

Vibration-based damage localisation: Impulse response
identification and model updating methods

Von der Fakultät für Bauingenieurwesen und Geodäsie
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des Grades

DOKTOR-INGENIEUR
DR.-ING.

genehmigte Dissertation von
BENEDIKT HOFMEISTER, M.ENG

2023

Referent: Prof. Dr.-Ing. habil. Raimund Rolfes
Korreferent: Prof. Dr. Christof Devriendt
Tag der Promotion: 21.04.2023

Abstract

Structural health monitoring has gained more and more interest over the recent decades. As the technology has matured and monitoring systems are employed commercially, the development of more powerful and precise methods is the logical next step in this field. Especially vibration sensor networks with few measurement points combined with utilisation of ambient vibration sources are attractive for practical applications, as this approach promises to be cost-effective while requiring minimal modification to the monitored structures. Since efficient methods for damage detection have already been developed for such sensor networks, the research focus shifts towards extracting more information from the measurement data, in particular to the localisation and quantification of damage.

Two main concepts have produced promising results for damage localisation. The first approach involves a mechanical model of the structure, which is used in a model updating scheme to find the damaged areas of the structure. Second, there is a purely data-driven approach, which relies on residuals of vibration estimations to find regions where damage is probable. While much research has been conducted following these two concepts, different approaches are rarely directly compared using the same data sets. Therefore, this thesis presents advanced methods for vibration-based damage localisation using model updating as well as a data-driven method and provides a direct comparison using the same vibration measurement data.

The model updating approach presented in this thesis relies on multi-objective optimisation. Hence, the applied numerical optimisation algorithms are presented first. On this basis, the model updating parameterisation and objective function formulation is developed. The data-driven approach employs residuals from vibration estimations obtained using multiple-input finite impulse response filters. Both approaches are then verified using a simulated cantilever beam considering multiple damage scenarios. Finally, experimentally obtained data from an outdoor girder mast structure is used to validate the approaches. In summary, this thesis provides an assessment of model updating and residual-based damage localisation by means of verification and validation cases. It is found that the residual-based method exhibits numerical performance sufficient for real-time applications while providing a high sensitivity towards damage. However, the localisation accuracy is found to be superior using the model updating method.

KEYWORDS: structural health monitoring, damage localisation, finite impulse filter, model-updating, multi-objective optimisation

Zusammenfassung

Die Zustandsüberwachung von Bauwerken hat in den letzten Jahrzehnten zunehmendes Interesse in Forschung und Industrie geweckt. Nachdem erste Überwachungssysteme bereits kommerziell eingesetzt werden, ist die Entwicklung leistungsfähigerer und präziserer Methoden der logische nächste Schritt. Insbesondere Sensornetzwerke mit wenigen Schwingungs-Messpunkten und Nutzung ambienter Erregungsquellen sind für praktische Anwendungen attraktiv, da dieser Ansatz kostengünstig und mit minimalen Änderungen an den überwachten Strukturen umgesetzt werden kann. Da für solche Sensornetzwerke bereits effiziente Methoden zur Schadensdetektion entwickelt wurden, rückt die Extraktion weiterer Informationen zur Schadenslokalisierung und -quantifizierung in den Fokus der Forschung.

Insbesondere zwei Ansätze liefern vielversprechende Ergebnisse für die Schadenslokalisierung: Der erste Ansatz basiert auf einem mechanischen Modell der Struktur, welches mittels einer Modellanpassung zur Lokalisierung möglicher Schäden genutzt wird. Der zweite Ansatz ist rein datengestützt und nutzt Residuen von Schwingungsschätzungen, um Schäden zu lokalisieren. Diese beiden Ansätze wurden bereits ausgiebig erforscht, aber nur selten direkt auf Grundlage derselben Datensätze verglichen. In dieser Dissertation werden daher Methoden zur schwingungsbasierten Schadenslokalisierung mittels Modellanpassung sowie einer datengestützten Methode weiter entwickelt und ein direkter Vergleich unter Verwendung derselben Schwingungsmessdaten durchgeführt.

Der in dieser Arbeit vorgestellte Ansatz zur Modellanpassung basiert auf einer Mehrziel-Optimierung. Daher werden zunächst die verwendeten numerischen Optimierungsalgorithmen eingeführt. Die Parametrisierung der Modellanpassung und die Formulierung der Zielfunktion werden auf Grundlage der Optimierungsalgorithmen entwickelt. Der datengetriebene Ansatz verwendet Residuen aus Schätzungen, welche mithilfe von finiten Impulsantwortfiltern mit mehreren Eingängen ermittelt werden. Beide Ansätze wurden anhand eines simulierten Kragarms unter Berücksichtigung mehrerer Schadensszenarien verifiziert. Anschließend wurden experimentell erhobene Daten einer Maststruktur verwendet, um die Ansätze zu validieren. Zusammenfassend bietet diese Arbeit eine Bewertung der Modellanpassung und der residualbasierten Schadenslokalisierung anhand von Verifikations- und Validierungsfällen. Es wurde festgestellt, dass die residualbasierte Methode eine für Echtzeitanwendungen ausreichende numerische Effizienz aufweist und gleichzeitig eine hohe Empfindlichkeit gegenüber Schäden bietet. Mit der Modellanpassung lässt sich hingegen eine genauere Lokalisierung erreichen.

