assessment Program (FLAP). Though these studies again are only concerned with structural integrity monitoring in metal ship structures. [67–69]

The main conclusion that can be taken away from this is that there is prior art in AE structural integrity monitoring of ships, even to a very applied level. Event classification in ships, or in any in-situ application, may be particularly important because of the significant amount of noise. Though unfortunately experience of AE monitoring on ships is limited to metal structures, while composite structures are under-represented as of this time. The next section will go into more detail on signal processing and classification of AE events.

3.2.5 Acoustic Emission source classification

In this section, common current approaches to AE source classification in composites are discussed. The processing segment of AE monitoring generally consists of parametrization of recorded AE waveforms and subsequent data analysis on this simplified representation. Data analysis in many current state-of-the-art studies involve machine learning algorithms in order to process the large amounts of AE data produced by testing. Recording and parametrization of AE waveforms into "features" is discussed; and an overview of common approaches to data analysis is presented. But first, a few notes on the recording and representation of AE waveforms are given.

Before a signal can classified or interpreted, it must pass from source through medium to transducer to be detected. This flow of information is depicted in Figure 3.19. A damage event releases acoustic energy in the form of a wave, which has certain distinguishing features. The wave is manipulated by the medium or media through which it travels, before arriving at the sensor. The AE system must then decide whether this source must be recorded. Many current implementations of AE monitoring make use of a threshold value to ignore AE reflection hits or low-amplitude noise. When a signal crosses the threshold for detection, the waveform is recorded to the computer storage. This reduces the amount of stored (noise) data for processing, though a poorly set threshold obviously has potentially disastrous consequences for the measurement data. A floating threshold can also used, as a fraction of the maximum amplitude of an AE event. The recorded waveform can then be discretized, or summarized through parametrization.

When a signal is detected and recorded, efforts can be made to classify its source mechanism. Much like humans distinguish musical instruments and sounds based on their *timbre* [70], classification of an AE event may be based on its character, consisting of signal features like frequency content, amplitude and rise time. But it may also be recalled from the previous section that the acoustic signal is altered by the medium by propagation mechanics such as reflection, attenuation, etc. Consequently, the signal received on the far right of Figure 3.19 is potentially dissimilar from the source signal. Therefore, either of two approaches can be taken from Figure 3.19: First, to try and correct for the signal manipulation through modelling of the entire chain of information propagation. This approach may be regarded as tedious, as well as sensitive to changes in set-up or material. The second option is to more or less accept the altered signal and investigate whether any information is retained even in the altered signal. Whether this leads to a more robust solution remains an open question, as the interpretation of the signal as transferred by the medium may in fact also result in an AE system sensitive to changes in set-up or material. Nonetheless, this is chosen as the preferred approach for the current project.

As a side-note in the current project, signal processing of the detected signal can also include source location. In fact, the potential of AE sensors to monitor larger areas can be considered as one of the major advantages of this method, especially for larger structures like ships: even if classification of a source mechanism is inaccurate, it may hold more valuable information about the structure's maintenance to know where most of the AE activity originates. Additionally, source location information may potentially be used to correct for the path-dependency through the medium of the signal's character as recorded by the AE transducer [71]. However, since the current project is limited to coupon-like specimens, source location is not pursued here.

In this section, current methods for the classification of source mechanisms from AE monitoring data are examined. First, the use of waveform features to describe a transient AE signal is

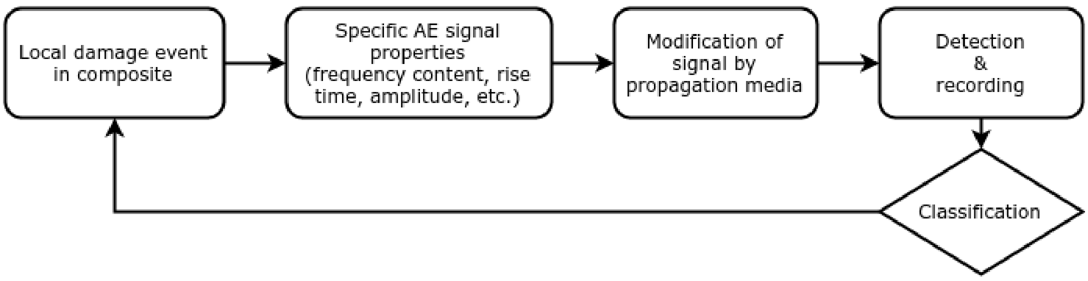

Figure 3.19: Schematic representation of the flow of AE information from creation (damage) to capture (AE detection by transducer), which ultimately is used for the "decision activity" of classification, which attempts to relate the recorded information back to its source.

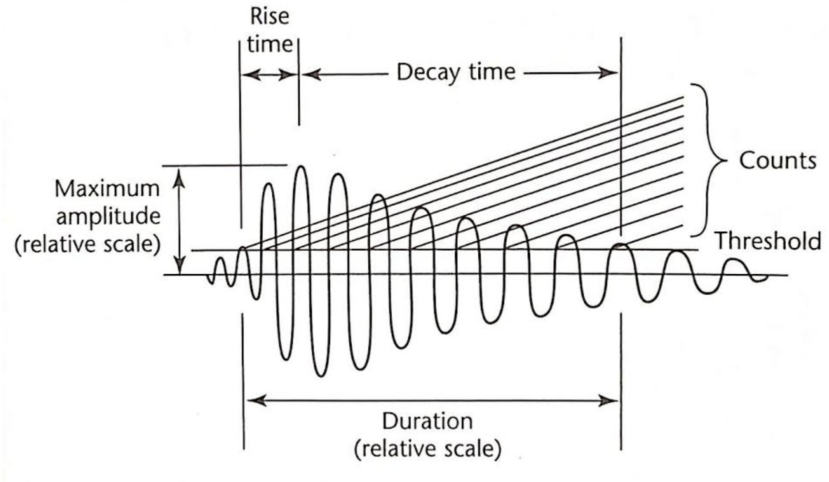

Figure 3.20: Commonly used waveform feature definitions [41]

discussed. Thereafter, current approaches to data analysis are discussed, involving the use of pattern recognition algorithms to enable processing of a plurality of waveform parameters. This section is concluded with an selection of AE monitoring studies based on pattern recognition and their parameters.

Transient waveform feature description

A common representation of AE events for AE monitoring is a description of the recorded waveform by a set of features such as rise time, maximum amplitude, counts, duration and energy [41]. Feature description first and foremost serves to convert the recorded waveforms into a data format more practical to subsequent analysis. The particular features used to describe a waveform are not set in stone, though particular features tend to be conventional. Some of these commonly used features are illustrated in Figure 3.20: amplitude, duration, rise time, decay time and counts. At this time it is important to realize that the threshold setting strongly influences the values and distributions of all of these waveform features. This stresses the fact that AE monitoring in its current implementation is primarily a configuration-dependent representation of the character of AE activity - not an absolute one.

In particular earlier studies on AE classification focus on one or several of these features, correlating damage mechanisms based on, for example, signal amplitude. It is found by some researchers that in GFRP samples, matrix cracking releases lower energy, lower amplitude signals while fibre pull-out or fracture is associated with higher amplitude signals. [15,18] However, because measured signal amplitude is sensitive to propagation effects it has since been established that amplitude alone is a poor classifier for source mechanisms [19]. The rising-slope criterion as proposed by Uenoya [72] presents an alternative approach, which might be considered more robust because it is not as dependent on the absolute signal amplitude itself.

The Fast Fourier Transform (FFT) is used to obtain frequency spectrum parameters of AE signals. A study by Giordano [73] identifies fibre breaking based on an acoustic fingerprint: a set of characteristic frequencies in single fibre composite tensile fracture. Analogously, Ramirez-Jimenez [74] discriminates fibre-matrix debonding, pull-out and fibre breaking by considering the primary frequency content of AE events in GFRP tensile specimens. Discrimination of damage mechanisms in specimens subject to an out-of-plane load based on partial power (spectral power density) of selected frequency intervals has also been shown to be effective [46,75]. Other

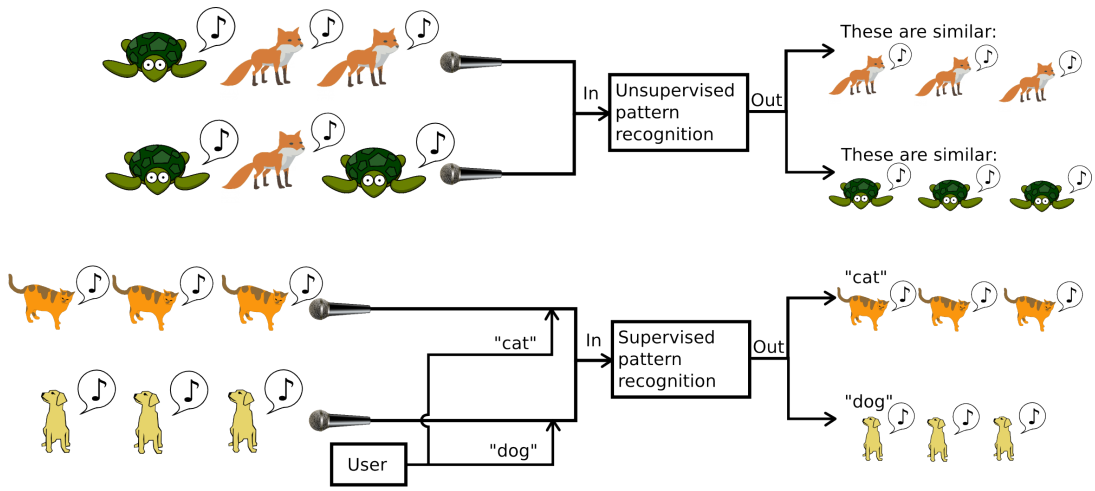

Figure 3.21: Conceptual depiction of the difference in the learning phase of the unsupervised (above) and supervised (below) pattern recognition algorithm. Images in this figure are from the public domain (www.publicdomainpictures.net)

studies have been able to discern damage mechanisms in tensile and flexural loaded GFRP specimens based on signal frequency content, revealing in this case fibre breaking to emit higher frequency peaks than transverse matrix cracking or fibre-matrix interface failure [76, 77]. In conclusion, the frequency spectrum may present one or several signal features worthwhile of including in classification efforts.

In many current AE-based damage classification strategies, typically six or more different waveform features are used as input to a pattern recognition algorithm [78] in order to describe a waveform with a more complete set of parameters compared to when only a single or a few waveform features are used. However, even in current implementations, AE source classification based on multiple waveform features has found limited success in distinguishing fibre breakage and matrix cracking [78, 79]. The use of pattern recognition for AE event classification is discussed in more detail in the following paragraphs.

Pattern recognition techniques

When dealing with the inherent volume that AE data presents, traditional analysis techniques inherently become impractical. Pattern recognition is therefore borrowed from the field of machine learning as a suitable tool to classify AE events based on multiple AE waveform features [15]. Although a signal feature in this context can be any scalar contributing to the signal's description, the features chosen for pattern recognition are often waveform features such as those described in the previous paragraphs. Similarity of vectors is decided based on (e.g. the Euclidian) "distance" between the vectors. A cluster of similar vectors is referred to as a class. Ideally, each waveform class represents a different damage mode or source mechanism, though especially when the algorithm is "unsupervised", this is not inherently guaranteed by the procedure. In the following paragraphs, supervised- and unsupervised learning are introduced. Subsequently, the use of these algorithms for the purpose of AE source classification is explained. [15, 41]

Pattern recognition algorithms are divided into two types: supervised and unsupervised algorithms. Unsupervised pattern recognition algorithms are used to cluster data into classes in the absence of prior knowledge about the nature of the classes. I.e. the term "unsupervised" refers to the ability of the algorithm to cluster data into classes without the user having to provide a learning dataset that includes information about which vector belongs to which class. This is both an important limitation and a feature of unsupervised pattern recognition algorithms because it also means that there is no guarantee that the clustering has a meaningful relation to the source mechanisms in the material. While some data-driven mathematical metrics are available to help alleviate this problem by providing a measure for the "quality" of clustering; such as the Davies-Bouldin or the Silhouette clustering performance indices [77, 80], the most scientifically rigorous method for verifying the classification would be to validate the clustering by some external means, if possible. However, this is not readily the case of AE monitoring of damage in composites.

In supervised pattern recognition, by contrast with the unsupervised algorithm, the user has prior knowledge about which vector corresponds to which class. This difference in the learning phase of the supervised vs the unsupervised method is conceptually depicted in Figure 3.21. In the supervised technique, a user with prior knowledge designates which vectors (here: sounds) belong to the clusters "cats" or "dogs". In the unsupervised algorithm, the user has no prior knowledge about what sounds foxes or tortoises make, so he depends on the unsupervised algorithm to make a meaningful distinction. In both cases, after the learning phase is completed, the algorithm can be presented with new animal sounds, which it classifies in either of the classes based on the distance of the feature vector to each of the clusters. [41]

Many recent studies use a Neural Network (NN) known as Kohonen's Self-Organizing Map (SOM) [81], an unsupervised algorithm, to develop a topologically reorganised map of the AE data feature-space whereby similar vectors are grouped together [7,15,22,77,78,80,82]. Clusters in the data are then identified using, e.g., k-means clustering, for which deductions can then be made pertaining to the source mechanisms based on the expected damage mechanisms in a particular laminate. Interestingly, this research group finds that it is found that for a two-class case (matrix cracking, fibre-matrix interface failure) the performance of a supervised k-Nearest-Neighbour (k-NN) classifier is comparable to the unsupervised Kohonen's map classifier, while the former is much easier to use [78]. This is especially interesting because k-NN is such an conceptually straightforward algorithm, determining the class of a "new" data point by majority vote of the first k neighbouring data points for which the class label is known.

In Table 3.2 an overview is given of a selection of AE classification studies that use machine learning algorithms for classification. Most of these use Kohonen's SOM with k-means clustering. A couple of observations may be listed from this table. Firstly, that most current AE studies do not use E-glass/vinyl-ester as a subject material; however, the methods used between studies on E-glass/polymer- and graphite/polymer-composites are roughly similar, leading to an expectation that these methods are likely to be applicable to E-glass/vinyl-ester epoxy. A second observation from this table is that k-means clustering is very popular. Thirdly and finally, it may be noted from Table 3.2 that both the distinguished damage mechanisms and the waveform features used for classification vary quite significantly across studies. With regard to the classification of AE sources this is taken to suggest that in the signal processing phase it is prudent to be careful drawing conclusions about the source mechanisms of the AE signals under examination. Furthermore, with regard to the range of waveform features used, it is concluded that their value in classification must be established by direct experimentation, while attention must be paid to their mutual interdependence. Finally, it is worth noting that practically all AE classification studies on composite materials are concerned with coupon specimens: Crivelli et al. [80] as well as Godin et al. [22] find AE classification based on machine learning to be robust against changes in layup, but none of the studies in this document publish robustness of their classification algorithms beyond these relatively modest changes in subject material. I.e.

the means to generalise AE classification approaches from coupon to substructure and structure (i.e. further steps in the building-block approach, so to speak) is largely not achieved or documented.

Table 3.2: Summary of classification parameters of a selection of AE classification studies, most of which implement the SOM for classification. *: A=amplitude; D=duration; R=rise time; CNTS=counts; CTP=counts to peak; E=(MARSE) energy of the waveform envelope; Eb=absolute energy; FCOG=FFT center of mass; FMXA=FFT peak freq.; FC=FFT central freq.; AR=amplitude ratio of peaks in FFT domain; FA=FFT amplitude; PP=partial power in FFT domain.

Study	**Subject material**	**Classification algorithm**	**Clustering**	**Damage mechanisms distinguished in signal clusters**	**Waveform features used***
McCrory (2015) [82]	graphite/epoxy	SOM; Unsupervised Waveform Clustering; based on ratio of modal content	k-means	matrix cracking; delamination	R, CNTS, Eb, D, A, FCOG, FC, FMXA
Crivelli (2015) [83]	graphite/epoxy	SOM	k-means	matrix cracking; impact-induced in-plane delamination	A, D, CNTS, E, FCOG, FMXA
Doan (2015) [84]	graphite/epoxy	P-clusters optimization based on Davies-Bouldin performance index	Gustafson-Kessel-based, with Mahalanobis distance	mechanical noise; fiber-related damage; internal friction noise; matrix-related damage; interface-related damage	Eb, A, E
Sause (2012) [85]	graphite/epoxy	P-clusters optimization based on clustering performance indices	k-means	matrix cracking; interface failure; fiber breakage	CTS/D, CTP/R, FMXA, FCOG, FMXA*FCOG, PP(6 domains)
Crivelli (2014) [80]	E-glass/polyester	SOM	k-means	breakage of short fiber mats; near-end failure modes; matrix degradation; in-plane delamination	A, D, R, CNTS, E, FCOG, FMXA
de Oliveira (2008) [77]	E-glass/polyester	SOM	k-means	transverse matrix cracking; crack initiation; longitudinal matrix splitting; fiber-matrix interface failure and/or edge delamination initiation; delamination; fiber cracking	A, R, D, E, CNTS, CTP, AR(1,2), E/F, CTS/D, FMXA(1,2), FA, FCOG
Godin (2006) [7]	E-glass/polyester	SOM	k-means	matrix cracking; fiber-matrix interface failure	R, D, CNTS, CTP, A, E
Godin (2005) [22]	E-glass/epoxy	SOM	k-means	matrix cracking; interface debonding; fiber failure; delamination	A, D, R, CNTS, CTP, E
Godin (2004) [78]	E-glass/polyester	SOM; (supervised) k-NN	k-means	matrix cracking; interface debonding; fiber failure	R, CNTS, E, D, A, CTP
Bhat (2003) [86]	graphite/epoxy	SOM	k-means	Hsu-Nielsen; scratch; spark; spike; tap; hydraulic	R, CNTS, E, A
Huguet (2002) [15]	E-glass/polyester	SOM	k-means	matrix micro-crack formation; fiber-matrix interface failure	R, CNTS, E, D, A, CTP

Limitations of current approaches and conclusions for the current project

Regardless of the type of pattern recognition algorithm used, any data-driven segment of the analysis in itself does not provide a scientifically rigorous methodology to classify damage mechanisms, or indeed any source mechanisms in AE monitoring of composites. This is because the data itself cannot be validated confidently using only the information contained in the AE waveforms. Or in terms of the supervised- and unsupervised algorithms: both processes require a secondary method to provide some level of external validation (supervision) of either labelling of data points (vectors), or of clusters as a whole. However, in many of current existing publications, this validation procedure is either exclusively data-driven, or based on post-mortem observations and expected damage mechanisms in the test specimen.

As an example, consider a study by Godin et al. [22], which extends the authors' previous work on E-glass/polyester to E-glass/epoxy. In this study, matrix cracking and interface failure are distinguished in UD samples based on amplitude; 70 [dB] being chosen somewhat arbitrarily based on past experience with E-glass/polypropylene [15]. The researchers find that tests on cross-ply laminates are still able to distinguish these same failure modes based on rise and decay behaviour and energy, signal amplitude is found to be of lesser value as a criterion. In this case, the use of multiple waveform features improves robustness of classification. However, it can be noted that there is a fundamental weakness in the two-class basic version of this type of supervised approach: the data is linearly separated along one feature-dimension (amplitude, at 70 [dB]) resulting in two classes in the dataset. This linear separation, whether explicit or implicit, means adding any other features to the data point's description (vector) would add superfluous information: by definition, a decision boundary can be drawn at 70 [dB]. Therefore, if damage mechanisms (e.g. matrix cracking and fibre breakage) can be separated along the single feature-dimension of amplitude, there is no need for pattern recognition algorithms in multidimensional feature-space. In fact, to a k-NN algorithm any added feature-dimensions would potentially reduce classification accuracy (—this readily follows from the definition of k-NN). Regardless, any of these potential limitations are moot if the classification cannot be validated in terms of damage mechanics. In this study, and many others, this is done by a combination of post-mortem inspection of test samples and expected damage mechanisms, while validity of clustering is supported by aforementioned data-driven mathematical metrics. These validation methods lack a connection with damage in a real-time capacity, however, while they provide no means to identify persistent noise sources, such as from friction of internal fractured surfaces.

Kohonen's SOM has been considered one of the most popular artificial neural network used for unsupervised learning applications in general [77], while pattern recognition and neural networks are a commonplace approach in current state of the art AE signal classification studies. These procedures are considered by some to be the most effective analysis tools available to solve the classification problem [80]. However, even with many studies asserting confidence these machine learning techniques, more than circumstantial evidence for the identification of, e.g., different damage mechanisms from AE activity, finds only limited external validation and thus remains largely a still sought-after achievement. In this way, assertions may be said to be echoing from the past when quoting Hamstad's writing about the state of AE applied to composites in 1986 [44]: "[...] fundamental errors which seem to be repeated over and over again [include] very questionably guessing about what the actual AE sources are". It may however be added that this "guessing" is not necessarily an attributable erroneous approach on the part of the researcher, but rather represents the challenge of validation inherent to AE monitoring of composites. The major challenge in the field of AE monitoring of composites is often cited to be that of classification [3]. But by considering the common definitions for classification and validation in machine learning and science in general, it can be seen that this is not the case: there is a plurality of machine learning and data analysis tools available, and many researchers find more than enough different ways of classifying the data. Therefore, it may be asserted that the major challenge in AE monitoring of composites lies not so much in classification, but in validation of source mechanisms.

Ultimately, the event classification techniques and tools described in this section are not to be seen as mutually exclusive, or even a conclusive answer to the classification problem. Rather, pattern recognition enables the use of any classifiers the researcher sees fit for the particular analysis case. The most important limitation in signal processing may then be the time and expertise available to the researcher. For the current project, it is asserted that developing new classification or clustering algorithms is not within the scope of the project goals, as many such algorithms already exist and the major challenge to advancing the field of AE monitoring of composites can be said to be not so much the classification of the data but validation of said classification.

In the current project therefore, it is considered to be more important to find which conclusions from the experiments can be made that are both congruent with existing knowledge on damage mechanics as well as following inductively exhaustive from the data. This touches on the inherent validation problem in AE monitoring: since direct observations of the source mechanics of AE events cannot be made, arguments for validation of AE sources must be made inductively, exhausting all probable explanations of a source mechanism until only the most likely answer remains. As it stands, conclusions on AE sources and therefore validation of any classification can seemingly not be made logically deductive; so what tools remain to draw conclusions are inductive exhaustion until a most likely conclusion can be plausibly accepted. In the current project, the approach to this inductive argument is based on the following facts: First, from the current state of documented knowledge on damage mechanics in composites, it can be asserted that low load level T-T fatigue gives rise to dominantly a single damage mechanism (matrix cracking) and that damage causes stiffness degradation[1]. Second, measurement of the strain in the test sample provides quantified evidence for stiffness degradation, and therefore a measure of existence of damage. And third, although it is not exactly known how damage grows on a per-cycle basis [87], from constitutive law it follows that there must be a relation between damage and the strain or stress in the material. This provides a basis for labelling hits based on the fatigue cycle phase segment they occur in. The practical aspects of the production-, testing- and data processing methodology in this project are explained in detail in the next chapter, whereas the analysis of testing results is presented in Chapter 5.

[1]Beware the *fallacy of the converse*! It is considered reasonable to assume that if there is stiffness degradation, there must be damage. But definitely not that all damage necessarily leads to stiffness degradation.

Chapter 4

Methodology

In this section the project methodology is explained in terms of the three main processes of coupon production, testing and data analysis. From the discussion of material mechanics in Section 3.1, it has been found that T-T fatigue-loaded composite plates loaded at low load levels (high-cycle fatigue) may be expected to be dominated by matrix cracking; whereas at higher load levels other damage mechanisms such as fibre-matrix interface failure, fibre breakage and delaminations start to occur more dominantly. If different damage mechanisms produce distinguishably different AE events, the first step in a building-block approach to AE monitoring is to investigate only one dominant damage mechanism, confirming it can be distinguished from any noise present in the system. From the preceding, the most straightforward method to produce a relatively isolated case of matrix cracking is to load coupons in low load level fatigue. This ensures the source mechanisms for acoustic activity are dominantly matrix cracking; but, analogously, this also means that the stiffness degradation in the material is attributable mainly to this damage mechanism. This allows validation of the AE monitoring data by measuring the material's stiffness degradation, which can be accomplished with DIC strain measurements. An additional piece of validation can be exploited: recent work by Pascoe et al. [87] suggests that the instant of crack growth and the emission of elastic energy are measurably inter-related, indicating a relation between the Kaiser effect and a threshold load for damage growth. Therefore, correlating the phase of the applied fatigue load with the occurrence of an AE hit is used as method of organising and labelling the AE hits with constitutively-informed "metadata". This approach is depicted by the blue path in Figure 4.1. In this figure, the current approach may be seen as adding a few extra validation checks between the AE data generation and -analysis, as compared to the data-centred approach, which is more common in existing research (see the previous chapter). The damage mechanics (Chapter 3) justify the use of stiffness degradation, which in turn can be used as validation for any applied filtering (in this case the fatigue cycle phase labelling). The underlying mechanics represented in any clusters of AE hits are in turn validated by the preceding steps.

With regard to validation it is also recalled to the reader's attention that post-mortem NDE evidently would provide a suitable and common means to validate both the presence and the extent of damage after testing, as well as potentially discerning different damage mechanisms. For the coupon specimens under consideration in the current projects, attempts have been made at post-mortem NDE through the methods of both phased-array ultrasonic imaging and optical coherence tomography. Both of these methods have not been able to reliably discern between

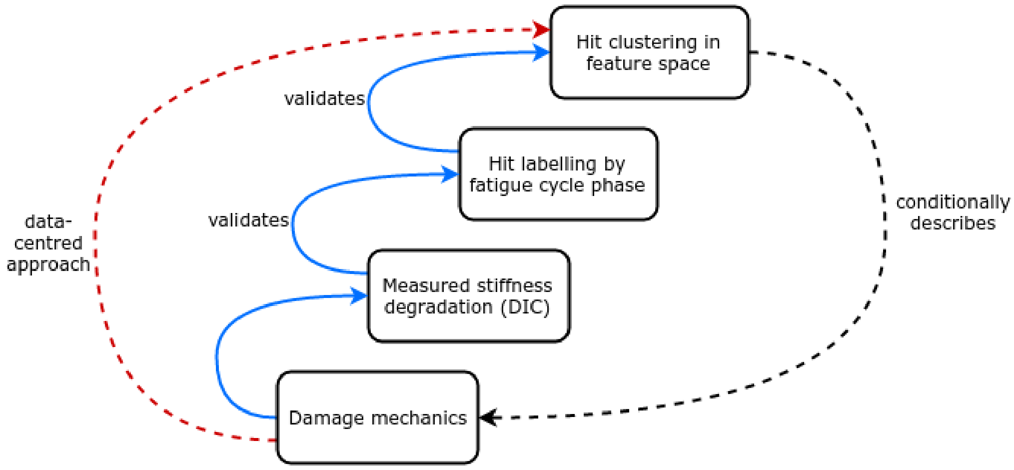

Figure 4.1: Schematic depiction of the project's approach (blue path) to AE classification and validation. Starting from damage mechanics, each step in the analysis is supported by the preceding ones. If all validation steps are sound and assertions arising from the data analysis can be made confidently enough, the AE data may be said to describe the damage mechanics (black dashed path). In contrast, the data-centred approach is depicted.

tested and untested samples, or indeed to discern recognisable features in the composite at all. Although this result could have a number of causes, it is considered most likely that for this particular composite, there are dominant features present in the material at both the ultrasonic wavelength ($O(1$ [mm]$)$) and the optical wavelength ($O(1$ [μm]$)$), such that both wavelengths are scattered enough to not produce an image in which features can be distinguished. The relatively large thickness of the coupons is likely to contribute to a greater amount of scattering as well. Regardless, the best guess of resolving this issue would be to use these or other imaging techniques at a different wavelength, which have not been available at the time.