SCHLAGWORTE: Zustandsüberwachung, Schadenslokalisierung, finite Impulsantwortfilter, Modellanpassung, Mehrzieloptimierung



Table of Contents

1	Introduction	1
1.1	Motivation	1
1.2	Structural health monitoring	2
1.3	Damage detection and localisation	4
1.4	Localisation using mechanical models	5
1.5	Localisation using data-driven methods	6
1.6	Research gap	8
1.7	Scope and structure of the work	9
2	Numerical optimisation algorithms	11
2.1	Global optimisation	11
2.1.1	Optimisation problems	12
2.1.2	Constrained optimisation	13
2.2	Classification of numerical optimisation algorithms	14
2.2.1	Algorithms based on random numbers	15
2.2.2	Deterministic algorithms	16
2.2.3	Research gap	17
2.3	Global Pattern Search	18
2.3.1	Search pattern	18
2.3.2	Algorithm description	19
2.3.3	Discretisation of the design variable space	22
2.3.4	Reuse of cached samples	22
2.3.5	Verification using test functions	23
2.3.6	Benchmark using test functions	27
2.4	Multi-Objective Global Pattern Search	32
2.4.1	Non-dominated sorting and Pareto frontiers	33
2.4.2	Multi-objective extension for Pattern Search	34
2.4.3	Details of the algorithm	35
2.4.4	Verification using test functions	37

3	Damage localisation using model updating	47
3.1	Fundamentals of model updating	47
3.1.1	Finite element model updating	47
3.1.2	Parameterisations	48
3.1.3	Objective functions	49
3.1.4	Modal analysis	51
3.2	Damage localisation method	53
3.2.1	Parameterisation of damage distribution	53
3.2.2	Objective function formulation	54
3.3	Verification using numerical simulation	55
4	Damage localisation using impulse response identification	65
4.1	Theory of impulse response filters	65
4.1.1	Data-driven damage localisation	66
4.1.2	Transmissibility functions and impulse response filters	67
4.1.3	System and control theory considerations	68
4.2	Damage localisation method	69
4.2.1	Single-input impulse response identification	69
4.2.2	Expansion to the multiple-input/multiple-output case	71
4.2.3	Interpolated finite impulse response filters	73
4.2.4	Computational performance	77
4.2.5	Damage-sensitive residuals	77
4.3	Verification using numerical simulation	79
5	Experimental validation	85
5.1	Girder-mast structure	85
5.1.1	Measurement campaign	86
5.1.2	Impact of structural damage	88
5.2	Model-based localisation of symmetrical damage	89
5.2.1	Modal analysis	89
5.2.2	Finite element model	90
5.2.3	Parameterisation	91
5.2.4	Damage identification and localisation	92
5.3	Data-driven localisation of symmetrical damage	96
5.3.1	Identification of impulse response filters	96
5.3.2	Damage identification and localisation	98
5.4	Comparison of localisation results	102
5.5	Unsymmetrical damage scenario	102
5.6	Comparison of numerical performance	107

6	Benefits and limitations	109
6.1	Numerical optimisation	109
6.2	Damage localisation based on model updating	109
6.3	Data-driven damage localisation	110
6.4	Environmental and operating conditions	111
7	Conclusion	113
7.1	Summary	113
7.2	Future research topics	114

Table of Abbreviations

NSGA-II Non-dominated Sorting Genetic Algorithm-II.

DIRECT DIviding RECTangles algorithm.

FDD Frequency-Domain Decomposition.

FE Finite Element.

FIR Finite Impulse Response.

GPS Global Pattern Search.

IEPE Integrated Electronics Piezo Electric.

IFIR Interpolated Finite Impulse Response.

MAC Modal Assurance Criterion.

MOGPS Multi-Objective Global Pattern Search.

SHM Structural Health Monitoring.

SSI Stochastic Subspace Identification.

Chapter 1

Introduction

1.1 Motivation

As infrastructure becomes increasingly complex and growing numbers of facilities need to be operational at all times, maintenance becomes more and more important. Traditionally, maintenance is carried out based on a predetermined schedule of inspections and planned overhauls. To reduce the workload of technicians and reduce cost, monitoring data can be leveraged to facilitate more flexible maintenance schemes. Based on this data, algorithms are continuously employed to observe the data and decide in real-time whether damage is likely. The monitoring then makes it possible to reduce down-time and enhance the maintenance planning by detection of damage in an early stage. A condition-based maintenance plan therefore enables deferring inspections to a point in time where overhauls are actually necessary and intermediate inspections can be largely skipped [95].

Because maintenance costs account for a large part of the operational expenditure of infrastructure facilities, Structural Health Monitoring (SHM) has received more and more attention in the past decades [60]. The anticipated optimisation of maintenance procedures and the associated cost savings as well as the potential benefits to public safety also lead to a high interest from industry in this area. A major drawback of monitoring solutions, however, is the cost of equipping structures and machinery with sensors, recording and continuously processing the data. Recently, miniaturisation and dropping costs of sensors and network-capable devices make industrial applications economically feasible [23].

This thesis therefore addresses the topic of vibration-based damage localisation in an effort to advance the algorithms that enable structural health monitoring. To this end, a damage localisation method based on a mechani-

cal model as well as a data-driven method are presented and evaluated using numerical simulations and measurement data.

1.2 Structural health monitoring

According to Farrar and Worden [36], the capabilities of SHM methods can be classified using five levels, as shown in Figure 1.1. These SHM levels are ordered and depend on each other. In this scheme, the most basic functionality, and thus the fundamental level, is damage detection. On top of the detection, the subsequent levels are localisation, classification, quantification and lifetime prognosis. The specificity and practical usefulness increases with the SHM level, however, achieving higher levels also requires the use of more complex methods and algorithms.




Figure 1.1: Levels of SHM according to Farrar and Worden [36].

In this work, the three levels detection, localisation and quantification are addressed. In this context, detection refers to identifying that a damage occurred somewhere in the monitored structure. A method is capable of localisation, when it can constrain the damage to specific geometric regions of the structure. Finally, quantification is the capability to determine the extent of the damage. Classification is less of an issue due to the unsupervised nature of the methods discussed in this thesis. Hence, damage is expected to manifest as a decrease in the structural stiffness and a detailed classification is therefore skipped implicitly. This is not only true for the experimental structures considered in this work, but also for more practical setups and construction materials [73].

SHM methods rely on sensors and measurement data from the monitored engineering structure. In structures such as wind turbines, bridges, skyscrapers or hydroelectric dams, there exist a multitude of possible physical quantities which can be measured. Strain, displacement, velocity, acceleration, tilt angle

and temperature are commonly used. The capabilities of a monitoring system are dependent on the measurement locations and the types of sensors. Depending on the measurement setup, local techniques which monitor material integrity on smaller length scales as well as global techniques, which can detect structural changes anywhere in the structure, can be used.

On the one hand, SHM systems can be implemented using a passive setup which only relies on measurement signals present due to the environmental and operational dynamic excitation experienced by the monitored structure. Such systems are also referred to as ‘output-only’ setups, as there is no active transfer of energy to the structure. On the other hand, an SHM system can also employ an active monitoring strategy, which involves actuators that can manipulate the structure in a controlled way to obtain an insight into the structural response [112]. Passive systems are usually cheaper and simpler, while active systems may provide better data quality due to the known signals that actuators can provide.

While in laboratory settings it is possible to apply predetermined mechanical excitation signals to a structure, such as active local monitoring of plate-like structures using Lamb waves [109], this is usually complex and expensive in a long-term monitoring setting on large structures. This thesis therefore concentrates on methods which extract information about the mechanical integrity of the monitored structure based on passive excitation using ambient vibration sources. The ambient nature of the exciting forces also means that they can usually not be directly measured [36]. As a result, vibration time series measured under ambient excitation cannot be directly employed for analysis. Rather, the data has to be processed using an identification method, which yields information about the underlying structural dynamic behaviour of the structure monitored. Ambient vibration-based monitoring approaches can be classified by the amount of information that is available in addition to the actual vibration measurement signals.

- Model-based methods rely on an accurate mechanical model of the structure to be monitored and judge upon its integrity by comparing the behaviour of the model to measured data
- Data-driven methods, also known as model-free methods, take a signal processing approach and deduce structural changes purely from differences in the characteristics of the measurement data

An issue, that affects many areas of monitoring, is the variability of environmental and operating conditions. Due to these fluctuations, e.g. ambient temperature, wind speeds or rotational speeds of machinery, the mechanical system properties may change over time. This poses a challenge for monitoring, as it is often hard to discern whether changes in the structure are due

to environmental factors or due to damage. Since environmental fluctuations happen on a rather long timescale compared to the vibration period of typical engineering structures, a distinction can be made between short-term and long-term dynamics. Here, short-term dynamics refers to the structural vibration in the lowest eigenfrequencies of the structure, while long-term dynamics refers to the changes in structural behaviour due to changes in the environmental and operating conditions. As this thesis focuses on the development of localisation methods using short-term dynamics, the treatment of long-term dynamics effects is intentionally omitted for the sake of clarity and brevity.

This thesis presents novel approaches for monitoring based on ambient vibration using both model-based and data-driven techniques. Specifically, a model updating approach based on multi-objective optimisation and damage distribution functions is pursued. The data-driven approach presented in this thesis is based on a residual power metric obtained using multiple-input finite impulse response filters. Verification and validation is carried out using simulations and measurement data of an experimental structure, respectively.

1.3 Damage detection and localisation

As implied by Figure 1.1, damage localisation poses a substantially harder challenge than damage detection. More precisely, a localisation method can also detect a damage, while a detection method cannot necessarily localise a damage. Hence, for damage detection, many structural dynamic identification approaches are suitable when combined with statistical methods. For example, Tsiapoki [117] presented advanced damage-sensitive features extracted using auto-regressive models. Haeckell et al. [50] and Penner [89] presented damage detection approaches based on the operational modal analysis methods Stochastic Subspace Identification and Frequency-Domain Decomposition, respectively.

Damage localisation requires a more specialised approach than damage detection. To achieve this, many authors follow a model-based approach using a finite element discretisation of the structure to be monitored. In this approach, a model updating scheme is used to fit the model to the measurement data, which thereby recovers the damage position as well as its severity. This approach was followed by Jahjouh [54] and Schröder et al. [104], who proposed different optimisation methods for the model updating problem. A more recently developed approach is the damage localisation using a purely data-driven approach. To this end, Wernitz et al. [123] presented a damage localisation method based on state-space projection and Kalman filtering.

The following sections discuss the model-based and data-driven localisation approaches in more detail.

1.4 Localisation using mechanical models

Model-based approaches rely on a mechanical model of the structure to be monitored. For vibration-based monitoring, the model has to be able to capture the vibration behaviour of the structure. In simple cases, this can be achieved by employing mass and spring models which form a multiple degree of freedom system. For more complex structures, the finite element method is employed to aid in the analysis.

Using the mechanical model, information captured by the measurement instrumentation can be interpreted in terms of changes in the structural properties. Therefore, the model is parameterised and these parameters are tuned numerically so that the model matches the measured data as closely as possible. This concept is known as model updating in literature [37, 82]. Some ways to implement the model update are

- Analytical gradients of stiffness and mass matrices with respect to modal parameters [71]
- Approaches based on proportional flexibility matrices [34]
- Finite element solver and a local optimisation algorithm [107]
- Finite element solver and a global optimisation algorithm [45, 104]
- Finite element-based multi-objective approaches [91]
- Finite element solver combined with neural networks [69]

A major problem that has to be dealt with in model updating using real-world measurement data arises due to the systematic error between the model and the actual structure. While analytical methods deliver correct results in verification settings where the damaged state is simulated in a mechanical model, they tend to break down when applied to measurement data. Hence, a trend towards more resilient and error-tolerant methods, which suppress systematic errors using advanced error metrics, can be observed in the literature [1]. The inevitable mismatch between the model and the actual structure as well as measurement uncertainty lead to an undetermined optimisation problem, which often has several local minima. In this regard, the methods can be ranked by their resilience, starting with the lowest: analytical, local optimisation, global optimisation, multi-objective optimisation. Neural networks and meta-model approaches impose additional limitations when compared to numerical optimisation, as they need to be trained before they can be used, which tends to increase the numerical cost [14].

The mechanical model can be parameterised using different strategies. A popular method is to assign parameters to the stiffness of individual elements or to the stiffness of regions of the model which are comprised of multiple elements [106]. A more rarely used method is the use of an analytical function with few parameters which prescribe the distribution of stiffness to the whole model [114]. In addition to the various methods of updating and parameterisation, there are also several ways to process the vibration measurement data prior to updating. For example, eigenmodes [113], eigenfrequencies [106], frequency response functions and transmissibility functions [75] can be identified. As a result, there are numerous ways in which updating algorithms, parameterisation and data processing can be combined to yield a model updating procedure.