Coupons of two different layups are chosen to be investigated: a ±45° and a 90°-dominated layup. These layups are chosen for two reasons: First, to allow for some amount of generalisation about the characteristics of acoustic activity originating from matrix cracking, independent of layup. And second, because both of these layups may exhibit different secondary damage mechanisms; namely mainly transverse fibre debonding for a 90°-dominated layup, and mainly interfacial damage for a ±45° layup. For the initial dataset, a similarity parameter between the coupon types is chosen to be a load level of 35%, where the load level is defined as the maximum fatigue load as a fraction of the respective coupon's Ultimate Tensile Strength (UTS). To expand the dataset and because a large difference in damage accumulation is found between the coupon types at this load level, two extra fatigue test types are added to the initial two: fatigue loading of a 90°-dominated coupon at a static strain-similar level to the ±45° coupon; and a ±45° coupon at a higher load level of 45%.

This chapter is divided into three sections, representing coupon production, coupon static- & fatigue testing; and data preprocessing & -filtering respectively. First, production of the ±45°- and 90°-dominated Glass-Fibre (GF)/vinyl-ester coupons is described in Section 4.1. Thereafter, testing of said coupons is described in Section 4.2. Both coupon types are quasi-statically loaded in tension to determine their respective UTSs, relative to which the fatigue F_{max} values are defined. The AE datasets, supported by DIC measurements, are the main objective of the testing campaign and are established through the T-T fatigue testing procedures presented in Section 4.2.3. Finally, analysis of the data obtained in fatigue testing is prepared by a series of data processing and -filtering steps, which are described in Section 4.3. This includes computation of the stiffness degradation from DIC data, as well as correlation- and labelling of AE hits with the instantaneous applied load. Analysis of the AE- and DIC data is presented in Chapter 5.

4.1 Coupon production

In this section, vacuum infusion production of the GF/vinyl-ester coupon batches is described. First, the production procedure is described step-by-step from preparation and vacuum infusion of the composite panels to sawing and tabbing of the coupons. This section concludes on discussing several notes on the production process. Measured final coupon dimensions are summarized in Table 4.1.

Coupon production is divided into three vacuum infusion batches of one panel each. The coupons are based on ASTM D3039 [88], though it may be noted that the accurate determination of the mechanical behaviour of the material is not a focus of the current study. Each panel is laid up out of 8 UD 1200 $[g/m^2]$ E-glass fibre plies, infused with the Dion VE1057 resin and mixing in the accelerator components in the ratios specified by the material data sheet. The material and laminate thicknesses are chosen to be compatible with DSNS's production methods outlined in Chapter 2. The first panel has layup $[90_3/0]_s$, the second panel has layup $[\pm 45_4]_s$ and the third panel has layup $[0/90_3]_s$. A set of 0°-layers is included in the 90°-dominated coupons because it is anticipated that a composite layup consisting of only 90°-layers would not be representative for any real-world engineering composite application. Tabs are applied as strips to sub-panels and the coupons of dimensions 175×25 [mm] are cut from these sub-panels. Handling of raw fibres is always done while wearing nitrile gloves. The production procedure of each of the panels is identical in methodology and is described by the following procedure:

1. Plies are cut from the roll of UD E-glass using scissors. Cuts perpendicular to the fibres are cut between the stitches (Figure 4.2a).
2. The surface of an aluminium plate is prepared as the laminating surface:
    i. A perimeter of approximately 2 [cm] is cleaned with ethanol, which will be used to tape the vacuum bag to.
    ii. The outer perimeter is scoured with a scotch-brite pad.
    iii. The outer perimeter is cleaned with ethanol again and covered with tape to protect it from the release agent.
    iv. Release agent is applied to the exposed plate surface.
3. The infusion setup is built (Figure 4.2b):
    i. Tape is removed from the outer perimeter and tacky tape is applied, not removing the protective layer from the tacky tape so as to avoid glass fibres adhering to it.
    ii. The GF laminate (e.g. $[0/90_3]_s$) is stacked and aligned on the plate surface.
    iii. The laminate is covered with a peel ply layer, overlapping the laminate edge.
    iv. The peel ply layer is covered with release film.
    v. The release film is covered with the infusion mesh, leaving about 2 [cm] clear of the infusion mesh at three of the four laminate edges, while overlapping the fourth edge by a few centimeters.

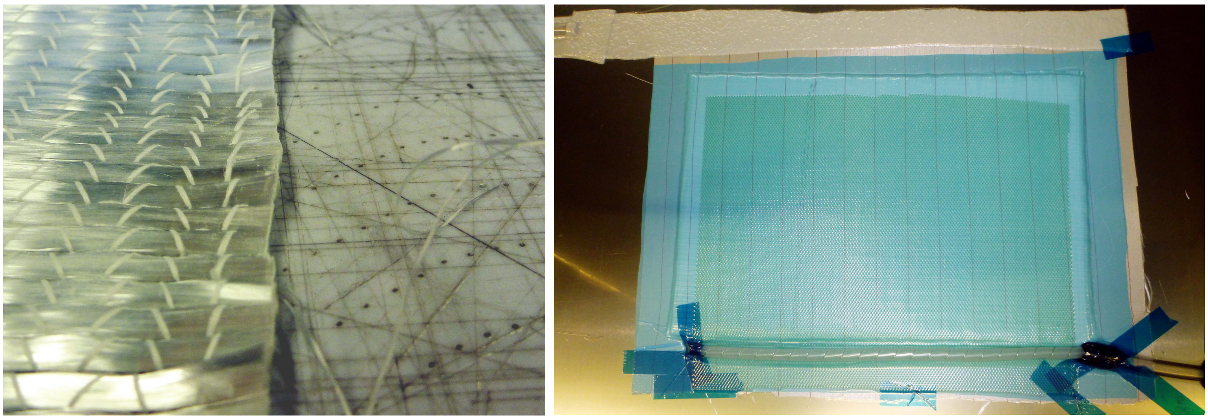

(a) **(b)**

Figure 4.2: **(a)** Close-up of a cut UD E-glass sheet showing the cut between the stitches; **(b)** The infusion layup prior to resin infusion, showing the vacuum air tube in the top left and the resin supply tube in the bottom right.

vi. Corners of the ply stack are taped under slight tension to the aluminium plate to keep them in place.

vii. A runner is added to the edge with the overlapping infusion mesh (bottom of Figure 4.2b), taping it in tension to the stack in order to open up the runner. Tape one end of the runner closed. The other end is connected to the resin supply tube and the connection is taped with both tacky tape and regular tape in order to protect the vacuum bag from any protruding sharp edges on the runner or supply tube.

viii. The vacuum air tube is positioned such that its path raises from the laminate surface about 1 [m] straight up from the laminate before it returns down to the vacuum pump. This is to ensure no resin will flow in the pump.

ix. Add breather material on the side opposite to the runner (top of Figure 4.2b), taking care that the breather material contacts the peel ply layer. The breather material forms a path to the vacuum air tube.

x. Both the vacuum air tube and the resin supply tube are secured to the tacky tape on the plate edge using an extra circumferentially applied layer of tacky tape.

xi. The vacuum bag is applied to the stack, secured by the previously applied tacky taped plate perimeter.

xii. The resin supply tube is clamped shut and the pressure inside the vacuum bag is brought to below 10 [mbar] for at least 30 *min* to clear the stack of any trapped air.

4. The pressure inside the vacuum bag is set to 150 [mbar] (i.e. higher than the resin degassing pressure - see next step).

5. Infusion is prepared and started:

   i. The resin is mixed in the ratios supplied by the manufacturer. The NORPOL PER.24 activating component is added last. The mixture is gently stirred for about a minute using a wooden stirring implement.

   ii. The resin mixture is degassed at 120 [mbar] for 15 minutes. The degassing pressure is chosen to be lower than the vacuum infusion pressure, but sufficiently above the 20°C vapour pressure of styrene (6 [mbar]), which is present in the resin mixture.

iii. Resin infusion is started by easing pressure on the clamp that is installed on the resin supply tube. The resin flow is held for a few minutes at a few centimeters beyond the tube clamp, in order to allow any gas still present in the tube to escape. Subsequently, the clamp is eased again to start the infusion.

6. The resin supply tube is clamped shut again after resin has partially penetrated the breather material at the vacuum end of the laminate.

7. The laminate is left to cure at room temperature for approximately 14 hours.

8. After curing, the vacuum is relieved and the laminate is de-bagged.

After production of the panel, it is sawed into sub-panels and tabbed with separately produced $[\pm 45]_8$ plain weave GF/epoxy tabs. The production process of the tabs is similar to that of the coupons themselves, with the notable two exceptions that the epoxy resin is degassed without setting a minimum pressure, and that its reaction is exothermic, requiring cooling by immersion of the resin container in water. For this reason, the tabs' production procedure is not repeated here. After E-glass/vinyl-ester panel production, subsequent sawing, tabbing and sawing again of the material into the final coupons is described by the following steps:

9. The panel is cut into four sub-panels (Figure 4.3a and 4.3b). A sample strip of 1.5 [cm] width through the middle of the panel is left clear to use for the determination of fibre volume fraction. Cutting is done using the CARAT diamond saw (blade width 3 [mm], Figure 4.3c). This relatively heavy saw is chosen because it is found to produce cuts that are more straight and consistent over the length of the relatively sizeable composite panels.

10. Tabs are applied to the coupons (sub-panels):

    i. The gauge surface of the coupon is taped off to protect it from subsequent abrasion and adhesive.

    ii. The to-be-tabbed surface of the coupon is sanded manually with P150 abrasive paper and degreased with ethanol. The product of this step is shown in Figure 4.3b.

    iii. The adhesive (3M EC9323) is prepared according the the manufacturer's instructions, adding $7wt\%$ glass pearls of 100 [μm] diameter to the mixture, to control for the bond line thickness.

    iv. Adhesive is applied to both surfaces and the adherents are brought together under light manual pressure.

    v. The tabbed sub-panels are enclosed in a vacuum bag, which is brought to an internal pressure of 800 [mbar].

    vi. The tabbed sub-panels are left to cure in an oven at $40°C$ for about 12 hours.

11. Tabbed coupon sub-panels are sawed into single coupons of 175×25 [mm] with the diamond saw (Figures 4.3c and 4.3d).

12. The sawed ends all of coupons are sanded manually with P150 abrasive paper (Figure 4.3e).

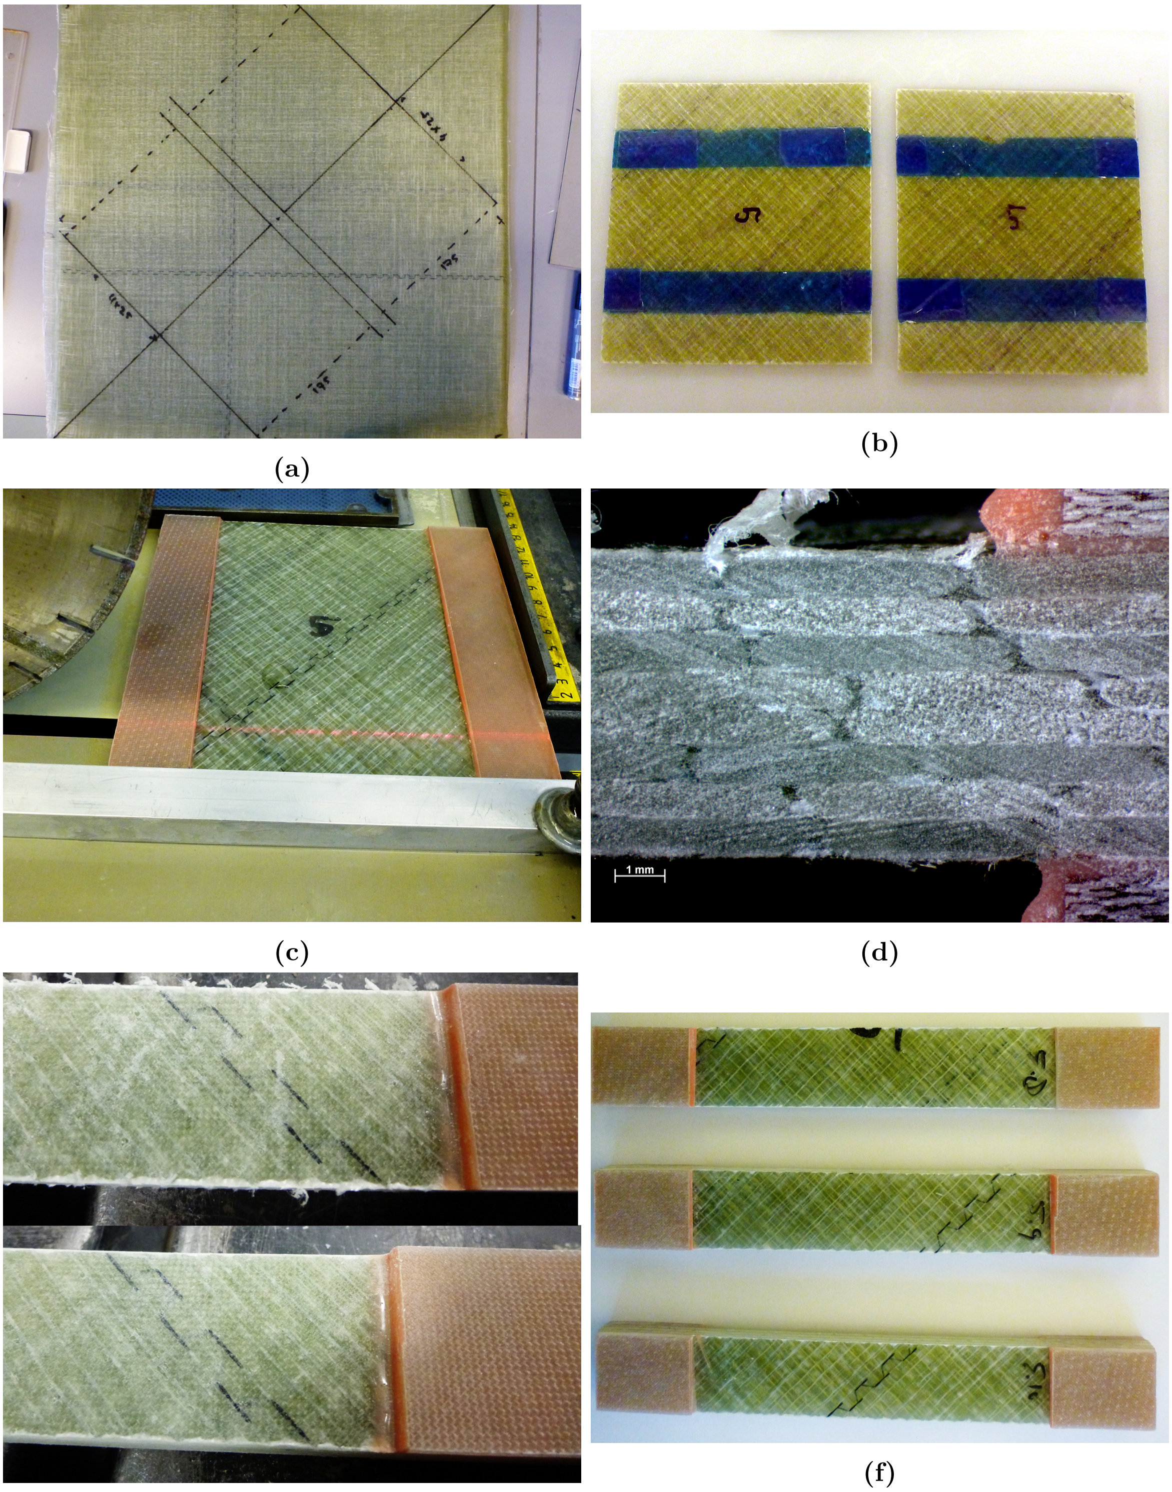

(a) (b) (c) (d) (e) (f)

Figure 4.3: **(a)** Cutting plan drawn on a ±45° panel. **(b)** Two ±45° sub-panels, taped and surface-treated ready for tabbing. **(c)** Sawing setup for a ±45° sub-panel. **(d)** Macroscope image of the long edge of coupon 5.10, after cutting but before sanding (Zeiss Discovery V8 at 2× magnification). Note the tabs and (orange) tab adhesive in the top right. **(e)** Coupon sanding before (top) and after (bottom) images. **(f)** "front" view of finished coupons 5.8, 5.9 and 5.10

Table 4.1: Averaged coupon dimensions and coefficients of variation (C_V) over the coupon populations.

Laminate	#	thickness [mm]	C_V	width [mm]	C_V	gauge length [mm]	C_V
$[90_3/0]_s$	14	6.28	0.57%	25.20	0.40%	108.9	0.52%
$[\pm45_4]_s$	20	6.18	0.44%	25.02	1.11%	109.4	0.89%
$[0/90_3]_s$	8	6.11	0.35%	24.67	0.70%	81.6	0.14%

This concludes the coupon production process. Figure 4.3f depicts a set of finished coupons. What follows are a few notes notes about the production process. In step (1.), a sample strip is cut out of the middle of the panel. From this strip, random pieces of 1.5×2.5 [cm] are cut for the determination of fibre volume fraction V_f. V_f is calculated based on the resin density as specified by the technical data sheet, and found to be approximately 57% for all panels. Note also in step (10.vi.), the tabbed coupons are cured at $40°C$. The technical data sheet of the adhesive specifies full curing after two hours at $80°C$. However, this temperature is decided to be avoided because of the low glass-transition temperature of vinyl-ester resins. Therefore, a longer curing cycle at lower temperature is chosen, based on the rule of thumb that the chemical reaction rate roughly halves for every $10°C$ decrease in temperature, as well as considering the fact that full curing occurs at room temperature in no more than 14 days, as specified by the 3M EC9323 adhesive's data sheet.

To conclude, an issue that is encountered during production is discussed: During initial vacuum infusion batches, an unexpected amount of air bubbles is found to propagate through the resin from the resin supply- to the vacuum pump side. Upon inspection, it is found that after resin exits the resin supply tube it flows back along the outside of this tube (but under the vacuum bag). This type of behaviour can be observed in the circled area of Figure 4.4. More significantly, the resin flow is not halted by the tacky tape securing the vacuum bag to the aluminium layup surface. No definite conclusion is reached as to the mechanics or chemistry of this defect, but it is observed that as the resin reaches the tacky tape, air bubbles start to enter the resin, deteriorating the quality of resin penetration in the laminate. It seems that the resin deteriorates, or possibly dissolves, the adhesive joint between the tacky tape and the aluminium plate, locally compromising the vacuum seal. This issue is resolved for the final production batches by switching out the standard Cytek tacky tape for a Terostat-81 one. Another tacky tape type, LTY900, was also tried with negative results. In Figure 4.4, an interim solution is depicted wherein three (two visible) tacky tape barriers are applied in an effort to prevent the resin from flowing back along the supply tube. Note from the figure that the resin does not flow beyond the first (Terostat-81) barrier, providing a solution to the back-flow problem. Evidently, it must be noted that forming of air bubbles in the resin is also influenced by the degassing depth and -duration prior to infusion. On the other hand, with a resin gel-time of 4 hours specified by the material technical data sheet, the resin should also not spend an excessive amount of time between mixing and infusion. A degassing time of 15 minutes at 120 [mbar] is found to suffice.

After production, the coupon dimensions are measured: thickness (6 measurements per coupon), width (3 measurements per coupon) and gauge length (2 measurements per coupon). These three dimensions are averaged per coupon, and the distribution of the average dimensions (C_V) is presented in Table 4.1. The count on the number of coupons is that of finished products. Usable material or half-products (including produced but not tabbed coupons) are not included. The quasi-static- and fatigue testing procedures of these coupons are presented in the next section.

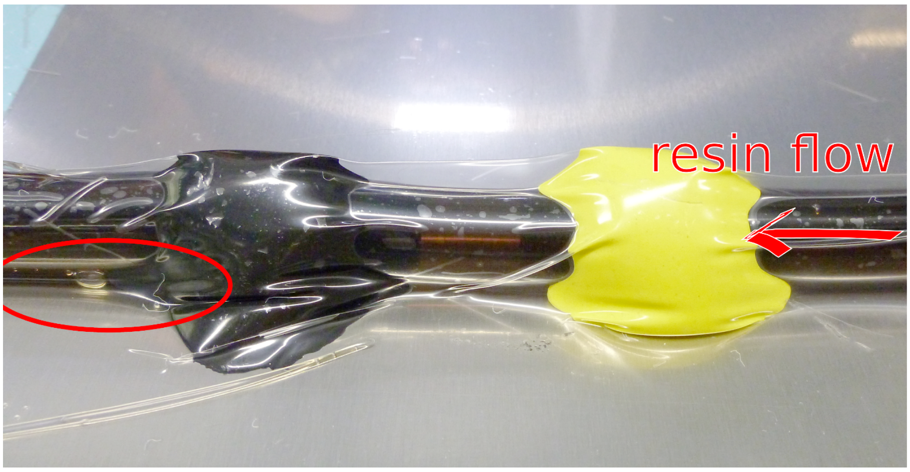

Figure 4.4: Interim solution to the resin back-flow production issue. Resin flowing from the right flows back along the outside of the tube (circled) after exiting it from the left. The first tacky tape barrier (Terostat-81; left and black) effectively blocks the resin back-flow, whereas the second barrier (LTY900; right, yellow and superfluous in this example) is found to be ineffective at achieving this feat. The final solution consists of only the Terostat-81 barrier.

4.2 Testing procedures

In this section, the testing procedures for the coupons are detailed. Static tests are performed mainly to establish the ultimate strength of the material, so that the maximum load for fatigue tests can be set relative to this value. The fatigue tests are the main source of data used in later analysis. Section 4.2.1 details the general test set-up and common parameters for both test types. Thereafter, Sections 4.2.2 and 4.2.3 respectively provide a detailed description of the static- and fatigue testing procedures.

4.2.1 Test set-up

In this section, the test set-ups and preparation for both static- and fatigue tests are described. First the physical set-up is summarized, after which the set-up preparation and parameters are given. The physical set-up for static and fatigue tests are identical. The static- and fatigue testing procedures are described in detail in the sections that follow.

In Figure 4.5(a) and (b), a schematic depiction of the test set-up and a close-up of a mounted coupon with attached AE sensor are shown respectively. The general test set-up is described referring to Figure 4.5a, starting at the test specimen and moving clockwise: The test specimen (coupon) is mounted in the grips of the 100 kN MTS load cell and an AE sensor is clamped to the coupon at the “bottom” end, as depicted in Figure 4.5b. The sensor is connected to the Vallen AMSY-6 system via a 34 [dB] amplifier; the AE acquisition system in turn is controlled and configured by a dedicated workstation, on which the AE data is also stored. The VIC3D DIC system is set-up with two cameras to register the displacement field on the front of the coupon (see DIC pattern in Figure 4.5b). An LED matrix light is set-up to provide consistent lighting of the coupon surface during testing. The DIC workstation is connected to the load cell system by a trigger, allowing the MTS workstation to trigger a set DIC images every 500 cycles in the fatigue tests. Note finally the $0-10$V analog load- and displacement signal in Figure 4.5a, which connects from the MTS load cell to both the DIC- and the AE workstation. Most notably, this allows load-displacement data to also be recorded with the AE data, enabling the correlation of AE hits with the fatigue cycle phase in which they occur. This correlation is discussed in detail in Section 4.3.2. It is found that care must be taken to set the gain on these analog outputs, so that the range of $0-10$V is effectively used; when signal voltages below roughly 0.1V are used, the signals are found to be dominated more and more by noise. It is further recognised that noise is introduced both in the control system of the load cell (introducing noise in the actual noise applied) and in the analog signal logged by the DIC system. It is found for the fatigue tests that the error in the load applied as logged by the MTS load cell can typically be $0.05-1\%$; whereas the analog signal recorded by the DIC system can be up to 15% for F_{min}, but typically 1.5% for F_{max}. For this reason, nominal load values, rather than analog readings, are used in the strain computations in Section 4.3.1.

From Figure 4.5a, an issue can be brought to attention: In hindsight, the current set-up is missing a clock synchronisation. While the communication of analog data between MTS load cell and both AE and DIC workstations is found to be highly functional for subsequent analysis, the lack of a clock synchronisation limits effective an effective comparison between DIC- and

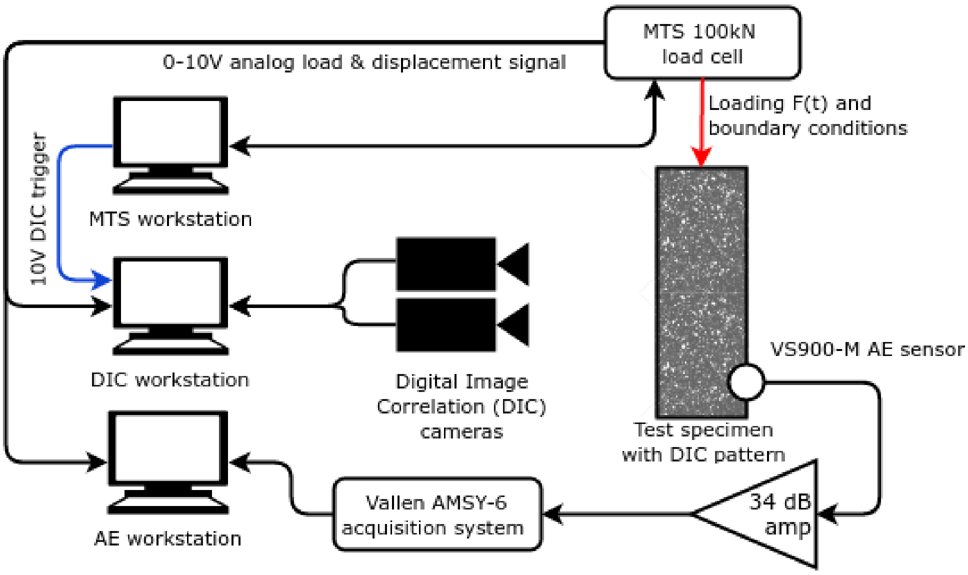

(a)

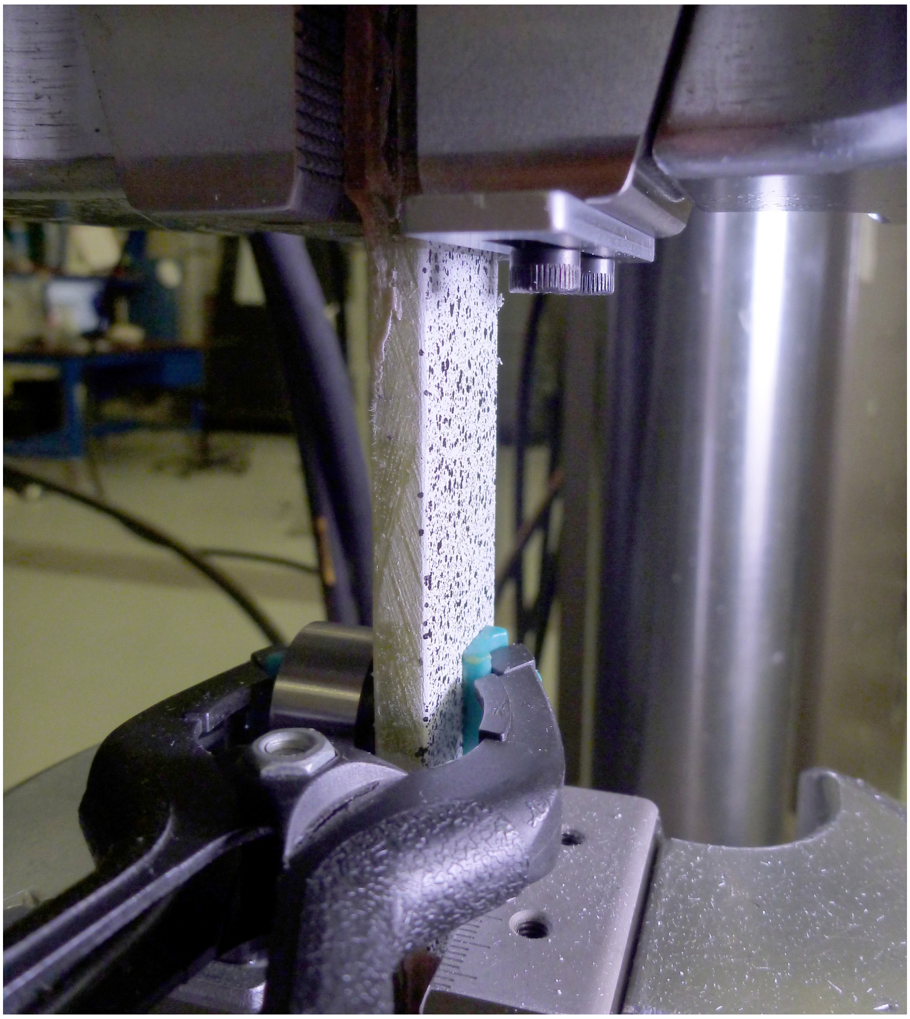

(b)

Figure 4.5: **(a)** Test set-up depicted schematically, and **(b)** close-up of a coupon loaded in the load cell, with DIC pattern on the front and clamped AE sensor on the rear of the specimen.