The choice of the parameterisation has to be considered together with the updating method, since the resulting computational effort is dependent on both. For example, to solve a model updating problem using global optimisation, it has to have only few unknown variables, hence it should not be parameterised using individual stiffness values of several hundred elements. Conversely, a parameterisation with few unknown variables tends to generate a non-convex optimisation problem, which can not be efficiently solved using a local optimiser [113].

This thesis pursues a global optimisation approach using modal parameters and especially focuses on a parameterisation which yields unique identification results. Ambiguities often arise as a result of too many parameters which lead to an under-determined mathematical problem with multiple non-physical candidate solutions. Therefore, in this work, damage distribution functions with only three design variables are used to increase the chance of obtaining unique identification results [45].

Furthermore, efforts towards dealing with uncertainty in the model updating process were made by various researchers [61, 65, 106]. This thesis, however, only deals with deterministic methods and intentionally omits statistics in favour of a more in-depth discussion concentrating on the core concepts of damage localisation.

1.5 Localisation using data-driven methods

Data-driven output-only monitoring methods allow for damage detection and localisation without a model of the physical structure. Since only the measurement data and no further knowledge of the structure is required, data-driven methods can be fully automated. Hence, this approach is particularly attractive when accurate modelling is too time-consuming to be practical or no documentation of the structure exists. Over the last decades, many different methods were proposed, most of them founded on statistical and signal processing the-

ory [36]. Due to the higher mathematical complexity compared to model-based monitoring methods, data-driven localisation methods are often more abstract and the results harder to interpret. Some prominent examples for such methods are

- Damage localisation using mode shape curvature [86]
- Power of difference processes obtained using state estimators [68, 125]
- Damage-sensitive features based on transmissibility functions [25, 75]

Modal parameters are widely used for structural dynamic identification and are therefore an obvious basis for monitoring methods. Dedicated modal analysis algorithms focus on this task only, while some parametric and non-parametric methods provide modal parameters as a by-product. Prominent examples of dedicated operational modal analysis methods are the frequency domain decomposition [18] and an improved Bayesian approach known as BayOMA [4]. The mode shape curvature method [86] directly employs identified mode shapes to identify kinks that relate to damage positions. However, this method is only applicable to beam-like structures and depends on a high sensitivity of the underlying modal analysis technique, which may not be the case under real-world conditions.

Over the last decades, several parametric identification methods were proposed. Notably, vector auto-regressive identification [92] as well as stochastic subspace identification [88] were applied successfully for structural health monitoring. Both of these approaches are based on identifying a limited number of structural modes, hence the name ‘parametric’. Mathematically, the structural modes are treated as the damped poles of the system. Based on such system identifications, damage localisation schemes can be implemented, for example by employing residuals [68]. A typical drawback of methods based on parametric identification is associated with high model orders, which increases the risk of obtaining non-structural modes. This makes it difficult to obtain models for a wide spectral range, hence, usually the model order is determined to be as low as possible and only as high as necessary [97, 131]. Concerning the computational effort associated with parametric methods, the resulting relatively low model orders are an advantage, since data can be processed rapidly.

Non-parametric identification methods do not consider individual modes of the underlying mechanical system, but rather the full vibration spectrum. However, the term ‘non-parametric’ is a bit misleading, as in the usual case of discrete measurement data, these techniques still employ a finite model order. Rather, the model order of non-parametric approaches is typically orders of magnitude higher than in a parametric approach. Since non-parametric methods are concerned with identifying the full spectral behaviour of structures,

many methods employ the Fourier transformation to obtain spectral responses. A popular approach considers transmissibility functions, which are obtained by the division of the spectra of two different sensor signals [75]. Since the vibration behaviour at a given point in the structure is dependent on the vibration of all other points in the structure, multiple input signals should be considered to obtain more accurate identification results. Such multiple-input transfer functions are enabled by a H_1 or H_2 estimation technique [3] and were employed for structural health monitoring [121]. However, the frequency-domain multiple-input identification is only possible if several uncorrelated excitation sources are present [74], which is not generally the case. Albeit, a major advantage of non-parametric approaches is the numerical stability of the identified systems. Further, post-processing is usually not needed after the identification to obtain a functioning model.

In this work, a variant of the transmissibility function method is pursued. An advantage of this approach is that it uses a non-parametric identification which is relatively insensitive towards its tuning parameters, making it well-suited for dealing with real-world measurement data. The proposed method uses an identification of the transfer characteristics in terms of finite impulse response filters. Using the identified characteristics, a vibration estimation is carried out and the residuals between the estimation and the measured signal are calculated. By analysing the residuals, structural changes linked to damage can be uncovered.

1.6 Research gap

The overview given in the preceding sections shows that there are some aspects of damage localisation methods which can still be improved. In case of model-based approaches, not many efforts have been made towards rugged multi-objective formulations. While objective functions based on modal parameters are well established, research by Perera and Ruiz [90] and Jin et al. [56] shows, that there is potential for improvement in the parameterisation of the finite element model. Additionally, the reciprocity between the model updating formulation, the optimiser and the numerical performance is often neglected. To tackle these issues, an efficient multi-objective approach using damage distribution functions and a relative error metric based on modal parameters is investigated in this thesis.

Regarding data-driven approaches, residual-based damage localisation has been a relatively recent development when compared to model updating. As such, well-established numerical methods have not yet emerged in this field. In this thesis, the aspects of rugged system identification and numerically efficient methods are addressed.

In addition to improvements of the individual methods, this thesis also provides a comprehensive comparison of model-based and data-driven approaches using verification and validation data sets. As most research articles are focused on a single method, to date there have not been many attempts at direct comparisons [7]. In this regard, this thesis especially examines issues related to the damage localisation.

1.7 Scope and structure of the work

The objective of this thesis is to develop advanced methods for damage localisation using both a data-driven approach as well as an approach based on mechanical models. The numerical solution of the model updating problem requires efficient optimisation algorithms. Hence, the development of numerical optimisation algorithms is also a topic this thesis focuses on.

The optimisation algorithm employed for the model updating is detailed in Chapter 2 while the overall updating method is described in Chapter 3. The details of the proposed data-driven damage localisation method are explained in Chapter 4. A validation and comparison of both damage localisation methods is given in Chapter 5. Finally, Chapters 6 and 7 point out the benefits and limitations of the proposed methods and summarise the findings of this thesis.