AE data somewhat. This issue is also illustrated in Section 5.4. Though it may be possible to alleviate a lack of synchronisation in post-processing, it is recommended for future tests that this issue is accommodated for from the beginning.

In addition to the general set-up described above, the test set-up of both static- and fatigue tests is prepared with the following steps: First, the coupon to be tested is degreased using ethanol and painted with a DIC speckle pattern using spray paint, following the common procedure used in the Delft Aerospace Structures and Materials Laboratory. The VIC3D DIC system is set-up and calibrated, first focusing the cameras on a coupon loaded in the load cell, then using a calibration pattern (typ. 14×10 pattern with 7 [mm] pitch). At least 15 images are shot of the calibration pattern in different orientations, after which the software calibration is run and outliers are discarded. A typical value found for the software's reported "projection error" is found to be 0.05, while typical in-plane strains of $O(10^{-4})$ are found when sampling 10 images of an undeformed coupon. The AE acquisition is set up in the following manner: the threshold is set such that no events are recorded while the sample is fixed in the load cell clamps (with zero applied load). Generally this is the case for a threshold of 57 [dB]. The rest of the AE parameters are determined empirically, based on Hsu-Nielsen sources and inspecting the recorded waveforms. A sample rate of 2 [MHz] is chosen, with a pretrigger length of 400 samples, for a total page length of 2048 samples (i.e. approx. 1 [ms] per recorded hit). A band-pass filter is present at 95-850 [kHz]. The re-arming time (recording cool-down) is 3 [μs]. After this set-up is completed, the static- or fatigue tests are described by their respective test procedures in the following two sections.

4.2.2 Static tests

After the test set-up is prepared as described in Section 4.2.1, a static tensile test consists of the following step-wise procedure:

1. The coupon is mounted in the MTS load cell:
   i. The coupon is secured in the bottom and top clamps, using the guiding implements on the clamps to ensure straight mounting.
   ii. The clamps are closed at $160 - 180$ [bar] hydraulic pressure.
   iii. The load cell is manipulated to apply 0 [kN] load, then control is switched to displacement and the measured displacement is nulled.
   iv. The VS900-M AE sensor is clamped to the coupon as depicted in Figure 4.5b. Apiezon M grease is used as the coupling agent.
2. A dark-coloured screen is placed behind the load cell, aiming to reduce sources of inconsistent lighting and glare for the DIC cameras.
3. DIC initial set-up and accuracy is checked by taking 10 pictures of the undeformed coupon and calculating strains.
4. The analog load-displacement signals are checked by inspecting the live read-out at the DIC workstation.
5. The MTS load cell procedure is configured: displacement controlled loading at 1.0 [mm/min]
6. AE set-up is verified:

i. Check for correct settings (see Section 4.2.1)
ii. AE recording is activated and three pencil lead breaks are performed on the rear end of the coupon (top, centre and bottom end of the gauge length) to verify that the system is correctly recording events.
iii. Recording is left running for $1 - 2$ minutes to validate if the threshold setting does not record any hits other than the pencil lead breaks.

7. Run the static test:

    i. Start DIC timed capture at once per 2000 [ms] and wait for about five pictures to be taken.
    ii. Start the MTS test procedure.
    iii. Check if no significant slip occurs.
    iv. Halt the MTS test procedure when load approaches zero (i.e. stiffness approaches zero).
    v. Manipulate the load cell displacement to zero, halt DIC- and AE capture.

The results of static testing are summarized in Figure 4.6 and Table 4.2. In addition, Figure 4.7 shows representative failure modes of statically tested coupons. Especially S90 and S090 types tend to fail near the clamped zone as depicted in Figure 4.7(a), while coupons of the D45 type tend to fail more towards the centre of the gauge length. Note also that the coefficient of variation C_V is higher the more a coupon type's failure is governed by failure of the matrix. From the load-displacement curves it can be seen that the tensile strength of S90 and S090 is higher than that of S45 by roughly a factor of 4, while the ductility is higher for coupons of type S45 than that of S90 and S090 by roughy a factor of 2. However, since these static tests serve mainly to establish a reference UTS for the coupon types, results are not discussed in further detail here.

Do note, however, test S5.12 in Figure 4.6c, which is not included in Table 4.2. Test S5.12 is tested 3 months after the other coupons in S45, indicating that the resin is still post-curing at room temperature after production. This can be seen most notably from the increased strength of S5.12 in Figure 4.6c. And indeed, the 10.2 [kN] UTS of S5.12 is 7% higher ($p = 2.5 \cdot 10^{-4}$) than the average of the S45 population. The physical significance of this curing effect is not the focus of the current project, but it is expected to have an effect on acoustic activity, which is shortly discussed in Section 5.1. It is further of note that in the static tensile tests of set S090 (Figure 4.6a), either some slipping in the grips occurs, or the load/displacement signal contains more noise than in other tests. Regardless, since the variance between tests is low and the main objective of these tests is the UTS determination, this is not considered a cause for concern.

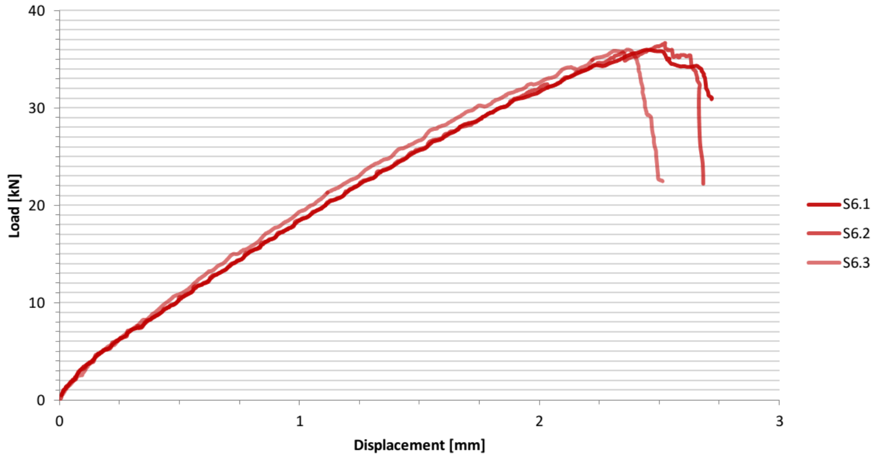

(a)

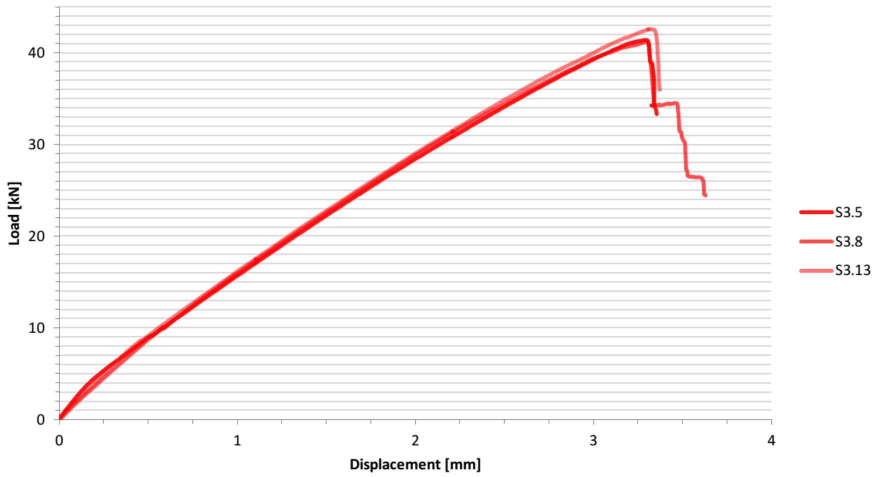

(b)

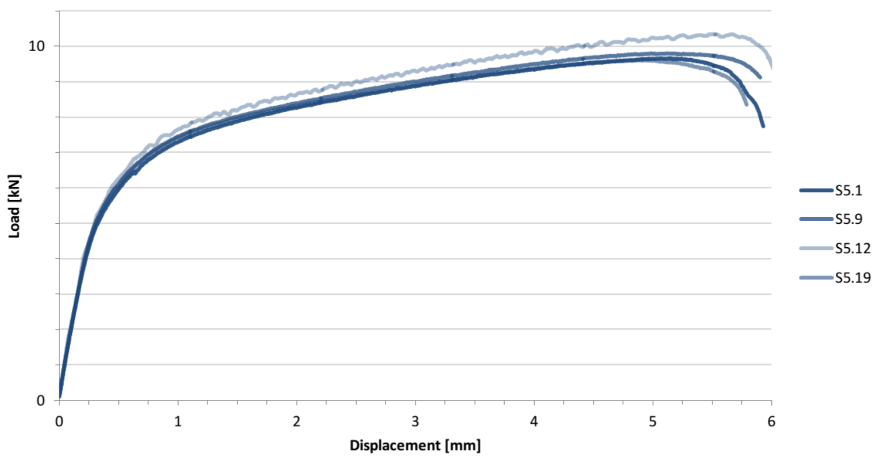

(c)

Figure 4.6: Force-displacement curves for the static tests, taken from logged MTS data.**(a)** Set S090 with layup $[0/90_3]_s$, **(b)** Set S90 with layup $[90_3/0]_s$; and **(c)** set S45 with layup $[\pm45_4]_s$.

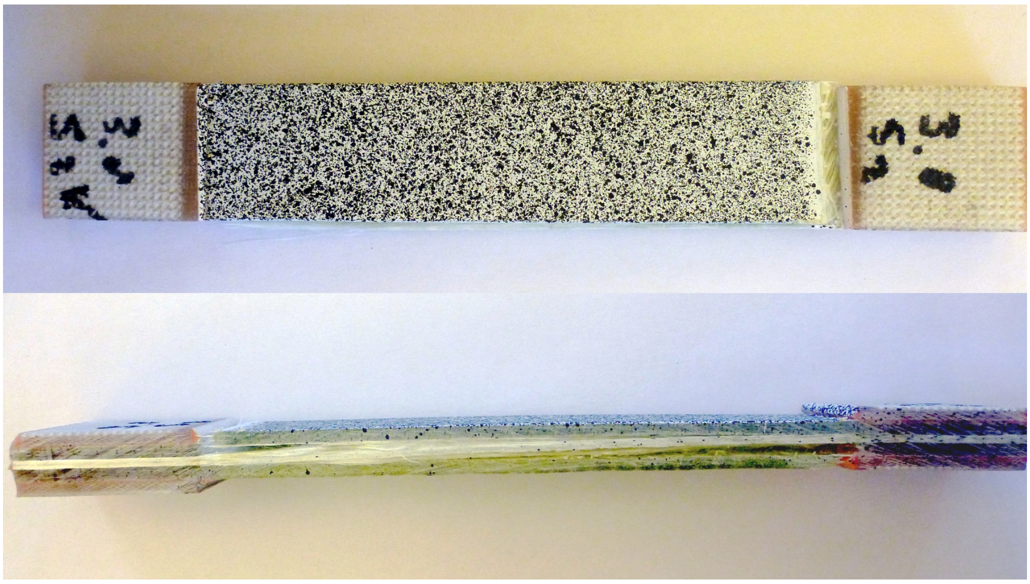

(a)

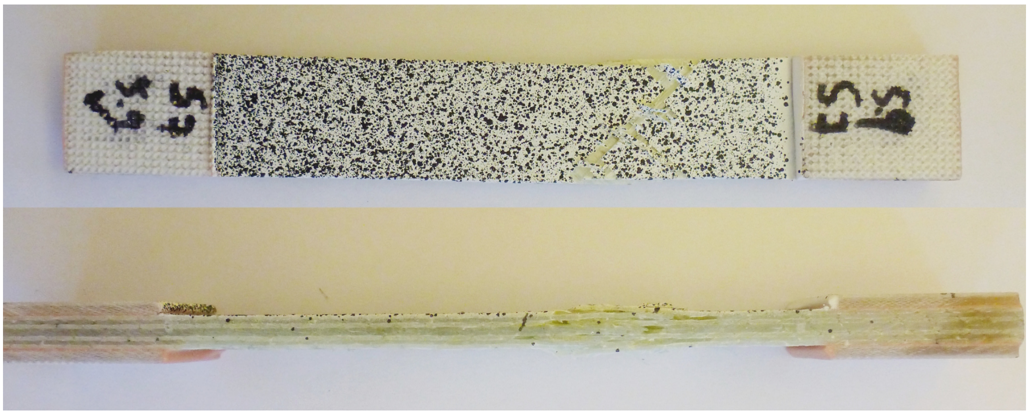

(b)

Figure 4.7: Front and side view of coupons in representative failure mode of **(a)** S90 and **(b)** S45 type coupons after static tensile testing until failure.

Table 4.2: Summary of the failure strengths of static tensile tests, determined from logged MTS data.

Set Name	Coupon layup	Test ID		Sample UTS [kN]
S090	$[0/90_3]_s$	S6.1		36.0
		S6.2		36.7
		S6.3		36.2
			μ	36.3
			C_V	0.009
S90	$[90_3/0]_s$	S3.5		41.4
		S3.8		41.2
		S3.13		42.6
			μ	41.7
			C_V	0.018
S45	$[\pm 45_4]_s$	S5.1		9.67
		S5.9		9.81
		S5.19		9.61
			μ	9.7
			C_V	0.10

4.2.3 Fatigue tests

After the test set-up is prepared as described in Section 4.2.1, a fatigue test consists of the following step-wise procedure:

1. Steps 1-4 are performed identical to those in Section 4.2.2.

5 The gain levels on the load-displacement signals are set to the highest power of 10 where both values are still expected to stay within 10 V (typically both gains are set to 10).

2. The MTS load cell procedure is configured, setting values F_{min}, F_{max} appropriately for the particular test. For all fatigue tests, $R = 0.1$. The fatigue spectrum is run under load control at 3 [Hz], with a DIC capture cycle every 500 cycles, as illustrated by Figure 4.8.

3. AE set-up is verified, identically to step 6 in Section 4.2.2.

4. Run the fatigue test:

   i. Arm DIC triggered capture (single shot) and test response by sending one trigger pulse from the MTS workstation manually.
   ii. Start the MTS test procedure.
   iii. Check if no significant slip occurs and the error of the measured load cell response relative to the driving signal is low.
   iv. Halt the MTS test procedure after sufficient fatigue cycles are completed or the coupon has failed.
   v. Manipulate the load cell displacement to zero, halt DIC- and AE capture.

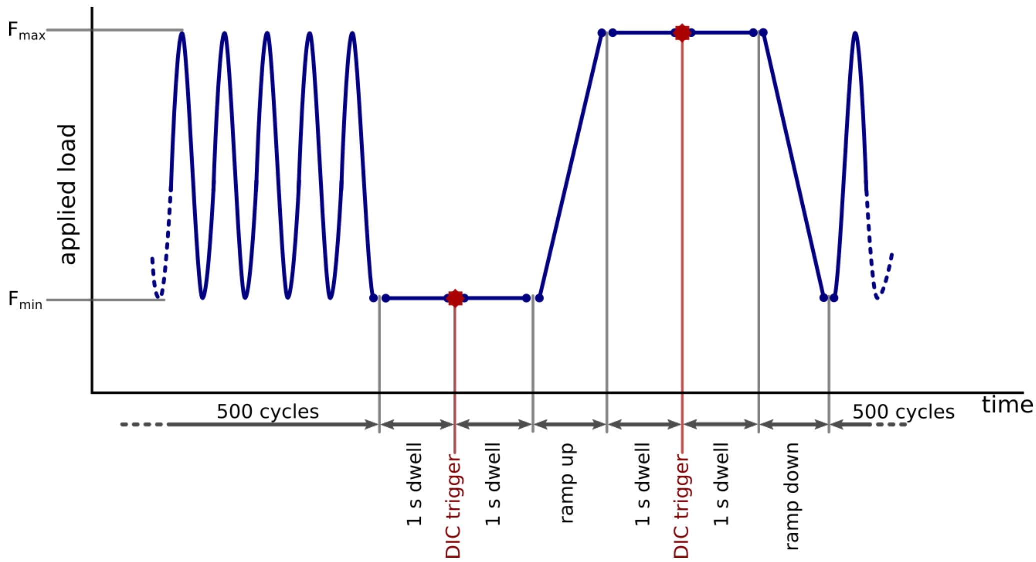

Figure 4.8: Schematic depiction of the fatigue loading procedure. 500 cycles are alternated with a pattern of holding the load at F_{min} and F_{max} to capture DIC images. R= 0.1 for all tests, so that in the current project a fatigue spectrum is defined only by F_{max}; which is set relative to a coupon type's UTS.

For the initial set of fatigue tests, the maximum load in the fatigue cycle is chosen relative to the UTS values determined in the previous section. A fatigue load level of 35% is chosen based on the discussion in Section 3.1, where load level in this project is understood as F_{max}/F_{UTS}, where F_{max} the maximum fatigue load and F_{UTS} the ultimate tensile strength. Sets D90 and D45 are therefore loaded in fatigue at a load level of 35%. However, it is found that at this level, set D90 coupons fail in relatively few cycles compared to coupons of set D45, which do not fail at all up to 700 kcycles. More importantly, AE activity is an order of magnitude higher in the D90 coupons compared to the D45 coupons, complicating analysis with the risk of significant AE sensor saturation. It is therefore decided to produce the S090 coupons, which, with an outer 0° layer instead of three 90° ones, may suffer less from the stress concentration at the tabs. This is confirmed to some degree by the increased strength of these coupons, established in the previous section. In addition, a lower load level of 11% is selected for the D090 tests, based on a static strain similarity with the 35% loaded D45 set. Finally, to investigate sensitivity of acoustic activity to the load level, a final fatigue category is established by cycling two D45 coupons at a 45% load level. Table 5.2 in Chapter 5 provides a full summary of the fatigue tests and their loading parameters.

4.3 Data processing procedures

After the AE, DIC and MTS data is collected from the fatigue tests, the data is filtered and processed so that the load data can be appended to AE data; and stiffness and stiffness degradation metrics can be extracted from DIC data. These procedures of data extraction, filtering and processing are described in this section.

Table 4.3 lists the main software packages used in the project. Data processing, filtering and analysis is performed mainly making use of open source Python libraries such as Numpy [89] (scientific computing), Pandas [90] (data handling and -analysis) and Scikit-learn [91] (data analysis and machine learning). Recording of AE and DIC data, as well as computation of (average) strains from DIC data is done through their respective proprietary software packages.

An overview of the input- and output products of the data processing steps is depicted in Figure 4.9. It shows how two main products are distilled from the AE and DIC data, allowing comparison of AE activity to the relative stiffness degradation (damage parameter) of a test specimen. The DIC data is processed first by calculating an average field strain in the center of the coupon surface, relative to the unloaded state. From this average strain, the fatigue modulus and the stiffness degradation can be computed. This processing procedure is detailed in Section 4.3.1.

The second data processing line concerns the AE data, which consists of the database of recorded waveforms (Vallen "tradb" data) and the database of parameterized hits ("pridb" data). The latter contains the waveform features of amplitude, energy, rise time, counts, duration and Root Mean Square (RMS) noise for each recorded waveforms, as well as the load and displacement measured by the MTS load cell, which can be correlated to the instant a hit is recorded. The partial power of five intervals in the frequency domain of the AE hits are appended as additional AE features derived from the FFT of the recorded waveforms. Additionally, the AE data processing includes the labelling of AE hits by the phase of the cycle they occur in. These processing procedures are detailed in Section 4.3.2. Analysis of the test data and discussion of its results is presented in Chapter 5. Except where noted otherwise, the data discussed in this document is preprocessed with the filtering and processing steps described in the sections below.

4.3.1 Digital Image Correlation data processing

The DIC data provides an approximation of coupon strain as a function of time. From the strain values, the fatigue modulus is determined, which is in turn used to determine the damage parameter. This section describes each of these steps.

First, the DIC data is processed using the VIC3D software. Surface strains are calculated relative to the unloaded condition for an area in the centre of the coupon face, covering at least 50% of the gauge length. The output strains used for analysis are then the averaged longitudinal surface strain over the area, each data point representing the average surface strain in the coupon alternatingly at F_{max} and F_{min} of the fatigue spectrum.

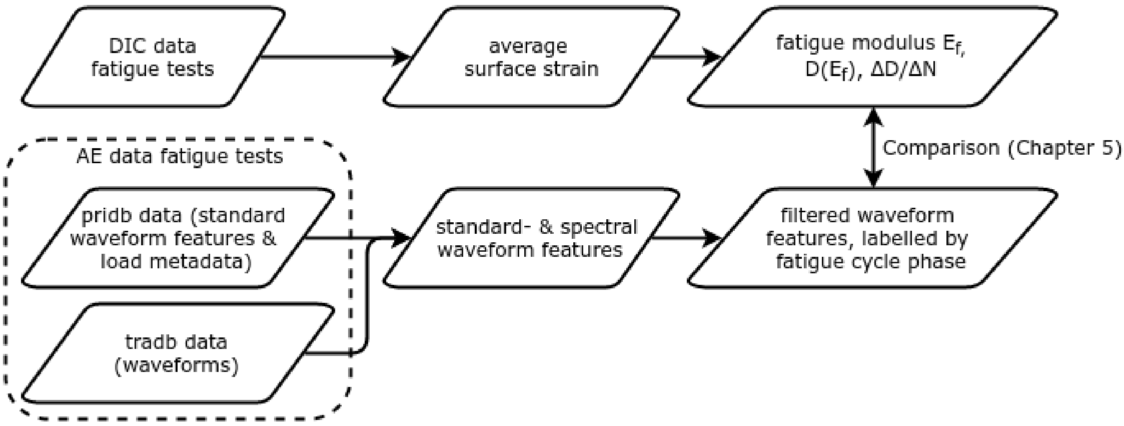

Figure 4.9: Schematic of the data processing scheme of both DIC and AE data, simplified to input- and output products only. Two processing 'lines' are discerned: (top row) the extraction of surface strain and subsequently the fatigue modulus, damage and damage parameter derivative from DIC data; and (bottom row) the processing of AE data, such that the cycle phase label and the spectral power features are appended to the AE waveform feature data.

Table 4.3: List of main software packages used for data processing. [a]Most important packages are listed. Anaconda's full package list can be found at http://docs.continuum.io/anaconda/pkg-docs

Software/package	Version
Vallen AE system software	R2014.0414.1
Correlated Solutions VIC3D DIC software	7
Pycharm professional IDE	2016.1
Anaconda[a] (Python distribution)	4.0.0 (x64)
Python	2.7.11
Cython	0.23.4
Numpy	1.10.4
Scipy	0.17.0
Pandas	0.18.0
Sqlite	3.9.2
Scikit-learn	0.17.1
Seaborn	0.7.0
Matplotlib	1.5.1

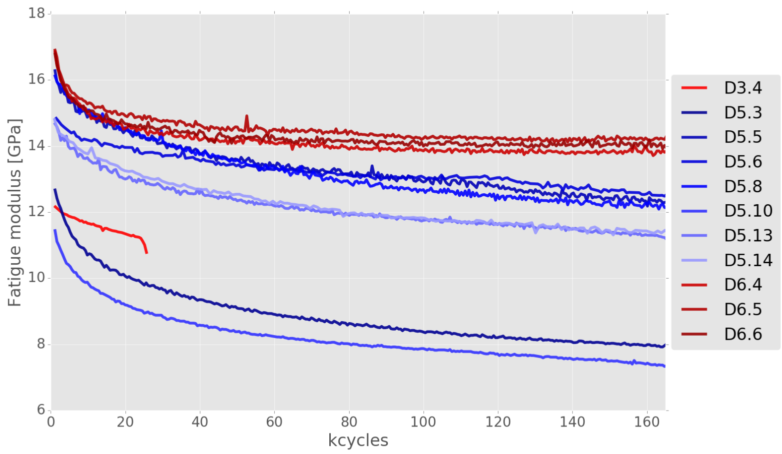

Figure 4.10: Fatigue modulus determined from raw DIC average surface strain data, for each of the tests under consideration.

From these surface strains, the fatigue modulus [92] is calculated. From material mechanics, the general stress-strain relation is given as $\sigma = E\varepsilon$, where σ stress, E stiffness and ε strain. So, if the surface strains calculated from DIC data are assumed to represent engineering strains, the fatigue modulus E_f can be defined by Eq. 4.1:

$$E_f = \frac{F_{max}}{A}\frac{1}{\varepsilon_{yy}} \tag{4.1}$$

Here, F_{max} is taken to be the fatigue cycle load maximum as set in the MTS program, ε_{yy} is the average DIC surface strain in the coupon longitudinal direction at the time of application of F_{max}; and A is the measured average cross-sectional area of the coupon. It is noted that ε_{yy} does not represent the engineering strain in the coupon, contrary to its usage here. Moreover, it represents a surface strain in the speckle pattern paint layer which may not be an accurate representation of the (engineering) strain in the sample. However, since in the current study the main use of this data is to compare among test samples, it is asserted that these simplifications can safely be used. The unfiltered fatigue modulus calculated for each of the fatigue coupons under consideration is depicted in Figure 4.10. Note the outlying behaviour of test D3.4, which is due to coupons in the D3 series failing after only 25-39 [kcycles]. Some differences between the different groups of tests can already be seen from this figure. For analysis of the damage parameter derived from the DIC data refer to the discussion in Section 5.1.

The damage parameter D is determined based on the fatigue modulus, as defined by Eq. 4.2 [92]. In this equation, E_0 is the initial stiffness, which is taken to be the coupon's highest (i.e. first) fatigue modulus. It is of note that the damage parameter is misleadingly named: it does not directly represent damage in the coupon - especially when multiple damage mechanisms are at play, but is primarily a normalized measure of stiffness degradation which in turn is definitely caused by damage.

$$D = 1 - E_f/E_0 \tag{4.2}$$

The most important limitation of this metric as a measure of fatigue damage in a composite is that it poorly represents the occurrence of multiple damage mechanisms, especially when they interact [92]. However, for the case of low load level fatigue testing in the current project, the test coupons can be expected to be dominated by only one damage mechanism: matrix cracking. Consequentially, and confirmed by analysis in Chapter 5, the damage parameter provides a relatively effective measure to validate the data obtained through AE monitoring.

4.3.2 Acoustic Emission data processing

Processing of the generated AE data aims to fulfil a series of purposes. First, the data need to be inspected: AE recording is manually initiated and ended, so in order to distil only the AE hits that occur during fatigue loading, the data must be truncated based on manually set time bounds. Additionally, some hits are considered outliers by inspection and are filtered based on their energy, amplitude or duration features. Furthermore, the "PA0" and "PA1" columns in the database file containing AE feature-described hits are also manually checked for the right conversion factors from Volts to [kN] or [mm] respectively. The second purpose of AE data processing in the current project is to label hits by the phase segment of the fatigue cycle they occur in. This is an automated process based in the aforementioned load (here: "PA0") data. Thirdly, the set of features natively used by the Vallen AE system software to describe AE hits is expanded by adding five partial spectral power features. Finally, processed AE data is saved per test to a serialized binary format, referred to as "pickling" [93] in the Python programming environment, so that it may be efficiently recalled by subsequent analysis steps. The processing steps described here are schematically depicted in Figure 4.11 and described in technical detail in the sections below. The Python code for the data retrieval and -filtering functions is hosted at `https://github.com/Unpluralized/PyAE`.

Data retrieval and inspection

Data from the Vallen system is organized into two file types: a file containing the recorded transient waveforms (tradb file) and a file containing the feature-described AE hits (pridb file). The feature set calculated by the Vallen system in the current project consists of six features: amplitude, energy, duration, counts, rise time and RMS noise. Both the pridb and tradb files are an instance of the popular database format SQLite, enabling retrieval and manipulation of the AE data by a variety of software. Amplitudes are converted to $[dB]$ as $Amp[dB] = 20\log_{10}(Amp[\mu V])$, to be consistent with their representation in the Vallen AE software.

Referring to the work flow depicted in Figure 4.11, the AE data of each test is inspected manually to determine filter parameters, which are then stored in a filter configuration file so they may be readily recalled. The types of filters mentioned here is non-exhaustive; only the filters that are ultimately used in service of the analysis in Chapter 5 are depicted. First, the time bounds of the test are determined by inspection: start and end of the test are decided relative to the loading start and end, with an added error margin of approximately a minute or two. Second, it is determined whether an upper bound for the amplitude or energy is needed to filter extreme outliers. This filter is only applied very sparingly to filter hits with energies of 10^6 [eu] and higher[1]. For all tests, all hits with duration < 1 [μs] are filtered out. These hits are found to constitute a very sizeable fraction of the total number of hits ($18 - 94\%$), while they can only be poorly described by the waveform features. Lastly, the conversion factor for

[1]Note: *eu* may be an obscure unit to the reader; it is merely an arbitrary *energy unit* defined in the Vallen AE system, resulting from numerical integration of the recorded signal (μV) over time (μs). Recall that the recorded waveform and even more so its parametrisation is only a representation (a consequence, if this term may be preferred) of the original acoustic energy release. Therefore, the nominal value for the energy of a waveform is of very limited significance anyway.

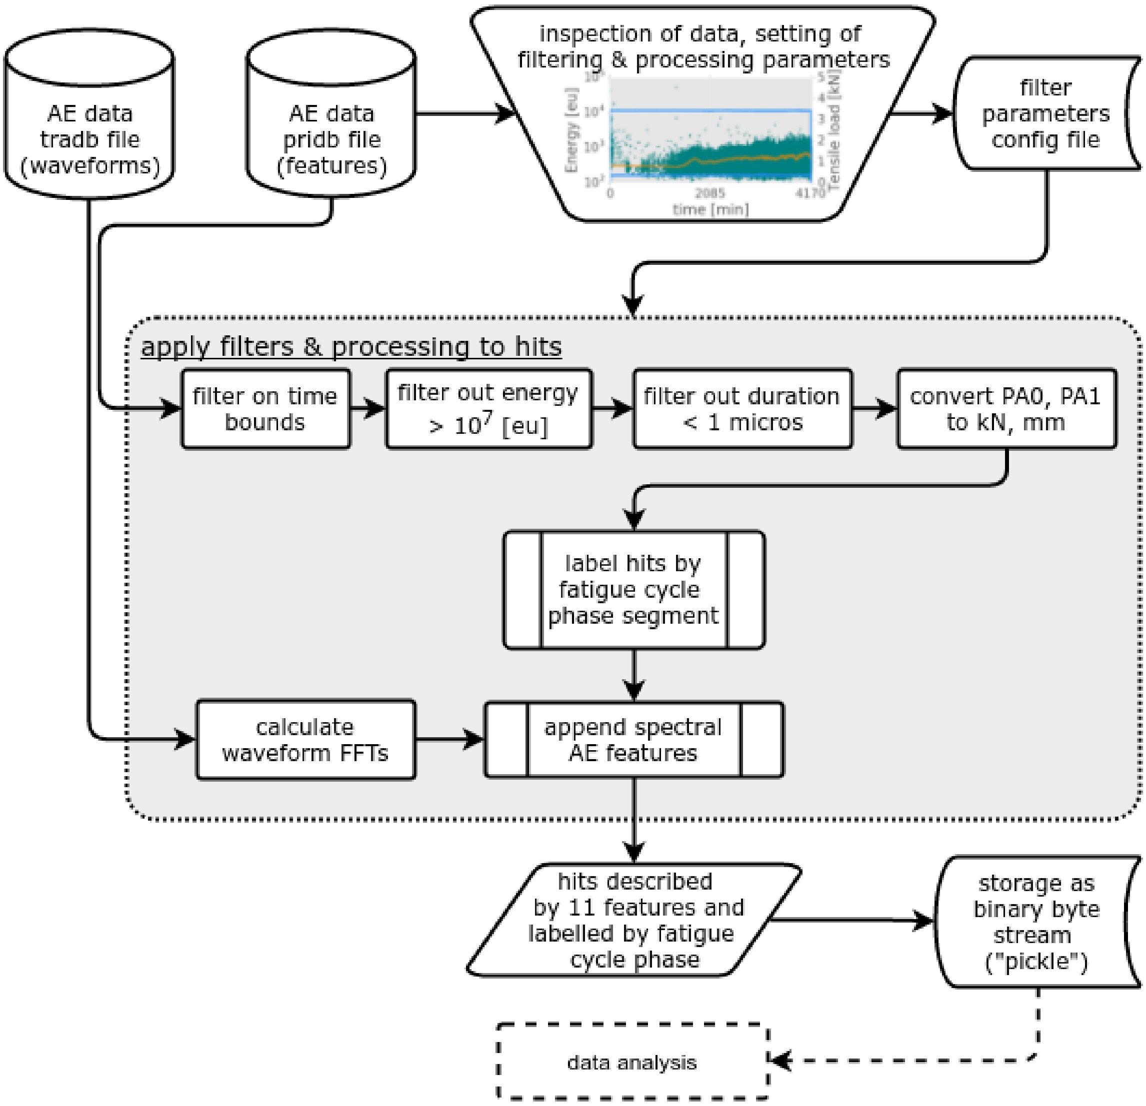

Figure 4.11: Flowchart depicting AE filtering and processing procedures.

the load- and displacement signal is set, by correlating it with equivalent MTS log data. This step can be done by inspection because only powers of 10 (i.e. 1, 10, 100) are used for the signal gain. The reason this gain value may vary between tests is because the gain is set as high as possible to reduce noise on the signal, while not being so high that the maximum expected load or displacement exceeds the analog signal scale of $0 - 10$ [V]. This "bug" in fact did occur for test D3.4, invalidating the load data above 10 [kN].

Appending the cycle phase as hit metadata

Because the database of feature-described AE hits also contains load information, it is possible to correlate each AE hit with the instantaneous load applied in the fatigue cycle, adding it as "metadata" to the hits. To this end, note that the Vallen AE system logs the analog input data representing the load and displacement to the pridb file with a frequency of 20 [Hz], but leaves this information as separate row entries from the AE hits. Therefore, to label the hits with the approximate instantaneous load and displacement a linear interpolation function native to the Pandas library [94] is applied to the load and displacement columns, using the time column as the index. This gives a reasonably accurate approximation of the instantaneous load and displacement, because at a fatigue frequency of 3 [Hz], this means there are still 3-4 load- and displacement measurements between each peak and valley of the fatigue cycle. The difference in load between each set of data points, ΔF, is now also calculated by means of Numpy's `gradient` function [94]. The accuracy of this "derivative" is low, but it does not need to be very accurate, as the next paragraph will show.

The fatigue cycle metadata can be simplified from continuous but not very accurate F and ΔF values to a more robust and discrete four-phase labelling. The four-phase labelling procedure is illustrated by Figure 4.12. Because the fatigue cycle is a sine wave, it can be divided into four equal parts based on the instantaneous load metadata (Figure 4.12a). In other words, the labelling boundaries are at $|\sin\theta| = |\cos\theta|$, which gives $\theta = \frac{\pi}{4} + n\frac{\pi}{2}$ (where $n \in \mathbb{Z}$) and therefore divides the sinusoidal phase in four equal segments: "up" (orange), "top" (red), "down" (light blue) and "bottom" (dark blue). "Top" and "bottom" are distinguished based on $F_{mean} \pm \frac{1}{2}\sqrt{2}F_a$, while a distinction between "up" and "down" cycle segments can be made based on the ΔF calculated in the previous paragraph. Finally note that by following the labelling decision tree in Figure 4.12b, ΔF in fact only needs to provide the sign (negative or positive) of the instantaneous load for labelling to be accurate. In subsequent analysis presented in Chapter 5, the labelling described here is found to provide a meaningful filter by providing a constitutively informed classification of hits.

Generating and appending spectral partial power features

To augment the set of AE features used to describe the AE hits in the analysis, five spectral power fraction ranges are derived as extra features. The frequency ranges are chosen somewhat arbitrarily based on prior use in work by Fotouhi [95]. They are $95 - 150$ [kHz], $150 - 250$ [kHz], $250 - 350$ [kHz], $350 - 500$ [kHz] and $500 - 850$ [kHz]. The FFT is computed using the `rfft` function in the Numpy library [94] and results are validated by comparing against FFT spectra of selected hits computed by the Vallen AE software. The five spectral features are then computed as fractions of the "full" spectrum between $95 - 850$ [kHz].

SQLite and its implementation in Python are generally fast, especially when only particular selections of a database are requested from the file. However, loading the transient waveforms embedded as "binary blobs" in the tradb files is found to be one of three most time-consuming

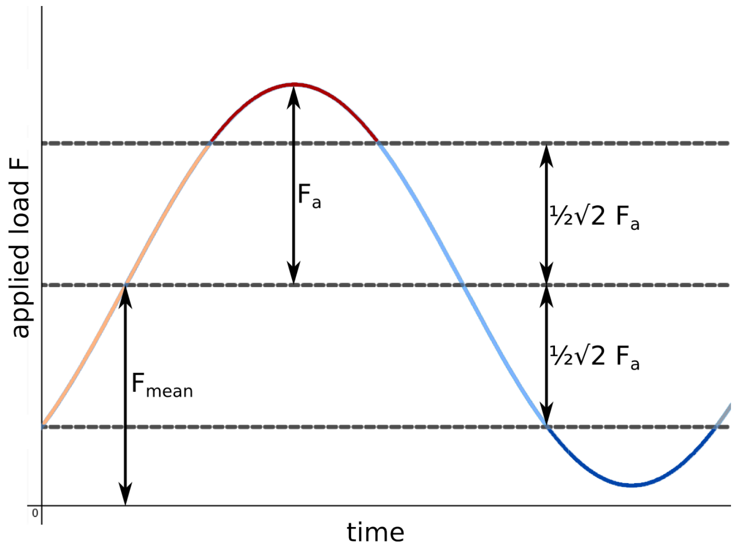

(a)

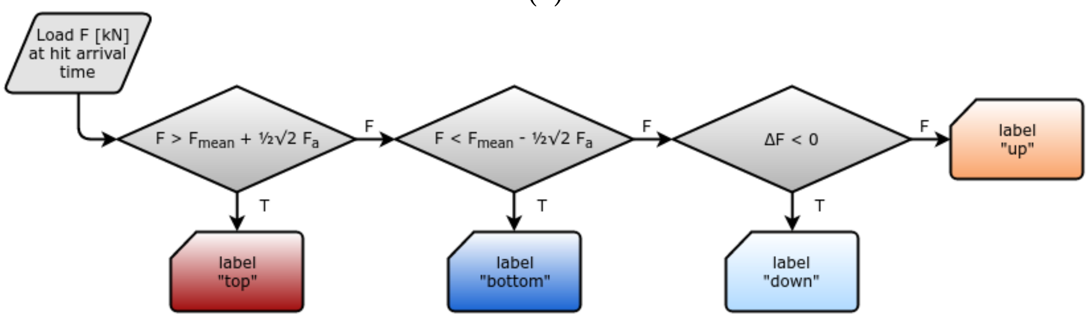

(b)

Figure 4.12: Labelling of hits by the fatigue cycle phase segment in which they occur, illustrated by **(a)** a single arbitrary fatigue cycle with relevant definitions and cycle segments coloured consistently with the labels in figure (b); and **(b)** the decision tree used by the hit labelling algorithm, where T=True, F=False.

```python
import pandas as pd
import numpy as np
import sqlite3

conn = sqlite3.connect(path_to_tradb_file)
ae_waveforms = pd.read_sql("SELECT Data FROM view_tr_data WHERE TRAI > 0",
        conn,
        coerce_float=False)
ae_waveforms = ae_waveforms["Data"].apply(
        lambda x: np.frombuffer(x, dtype=np.short))
```

Listing 4.1: Python (and a single line of SQL) code snippet to efficiently load all waveforms from a tradb file. After this operation, `ae_waveforms` is a Numpy array with a waveform on each row.

data processing steps - the other two being the computation of the waveform FFT and partial powers themselves, and the cycle-phase labelling of hits. However, using the method detailed in Listing 4.1 below, this computational step is found to be a factor of at least 200 faster (nominally 10^6 waveforms in 10 [s]) than any other tried method.

A second procedure for which performance is found to be very sensitive to implementation is the computation of partial powers and their insertion into the table of hit features. This procedure is sped up from a computation time of 16.7 hours for 10^6 waveforms when using Pandas' `DataFrame.apply()` function on a Pandas DataFrame (i.e. table) containing the waveform FFTs, to a computation time of just 2.5 minutes when calculating the partial power fractions first to a list of Numpy arrays, before appending the newly calculated spectral feature columns to the AE hit features DataFrame. In other words, calculation in-place in a DataFrame is found to be much slower than first calculating the features separately and then appending them to the table of AE hits as feature columns in a single step.

This concludes the description of the AE data processing steps as presented in Figure 4.11. After the steps described here are completed for each test, they are saved to a serialized binary format ("pickling" [93]) so that they can be quickly recalled for any analysis steps. In the next chapter, the analysis results are discussed. Except where noted otherwise, the data discussed in Chapter 5 is subjected to the filtering and processing described in these past sections.

Chapter 5

Discussion of results

Findings arising from data analysis of the acoustic emission data and DIC measurements are discussed in this chapter. The sections in this chapter are roughly organised from high-level description of the data to lower-level descriptions. First, Section 5.1 summarizes all static and fatigue tests and their organisation into four sets, as well as presenting an overview of the most relevant information obtained from the DIC measurements. A few general notes on the test results are also discussed. Section 5.2 presents the outcome of- and empirical basis for the labelling of hits by the phase of the fatigue cycle they occur in (procedurally described in Section 4.3). Thereafter, the test types are compared from the perspective of basic statistical descriptors such as mean, median and variance in Section 5.3. The spread and clustering in the feature-space is evaluated and discussed. Finally, Section 5.4 investigates trends in the AE data in more detail from the perspective of energy. The existence of a correlation between the stiffness degradation and AE hits occurring in different cycle phases is evaluated, thereby also validating some aspects of the spread and choice of clustering of AE data. Conclusions to the findings discussed in this chapter are presented in Chapter 6.

In the analysis of AE data various types of noise may be found. It is recognised that noise is in the eye of the beholder: whether a source classifies as noise depends on the interests and goals of the operator. For example, AE activity arising from friction between internally fractured surfaces in the coupon material may be regarded as noise from the perspective of strictly detecting damage mechanisms; however, for the purposes of structural integrity monitoring, this activity may contribute valuable information. For the purposes of discussion in the current analysis, three types of noise are defined, based their locations of origin: (1) noise external to the coupon, (2) noise external to the coupon material; and (3) noise internal to the coupon material. Noise in general then is considered any AE activity that does not originate from a damage mechanism in the coupon material. Noise type (1) constitutes any noise sources that do not originate in the tabbed coupon or its interfaces. This noise source includes, for example, load cell machinery noise, but does not include friction between the tabs and the load cell clamps, damage in the tabs, or friction inside the coupon itself. Noise type (2) constitutes noise occurring outside the E-glass/vinyl-ester coupon material. This includes the noise from type (1), but also all sources occurring in the tabs or their interfaces, such as damage in the tab-coupon adhesive bond. Finally, noise type (3) constitutes noise that occurs strictly within the E-glass/vinyl-ester coupon material. This may include AE activity from friction between internal fractured surfaces.

5.1 Tests summary and Digital Image Correlation results

A summary of all static and fatigue tests included in the discussion is given in Tables 5.1 and 5.2 respectively. Examining the table of fatigue tests, it is noted that the time between production and testing seems to have an influence on the acoustic activity: compare coupons D5.3, D5.5, D5.6, D5.8 (tested 4 months after production) to coupons D5.10, D5.13 (tested 2 months after production). This effect may be attributed to room temperature curing of the resin, particularly since save for the curing cycle for the coupon tab adhesive, no post-curing has been applied to the coupons. Indeed, it is confirmed by test S5.12 (Table 5.1) that over a period of 3 months, the ultimate strength during static testing of a single coupon is found to be increased by 7% relative to the initial set of static tests. This value is 7.2σ above the mean of the initial set of static tests ($p = 2.5 \cdot 10^{-4}$). Analogously, coupon static stiffness in the 0-3.4 [kN] load region is increased by 6%. It is theorized that both the increase in strength and stiffness and the decrease in acoustic activity are caused by continued cross-linking in the thermoset resin over this period of time. Interestingly, while this effect can be observed in some other studies [96], some researchers also document an opposite effect [44]. Evidently, this observation is of importance for the in-situ application of AE monitoring, at least for the beginning period of the lifetime of a structure. However, in the current analysis, this effect is not the main focus.

As a second observation, it may be noted that not all tests are run at the same AE threshold. This is a direct consequence of the chosen set-up procedure, where the threshold is chosen during the set-up as a minimum value at which no noise is being recorded by the AE system before the test starts. A dilemma surfaces when observing the distribution of ΔT between consecutive hits in Figure 5.1: noting the cut-off at $\Delta T = 3$ [ms] introduced by hit capture cool-down, it is judged that a limiting bias in the recorded acoustic activity cannot be ruled out. It is possible to remedy this bias by lowering the threshold, lowering the cool-down time after a threshold crossing; and/or reducing the length of recording per threshold crossing. Lowering the threshold parameter is undesired as it introduces extra noise to the data. Lowering the cool-down or length of recording parameters significantly may also be undesirable, because a limiting bias on the recorded acoustic activity also has the merit of filtering reflections of the same event in the material. However, an order of magnitude calculation reveals that even at the minimum ΔT of 3 [ms], a wave would travel $O(10\ [\mathrm{m}])$ in the material. This intuitively suggests reflections are less of a concern with the current parameter set, while some undesired limiting bias on acoustic activity may feasibly be expected. It is concluded that the margins set up to guard against reflections are higher than conceivably needed.

The AE threshold being not constant between all test sets has one other important consequence. It is readily apparent that all of the AE features amplitude, energy, RMS noise, rise time, counts and duration are calculated relative to the signal threshold crossing in some way or another. The effects of a different threshold setting may be cautiously called linear for some features such as amplitude or rise time, but most definitely non-linear for parameters like hit energy and hit intensity. Therefore, it is noted that for the following sections, conclusions from comparison of AE data between e.g. set D90 with other fatigue test sets must reserve an appropriate amount of caution.

Table 5.1: Descriptive summary of all static tests. [a]Specimen 5.12 has had 3 months more RT curing time than the other S45 coupons, causing its ultimate strength to be 7.2σ above the mean of the other S45 coupons. [b]Hits after filtering, between the time loading starts and load reaching its maximum.

Set name	Coupon laminate	Test ID	$F_{max,static}$ [kN]	AE threshold [dB]	Hits[b]
S090	$[0/90_3]_s$				
		S6.1	36.0	55	9681
		S6.2	36.7	55	10088
		S6.3	36.2	55	8814
		μ	36.3		
		C_V	0.009		
S90	$[90_3/0]_s$				
		S3.5	41.4	58	13906
		S3.8	41.2	55	11070
		S3.13	42.6	58	13949
		μ	41.7		
		C_V	0.018		
S45	$[\pm45_2]_s$				
		S5.1	9.7	58	32962
		S5.9	9.8	58	33691
		S5.19	9.6	58	26240
		μ	9.7		
		C_V	0.010		
		S5.12[a]	10.4	58	25147

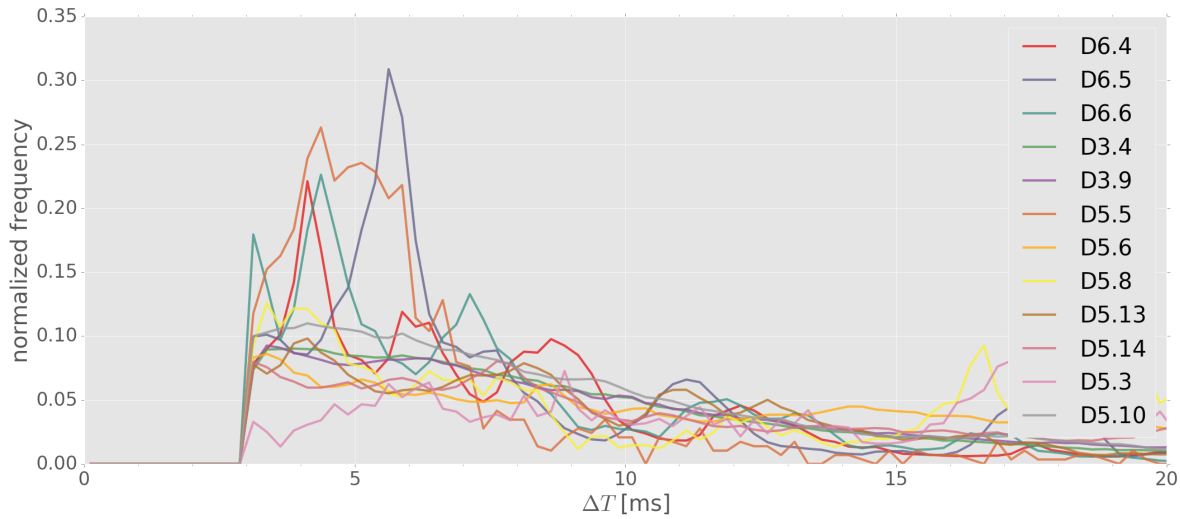

Figure 5.1: Section of the distribution of the time difference ΔT between consecutive hits. This graph depicts an unfiltered version of the dataset.

Table 5.2: Descriptive summary all of fatigue tests taken into consideration in this document.
[a]Average from static tests in the interval 0-F_{max}.
[b]The load level (= F_{max} in the fatigue spectrum) is given as a percentage of the average maximum load in the corresponding static tests (see table 5.1). Similarly, the strain level is the average engineering strain of the corresponding static tests at the load level that corresponds to F_{max} in the fatigue cycle.
[c]Hits after filtering, in roughly the first 930 minutes of each test.
[d]To failure.
[e]Alternating 5kcycles at 3Hz, .5kcycles at .5Hz, instead of 3Hz continuously.

Set name	**Coupon laminate**	**E [GPa]**[a]	$\mathbf{F_{max}}$ **[kN]**	**Load level**[b] **[%**$\mathbf{F_{max,static}}$**]**	ε**-level**[b] **[%]**	**Test ID**	**Cycles completed**	**AE threshold [dB]**	**Hits**[c]
D090	$[0/90_3]_s$	20.5	3.9	11%	0.16%	D6.4	244200	57	302512
						D6.5	239500	57	155902
						D6.6	205000	57	241925
D90	$[90_3/0]_s$	14.0	14.6	35%	0.83%	D3.4	25300[d]	57	340085
						D3.9	38700[d]	58	363035
D45L35	$[\pm 45_2]_s$	15.8	3.4	35%	0.16%	D5.5	164900	57	1013
						D5.6[e]	503200	57	129961
						D5.8	186000	57	15037
						D5.13	173700	57	87442
						D5.14	732700	57	851
D45L45	$[\pm 45_2]_s$	15.8	4.4	45%	0.23%	D5.3	232000	57	1326
						D5.10	175500	57	616669

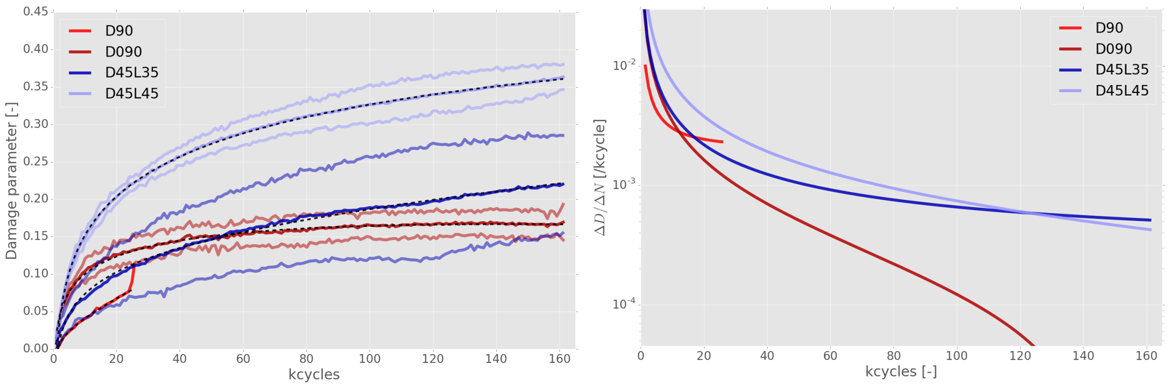

(a) **(b)**

Figure 5.2: Fatigue damage parameter and its derivative, computed from the DIC data. **(a)** Mean damage parameter and 95% confidence interval for each of the test sets. The best fit of Eq. 5.2 is shown (dashed line). Due to corruption of the DIC data for test D3.9, the D90 set only contains data from one test. **(b)** derivative of the damage parameter (analytical derivative of the best fit curve). The damage derivative of D090 drops below the y-axis minimum at 120 kcycles because of the representation of values close to zero in the log-scale.