Chapter 2

Numerical optimisation algorithms

In this thesis, the model updating problem is treated as a numerical optimisation problem and a global optimisation method is applied to solve it. The latter is required, because the objective function space arising from complex engineering optimisation problems usually contains multiple local minima. In addition, the objective space also contains numerical noise which further necessitates the usage of global methods. In the following, deterministic global optimisation algorithms based on the Pattern Search approach [33] are presented and the numerical performance is evaluated. The benchmark results show that the presented algorithms are well-suited to the aforementioned requirements. The EngiO framework [10], which was co-developed by the author of this thesis, is used to implement and run the presented algorithms.

This chapter contains parts of research published by the author of this thesis as well as submitted articles [10, 44, 45, 51].

2.1 Global optimisation

Global optimisation algorithms are used when the objective function space of an optimisation problem features many local minima. In such cases, local optimisation algorithms can find the true global optimum only by chance. Further, the use of local optimisation algorithms is hindered when no analytical derivatives of the objective function are available, which is the case for the model updating formulation considered in this thesis. Derivative-free global optimisation algorithms treat the objective function as a black box without knowledge of the internal structure of the optimisation problem. These derivative-free algorithms generally have a lower numerical performance than derivative-based methods, however, they impose no restrictions to the type of objective function

and are more robust. Thus, formulating the model updating problem as a global derivative-free problem makes it adaptable and readily applicable.

Depending on the engineering application at hand, the global optimisation problem comprises one or more objectives and may involve constraint equations. In the following sections, the single- and multi-objective formulations are introduced and constraint handling strategies are discussed.

2.1.1 Optimisation problems

Single-objective optimisation is concerned with minimising or maximising a scalar function. In the following, only minimisation is discussed, however maximisation can readily be achieved using a negative sign. Numerous basic engineering optimisation problems fall in the category of scalar, bounded, unconstrained, non-linear and derivative-free optimisation problems [105]

$$\text{minimise } f(\mathbf{x}) \quad \text{for } f \in \mathbb{R}, \mathbf{x} \in \mathbb{R}^n, \quad (2.1)$$

where f is the scalar objective function and \mathbf{x} is the n -dimensional vector of design variables. In essence, the goal is to find the design variable space point with the most extreme associated objective function value. The space of the design variables is bounded to the volume of a hypercube

$$\mathbf{x}_{\text{lb}} \leq \mathbf{x} \leq \mathbf{x}_{\text{ub}}, \quad (2.2)$$

where \mathbf{x}_{lb} and \mathbf{x}_{ub} are the lower and upper bounding vectors, respectively.

Multi-objective optimisation differs from single-objective optimisation by considering a vector-valued objective function instead of a scalar one. In multi-objective optimisation, the bounded, unconstrained, non-linear and derivative-free optimisation problem is solved

$$\text{minimise } \mathbf{f}(\mathbf{x}) \quad \text{for } \mathbf{f} \in \mathbb{R}^m, \mathbf{x} \in \mathbb{R}^n, \quad (2.3)$$

where \mathbf{f} , in contrast to Equation 2.2, is an m -dimensional objective function vector.

As a consequence, usually no single most extreme point results from the optimisation. Instead, there is an infinite number of compromise solutions which favour some of the objectives while disregarding others. The compromise solutions are situated on a hyper-dimensional surface in the objective space, which is referred to as the Pareto frontier, named after Vilfredo Pareto [81]. In a given set of objective value vectors, the compromise solutions can be identified using the Pareto non-dominance criterion and a corresponding non-dominated sorting algorithm. The non-dominance criterion mandates that a point $\hat{\mathbf{x}}$ is on the Pareto frontier, if there exists no point \mathbf{x} such that $\mathbf{f}(\mathbf{x}) \preceq \mathbf{f}(\hat{\mathbf{x}})$ [76]. This means that the solutions on the frontier are better than others at least in one

objective. A sketch of the objective value space of a multi-objective problem is shown in Figure 2.1.




Figure 2.1: Sketch of objective value space in multi-objective optimisation and Pareto dominance.

Multi-objective optimisation algorithms recover an approximation of the Pareto frontier, which can be extracted using the Pareto non-dominance criterion. In Figure 2.1, the points marked by squares constitute the non-dominated set. The points marked as dominated are Pareto-dominated by several other points and are therefore not on the Pareto frontier. For the example of a two-objective problem this means that there must be no point below or to the left of a point on the frontier, as illustrated by the arrows in Figure 2.1.

2.1.2 Constrained optimisation

Some engineering problems require the use of constraints, which define a feasible region inside the bounded design variable space. The constrained multi-objective optimisation problem can be stated similar to Equations 2.2 and 2.3

$$\begin{aligned}
 &\text{minimise } \mathbf{f}(\mathbf{x}) && \text{s.t. } \mathbf{g}(\mathbf{x}) \geq \mathbf{0} \\
 & && \mathbf{h}(\mathbf{x}) = \mathbf{0} \\
 &\text{for } \mathbf{f} \in \mathbb{R}^m, \mathbf{g} \in \mathbb{R}^p, \mathbf{h} \in \mathbb{R}^q, \mathbf{x} \in \mathbb{R}^n,
 \end{aligned} \tag{2.4}$$

where $\mathbf{g}(\mathbf{x})$ consists of one-sided inequality constraints and $\mathbf{h}(\mathbf{x})$ contains equality constraints with p and q constraint equations, respectively. Since Equation 2.4 poses a fundamental maths problem, many methods have been proposed to deal with constraints in numerical optimisation tasks. A popular and simple way to handle such problems is the expansion of the objective function using static penalty techniques. Following this approach, the constrained optimisation

problem is redefined as

$$\begin{aligned} & \text{minimise } \mathbf{f}(\mathbf{x}) + \Phi(\mathbf{g}(\mathbf{x}), \mathbf{h}(\mathbf{x})) \\ & \text{for } \mathbf{f}, \Phi \in \mathbb{R}^m, \mathbf{g} \in \mathbb{R}^p, \mathbf{h} \in \mathbb{R}^q, \mathbf{x} \in \mathbb{R}^n, \end{aligned} \quad (2.5)$$

where Φ is the vector-valued penalty function.