The fatigue damage parameter and its derivative is derived from the DIC measurements (see Section 4.3) and presented in Figure 5.2. The damage parameter is repeated in Eq. 5.1, with E_f the fatigue modulus and E_0 the initial stiffness. In the current analysis, the highest fatigue modulus recorded through the DIC data is taken as E_0. To show trends in the damage parameter and $\Delta D/\Delta N$ more clearly, the damage parameter is fitted with the generalized function given in Eq. 5.2; its analytical derivative given in Eq. 5.3 and with a, b and c fitting parameters. Fitting is done using the curve fit function from the SciPy library [94]. The function in Eq. 5.3 is chosen particularly because it is found to improve the representation of the asymptotic behaviour of the damage parameter of set D090: the mean of the mean standard errors for the fit of this function is lower than for $y = ax^{\frac{1}{b}} + c$ by a factor of $2.3 \cdot 10^5$. As a final note on the damage parameter it must be brought to the attention of the reader that the metrics of D and $\Delta D/\Delta N$ are not necessarily a direct measure of damage accumulation, but primarily represent a measured stiffness degradation (and its derivative). In other words: the damage parameter measures only damage that causes a stiffness degradation.

$$D = 1 - E_f/E_0 \tag{5.1}$$

$$y(x) = a^x x^{\frac{1}{b}} + c \tag{5.2}$$

$$\frac{dy}{dx} = \frac{1}{b}(a^x x^{\frac{1}{b}-1}(bx\ln(a) + 1)) \tag{5.3}$$

From the depiction of the damage parameter in Figure 5.2, several observations are brought to the attention of the reader. Firstly, that the variance in fatigue modulus (and therefore damage parameter) in set D45L35 is relatively high. Examining the fatigue modulus of the individual tests reveals this is at least in part attributed to the room-temperature curing effect discussed earlier in this section: coupons that are tested later tend to start out at a higher stiffness. However, the derivative of the fatigue modulus remains relatively constant between D45L35 tests. A second observation lies in the comparison between sets D45L35 and D45L45: D45L45 starts out with a higher stiffness reduction per cycle, causing an ultimately more severe reduction in stiffness. However, from Figure 5.2(b) it can be seen that $\Delta D/\Delta N$ of D45L35 and D45L45 ultimately converge towards a similar degradation per cycle of approximately $5 \cdot 10^{-4}$/kcycle. Conversely and thirdly, set D090, though starting out with a relative degradation per cycle that is most comparable to set D45L45, can be seen to asymptotically grow to a constant damage (approx. 0.32) and zero (or unmeasurably small) degradation per cycle. Finally, set D90 (test D3.4), failing after 25 kcycles, counter-intuitively exhibits the least relative degradation per cycle, while sharply increasing approaching failure.

The interpretation of these degradation results is not completely unambiguous. However, (applying Occam's razor) an explanation fitting the expected behaviour of damage progression as discussed in Section 3.1 can be formulated: in set D090, damage builds up until matrix crack saturation is reached, but not enough mechanical work is supplied to progress beyond this damage mechanism. In contrast, both D45 sets continuously keep degrading. Since stiffness is continuously degrading, damage can be said to continue progressing. However, this does not definitively prove the progression from one damage mechanism to another: it is theoretically possible that for the D45 set, matrix crack saturation is not reached over the span of the fatigue tests. In addition, the observed behaviours can also be explained from the layup types: initially, there is 'slack' in the longitudinal fibres in the D090 coupon, while the resin itself carries most of of the load. Only after sufficient strain has built up in the resin, the fibres start being loaded in tension. Therefore, the strain limit of set D090 in Figure 5.2 is set by the stiffness of the fibres. In contrast, comparing sets D45, there are no fibres spanning the complete gage length of the coupon; therefore, the stiffness degradation is not limited by the fibres and is allowed to continuously progress. This result is in accordance with the comparison of matrix crack saturation under strain- versus load control in Section 3.1 as documented by (for example) Li et al. [2]. Nonetheless, the fact that stiffness degradation of the D090 tests is relatively constrained compared to the D45 samples may provide the simplest valuable piece of validation data for subsequent AE data analysis.

5.2 Hit labelling by fatigue cycle phase

Using the load-time data recorded alongside the recorded AE hits, the AE hits can be correlated to the instantaneous load applied to the coupon. The hits are labelled by cycle phase segment based on the instantaneous load and $\Delta Load$ values, as described in Section 4.3. The labelled hits are referred to as 'top'-, 'up'-, 'bottom'- or 'down'-labelled hits. Additionally, 'up'- and 'top'-labelled hits together referred to as 'loading'-labelled, while 'down'- and 'bottom'-labelled are together referred to as 'unloading'-labelled. A visualised outcome of this classification is presented in Figure 5.3. This figure depicts a roughly representative result for tests in sets D090 and D45, with the notable exception of test D5.10, which makes for a significantly more noisy classification than the other tests. That is, in test D5.10 the boundaries for this classification-by-inspection are less clear, and in setting the class boundaries it cannot be avoided to intersect large amounts of hit groupings which seem to belong together when judged by the visualisation type of Figure 5.3.

Disregarding the labels for a moment, a few observations are made on the AE activity in general, referring to the example of Figure 5.3. One type of AE activity is distinctly found throughout the test above roughly 3 kN in the fatigue cycle. This type of load activity is relatively abundant in the beginning of the test, but reduces sharply over a period in the order of 25 kcycles for all tests. At approximately 230 kcycles (in D5.3) an anomalous concentration of hit activity is observed. Conversely, in this test hits below 3 kN follow a different pattern: initially, activity is non-existent or undetectable, while it clearly grows as the test progresses. Inspection of the acoustic activity reveals that beyond roughly 150 kcycles, the unloading-labelled activity tends

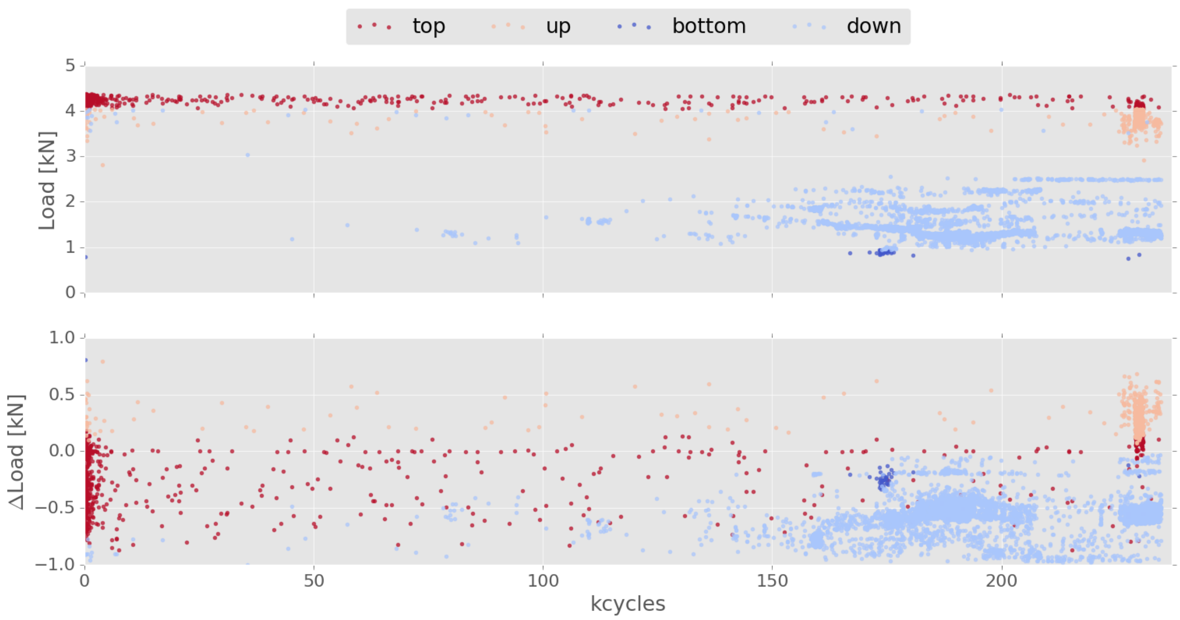

Figure 5.3: Load and ΔLoad history for hits in test D5.3. The hits are labelled based on the phase of the fatigue cycle they occur in (sinusoidal period 2π divided into four equal phase segments).

to settle at an average of $O(10^4)$ hits per 5 kcycles; this is in contrast to the ‘loading’-labelled hits, which tend to stay within the range of $O(1-100)$ hits per 5 kcycles on average. Comparing between different test types, it is further of note that in set D45L45, the load level at which top-labelled activity occurs has a higher variance compared to set D45L35. The top-labelled activity is also nominally higher and more constant throughout the test than in set D45L35, whereas no clear trend is detected for the unloading-labelled activity. Top-labelled activity of set D090 is roughly similar to set D45L35, tending to diminish as the test progresses. The latter behaviour is conceptually consistent with the Kaiser effect, where repeatedly loading the material at the same load causes less and less acoustic activity. By examining the top-labelled hits as a function of displacement instead of load, it is revealed that the felicity effect applies most notably as a function of the applied work and in the early stages of testing ($N < 50$ kcycles).

Finally, a few remarks about the characteristics of the unloading-labelled hits are brought to the attention: The unloading-labelled hits are found mainly in the bottom half of the unloading phase of the cycle. In tests or periods of tests with high down-labelled AE activity, hits start to coalesce along equally separated bands in ΔLoad (e.g. visible as 4-5 horizontal bands starting at approx. 175 kcycles in the figure). Moreover, some of these bands do not appear at a constant load, but exhibit ‘walk’ during testing: the activity starts out at one load level, but later gradually shifts to either a lower or a higher level. This phenomenon is most apparent over the longer time scale of tests D5.6 and D5.14. It is theorized that this phenomenon is not rooted in material mechanics. Rather, it is possible that this phenomenon is a consequence of hit intensity saturation of the AE system. This assertion is substantiated by the discussion of Figure 5.1.

From these findings it is asserted that the unloading-labelled and loading-labelled AE hits may come from distinctly different mechanical sources. Therefore, the following section explores the (dis)similarity of these source types by presenting a description and an analysis of the feature-space. Subsequently, Section 5.4 investigates the predictive value of the top-labelled hits in an energy-based approach.

5.3 Distribution and clustering in the feature-space

In this section, an overview is given of the statistical distribution and mutual similarity of AE data. This is done mainly from the perspective of the test sets as a whole. The most dominant clusters in AE feature-space are also presented, confirming that the fatigue cycle phase-based labelling of hits correlates with mutual (dis)similarity in AE feature-space.

5.3.1 Descriptive statistics and distributions

Some of the waveform features are first discussed in overview using basic descriptive statistical parameters. This representation does not accurately represent the finer behaviour present in the AE data, particularly because it implicates an assumed Gaussian distribution, which is considered a simplifying assumption for many of the waveform features. However, using statistical descriptors does provide an overview of a relatively large dataset with relatively few meaningful parameters. The statistical descriptors used here are mean, median, coefficient of variation (C_V) and the unbiased skew (normalized by $N - 1$) [94]. Table 5.3 lists the nominal values of the statistical descriptors of amplitude, energy, RMS noise, counts, duration and rise time. In addition, Figure 5.4 visualizes the complete distribution of all of the AE features used by means of Kernel Density Estimation (KDE)s, separately for each of the tests in sets D090 and D45. In this figure, the data is also separated into the loading- and unloading-labelled hits. For this figure, it is useful to note that the loading-labelled hits represent a small minority of the total amount of hits - therefore, the unloading-labelled distributions are also a fair approximation of the complete collection of hits. Note set D90 is not included because of corruption of the DIC data for test D3.9 and of load data for test D3.4. Further note that in Table 5.3 that the coefficient of variation is not given for the skew values. This is because values of C_V are of misleading significance when describing distributions that approach zero.

Of the statistical descriptors under consideration, firstly consider the skewness. The nominal skew values for the AE features under consideration are generally to the left and positive (0.01 to 0.686). For the AE features of amplitude and energy, the highest skew values tend to correspond to the highest-loaded test types, while the lowest skew values and even some mildly negative ones tend to correspond to the lower loaded test types. On the other hand, particularly for features RMS noise, counts and rise time, the skewness for type D090 rises significantly above that of other test types. Recalling a positive skew indicates a ‘skew to the left’ or a ‘right-tailed distribution’, this suggests that RMS noise, counts and rise time for D090 are centred more towards lower values or have more of a tail to the right, while for energy and amplitude, the distribution is more symmetric.

The observations about skew become more clear when including the coefficient of variation and the mean and/or median values (centre values) in the discussion. The generally lower (relative) variance and centre values of the AE features for the test types with lower load level compared to higher load level complements the observations about rise time: the lower loaded test types are centred around a significantly lower value, with less of a tail to towards higher values. This implies that a higher load level generally causes the centre of the distribution to shift to a higher level, but also increases the tail to the right. The high skew and relatively high variation of

Table 5.3: Descriptive statistics for the distribution of AE features amplitude, log_{10}(energy), RMS noise, counts, duration and rise time.

	mean	**median**	**skew**	$\mathbf{C_V}$
Amp [dB]				
D45L35	60.6	60.6	0.010	0.035
D45L45	60.9	60.4	0.171	0.046
D90	63.2	62.1	0.276	0.066
D090	60.1	60.1	-0.019	0.023
$\mathbf{C_V}$	0.023	0.015		0.426
$\log_{10}$**(Eny [eu])**				
D45L35	2.80	2.80	0.000	0.077
D45L45	2.93	2.87	0.191	0.104
D90	3.37	3.30	0.160	0.137
D090	2.70	2.70	-0.003	0.070
$\mathbf{C_V}$	0.101	0.091		0.315
RMS [μ**V**]				
D45L35	63.4	63.7	-0.049	0.100
D45L45	80.1	79.7	0.024	0.219
D90	86.2	85.7	0.027	0.218
D090	76.7	75.6	0.146	0.103
$\mathbf{C_V}$	0.126	0.122		0.420
Counts [-]				
D45L35	2.91	3	-0.050	0.596
D45L45	2.86	2	0.281	1.066
D90	5.94	4	0.372	0.877
D090	1.83	1	0.686	0.660
$\mathbf{C_V}$	0.525	0.516		0.268
Dur [μ**s**]				
D45L35	13.2	15.9	-0.309	0.666
D45L45	23.8	15.1	0.391	0.934
D90	51.7	39.5	0.273	0.865
D090	10.0	7.80	0.219	0.995
$\mathbf{C_V}$	0.769	0.704		0.165
RiseT [μ**s**]				
D45L35	2.91	3.00	-0.050	0.596
D45L45	2.86	2.00	0.281	1.066
D90	5.94	4.00	0.372	0.877
D090	1.83	1.00	0.686	0.660
$\mathbf{C_V}$	0.525	0.516		0.268

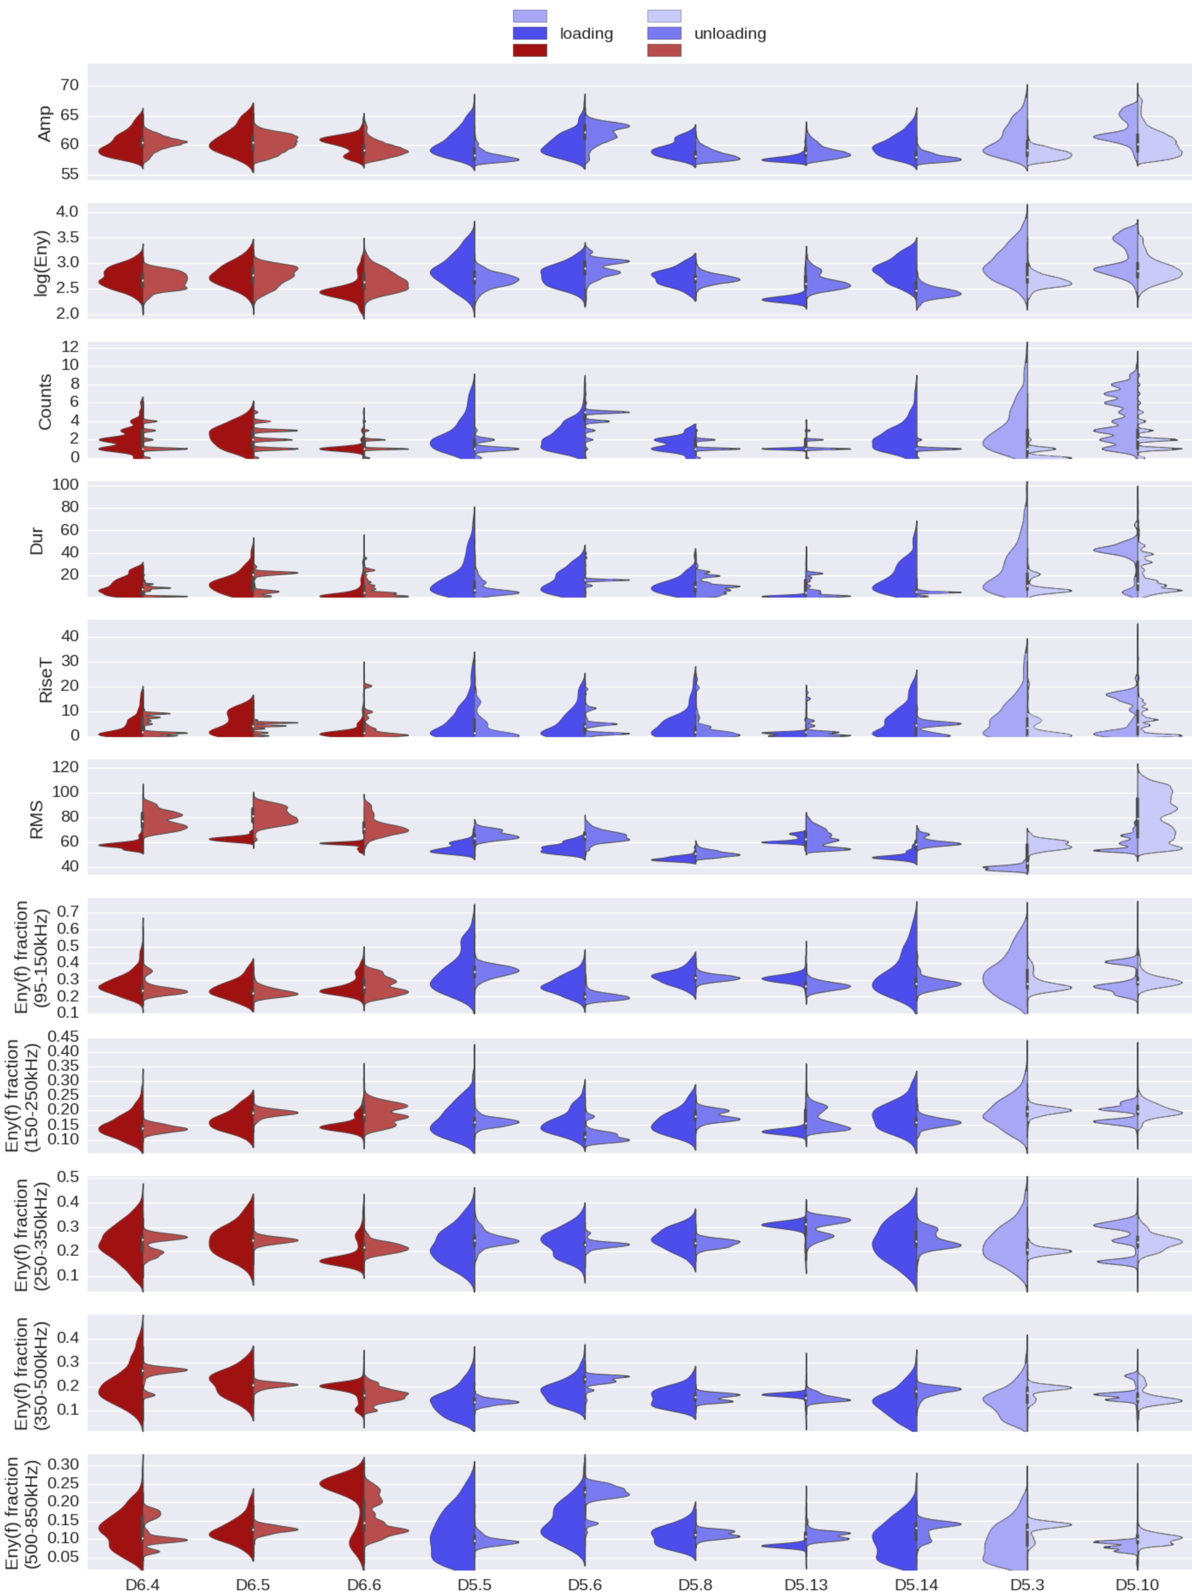

Figure 5.4: KDE distribution of AE features for each of the fatigue tests in sets D090 and D45, separated into loading- (left) and unloading-labelled (right) hits. The KDEs are normalised by maximum width.

counts and rise time in set D090 is offsetting the significantly lower centre values for this test set. Finally, note the AE energy shows a positive linear correlation with load levels ($R^2 = 0.95$) and strain level ($R^2 = 0.97$). A linear correlation between amplitude and load level also exists ($R^2 = 0.87$).

Considering the relative spread of statistical metrics, it is observed that energy and amplitude are most readily explainable: with increasing load level, the centres of the distributions shift to higher values, while the right-tail and variance increase in magnitude. The C_V of amplitude and energy is also most sensitive to the test type. While some trends can be observed in counts duration and rise time - e.g. a similar rise time and increasing C_V for counts, duration and rise time from set D45L35 to D45L45, the trends across all test sets are less pronounced.

Finally considering the complete distributions of waveform features presented in Figure 5.4, several observations can be added to the discussion. In this figure, some agreement between tests or test sets can be seen for several of the waveform features. However, there is also a significant amount of 'outlying' information. For instance, considering the energy, the loading-labelled hits seem to be centred roughly at the same energy level for most tests; while tests D6.6, D5.13 and D5.10 most notably show outlying behaviour. It is also notable that for test D5.6, some features behave differently: especially the partial power of the higher frequencies is centred at a higher value for the unloading-labelled hits than for other tests. Considering this test is ran at a different fatigue cycle frequency (see Table 5.2), this may suggest a strong rate-dependency for some AE features. However, since only one test is performed with this configuration, conclusions from this observation are limited.

It is evident that after separating loading- and unloading-labelled hits, the data is still not in a form where class boundaries might readily and confidently be determined based on one or a few features. This complexity hints at the appropriateness for machine learning tools in deepening investigation of the data. In the following sections, the data is therefore investigated by two other approaches. First, the similarity of hits in feature-space is analysed in Section 5.3.2, which includes a metric based on the overlap of distributions, as well as visualisation of clustering in the data. The second method to deal with the data complexity is by simplification, which is done in Section 5.4 by considering only the waveform feature of energy feature.

5.3.2 Similarity in feature-space and clustering

In this section, statistical exploration of the data is expanded by using the overlap of two feature value distributions is as a similarity parameter between test sets. Thereafter, clustering in the data is investigated in order to reveal relative similarity or dissimilarity of hits in the dataset. The AE hits are separated into hits occurring in either the unloading or the loading phase of the fatigue cycle, as described in the previous section. In order to clarify the current analysis, the concept of using the overlap between two feature distributions as a similarity parameter is first explained. The distributions of feature values for two test sets are determined by integral-normalized multi-modal Gaussian KDEs on a common axis. The multi-modal KDE from the SciPy library is used with default parameters [94]. The overlap fraction is then computed as the overlapped area under the curve as a fraction of the total area under the curve of the KDE probability density function of both test sets. Therefore, an overlap value of close to unity can be said to suggest high confidence of one subset being a member of the other. Low overlap values on the other hand signify a likely difference between the waveform features of two test sets. In the current analysis, each overlap fraction is a metric derived from the comparison of one feature between two test sets. Each test is split by their cycle phase labels ('loading' and 'unloading'), so with 4 test sets, there are $6+1$ unique combinations (D45L35 against D45L35T

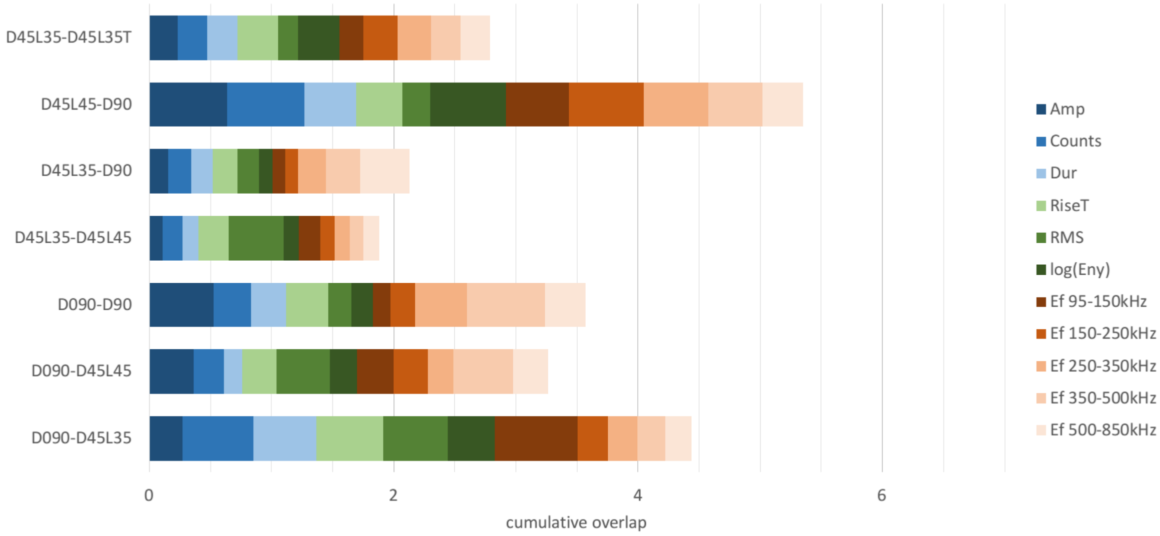

(a) Loading-labelled hits.