According to Coello Coello [26], penalty techniques are divided into interior and exterior approaches. Interior methods restrict the design space to feasible solutions only. On the one hand, if there is only a very small feasible region, an application of these approaches can result in entirely missing the feasible region. On the other hand, in case of success they will definitively provide feasible solutions only. In contrast, exterior methods do not restrict the design space and thus are suited for problems with small or even non-existent feasible regions. However, the optimal solution obtained using an exterior approach will often not be inside the feasible region. The choice whether to employ an interior or exterior penalty function has to be made before actually running the optimisation. In practice, numerical experiments are usually required in order to find a penalty formulation that delivers satisfying results for a given objective function.

As an example, using the exterior linear penalty technique, the violation of a single inequality constraint $g(\mathbf{x})$ is penalised by

$$\Phi_i^g = \begin{cases} -r_i^g \cdot g(\mathbf{x}) & g(\mathbf{x}) < 0 \\ 0 & g(\mathbf{x}) \geq 0, \end{cases} \quad (2.6)$$

where r_i^g are the corresponding penalty vector entries and i denotes the objective function index. In the case of a single equality constraint $h(\mathbf{x})$, the linear penalty function

$$\Phi_i^h = r_i^h \cdot |h(\mathbf{x})| \quad (2.7)$$

evaluates the constraint violation weighted by penalties r_i^h specified for each objective i . This vector formulation of the penalty functions enables imposing constraints on multi-objective optimisation problems as well. The final penalty added to the objective function is expressed by the sum over all equality as well as inequality penalty values

$$\Phi_i = \sum_g \Phi_i^g + \sum_h \Phi_i^h. \quad (2.8)$$

2.2 Classification of numerical optimisation algorithms

Since general-purpose computers were invented, they have been used to solve numerical optimisation problems. This led to the development of an immense

number of algorithms dedicated to this topic. To contextualise the method outlined in the following, this section provides a classification of the currently predominating algorithm types. The formulation chosen for the model updating problem in this thesis requires the usage of single-objective or multi-objective global algorithms. Convex or local optimisation approaches are therefore not discussed further.

The sample generation of global optimisation algorithms can be either based on random numbers or based on a deterministic approach. Derivatives of the objective function are not employed by the currently predominating approaches. Additional features make some algorithms distinctive, such as special constraint handling strategies. However, since constraint handling can also be incorporated into the objective function via penalties, as outlined in Section 2.1.2, these intricacies are not regarded in the following.

Hence, for suitable global algorithms, the main differences lie in the employed sample generation schemes. The following sections outline the differences between these sample generation schemes as well as the state of research in this field. As the single and multi-objective approaches only differ in the handling of objective function evaluations, no further distinctions are made in the discussion below.

2.2.1 Algorithms based on random numbers

The most commonly used global derivative-free algorithms are based on random-number sample generation schemes. The random numbers are obtained numerically using a pseudo-random number generating algorithm. Thereby, a stochastic exploration of the design variable space of the optimisation problem is achieved. Many optimisation algorithms that utilise random numbers are inspired by natural processes like biological or chemical phenomena [130]. These algorithms are also referred to as metaheuristic algorithms.

Many optimisers based on random numbers belong to the class of evolutionary algorithms. Two very popular approaches are Evolution Strategies [105] and Genetic Algorithms [40]. The essence of both approaches is the abstraction of the Darwinian evolution theory as well as natural selection of biological systems. The natural process “survival of the fittest” is represented in terms of mathematical operators like crossover, recombination, mutation and fitness evaluation. Other examples of metaheuristic algorithms based on random numbers are Simulated Annealing [64], inspired by the annealing process of metals, Particle Swarm optimisation [59], inspired by swarm intelligence of fish and bird behaviour, and Harmony Search [39], inspired by the arrangement of musical harmonies.

The Non-dominated Sorting Genetic Algorithm-II (NSGA-II) [28, 29] is a multi-objective expansion of the evolution-inspired Genetic Algorithm and

known to perform well on most problems [20]. In this algorithm, solutions are ranked by the computation of non-dominated sets and a crowding distance is assigned to find members of the new population. To date, it is arguably the most popular multi-objective optimisation algorithm.

More recent global optimisation methods include the Grey Wolf optimiser [79], Spotted Hyena optimiser [31] and the Whale Optimisation Algorithm [78]. As the naming suggests, the ideas behind these algorithms are motivated by animal behaviour. More generally, these motivations and inspirations can be referred to as metaphors. The use of such metaphors, however, has been criticised recently, since it may obscure the mathematical nature of the actual algorithm. For example, Sörensen [110] showed that some algorithms published under different names and using different metaphors were actually mathematically identical. In particular, Harmony Search was shown to be a variant of the Evolution Strategy algorithm. For this reason, the author of this thesis tries to avoid metaphors and instead focuses on the mathematical concepts behind the algorithms.

A property often associated with metaphor-based algorithms is a high number of parameters. Since e.g. animal behaviour is hard to put into mathematical formulae, the numerical modelling needs to be tuned to obtain the desired effect. This mathematical tuning in turn leads to numerous free parameters which have to be set in a problem-specific way to obtain optimal performance. While very good numerical performance can be obtained when tuning the algorithm to specific benchmark problems, the performance on real-world problems using default parameter settings is often sub-optimal. This is exacerbated by complex interactions between the parameters, which may lead to parameter tuning by a ‘trial-and-error’ approach and high computational effort to obtain satisfactory solutions. To tackle this issue, metaphor-free algorithms were proposed recently by Rao [99], which have as little as a single parameter, while retaining an acceptable numerical performance. The numerical benchmark in Section 2.3.6 discusses several currently popular random number-based approaches.

2.2.2 Deterministic algorithms

Contrary to algorithms based on random numbers, deterministic methods generate samples in a predetermined and predictable way. Hence, deterministic algorithms do not rely on random permutations in order to converge to an optimal solution. The design space trajectory of a deterministic algorithm is only dependent on the objective function. Thus, the result for a given optimisation problem is always the same and does not vary for different runs as is the case for algorithms based on random numbers.

Deterministic approaches are most often applied for convex and gradient-based optimisation. Popular methods from this class are Coordinate Descent

[128], Pattern Search [48] and Sequential Quadratic Programming [16]. However, these convex optimisers can not reliably find the global optimum in a problem with several local minima. Only few deterministic global optimisation algorithms were proposed thus far. A popular one is the DIviding RECT-angles algorithm (DIRECT) algorithm [57], which exploits geometric patterns and shapes to traverse the design variable space. Evtushenko and Posypkin [35] proposed an approach which employs box-constraints and a branch-and-bound optimisation scheme to tackle multi-objective problems. However, some of these approaches lack numerical robustness and there have not been many attempts made at comparing them to more mainstream methods, like NSGA-II.

2.2.3 Research gap

Even though the currently predominating global optimisation algorithms are based on random numbers, deterministic algorithms have interesting numerical properties which can be advantageous for some applications. The numerical performance in many engineering optimisation problems is dependent on the number of samples needed to achieve results sufficiently close to the global optimum. Deterministic algorithms are interesting in this regard, since the positioning of sampling points is predetermined in such approaches, which can be exploited to speed up computations. One such application is the fuzzy α -level optimisation, where a deterministic algorithm proposed by Hübler and the author of this thesis [51] is able to provide similar or better results compared to state of the art methods by employing a sample reusing approach. A reason for the relative unpopularity of deterministic global optimisation algorithms lies in their low suitability for problems with high numbers of design variables. This, however, can be countered by employing a formulation of the objective function with a low number of variables.

In multi-objective optimisation, there aren't many successful algorithms to begin with. The current literature is mostly dominated by the algorithm NSGA-II [28] and attempts to create equally versatile and high-performance approaches have mostly failed [20]. There have been some contributions considering deterministic multi-objective approaches [35], however these have mostly not been employed to solve practical engineering problems. Hence, novel methods with a different set of strengths and weaknesses compared to NSGA-II are needed to enhance the number of options to choose from.

This thesis thus follows a deterministic approach which is outlined in the following sections.

2.3 Global Pattern Search

The Global Pattern Search (GPS) algorithm was introduced by the author et al. [45] and is a core innovation of this thesis. It is based on the local Pattern Search algorithm [33] and augmented by the simultaneous tracking of multiple local minima in order to yield a global technique. The resulting global algorithm belongs to the deterministic class of optimisation algorithms and thus does not rely on pseudo-random numbers to generate its samples. The GPS algorithm is also related to the concepts of Coordinate Descent [128] and Direct Search [43], as it employs the round-robin design variable variation applied by these methods. The GPS algorithm was published as an open-source code written in MATLAB syntax by Berger et al. [10].

2.3.1 Search pattern

As its name implies, the Global Pattern Search approach is based on a search pattern, which is used to sample the design variable space. This pattern is centred around a current best sample \mathbf{b} , as depicted in Figure 2.2. The distance between the current best sample and the generated samples \mathbf{s}_j is controlled by the step width vector \mathbf{w} , which comprises the step widths associated with the design variables. The search pattern is aligned with the design variable axes, which leads to a cruciform shape. Additionally, all points generated by the algorithm are located on a regular grid. The grid spacing is defined by the parameter N , which is discussed in detail in Section 2.3.3.

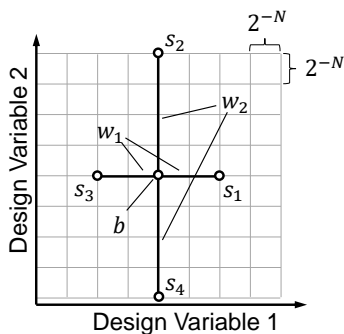


Figure 2.2: Sample generation scheme of the Pattern Search algorithm. A variation along the axes is performed around the base vector \mathbf{b} using the step width vector \mathbf{w} , resulting in the sample coordinates \mathbf{s}_j . The integer grid is indicated using grey lines.

In each iteration of the algorithm, the sampling points \mathbf{s}_j are evaluated using the objective function. The results of these evaluations are then used to determine whether the sampling points led to improved objective values compared to the currently best sample \mathbf{b} . When a sampling point with an improved objective value is found, the current best sample is replaced with the improved sampling point. This way, the search pattern is used to traverse the search domain in the direction of improved objective function values. When no more sampling points with improved objective function values can be found, the algorithm reduces the step width by half its original length. This way, the search pattern is spatially refined to allow for convergence of the optimisation scheme.

The main difference between the local Pattern Search algorithm and GPS is that the latter employs a set of currently best samples instead of a single one. The number of simultaneously tracked best samples is represented by the parameter T , which is the only parameter that influences the trajectory of the GPS algorithm during the optimisation run. When choosing $T = 1$, GPS becomes a local optimisation method which resembles the local Pattern Search approach. Higher values of T lead to increasingly more global search domain coverage. The GPS approach can be thought of as a parallelised local optimiser, where multiple local minima are tracked and refined. In contrast to a multi-start approach, however, GPS only tracks local minima until all T points are close to the global minimum, after which it only refines the global minimum solution. This behaviour is further discussed in Section 2.3.5.

While this thesis does not attempt the mathematical proof of convergence properties, a proof for the local Pattern Search algorithm, essentially the $T = 1$ case, was published by Torczon [116]. When approaching $T \rightarrow \infty$, the GPS algorithm becomes a full-factorial grid sampling approach, where convergence is also guaranteed. Therefore, a proof for the cases $1 < T < \infty$ remains a topic for future research.

2.3.2 Algorithm description

The pseudo-code of GPS is given in Algorithm 1. In this algorithm, the currently tracked global minima are stored in the matrix \mathbf{H} , which has T rows to store their sample coordinates. Since \mathbf{H} contains the best samples, it is also referred to as the ‘hall of fame’ in the following.