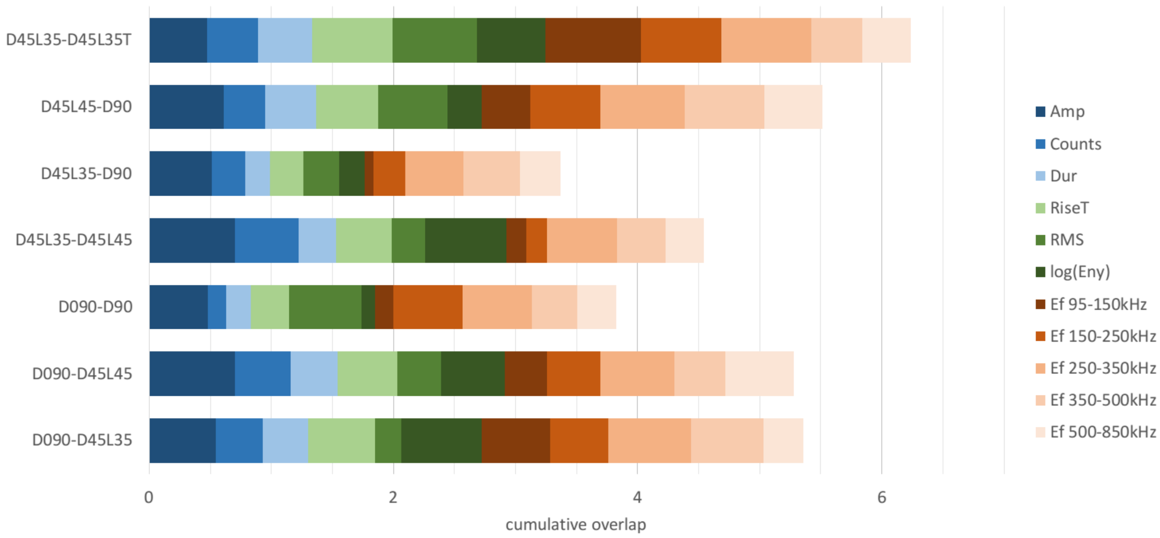

(b) Unloading-labelled hits.

Figure 5.5: Cumulative overlap fractions representing similarity of AE feature value distributions between test types, broken down into contributions from each of the AE features.

is added to compare AE hits early in fatigue testing to those later in fatigue testing). Testing each combination for all of 11 features, Figure 5.5 depicts a summary of 154 overlap fractions, which are organised cumulatively per combination of test sets. If the cumulative overlap is high, test sets can be considered relatively similar by this metric. Though one important reservation must be mentioned: because the waveform features are not independent, the nominal overlap fraction values and particularly their sum is of limited analytical value.

Several observations are made from this figure. Firstly, that it is not readily possible to single out one or a few AE features as distinctly contributing in linearly separating the data in AE feature-space. Most AE features contribute in a varying but significant degree to the similarity or difference of two test sets. Secondly, it is noted that the loading-labelled hits exhibit lower and more varied cumulative overlap values on average, compared to the unloading-labelled hits. Additionally, from a damage mechanics perspective, the loading-labelled hits seem to provide more meaningful information than the unloading-labelled hits: e.g. in both Figure 5.5(a) and (b), D45L45 is more similar to D90 than D090; D090 is more similar to D35L35 than to D35L45. However, the differences are more pronounced for the loading-labelled hits. In addition, the most "disagreement" between the loading- and unloading-labelled hits is found for the comparisons D45L35-D45L35T and D45L35-D45L45. The unloading-labelled hits score these set comparisons as relatively similar, while the loading-labelled hits score them as relatively dissimilar.

The analysis of Figure 5.5 only weakly suggests that the labelling of hits by their cycle phase may be a meaningful one: some differences between the two categories may be explainable based on mechanics or layup, while any existing trends are certainly obscured in this analysis. On the other hand it is also clear that significant similarity exists between the AE hits of different test types. Two analytical tools can be added into the discussion to investigate the data more closely. The first is Independent Component Analysis (ICA) as implemented in the Scikit-learn library [97, 98]. This is a projection that is often described in contrast with Principal Component Analysis (PCA), which aligns data along the axes that maximize variance. Conversely, ICA projects the data along axes that maximize mutual independence of these axes. This transformation tool is chosen because the AE data is described by a set of highly interdependent features in a non-Gaussian feature-space. By comparison, PCA is more suited to describe highly Gaussian data. At the same time, ICA can also be applied with an N by $(m < N)$ mixing matrix, allowing for a projection onto 2D space. Evidently, in the projection to lower dimensions, the number of independent components that can be identified by the ICA algorithm is also reduced. An ICA transform of all hits in sets D090, D45L35 and D45L35T is depicted in Figure 5.6.

The second projection method that is introduced is a stochastically-informed projection of multidimensional data into lower-dimensional space, named t-SNE [99, 100]. This projection is developed at the TU Delft *Pattern Recognition and Bioinformatics Group* specifically to project large datasets of high-dimensional data into a 2- or 3D space, while aiming to represent the relative mutual distances of data points in N-dimensional space as well as possible. Therefore, the potential of the projection offered by t-SNE is recognised for representatively visualising similarity between individual hits in the current data. This is anticipated to be useful, first, because it allows manual inspection of the multidimensional data and its clustering. And second, because this algorithm depicts a representation of the distance between points, which is a governing metric for many classification algorithms, such as k-NN. Figure 5.7 depicts two t-SNE visualisations of test D5.3, where all hits are labelled by their cycle phase (a); and only the top-labelled hits, labelled by the number of kcycles completed (b). These figures are produced using the t-SNE implementation from the Scikit-learn library [97]. An ICA with an 11×8 mixing matrix is used as a preprocessing step. By varying the number of output

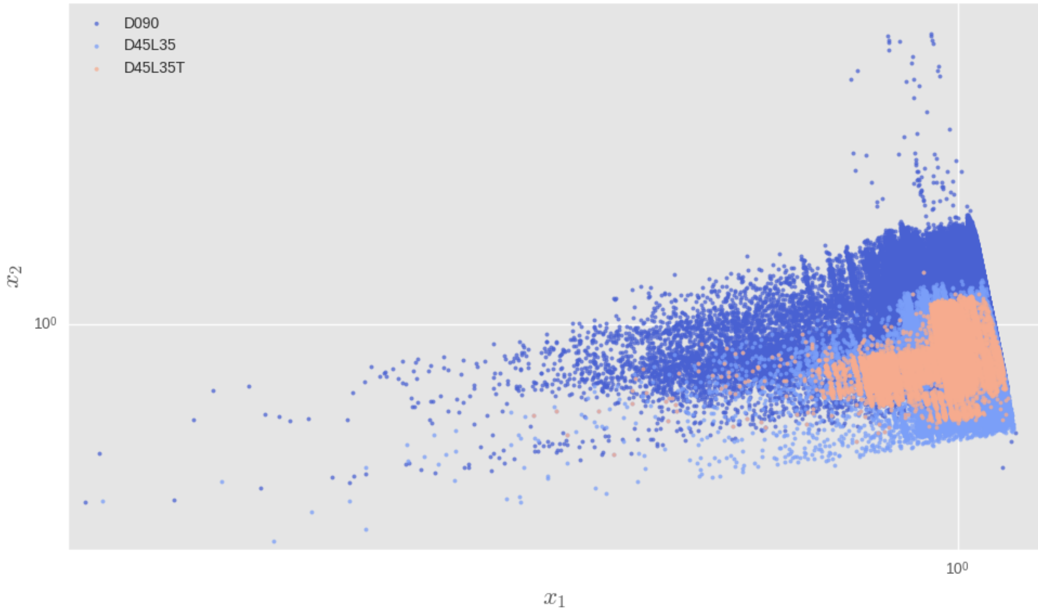

Figure 5.6: Two-dimensional ICA projection of the 11 waveform features for complete test sets D090, D45L35 and D45L35T illustrating the similarity between the the complete sets of AE hits from different test types.

dimensions m in the ICA preprocessing step between $2 - 11$, it seems to be the case that for approximately $7 \leq m \leq 9$, test D5.3 tends to be clustered more clearly and consistently in two linearly separable clusters (mostly 'down' and mostly 'top'). But since the t-SNE operation has a non-convex cost function [97] and the shape of the visualisation is therefore non-deterministic, confidence on finding consistent trends in this dimensionality dependency is low. To account for the non-convex nature of t-SNE, a given analysis is ran three times with random initialisation parameters and the output is inspected for stability. Results of the t-SNE are found not to be particularly sensitive to variation of the algorithm parameters, so default algorithm parameters are used.

Figure 5.6 offers several insights of note. Most notably, it revisits the suspicion arising from the discussion of Figure 5.5: the AE hit feature-space of different test sets are partial subsets of each other. Most notably this is true for D45L35T relative to D45L35. In other words, this figure hints once more that the test sets are not readily linearly separable. This can be seen as a validation for the application of machine learning to AE monitoring data, because this field offers tools to separate data that is not readily linearly separable. From a material mechanics perspective on the other hand, the fact that test sets may not be linearly separable may be considered a necessary but insufficient condition for the potential use of AE data as a means to discern damage progression: If all tests start out with similar damage mechanisms (e.g. matrix cracking) at a damage growth rate in the same order of magnitude, and acoustic emission hits contain information about the damage progression, it may be expected that the tests under consideration exhibit AE activity of the same character, at least in early stages of testing. Though it is also apparent that the variance of AE hit features varies significantly between test types - this particular observation is readily analogous to the findings in Section 5.3.1.

Whereas the ICA plot presents a comparison between test sets, the t-SNE plots in Figure 5.7 compares hits within a test set. An important conclusion can immediately be drawn from Figure

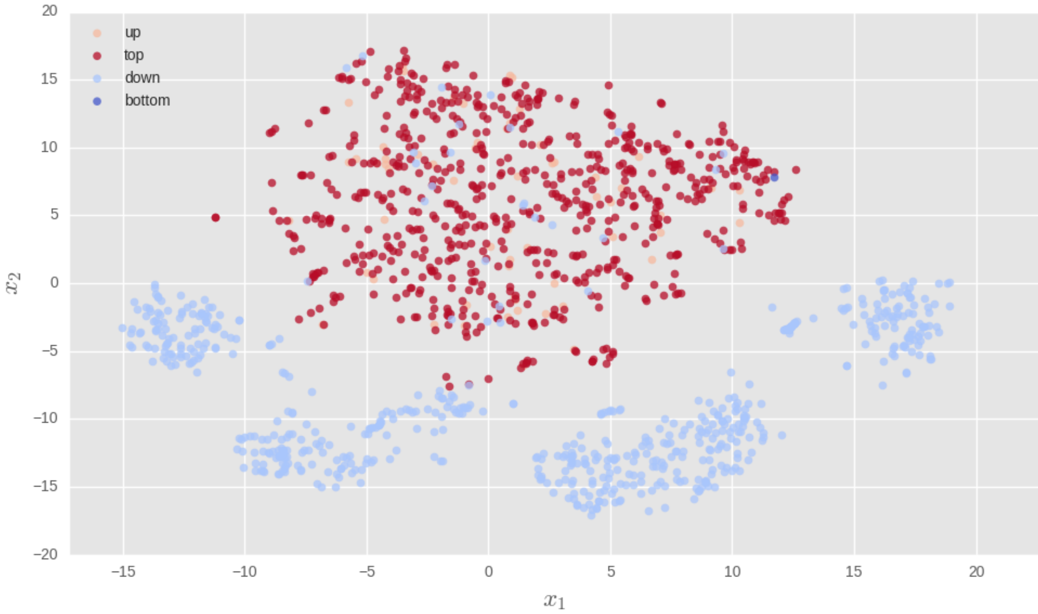

(a) t-SNE projection of all hits in test D5.3.

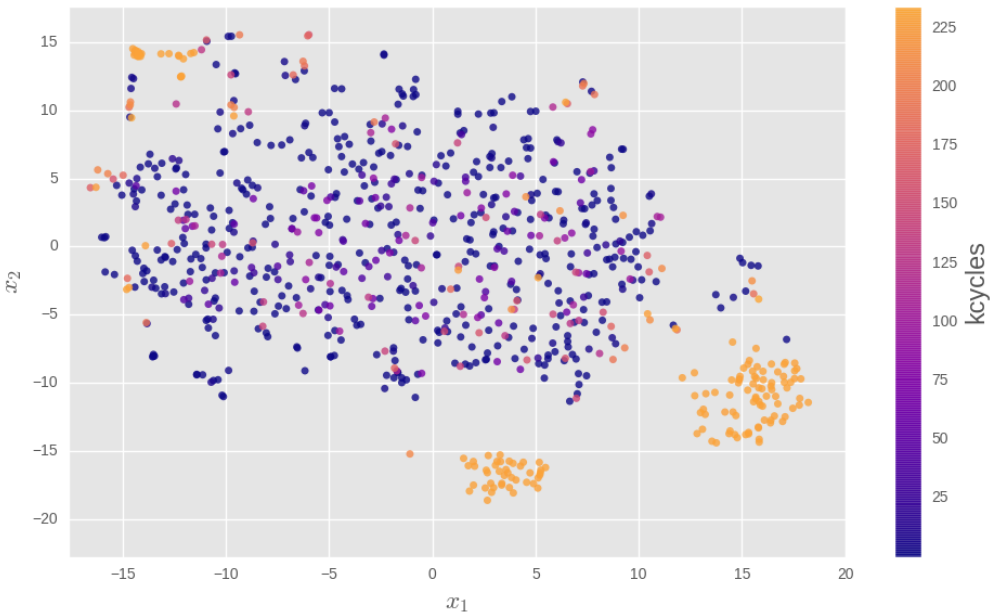

(b) t-SNE projection of top-labelled hits in test D5.3.

Figure 5.7: Two t-SNE projections of test D5.3. This projection aims to correctly represent distance between data points in the transformation from 11- to 2-dimensional space. **(a)** Two-dimensional t-SNE projection of the 11 waveform features of test D5.3 with hits labelled by cycle phase. **(b)** Similar t-SNE projection for only the top-labelled hits, labelled by the number of kcycles completed at the time of occurrence.

5.7a: the labelling of hits by cycle-phase correlates with the most major classification that can be made from the perspective of similarity between waveforms. Note however also that the correlation between labelling and clustering is not perfect: several down-labelled hits are found to be similar to the top-labelled hits in the feature-space. Moreover, the up-labelled hits tend to be similar to the top-labelled hits, but not to the down-labelled hits. From examining the other tests, the down- and bottom-labelled hits tend to be grouped together; while there are also cases of down- and top-labelled hit clusters connecting or overlapping. The image of Figure 5.7a is roughly representative of the results in both the D090 and D45 sets. However, the D45 sets produce the most clearly separable clusters (mostly 'down' vs mostly 'top') compared to the D090. Finally, it is of note that the down-labelled hits are significantly more prevalent in the coupons that have experienced less curing time before testing, making clusters less easily identifiable. Furthermore, comparing Figure 5.7b to Figure 5.3, it can be seen that the clustering within the top-labelled hits also reveals some dissimilarity between the main body of hits and the anomalous burst of hits identified in Figure 5.3 at 230 kcycles. This type of anomaly is only distinguishable in the D45L45 and D090 sets (i.e. not in the D45L35 set). Of course, it remains to be seen if this distinct difference also has a basis in damage mechanics. This question is further investigated in the next section, where the classification between loading- and unloading-labelled hits is analysed from the perspective of energy intensity. More importantly, the validity of the current choice of hit labelling is shown through correlation of the energy intensity of the loading-labelled hits with the derivative of the damage parameter.

5.4 Examining hit energy

In this section, the discussion on AE activity is deepened by simplifying the discussion of the AE waveform to the hit energy only. This step has a couple of reasons: First, the analysis is simplified by only considering a single feature. Second, this eliminates any concerns about interpreting mutual interdependence of AE features; and third, the waveform energy is considered a good candidate for closer inspection, because the waveform energy relates to the actual AE energy, which in turn relates to the energy released by any source mechanisms. First, trends in the AE energy distribution as a function of time are discussed. Thereafter, the energy intensity is correlated with the derivative of the damage parameter obtained from the DIC data.

5.4.1 Energy distributions

In this section the distribution of the energy is examined as a function of time and load. First, as a function of both time and load applied to the coupon at the instant the AE hit is recorded; and second in time by means of cumulative spectral energy plots. The first representation of the AE energy can be understood as adding a third dimension to Figure 5.3 by scaling the z-axis with hit energy. Figure 5.8 shows this in cumulative fashion for test D5.14. Load-time space is divided into bins and each bin is scaled by the energy: each discrete time increment is in fact a histogram using hit energy as bin weights. Note test D5.14 and not D5.3 is presented in this figure, because D5.14 exemplifies the behaviour over the longer term better, while still keeping the activity in the top of the fatigue cycle marginally intelligible in the figure. The second representation is presented in Figure 5.9 and is referred to as a cumulative time-sectioned KDE. Each line in this figure is a Probability Density Function (PDF) of the AE energy spectrum determined by KDE for that increment (bin) of time. The PDF is then normalized by the area under the curve and scaled by the cumulative energy in the time bin. This type of representation is meant to give an insight into the energy of the top-labelled hits relative to the complete collection of hits.

The results in Figure 5.8 are considered an expansion of the discussion in Section 5.2. Note therefore that the features distinguished in Figure 5.2 are also present in Figure 5.8: Recall from previous discussion that the loading-labelled hits are initially active but reduce in activity as the test progresses, while the unloading-labelled hits are initially absent, but grow steadily after the initial stages of testing. From 5.8, it can be seen that hits at the high end of the load spectrum constitute only a marginal fraction of the total amount of energy carried by all of the AE hits. Now, also consider Figure 5.9, which shows representative behaviour in that the very early stages of testing are again dominated by the top-labelled hits, as well as the unloading-labelled hits appearing later in testing, but at a lower energy level. As an aside, note that the sudden increase at in the last time-section at the 10^3 [eu] energy level correlates with the anomaly found for the top-labelled hits at the end of Figure 5.3.

From examining these figures, the most obvious distinguishable feature is that even though the higher energy hits occur near the top of the load cycle, the vast majority of AE energy is carried by the hits arriving during unloading of the sample. This immediately presents an interesting problem for any constitutive interpretation of this data: assuming that the hits arriving during loading and those arriving during unloading are in fact caused by two different

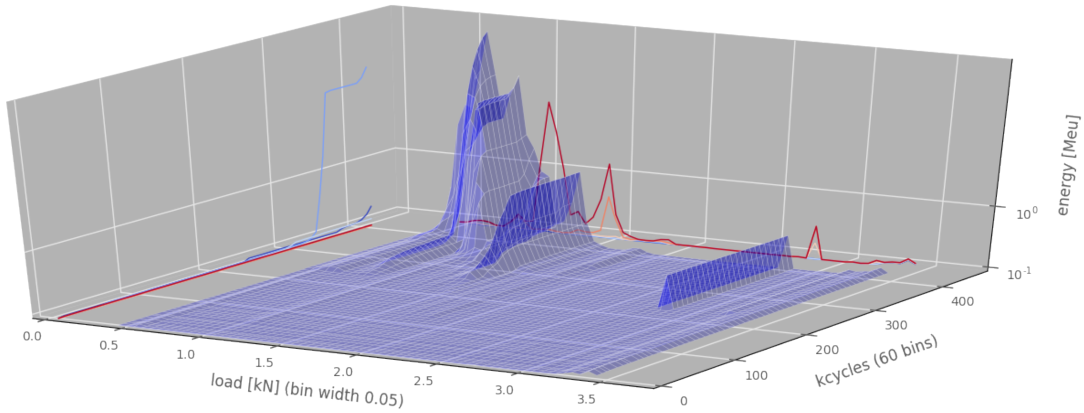

Figure 5.8: Cumulative AE hit energy spectrum in time-load space for test D5.14. Note the projections on the vertical planes are cross-sectional (non-cumulative).

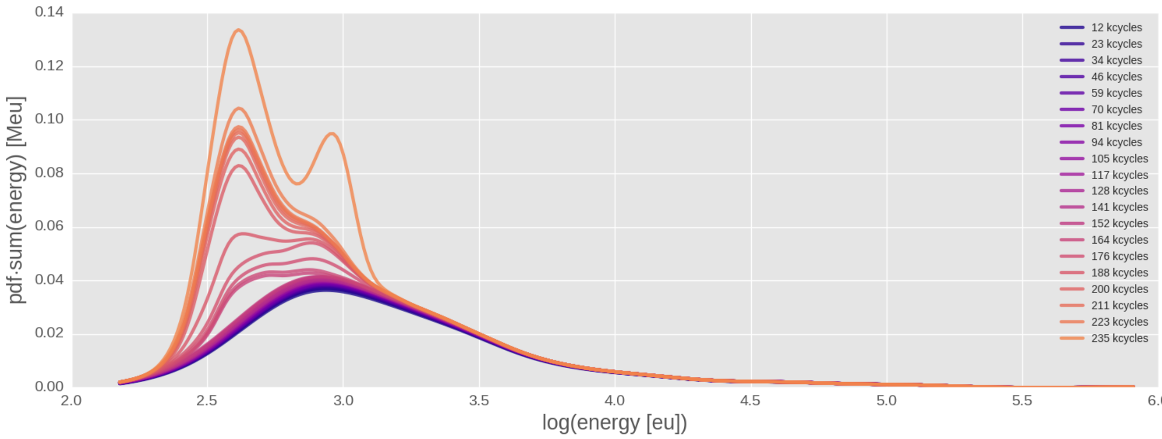

(a) Time-sectioned cumulative KDE for test D5.3 (all hits).

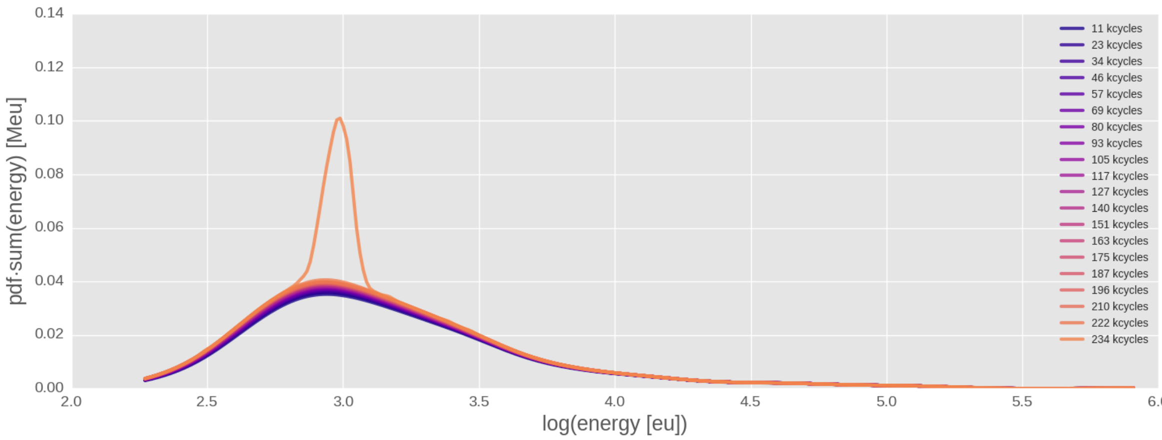

(b) Time-sectioned cumulative KDE for test D5.3 (top-labelled hits).

Figure 5.9: Cumulative AE hit energy spectrum for test D5.3, sectioned in time. Each line is a KDE representing the PDF of the energy level of the hits arriving in that interval in time.

types of source types, it is not readily deducible what these two (or more) source types are. A set of hypotheses is formulated, which can be evaluated for either of the labelled classes (loading- and unloading-labelled):

Hits in the given class represent mainly...

1. Noise external to the coupon material. This can include machine noise, but also noise from damage in the tabs or adhesive.
2. Noise internal to the material. This type of activity can be attributed to friction of internal fractured surfaces.
3. Damage (i.e. dominantly matrix cracking) developing or initiating in the material.
4. Progression of damage into mechanisms other than matrix cracking.

It is interesting to note that from a deductive perspective, it does not seem conclusively possible to falsify or confirm any of the above statements. However, several assertions can still plausibly be made about the relative expected contribution from these hypothetical behaviours. Firstly, any assertions that a significant amount of AE activity is caused by shear damage in the adhesive (part of (1)) are carefully considered dismissible, as the adhesive is only loaded up to 10% (compared to 45% tensile in the coupons) of its failure strength. However, damage mechanisms under compression cannot be reasonably ruled out. The second component of (1) is that of external mechanical noise from the load cell machinery. An expectable characteristic of machine noise is that is most likely to have an approximately constant intensity throughout testing, which neither of the currently identified classes exhibit. From this, it is concluded that mechanical noise is a possible, but not likely source of AE activity. To this discussion it may be added that repeating any of the fatigue tests with an untabbed coupon could help rule out compressive damage sources in the tab or adhesive. Likewise, insight into the influence of compressive damage mechanisms in the clamped area of the coupon on AE activity could be gained by running a test with significantly higher clamping pressure on an untabbed or tabbed coupon. Repeating a test in a different load cell could help confirm or invalidate the assertion that external mechanical noise is not appreciably present in the AE data.

Another important assertion follows from material mechanics, existing work on AE; and the measured reduction in stiffness: it is quite certain that material damage is developing and dissipating energy as AE hits. It follows that at least one of the labelled categories (unloading or loading) must contain AE hits from material damage sources; i.e. statement (2) must be true for at least one of the labelled classes. Again, from the stiffness reduction trends, damage per time interval (or per cycle) must be high initially, reducing over time. This strongly suggests that the unloading-labelled hits are not related to the initial stiffness degradation, because this type of activity follows an opposite trend. It is interesting to note, however, that the unloading-labelled hits follow a trend correlating roughly with the stiffness degradation itself. The loading-labelled hits on the other hand do follow a pattern that is consistent with the stiffness reduction in the sample, making this class of hits the more promising candidate for further investigation. The next section investigates this correlation in more quantitative detail.

If the unloading-labelled hits do not represent external noise, the first candidate hypothesis to explain the source mechanism of these hits might be internal noise (2) caused by friction between internal fractured surfaces. To investigate the internal friction hypothesis, consider the following piece of data: the activity of unloading-labelled hits for set D090 starts early and stays roughly constant at $2 \cdot 10^3$ hits per kcycle on average, whereas the same type of activity in

D45 sets generally starts out as non-existent and can take roughly $50-400$ kcycles to reach this level. The noise hypothesis for the unloading-labelled hits is in this sense supported by the fact that the unloading-labelled activity never decreases: because the rate of stiffness degradation is reducing throughout testing, it may be reasonably expected that the acoustic activity relating to this degradation follows the same trend. Furthermore, from the expected damage mechanisms in both coupon types (ref. Section 3.1), it can be reasonably assumed that matrix cracking is the first damage mechanism in both test types. But since the fatigue cycling is always in tension, friction between fractured surfaces of transverse cracks cannot readily be expected; although it cannot be completely ruled out either. Another way for hypothesis (2) to hold would be if friction between longitudinal cracks is the cause of this AE source. Longitudinal cracks could be manifested, for example, in inter-ply delamination as a damage mechanism progressing from matrix cracking (hypothesis (4)). This theory can be considered plausible based on the fact that coupons in set D090 will exhibit a more significant stiffness incompatibility between 0°- and 90°-plies, compared to the ±45° plies in the D45 sets, which could explain the difference in unloading-labelled hit activity between sets D090 and D45. But this theory also leaves a rather significant question unanswered: why is this type of activity almost exclusively abundant in the lower half of the unloading cycle and not in the upper half of the unloading- or the lower half of the loading cycle? No unambiguous conclusion can therefore be drawn on the noise-based explanation of the unloading-labelled hits' source mechanism. With regard to the hypothesis that these hits are not noise, however, it can be noted that because these hits occur during unloading, the only conceivable way they could originate not from any kind of noise but from a damage mechanism, would be if one more more (fatigue) damage process of this particular material are governed by some undocumented viscoelastic behaviour.

Finally, the question of whether any of the identified AE event classes are related specifically to either crack initiation or propagation is evaluated (implied in hypothesis (3)). Other researchers have found that transverse cracks tend to propagate across the full width of the specimen (almost) instantaneously [14, 77, 101]. It is therefore concluded that at the current level of filtering/clustering of the data, no confident assertions can be made on whether propagation and initiation produce distinctly different types of AE events.

Evaluating relative plausibility of the stated hypotheses, it is concluded that the loading- or top-labelled hits are a highly likely candidate to exhibit a meaningful correlation with stiffness degradation. For the unloading-labelled hits it cannot be definitely concluded whether their origins are in internal friction or an undocumented viscoelastic material damage process. For all hit types it is judged that external noise, while not definitely impossible, is unlikely to be a dominantly contributing source mechanism. Additionally, it must be noted that all of the assertions in this section are concerned with attributable *trends*; any or all of the identified classes of AE hits may still consist of multiple different source types. Finally, ending this discussion on somewhat of a philosophical note: consider that from the current discussion as well as the analysis presented in the next section, it is apparent that the activity of the top-labelled hits bear some constitutive connection to the stiffness reduction in a coupon. Even so, this apparent correlation does not guarantee a relation to damage in the material due to tensile cycling. For example, AE events could originate from a different source that also releases energy at the top of the cycle, such as compressive damage in the tabbed area, the adhesive or the tabs themselves. This touches on an inherent "problem" of any interpretation of AE data based on hit energy that is simply based on conservation of energy: any source that emits AE energy during testing may be expected to emit more energy at the top of the cycle, where most work is applied to the whole system, than at a lower load level.

5.4.2 Energy intensity history and correlation with the damage derivative

The significance of the energy intensity history is the main focus of this section. In the current project, any "intensity" is understood as a summed quantity per interval of time; i.e. energy intensity is the sum of the hit energies per an interval of, for example, 1000 fatigue cycles. Inspecting the data, it is revealed that all tests start out with a short period (no more than roughly 2 hours, or 20 kcycles) where the acoustic activity is high and strongly decreasing, with energy levels and amplitudes of larger variance and magnitude. After this period, activity settles at an energy level 1-2 orders of magnitude lower. The initial short period of high activity is referred to as the settling phase. The D90 set is regarded as an exception to this case because in these tests acoustic activity is so high that discernibility of the settling phase is drowned out by the base level acoustic activity. Furthermore, both the energy level and amplitude of the hits and the variance of the energy level are on average higher in set D90 compared to the other test sets: about one order of magnitude higher for the nominal energy; and up to a factor two for the variance. By filtering out lower energy hits, it can be revealed that the settling phase also appears in the D90 set.

For all tests performed at a low load level, a sequence or part of a sequence can be observed in time. This is most apparent by plotting the energy intensity as a function of time (e.g. with a bin size of 1000 cycles; see Figure 5.11), but it can also be seen in hit intensity, hit energy or hit amplitude plots. The observed sequence is illustrated schematically in Figure 5.10. Note that this figure essentially shows the trends discussed in the previous sections. After the settling phase (1), the acoustic activity will go through a valley phase (2). In this phase, AE hits do arrive, but their arrival is relatively sporadic and the variance of their energy and amplitude is relatively large. Finally, AE activity ramps up in phase (3) where the energy intensity enters a phase of constrained growth. The sequence described here is not observed completely in each of the tests; however, it can be theorised from, for example, comparing test D5.3 to tests D5.14 and D5.13 that some coupons were not loaded long enough to have entered phase (3), while in others phase (2) is so short that it does not seem to appear at all. Notably, phase (3) activity tends to start almost immediately for set D090, while in D45 tests this type of activity only becomes discernible after $50 - 150$ kcycles. Because there is significant variation within the D45 category, it is unclear whether this is effect is attributable to the difference in damage mechanics between these two coupon types, or whether it may be related to the curing effects discussed in the beginning of this chapter.

By once more separating the data into hits arriving during unloading and loading, it is revealed that the energy intensity of the loading-labelled hits tends to contain information that correlates more accurately with the change in the damage parameter than the complete set of AE hits would. Moreover, phases (1) and (3) are found to be roughly dominated by the loading- and the unloading-labelled hits respectively. This is depicted using empirical data in Figure 5.11 and idealized in Figure 5.10. A practical limitation of the visualisation in Figure 5.11 is brought to the attention of the reader: The data from AE and DIC measurements are not mutually synchronised. This has three causes: first, the DIC and AE systems are physically separate systems that do not communicate a mutual clock between them in the current set-up. Second, the fatigue cycle has frequency 3.0 [Hz], but is interrupted every 500 cycles to capture DIC images. This causes a variable offset in the timing (in kcycles) relative to AE data, as the number of cycles completed at each point in time is calculated based on the averaged loading frequency over the whole test. This average frequency is approximately 2.9 [Hz] for most tests (2.0 [Hz] for D5.6). Third and finally, the DIC measurements are not "as actual" as the AE measurements are: DIC measurements are only taken every 500 fatigue cycles, whereas AE hits practically instantaneous. Some of these causes for a lack of synchronisation in the data may

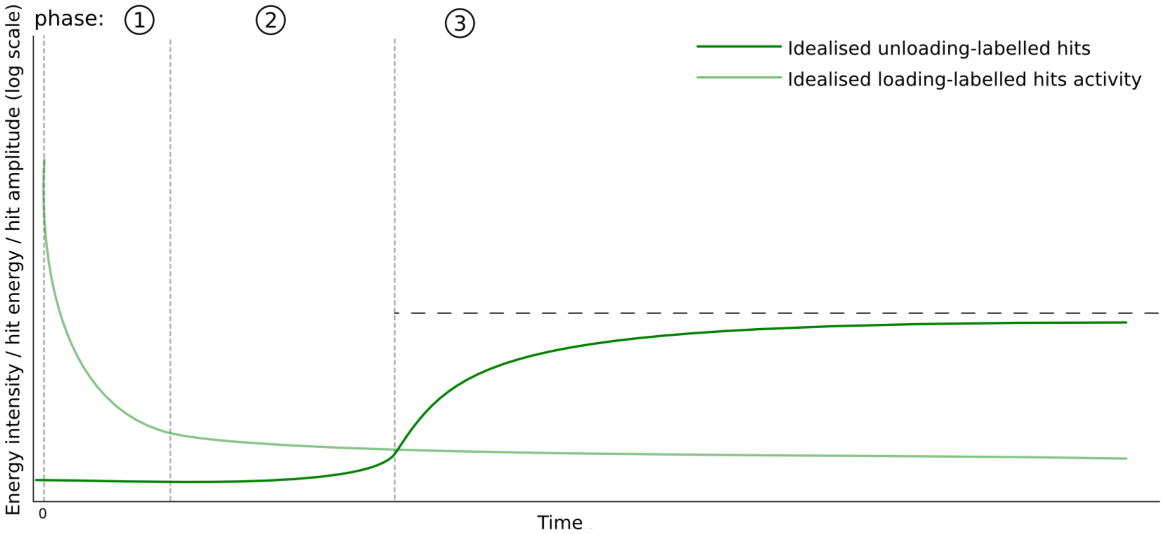

Figure 5.10: Schematic depiction of the observed hit energy (intensity) trend. Phases (1) settling, (2) valley and (3) constrained growth are indicated. An idealized separation of the behaviour of hits arriving during loading (light green) and unloading (dark green) is depicted.

be mitigable by more elaborate post-processing, but in the current analysis they are left as-is.

Closely inspecting Figure 5.11, it can be seen that even though detailed behaviour may be correlated between the change in the damage parameter and the energy intensity of the loading-labelled cycles for both depicted test types, the correlation is fairly rough. It is also evident that the correlation for tests in set D090 is less clear. This may be in part due to stiffness degradation being limited by the 0°-layers, while actual damage is expected to continue developing. One of the most notable features in Figure 5.11a is actually where correlation is lost - exactly at the anomalous cluster in Figure 5.3. However, the observed correlation of these two metrics is also rather rough. And indeed, while for most tests, correlation between energy intensity and ΔD can mostly be attributed to the top-labelled cycles, although for some tests it is not completely clear whether some information might also be found under other labels, while the up- or top-labelled hits may contain a fair amount of noise. The down-labelled hits as a class are not found to positively correlate with the evolution of ΔD in any of the tests. Given these observations, it may be feasible to expect a progression in the phase of the fatigue cycle in which AE hits occur. For example, the AE hits correlating with the evolution of ΔD may occur at a lower and lower load in the fatigue spectrum, as illustrated in Figure 5.12a. However, referring back to Figure 5.3, it can be seen that the separation between top-labelled and all the other hits is and remains high throughout testing. Therefore, the hits at the top of the cycle never shift out of the top-labelled cycle segment and the hypothetical behaviour in Figure 5.12a cannot be confirmed by the current data. Nonetheless it is possible, but not readily defensible, that this behaviour still exists, but is manifested in lower energy hits that are truncated by the AE threshold. The amplitudes of top-labelled hits tend to be centred around values of approximately $59 - 61$ dB; this centre remaining relatively constant throughout testing. It cannot however be confidently asserted from the data that there is not, for example, a second or broader peak at a lower level in the distribution of AE hit amplitude. Therefore, truncation of meaningful data by the AE threshold is considered a possible outcome of the current configuration. It is on the other hand notable that this finding agrees with other works in literature, as discussed in Section 3.2.5, in that the damage mechanism of matrix cracking is manifested in lower amplitude AE hits.

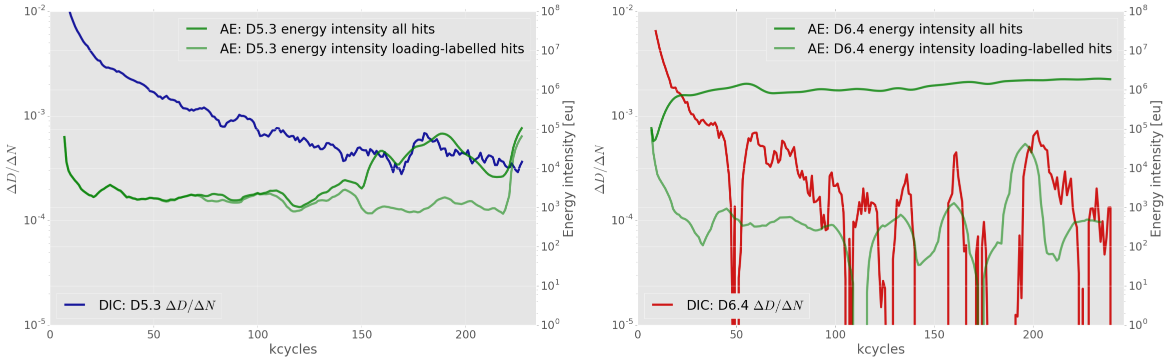

(a) test D5.3 **(b)** test D6.4

Figure 5.11: Energy intensity history (bin size of 1000 cycles) of tests **(a)** D5.3; and **(b)** D6.4. Both graphs also contain the damage parameter derivative, as well as a filtered energy intensity containing only the loading-labelled hits. The damage derivative of D6.4 only drops below the x-axis because of the representation of values close to zero in the log scale.

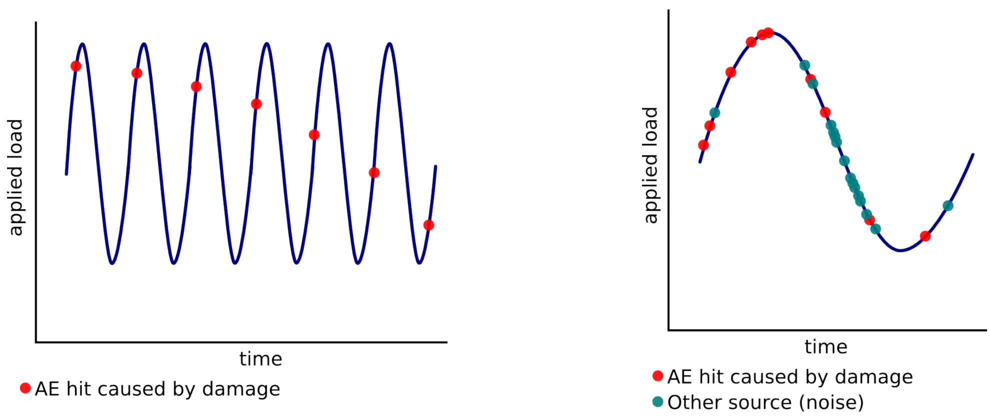

(a) **(b)**

Figure 5.12: Two hypothetical (not necessarily mutually exclusive) trends in AE activity. Both of these cannot be confirmed in the current study. **(a)** Activity caused by material damage occurs at a lower and lower load, as progressing damage creates stress concentrations lowering the threshold for damage growth through successive fatigue cycles; **(b)** Activity caused by material damage is present during all phases of the fatigue cycle, but is not readily recoverable from the unloading phase due to high activity of other AE sources.

Table 5.4: Mean squared errors (of normalized values in $\log_{10}$ scale) between the $\Delta D/\Delta N$ parameter and energy intensity of selections of hits by fatigue cycle phase, for the first 930 minutes of each test. The change relative to the mean squared error for "all hits" is given in parentheses, with improvements emphasized in bold text. Note the exclusion of test D5.6, because its behaviour is outlying (see note **(e)** in Table 5.2).

Set name	Test ID	all hits	top-labelled hits	loading-labelled hits
D090	D6.4	0.58	0.27 (**-53%**)	0.27 (-53%)
	D6.5	1.57	0.27 (**-83%**)	0.27 (-83%)
	D6.6	0.69	0.39 (**-43%**)	0.52 (-25%)
D45L35	D5.5	0.56	0.15 (**-73%**)	0.15 (-73%)
	D5.8	1.94	0.18 (**-91%**)	0.18 (-91%)
	D5.13	1.37	0.08 (**-94%**)	0.30 (-78%)
	D5.14	0.31	0.17 (**-45%**)	0.18 (-42%)
D45L45	D5.3	0.19	0.08 (-58%)	0.07 (**-63%**)
	D5.10	0.94	1.52 (+62%)	1.24 (+32%)

One open question these observations leave somewhat open is whether the hits descriptive of damage (and/or stiffness degradation) remain to be found at the top of the fatigue cycle throughout fatigue loading. I.e., it is possible that the "meaningful" AE hits in fact do arrive in the top of the cycle at the beginning of testing but as they start arriving at lower loads, they start occurring at lower amplitudes and are more drowned out by the activity of other types of AE sources; namely the dominant activity identified in the unloading phase. This situation is illustrated with Figure 5.12. As an explorative trial, a k-NN (k 3 to 5) clustering algorithm is used, preprocessed by ICA, with the aim to classify hits similar to selected sections of top-labelled hits. k-NN is chosen for this exploratory operation, because of its conceptual symmetry with the t-SNE projection, as well as being a common choice for AE studies in general. But based on the correlation of the outcome with the ΔD parameter, predictions are not found to provide any new information. This is asserted to be mainly caused by the sensitivity to accuracy in the current approach: because the top-labelled hits are only a relatively small number of hits, misclassification of $1-2\%$ of hits can already significantly influence the derived energy intensity. Therefore, the behaviour represented by Figure 5.12b cannot be readily confirmed or falsified at this time.

Finally, one other possible interpretation of the outcome of the correlation between energy intensity and ΔD must be considered: the DIC data is not a perfect representation of damage in the specimen. It is possible that the difference between the information in the loading-labelled energy intensity and in the DIC data represents extra information present in the AE data, rather than noise. If this is true, damage must be developing which does cause AE events, but does not cause (appreciable) stiffness degradation. A candidate for this type of AE source could, for example, be compressive damage in the coupon tabs or tab adhesive, or possibly even some types of longitudinal cracking. Regardless, this hypothesis cannot be evaluated with the current data, so it must be left as an open question.

In Table 5.4, the correlation between the stiffness degradation and energy intensity of loading- or top-labelled hits for test sets D45 and D090 is quantified by means of the mean squared error. This table is established by normalizing the type of curves shown in Figure 5.11 by first their variance and then their median values. This normalization is chosen because, by inspection, it leads to a reasonable alignment of the three curves. The mean squared error between respective curves is then calculated as a metric of correlation (lower is better), where the squared error is calculated for each instance (every 1000 fatigue cycles) and averaged over the 930 minutes used as a standard period in this project. In addition, Table 5.4 also shows the relative change of the error when only top- or loading-labelled hits are considered. Although this representation strongly suggests that selecting only the top-labelled hits is generally a significant improvement of the correlation between energy intensity and stiffness degradation, while adding the up-labelled hits yield little or no improvement, it must be stressed that the simplification of this representation is anything but unambiguous. As can be seen in Figure 5.11, the quality of correlation varies for certain sections of the stiffness degradation curve; this is true for all tests. This, coupled with a higher activity of the higher loaded coupons overall, also explains why the correlation is less convincing for the D45L45 test set (particularly D5.10). It may be concluded that the labelling of AE hits based on the fatigue cycle phase seems to hold promise, while the currently applied algorithm of simply dividing the fatigue cycle phase in four equal segments may lack the sophistication that may be required to reveal the correlation of stiffness degradation with hit energy intensity in appropriate detail. By extension, it is also noted that any normalization approach, such as the one chosen here, will be somewhat arbitrary: because of the sporadic quality of the sought after correlation, significant "noise" still present in the data—namely unexplained AE behaviour—prevents metrics such as the initial-, mean-, median-, maximum-, or long-term trend to be meaningfully defined across different tests. Future study will need to investigate whether a more sophisticated labelling might reveal the trends found in this document with higher fidelity, enabling a more robust comparison to stiffness degradation.

The correlation of hit intensity with stiffness degradation must be concluded with a somewhat limited set of definitive answers. Said correlation can be found and conforms to the expected behaviour of reducing over time in this type of test and being consistently active near the top of the fatigue cycle. No distinct progression from one damage mechanism to the next can be readily identified; though with the current data this neither confirms or falsifies the hypothesis that the AE data contains this kind of information. In these particular testing conditions, it is found that the source mechanism of the unloading-labelled hits is particularly uncertain. The findings in this section further strongly suggest that the clusters found in Section 5.3.2 are validated by a connection to a constitutive basis. This is considered an essential finding, as it proves that distinctly different source mechanisms can be correlated to distinctly different clusters of hits in AE feature-space.

Chapter 6

Conclusions

In this chapter, conclusions are presented as derived from the analysis of results in Chapter 5. Recommendations for future work are discussed in Chapter 7. In the current project, a constitutively-informed approach to the classification of AE data is developed and validated. The main research question is posed as: *How can in-situ AE measurements on a GFRP structure be used to meaningfully monitor damage development?*. This question is investigated by defining damage development as primarily based on stiffness degradation, which in turn provides the validation data for the experiments. An AE dataset is established from fatigue-loaded ($R = 0.1$) 90°- and ±45°-dominated GFRP coupons, such that damage in the samples is dominated by matrix cracking. DIC is used to derive the relative stiffness degradation measured on the surface of the samples over the duration of the fatigue test. AE hits are labelled based on the load applied to the sample at the instant an AE hit is recorded. This labelling (classification) is validated with the stiffness degradation data. In addition, the classification made based on the cycle phase segments is correlated against the clustering of the AE hits in 11-dimensional waveform feature-space. At least two classes of hits are identified in the feature-space: the loading-labelled hits start with high hit intensity, but reducing as the test progresses; whereas the unloading-labelled hits start at non-existent hit intensity, which increases as the test progresses. The unloading-labelled hits tend to have lower median energy- and amplitude values, but comprise the overwhelming majority of the total amount of AE hits over the duration of a test.

From the discussion of results in Chapter 5, three primary conclusions are identified. Firstly, using the fatigue cycle phase as a means to classify AE data is shown to lead to a meaningful distinction between at least two major AE sources found in the current dataset: For the material- and testing configuration in the current project, hits occurring during the loading segments of the fatigue cycle are shown to correlate better with the stiffness degradation than the complete set of hits does. This is true for both coupon types, while the differences in AE (energy) activity between D090 and D45 test sets are also correlated with the stiffness degradation to a first order. It is added that although some of the loading-labelled AE activity is explained by stiffness degradation, all tests exhibit bursts of this type of activity that cannot be explained in this way. Therefore, while current results are promising, the current quality of correlation is not consistent enough to base predictions of stiffness degradation due to matrix cracking on. It is interesting to note that in the current project, there is actually no particular reason to assume that the methodology used validates the distinction of one damage mechanism or another in the AE data. Rather, with the current validation method based on stiffness degradation, the classification that can be made in AE data is most accurately described as ‘classes of hits that contribute to stiffness degradation in the direction of loading’ and ‘classes of hits that do not’.

Secondly, the AE hit classification made based on the fatigue cycle phase segment is approximately reflected in the clustering found in the AE hits when they are described by the 11 waveform features of amplitude, energy, duration, rise time, counts, RMS noise; and 5 partial power fractions in the frequency spectrum. Because by extension, this clustering corresponds to the hits correlating with the stiffness degradation, it is concluded that for the major three source mechanisms present in the current tests, use of clustering algorithms to discern these sources is qualitatively validated. However, it is noted that by this definition, it is, for example, possible that the class represented by the clusters of anomalous loading-labelled hits may also be rooted in the same type of damage mechanism (i.e. matrix cracking) as the rest of the loading-labelling hits, while not contributing appreciably to stiffness degradation - see also the discussion in the next chapter.

Thirdly and finally, from the preceding, it is concluded that particularly the combination of validation of AE classification through stiffness degradation and the representation of the AE feature-space that is offered by t-SNE shows promise as an effective and consistently verifiable approach to the basis of the AE classification problem. Particularly t-SNE, in combination with ICA, is found to provide a much-needed means to inspect AE data in a manner representative of relative similarity between AE hits. Conveniently, this method is also conceptually compatible with some classification algorithms like k-NN. This enables the operator to manually label hits, or to verify clustering proposed by an unsupervised classification algorithms.

From these three main conclusions, it follows that the main research question to the project is answered, at least in a provisional form, by the approach taken in the current project: this approach provides a methodology that shows promise to develop AE-based structural integrity monitoring of composites, using the well-defined and meaningful validation metric of stiffness degradation; and an approach to filtering and classification that is rooted in mechanics and can be validated by the aforementioned metric. For completeness, it is noted that the finding of a correlation between stiffness degradation and particular forms of AE activity implies that the answer to the sub-question concerning the existence of meaningful information in the AE data is a conditional *yes*. The condition here is that in the current project, stiffness degradation has been chosen as the metric for damage development. Finally, the answer to the main research question can be said to be of a provisional form because it is only answered for this particular material and testing conditions, without the proposed methodology being proven to be extensible to, for example, higher fatigue loads, other damage mechanisms or more complex laminates. However, there is no apparent reason to expect the current methodology will not be valuable in extending the analysis to these variables.

In addition to the primary conclusions listed above, a series of secondary conclusions of note are listed. Firstly, it is noted that in the current project, the majority of AE energy and activity is found in the bottom half of the downward fatigue cycle. This finding gives rise to the dilemma that, if the hits in this phase of the fatigue cycle cannot be attributed to noise external to the coupon, at least one of the following statements must be true: (1) The AE activity found in the unloading phase of the cycle originates mainly from noise sources internal to the coupon. In this case, the assertion “AE hits corresponding to noise represent the majority of recorded AE hits but [an] insubstantial phenomenon in terms of energy” - Doan et al. [102], does not hold for this particular material and testing conditions. This type of activity could then be caused by internal friction. (2) The AE activity found in the unloading phase of the cycle mainly originates from sources internal to the coupon, but they are not noise. In this case, said activity signifies a type of unexplained material behaviour. This could, for example, be indicative of viscoelastic behaviour during cycling, and/or arise from Poisson's contraction in the material. This unresolved dilemma undoubtedly warrants further research into the underlying mechanism of this AE source: If this type of AE activity is attributed to friction noise internal to the coupon, it may be promising to investigate it as a metric indicative of existing damage, whereas if this type of activity is attributed to an undocumented material behaviour, it may lead to gains in the understanding of damage mechanisms in composites.

Second, it is noteworthy that the current approach is particularly valid for T-T fatigue at low load levels. The stiffness degradation used here as a validation metric has a valid constitutive basis only because matrix cracking is almost exclusively the only damage mechanism. The current approach has to be built up from this uni-modal damage process so as to not introduce too much 'damage mechanism noise'; and indeed, the bursts in loading-labelled AE-activity which cannot be explained from the stiffness degradation may well be rooted in damage mechanisms other than matrix cracking, though this is not definitively proven in the current analysis. Also, as noted in Chapter 5, hit intensity saturation of the AE sensor cannot be ruled out, particularly for test sets like D90 where the load level and AE activity is higher. Therefore, at least for research on AE-based structural integrity monitoring, it is recommended to perform tests at a low load- or strain level (also depending on the damage mechanism to be investigated). This is especially noteworthy because it cannot currently be ruled out if the AE threshold of 57 [dB] is truncating the distribution of AE hits originating from relevant damage mechanisms. This poses an especially interesting limitation, because if structural integrity should be a function of cumulative AE activity, the system could be sensitive to any 'missed' AE events and may therefore be unstable in its predictive value. Though imaginably, over large populations of AE hits it may be possible to estimate the complete distribution of hits based on a truncated one.

This chapter is concluded on a few general assertions about the project approach and its results. In summation, the techniques used in this project show promise in their efficacy to distil AE hits originating from matrix cracking and causing stiffness degradation, for this particular material- and testing configuration. It is therefore also asserted that, in this context, the results of the current project do not suggest any particular reason why the building-block would not provide a suitable blueprint for extending analysis by the current methods to other damage mechanisms, layups and/or materials. However, for this approach to be successful, it is of fundamental importance that each next 'block', consisting of added complexity, is small enough, such that confidence in the causes of this increment of complexity is high. This notion is so important because each block in the building-block approach inherits the confidence level of the assertions of the previous block(s). In this way, the current approach is contrasted against pure data-driven approaches: the mechanics-informed building-block approach inherently builds towards exhaustive classification more slowly and deliberately, but in doing so, this approach may provide insights on both structural integrity monitoring and damage mechanics at each of those deliberate steps, thereby contributing towards a better understanding of composites. For the goals of DSNS outlined in the first two chapters of this report, it is analogously concluded that the current approach would be a careful, but confident method to build up the understanding of AE-based structural integrity monitoring of composites, towards a tool that may provide an accurate 'log' of real-time damage. It may evidently be added that in an in-situ application, more elaborate filtering and/or classification may be needed to isolate the AE sources of interest. For more discussion on recommendations for the development of this subject, refer to the next chapter, where the discussion of the building block approach and perspective on future work is expanded in detail.

Chapter 7

Recommendations

In this final chapter, this document is concluded by presenting recommendations for future research. These recommendations are founded mainly on the conclusions presented in the previous chapter. First, several recommendations on the subject of the analysis methodology itself are given. Thereafter, recommendations for expansion of the current methodology in the context of the building-block approach are presented.

The current project purposefully relies heavily on existing Python libraries. Nonetheless, development of filtering and processing functions has been found to take up a significant amount of time and resources in the project. It is encouraged to build on the now existing Python functions, instead of reinventing the wheel for every following project on the current subject. It may even be prudent to expand the current tool set into a moderated open source library of AE filtering- and processing functions. In addition, it is judged that because the simple fact that visualisation of the large, multidimensional, datasets produced through AE monitoring is non-trivial, it presents an obstacle for the understanding of AE data. The tools used in the current project, such as fatigue cycle segment-based hit labelling; and t-SNE provide an answer to this challenge to some degree. However, the current project limits itself to only a selection of visualisations, while their use is often still limited by computational performance. With regard to data visualisation for the researcher, it is recommended to investigate graphics-accelerated and/or interactive (e.g. PyQt, VisPy or Bokeh for Python) visualisation libraries, so that this limitation may be overcome for larger subsets of AE data than is currently the case. By extension, it would also be interesting and potentially productive to build an interface based on the type of visualisation in Figure 5.3 to enable, quite literally, intuitive manual labelling of data. Evidently, the prospect of large datasets being produced by AE monitoring is particularly relevant for any envisioned in-situ AE monitoring applications. For this purpose it is possible to simplify relatively complex AE data to one or several meta-parameters such as energy per minute and/or the shape of the AE feature distributions.

Furthermore, with regard to the hit-labelling approach itself, it is asserted that the current approach of splitting the fatigue cycle into four equal phase segments may be a relatively blunt tool, even though in its current use, it is found to be rather robust. It is possible that a more sophisticated labelling procedure may produce a classification that has higher validity with regard to stiffness degradation and/or inherent clustering in the data. For example, the discussion in Section 5.4 on where in the fatigue cycle phase exactly damage occurs and if and how that changes over the duration of N cycles, remains largely unanswered. Therefore, the determination of in which phase in the fatigue cycle (matrix) damage occurs and AE energy is emitted is a recommendable field of study in itself. Answering these questions may yield results of fundamental significance to both AE monitoring of composites and the understanding of composite damage mechanics.

In the current data, three different sources are distinguished. The two main distinguishable sources are approximately organised in the unloading-labelled hits and the loading-labelled hits respectively. Although the loading-labelled hits are asserted to be related to instantaneous damage development with high confidence, the source mechanism of the unloading-labelled hits is left as a largely open question. As discussed in the previous chapter, identifying the source mechanism of this category of hits may lead to an increased understanding of damage mechanics in composites and/or imply a significance of internal friction noise for the application of AE-based structural integrity monitoring of composites. However, it is also possible that this source constitutes a form of noise external to the coupon. It is important to investigate the mechanism behind this source in more detail, ruling out or confirming its rooting in external noise: a set of tests can be repeated without tabs on the coupons, in a different load cell and/or with different clamping pressure to deduce whether this source is related to damage in the tab material/the adhesive, load cell machinery noise, or compressive damage in the clamped region of the coupon. Alternatively, if it cannot be related to the aforementioned mechanisms, this source may be related to internal friction of fractured surfaces, which may be evidenced by the rate dependence found in the AE data for test D5.6. In that case, it warrants investigation whether this type of AE activity can be related to the damage state, e.g. via the detection of existing delaminated surfaces.

Analogously, for the loading-labelled AE activity, bursts of anomalous activity are found in this category, which correlate with divergence from correlation with the stiffness degradation. In other words, the loading-labelled activity still contains a 'signal' and 'noise', with respect to stiffness degradation. Similar assertions as discussed for the unloading-labelled hits in the previous paragraph can be repeated: it is not definitely proven that this category of 'noise' is rooted in material mechanics or in noise sources external to the coupon material. A limited set of observations suggest that this anomalous activity may correlate with other damage mechanisms, or quite simply damage that does not contribute significantly to stiffness degradation under the applied loading. The set of proposed testing conditions mentioned in the previous paragraph can help rule out external noise sources, while loading of $\pm 45°$-dominated coupons at a slightly higher load level, or for a higher number of cycles, can contribute confidence to whether this anomalous activity indeed alludes to other damage mechanisms. The same can be done for the [0/90]-type coupons, because with the stiffness degradation in this coupon type being constrained by the 0°-layers, late-fatigue-stage AE-activity in these samples may represent inter-ply delamination, if noise sources can be ruled out. Finally, a more thorough validation of any of the underlying mechanisms of AE sources may be enabled if post-mortem inspection of the damage in the coupons were made possible. If no suitable NDE methods are available, it may still be worthwhile to switch to a lighter UD weave to a thickness range more compatible with (phased-array) backscatter ultrasonics, enabling inspection via these means. Evidently, this step would be provide less valid evidence for the current material configuration, but from a fundamental perspective on research of AE-based structural integrity monitoring, it can still provide contributions that may otherwise be out of reach.

For future steps in the building-block approach, a few recommendations are presented. First and foremost, it is once more emphasized that any extension of this approach to other materials, loading modes or structures needs to be done in careful and deliberate steps so that any new assertions can be substantiated with high confidence - i.e. building from simple load cases and low load levels. The main reasons for this are that a complex, highly-active (AE) dataset drastically reduces separability by the fatigue cycle labelling approach used in the current project; and the fact that sensor saturation in itself complicates any validation strategy unnecessarily. From this framework, future steps in the building-block approach can be explored. It is asserted that it may be advisable to repeat testing at a lower threshold (e.g. 54 instead of 57 [dB]). Activity of hits associated with stiffness degradation (and thereby matrix cracking) in the current test set is found to be concentrated at a relatively low amplitudes. Therefore, it is considered worthwhile to investigate whether the mode of the amplitude shifts to lower values in time, and if the effect(s) described in Figure 5.12 can be confirmed or falsified. Alternatively to performing extra tests at a higher threshold, it may suffice to instead re-process the current data with a *higher* threshold. In general, it must be noted that in any AE-based study, the threshold strongly influences the resulting AE data. The threshold can be set such that the majority of background AE activity is filtered out, though this will inevitably also filter activity from source mechanisms based in damage processes. There is therefore no universal recommendation for the threshold setting, save for that the researcher must be aware of its consequences both as a filter and to AE waveform feature values.

In addition, it is noted that in the current study no particular focus has been given to the influence of changing attenuation- and propagation properties of AE waves as the coupon becomes progressively more damaged. It may be expected that at any time, priorly developed damage influences the character of subsequent hits, although no explicit evidence for this has been examined in the current project: even though it is found that the energy- and amplitude of loading-labelled hits remain relatively constant for the current test parameters, it is recommendable to dedicate an analysis to quantifying these effects for a selection of waveform parameters. It is evident that a progressive change in the character of AE hits may have important consequences for the purposes of (in-situ) AE-based structural integrity monitoring of composites.

When the building-block is expanded to other damage mechanisms, it is important to note that the stiffness degradation becomes less and less of a representative metric for damage in the coupon as damage mechanisms become more multi-modal and self-interacting. This limitation may be managed through application of a mixed mechanics- and data-driven approach: first, the character of matrix cracking can be determined from strain-controlled T-T tests of pure 90° coupons. Using the character of these hits, it can be more confidently asserted if and how dominant the prevalence of other damage mechanisms than matrix cracking is in the [0/90]-type coupons in the current project. A series of double cantilever beam T-T fatigue tests, using the same hit labelling approach as in the current project, may also be used to increase confidence in whether or not the secondary AE activity found in the [0/90]-type coupons is in fact delamination, and if it can be universally detected at all. Additionally, a limited set of studies on environmental influences indicates that AE activity reduces with material ageing (also seen in the current project) and exposure to, e.g., water and saline environments. It should be prudent to quantify this effect for the current material configuration and verify whether it also holds after filtering. Expanding the building-block even further, to the structure level and/or in-situ trials, it is anticipated that effects like attenuation, dispersion, scattering and diffraction may influence the character of AE hits more significantly. The actual significance of these effects can be asserted in these types of trials, while source location in conjunction with parameter correction techniques [71], or localised measurements may pose a suitable approach to deal with these effects.

In conclusion, it is noted that any methodology outlined here to expand the building-block approach of AE-based structural integrity monitoring is merely one of many possible implementations: fully mapping out a methodology culminating in achieving the ultimate goal of universal AE-based structural integrity monitoring capabilities of composite structures would be unavoidably speculative in nature. The building-block approach provides a useful framework to apply the scientific method towards an incrementally complex goal, but the process itself must be iterative, building on the most reliable assertions of each successive step.

References

[1] private correspondence with Laurent Morel, technical specialist composite structures at Damen Schelde Naval Shipbuilding, 2015.

[2] C. Li, F. Ellyin, and A. Wharmby. On matrix crack saturation in composite laminates. *Composites Part B: Engineering*, 34(5):473–480, 2003.

[3] M.A. Hamstad. Thirty years of advances and some remaining challenges in the application of acoustic emission to composite materials. In *Acoustic Emission - Beyond the Millennium*, pages 77–91. Elsevier Science, 2000.

[4] Damen Shipyards Group website. http://www.damen.com/. [accessed 2015-06-16].

[5] P. Davies, D. Choqueuse, and H. Devaux. Failure of polymer matrix composites in marine and off-shore applications. In *Failure Mechanisms in Polymer Matrix Composites*, pages 53–78. Woodhead Publishing Limited, 2012.

[6] Y. Okumoto, Y. Takeda, M. Mano, and T. Okada. *Design of ship hull structures.* Springer, 1st edition, 2009.

[7] N. Godin, S. Huguet, and R. Gaertner. Influence of hydrolytic ageing on the acoustic emission signatures of damage mechanisms occurring during tensile tests on a polyester composite: Application of a Kohonen's map. *Composite Structures*, 72(1):79–85, 2006.

[8] E.P. Gellert and D.M. Turley. Seawater immersion ageing of glass-fibre reinforced polymer laminates for marine applications. *Composites Part A: Applied Science and Manufacturing*, 30(11):1259–1265, 1999.

[9] A.G. Gibson, J.T. Evans, G. Kotsikos, S.D. Speake, and J.M. Hale. Fatigue of marine laminates in aqueous environments. *Plastics, rubber and composites*, 29(10):533–538, 2000.

[10] E. Poodts, G. Minak, and A. Zucchelli. Impact of sea-water on the quasi static and fatigue flexural properties of GFRP. *Composite Structures*, 97:222–230, 2013.

[11] J.-C. Anifrani, C. Le Floc'h, D. Sornette, and B. Souillard. Universal Log-Periodic Correction to Renormalization Group Scaling for Rupture Stress Prediction From Acoustic Emissions. *Journal de Physique I*, 5(6):631–638, 1995.

[12] A. Johansen and D. Sornette. Critical Ruptures. *The European Physical Journal B*, 18(1):24, 2000.

[13] R. Talreja. Fatigue of Composite Materials: Damage Mechanisms and Fatigue-Life Diagrams. In *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, volume 378, pages 461–475, 1981.

[14] M.V. Cid Alfaro, A.S.J. Suiker, and R. de Borst. Transverse failure behavior of fiber-epoxy systems. *Journal of Composite Materials*, 44(12):1493–1515, 2010.

[15] S. Huguet, N. Godin, R. Gaertner, L. Salmon, and D. Villard. Use of acoustic emission to identify damage modes in glass fibre reinforced polyester. *Composites Science and Technology*, 62(10-11):1433–1444, aug 2002.

[16] M. Bhat and C. Murthy. Fatigue damage stages in unidirectional glass-fibre-epoxy composites: identification through acoustic emission technique. *International Journal of Fatigue*, 15(5):401–405, 1993.

[17] S. Koussios and R. Marissen. TUDelft course material for AE4634: Fracture of advanced materials, 2014.

[18] S. Barré and M.L. Benzeggagh. On the use of acoustic emission to investigate damage mechanisms in glass-fibre-reinforced polypropylene. *Composites Science and Technology*, 52(3):369–376, jan 1994.

[19] W.H. Prosser, K.E. Jackson, S. Kellas, B.T. Smith, J. Mckeon, and A. Friedman. Advanced , waveform based acoustic emission detection of matrix cracking in composites. *Materials Evaluation*, 53(9):1052–1058, 1995.

[20] E.K. Gamstedt and B.A Sjögren. Micromechanisms in tension-compression fatigue of composite laminates containing transverse plies. *Composites Science and Technology*, 59(2):167–178, 1999.

[21] L.E. Asp, L.A. Berglund, and R. Talreja. Prediction of matrix-initiated transverse failure in polymer composites. *Composites Science and Technology*, 56(9):1089–1097, 1996.

[22] N. Godin, S. Huguet, and R. Gaertner. Integration of the Kohonen's self-organising map and k-means algorithm for the segmentation of the AE data collected during tensile tests on cross-ply composites. *NDT and E International*, 38(4):299–309, 2005.

[23] P. Liu, R.M. Groves, and R. Benedictus. 3D monitoring of delamination growth in a wind turbine blade composite using optical coherence tomography. *NDT & E International*, 64:52–58, 2014.

[24] P. Liu, R.M. Groves, and R. Benedictus. Optical Coherence Elastography for measuring the deformation and strain within glass fibre composite. *Applied Optics*, 53(22):5070–5077, 2014.

[25] C. González and J. LLorca. Mechanical behavior of unidirectional fiber-reinforced polymers under transverse compression: Microscopic mechanisms and modeling. *Composites Science and Technology*, 67(13):2795–2806, 2007.

[26] S.T. Pinho, R. Gutkin, S. Pimenta, N.V. de Carvalho, and P. Robinson. Fibre-dominated compressive failure in polymer matrix composites. In *Failure Mechanisms in Polymer Matrix Composites*, pages 53–78. Woodhead Publishing Limited, 2012.

[27] K. Oguni, C. Y. Tan, and G. Ravichandran. Failure Mode Transition in Unidirectional E-Glass/Vinylester Composites under Multiaxial Compression. *Journal of Composite Materials*, 34(24):2081–2097, 2000.

[28] N. Kosmann, J.M. Karsten, M. Schuett, K. Schulte, and B. Fiedler. Determining the effect of voids in GFRP on the damage behaviour under compression loading using acoustic emission. *Composites Part B: Engineering*, 70:184–188, 2015.

[29] C. Kassapoglou. TUDelft course material for AE4509: Advanced design and optimization of composite structures, 2014.

[30] R. Marissen and H.R. Brouwer. The significance of fibre microbuckling for the flexural strength of a composites. *Composites science and technology*, 59:327–330, 1999.

[31] M. Fotouhi and M. Ahmadi Najafabadi. Investigation of the mixed-mode delamination in polymer-matrix composites using acoustic emission technique. *Journal of Reinforced Plastics and Composites*, 33(19):1767–1782, 2014.

[32] S. Benmedakhene, M. Kenane, and M. L. Benzeggagh. Initiation and growth of delamination in glass/epoxy composites subjected to static and dynamic loading by acoustic emission monitoring. *Composites Science and Technology*, 59(2):201–208, 1999.

[33] ASTM International. D5528 Standard Test Method for Mode I Interlaminar Fracture Toughness of Unidirectional Fiber-Reinforced Polymer Matrix Composites, 2013.

[34] O. Ceysson, M. Salvia, and L. Vincent. Damage mechanisms characterisation of carbon fibre/epoxy composite laminates by both electrical resistance measurements and acoustic emission analysis. *Scripta Materialia*, 34(8):1273–1280, 1996.

[35] M.O.W. Richardson and M.J. Wisheart. Review of low-velocity impact properties of composite materials. *Composites Part A: Applied Science and Manufacturing*, 27(12):1123–1133, 1996.

[36] Z. Aslan, R. Karakuzu, and B. Okutan. The response of laminated composite plates under low-velocity impact loading. *Composite Structures*, 59(1):119–127, 2003.

[37] T. Shyr and Y. Pan. Impact resistance and damage characteristics of composite laminates. *Composite Structures*, 62(2):193–203, 2003.

[38] A. Malhotra and F.J. Guild. Impact damage to composite laminates: Effect of impact location. *Applied Composite Materials*, 21(1):165–177, 2014.

[39] S. Sanchez-Saez, E. Barbero, R. Zaera, and C. Navarro. Compression after impact of thin composite laminates. *Composites Science and Technology*, 65(13):1911–1919, 2005.

[40] ASTM International. E1316 Standard Terminology for Nondestructive Examinations, 2014.

[41] R.K. Miller, E.v.K. Hill, and P.O. Moore. *Nondestructive Testing Handbook, vol.5, Acoustic Emission Testing*. ASNT, 3rd edition, 2005.

[42] The Collaboration for NDT Education. Introduction to Acoustic Emission Testing. `https://www.nde-ed.org/EducationResources/CommunityCollege/Other%20Methods/AE/AE_Index.htm`, 2014. [accessed 2015-06-02].

[43] R.W.B. Stephens and A.A. Pollock. Waveforms and Frequency Spectra of Acoustic Emissions. *The Journal of the Acoustical Society of America*, 50(3B):904, 1971.

[44] M.A. Hamstad. A review: Acoustic emission, a tool for composite-materials studies. *Experimental Mechanics*, 26(1):7–13, 1986.

[45] S. Pant, J. Laliberte, M. Martinez, and B. Rocha. Derivation and experimental validation of Lamb wave equations for an n-layered anisotropic composite laminate. *Composite Structures*, 111(1):566–579, 2014.

[46] J. Bohse. Acoustic emission characteristics of micro-failure processes in polymer blends and composites. *Composites Science and Technology*, 60(8):1213–1226, 2000.

[47] M.R. Gorman. Plate wave acoustic emission. *The Journal of the Acoustical Society of America*, 90(1):358, 1991.

[48] S.K. Datta. Wave propagation in laminated composite plates. *The Journal of the Acoustical Society of America*, 83(6):2020, 1988.

[49] A.H. Nayfeh. The general problem of elastic wave propagation in multilayered anisotropic media. *The Journal of the Acoustical Society of America*, 89(4):1521, 1991.

[50] M.R. Gorman. Plate waves produced by transverse matrix cracking. *Ultrasonics*, 29:245–251, 1991.

[51] A.A. Pollock. Critical AE Problems for the Researcher. *Journal of Acoustic Emission*, 9(2):140–141, 1990.

[52] W.H. Prosser. Applications of advanced, waveform based AE techniques for testing composite materials. In *Proc. SPIE*, volume 2944, pages 146–153, 1996.

[53] M. Surgeon and M. Wevers. Modal analysis of acoustic emission signals from CFRP laminates. *NDT & E International*, 32(6):311–322, sep 1999.

[54] W.H. Prosser, M.D. Seale, and B.T. Smith. Time-frequency analysis of the dispersion of Lamb modes. *The Journal of the Acoustical Society of America*, 105(5):2669–2676, 1999.

[55] P. Ochôa, V. Infante, J.M. Silva, and R.M. Groves. Detection of multiple low-energy impact damage in composite plates using Lamb wave techniques. *Composites Part B: Engineering*, 80:291–298, 2015.

[56] M.B. Barroso, D. Barazanchy, M. Martinez, R.M. Groves, and R. Benedictus. Time Reversal of Lamb Waves for Damage Detection in Thermoplastic Composites. In *International conference on adaptive structures and technologies*, 2013.

[57] J. Kaiser. *An investigation into the occurrence of noises in tensilte tests or a study of acoustic phenomena in tensilte tests.* Phd thesis, Technische Hochschule München, 1950.

[58] D. Zarouchas and D. van Hemelrijck. Mechanical characterization and damage assessment of thick adhesives for wind turbine blades using acoustic emission and digital image correlation techniques. *Journal of Adhesion Science and Technology*, 28:37–41, 2014.

[59] C.U. Grosse and M. Ohtsu. *Acoustic Emission Testing. Basics for Research - Applications in Civil Engineering.* Springer, 1st edition, 2008.

[60] ASTM International. E750 Standard practice for characterizing acoustic emission instrumentation, 2010.

[61] ASTM International. E976 Standard Guide for Determining the Reproducibility of Acoustic Emission Sensor Response, 2010.

[62] Vallen Systeme GmbH. Acoustic Emission Sensors specification, 2012.

[63] P.E. Mix. *Introduction to non-destructive testing.* Wiley, 2nd edition, 2005.

[64] G. Qi. Wavelet-based AE characterization of composite materials. *NDT and E International*, 33(3):133–144, apr 2000.

[65] A. Lee, G. Wang, S. Ternowchek, and S.F. Botten. Structural Health Monitoring on ships using Acoustic Emission testing. In *SSC symposium.* SSC, 2014.

[66] I. Baran, M. Nowak, A. Jagenbrein, and H. Bugłacki. Acoustic Emission Monitoring of Structural Elements of a Ship for Detec- tion of Fatigue and Corrosion Damages. In *30th European conference on Acoustic Emission Testing*, 2012.

[67] L. Rogers. Acoustic Emission Method: Remote Monitoring and VALID Trials - corporate presentation of VALID results. Technical report, Lloyd's Register, 2009.

[68] Ship Structure Committee. SSC-468: Development of a Structural Health Monitoring Prototype for Ship Structures. Technical report, Ship Structure Committee, 2013.

[69] K. Stambaugh, I. Drummen, C. Cleary, R. Sheinberg, and M. Kaminski. Structural Fatigue Life Assessment and Sustainment Implications for a New Class of US Coast Guard Cutters. Technical report, United States Coast Guard, 2014.

[70] A. Caclin, S. McAdams, B.K. Smith, and S. Winsberg. Acoustic correlates of timbre space dimensions: a confirmatory study using synthetic tones. *The Journal of the Acoustical Society of America*, 118(1):471–482, 2005.

[71] S.Kh. Al-Jumaili, K.M. Holford, M.J. Eaton, and R. Pullin. Parameter Correction Technique (PCT): a novel method for Acoustic Emission Characterisation in Large-Scale Composites. *Composites Part B: Engineering*, 75:336–344, 2015.

[72] T. Uenoya. Acoustic emission analysis of interfacial fracture of laminated fabric polymer matrix composite. *Journal of Acoustic Emission*, 13:95 – 102, 1995.

[73] M. Giordano, A. Calabro, C. Esposito, A. D'Amore, and L. Nicolais. An acoustic-emission characterization of the failure modes in polymer-composite materials. *Composites Science and Technology*, 58(12):1923–1928, 1998.

[74] C. R. Ramirez-Jimenez, N. Papadakis, N. Reynolds, T. H. Gan, P. Purnell, and M. Pharaoh. Identification of failure modes in glass/polypropylene composites by means of the primary frequency content of the acoustic emission event. *Composites Science and Technology*, 64(12):1819–1827, sep 2004.

[75] A.R. Oskouei and M. Ahmadi. Acoustic Emission Characteristics of Mode I Delamination in Glass/Polyester Composites. *Journal of Composite Materials*, 44(7):793–807, 2010.

[76] D.B.B. Ferreira, R.R. da Silva, J.M.A. Rebello, and M.H.S. Siqueira. Failure mechanism characterisation in composite materials using spectral analysis and the wavelet transform of acoustic emission signals. *Insight - Non-Destructive Testing and Condition Monitoring*, 46(5):282–289, 2004.

[77] R. de Oliveira and A.T. Marques. Health monitoring of FRP using acoustic emission and artificial neural networks. *Computers & Structures*, 86(3-5):367–373, 2008.

[78] N. Godin, S. Huguet, R. Gaertner, and L. Salmon. Clustering of acoustic emission signals collected during tensile tests on unidirectional glass/polyester composite using supervised and unsupervised classifiers. *NDT and E International*, 37(4):253–264, jun 2004.

[79] M.R. Gorman and W.H. Prosser. AE Source Orientation by Plate Wave Analysis. *Journal of Acoustic Emission*, 9(4):282–288, 1990.

[80] D. Crivelli, M. Guagliano, and A. Monici. Development of an artificial neural network processing technique for the analysis of damage evolution in pultruded composites with acoustic emission. *Composites Part B: Engineering*, 56:948–959, 2014.

[81] T. Kohonen. The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990.

[82] J.P. Mccrory, S.Kh. Al-jumaili, D. Crivelli, M.R. Pearson, M.J. Eaton, C.A. Featherston, M. Guagliano, K.M. Holford, and R. Pullin. Composites : Part B Damage classification in carbon fibre composites using acoustic emission : A comparison of three techniques. *Composites Part B*, 68:424–430, 2015.

[83] D. Crivelli, M. Guagliano, M.J. Eaton, M. Pearson, S.Kh. Al-Jumaili, K.M. Holford, and R. Pullin. Localisation and Identification of fatigue Matrix Cracking and Delamination in a Carbon Fibre Panel by Acoustic Emission. *Composites Part B: Engineering*, 74:1–12, 2015.

[84] D.D. Doan, E. Ramasso, V. Placet, S. Zhang, L. Boubakar, and N. Zerhouni. An unsupervised pattern recognition approach for AE data originating from fatigue tests on polymer composite materials. *Mechanical Systems and Signal Processing*, 64-65:465–478, 2015.

[85] M.G.R. Sause, A. Gribov, A. R. Unwin, and S. Horn. Pattern recognition approach to identify natural clusters of acoustic emission signals. *Pattern Recognition Letters*, 33(1):17–23, 2012.

[86] C. Bhat, M.R. Bhat, and C.R.L. Murthy. Acoustic emission characterization of failure modes in composites with ANN. *Composite Structures*, 61(3):213–220, 2003.

[87] Submitted: J.A. Pascoe, D. Zarouchas, R.C. Alderliesten, and R. Benedictus. Using Acoustic Emission to Understand Fatigue Crack Growth Within a Single Load Cycle. *International Journal of Fracture*, 2016.

[88] ASTM International. D3039/D3039M Standard test method for tensile properties of polymer matrix composite materials, 2014.

[89] S. van der Walt, S. C. Colbert, and G. Varoquaux. The numpy array: A structure for efficient numerical computation. *Computing in Science Engineering*, 13(2):22–30, March 2011.

[90] W. McKinney. Data structures for statistical computing in python. In S. van der Walt and J. Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 51 – 56, 2010.

[91] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[92] R.C. Alderliesten. TUDelft course material for AE4735: Fatigue, 2014.

[93] Python software foundation. Python language documentation, version 2.7. https://docs.python.org/2.7/. [accessed 2016-06-02].

[94] E. Jones, T. Oliphant, P. Peterson, et al. Scipy: Open source scientific tools for Python (documentation). http://www.scipy.org/. [accessed 2016-05-26].

[95] M. Fotouhi and M. A. Najafabadi. Acoustic emission-based study to characterize the initiation of delamination in composite materials. *Journal of Thermoplastic Composite Materials*, 29(4):519–537, 2016.

[96] J. Kakakasery, V. Arumugam, K. Abdul Rauf, D. Bull, A.R. Chambers, C. Scarponi, and C. Santulli. Cure cycle effect on impact resistance under elevated temperatures in carbon prepreg laminates investigated using acoustic emission. *Composites part B: Engineering*, 75:298–306, 2015.

[97] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn documentation. http://scikit-learn.org/. [accessed 2016-05-26].

[98] A. Hyvärinen and E. Oja. Independent component analysis: Algorithms and applications. *Neural Networks*, 13(4-5):411–430, 2000.

[99] L. van der Maaten. Visualizing high-dimensional data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

[100] L. van der Maaten. t-SNE github page. https://lvdmaaten.github.io/tsne/. [accessed 2016-05-26].

[101] W.H. Prosser. Advanced AE techniques in composite materials research. *Journal of Acoustic Emission*, 14:S1–S11, 1996.

[102] D. Doan, E. Ramasso, V. Placet, L. Boubakar, and N. Zerhouni. Application of an unsupervised pattern recognition approach for AE data originating from fatigue tests on CFRP. In *31th Conference of the European Working Group on Acoustic Emission*, pages 465–478, 2014.