In the first two lines of Algorithm 1, the coordinates of the initial sample and the initial step width are calculated depending on the resolution parameter N . As the initial sample is the best sample known to the algorithm upon initialisation, it is assigned to the first row of the ‘hall of fame’ \mathbf{H} as the currently best sample. The value 2^{N-1} is chosen so that the initial sample is exactly in the middle of the design variable space and the initial step width spans exactly half the design variable space. This leads to an initial search pattern

Algorithm 1 Global Pattern Search (GPS)

```
1:  $w_i \leftarrow 2^{N-1}$  {initialise step width vector}
2:  $H_{1,i} \leftarrow 2^{N-1}$  {initialise hall of fame}
3: loop
4:    $\mathbf{b}_k \leftarrow \mathbf{H}_k$  {take  $T$  base coordinates from hall of fame}
5:   for  $k = 1$  to  $T$  do
6:     generate  $2n$  sampling coordinates  $\mathbf{s}_j$  for each  $\mathbf{b}_k$ 
7:     deduplicate  $\mathbf{s}_j$ 
8:     clamp sampling coordinates  $\mathbf{s}_j$  to domain boundaries
9:     calculate  $\mathbf{x}_j$  from  $\mathbf{s}_j$ 
10:     $y_j \leftarrow f(\mathbf{x}_j)$  {sample using cache}
11:  end for
12:  update hall of fame  $\mathbf{H}$  using  $\mathbf{y}$ 
13:  if  $\mathbf{H}$  changed in update, then
14:    continue loop
15:  end if
16:  if every  $w_i$  is 1, then
17:    break loop
18:  end if
19:   $w_{\max} \leftarrow \frac{1}{2}w_{\max}$  {reduce largest step width}
20: end loop
```

which touches the boundaries of the design variable space and thus covers the maximum possible hypervolume. The algorithm always starts its optimisation run with exactly one point situated in the middle of the design variable space, making it a deterministic approach.

Lines 3 to 20 of Algorithm 1 comprise the iteration loop of the algorithm, which can be divided into sample generation using the search pattern, sampling and updating the ‘hall of fame’ as well as step width control. These parts of the iteration loop are addressed in the following.

The pattern-based sampling of GPS is described in lines 5 to 11 of Algorithm 1. In a loop over the current globally best samples taken from the ‘hall of fame’, the search pattern is applied to generate the sampling coordinates \mathbf{s}_j . If the number of samples available in \mathbf{H} is less than T , all samples in \mathbf{H} are used instead. In other words, the ‘hall of fame’ is filled up with points until it reaches the size T , at which point it maintains its size. For each of the base points \mathbf{b}_k taken from the ‘hall of fame’, variations along all n design variables are carried out. Since there is one sample in the positive as well as in the negative direction of each of the n axes of the search domain, $2n$ samples are generated per base coordinate \mathbf{b}_k . A graphical representation of the sample generation scheme

is shown in Figure 2.2. Since some of the samples may have been generated multiple times due to overlapping search patterns, these are deduplicated to remove redundant samples. In the next step, each sample is clamped to the search domain boundaries. If a sample is outside of the boundaries, it is moved onto the boundary. Since the sample coordinates are generated on an integer grid, the real-valued sample coordinates \mathbf{x}_j are calculated before the objective function can be evaluated. The integer grid is explained in more detail in Section 2.3.3.

The real-valued sample coordinates \mathbf{x}_j are used to evaluate the objective function using a caching mechanism. This caching mechanism is further specified in Section 2.3.4. The result of the objective function evaluation is then stored to the vector \mathbf{y} , which contains the objective values calculated during each iteration of the algorithm. The outlined sample generation scheme is only dependent on the samples in the ‘hall of fame’. Hence, all samples of an iteration can be evaluated independently, which means that the sample evaluation can be parallelised to increase numerical efficiency on many-core computers. After the evaluation of the samples, the ‘hall of fame’ matrix \mathbf{H} is updated using all samples generated in the iteration. This is facilitated by sorting the objective function values \mathbf{y} and assigning the best T samples to \mathbf{H} .

The step width control is described in lines 13 to 19 of Algorithm 1. If the updated ‘hall of fame’ changed during the update, the iteration loop is restarted using the same step width as before. This enables an exploratory search of the design space. Lines 16 to 18 of the pseudo-code contain a stopping criterion for the corner case where the step width is 1. When no change in the ‘hall of fame’ is achieved in an iteration, the step width is reduced. This is performed one axis at a time by determining the axis with the maximum step size and subsequently reducing it by half its length. The stopping criterion is encountered, when the step widths of all design variables reach 1. At this point, the algorithm can not produce any new candidate points on the integer grid. A further reduction of the step width on the integer grid would lead to a meaningless step width of zero and thus the algorithm is terminated. The algorithm is also aborted once a specified number of objective function evaluations is reached.

A major advantage of the GPS approach is that only one parameter controls the convergence behaviour. This is contrasted by the parameter count for traditional metaheuristics like Particle Swarm Optimisation or Genetic Algorithms, which employ four or more parameters. The parameter T defines the number of tracked points and thereby controls the balance between convergence rate and design variable space exploration. If T is chosen too low for a given problem, the search pattern fails to place samples close to the position of the global optimum, and convergence occurs at a local optimum instead. By choosing suf-

ficiently high values for T , the global optimum is discovered and the algorithm constricts its search pattern at that point.

Having only a single parameter means that GPS can be tuned to new optimisation tasks without much numerical experimentation. Furthermore, it prevents the emergence of complicated interactions between parameters. Another benefit stems from the deterministic nature of the algorithm: Since the grid is independent of the parameter choice, samples of previous runs can be reused to speed up the optimisation process of later runs. This strongly reduces the numerical effort when an increase of the parameter T is deemed necessary in a rerun.

2.3.3 Discretisation of the design variable space

The design variable space is discretised by introducing integer coordinates which form a Cartesian grid. This grid is defined using a resolution of 2^N , where N is only limited by the numerical capabilities of the programming language and computer architecture the algorithm is implemented in. This grid is depicted in Figure 2.2. New sampling points are generated in every iteration of the algorithm based on the coordinates of the T globally best points, as described in Section 2.3.1. The sample coordinates \mathbf{s} are located on the discrete integer grid and transformed onto the continuous design variable space by

$$x_i = x_{\text{lb},i} + s_i(x_{\text{ub},i} - x_{\text{lb},i}) \cdot 2^{-N}. \quad (2.9)$$

The definition of the resolution using powers of two stems from the way that the search pattern is modified during the flow of Algorithm 1: When the algorithm can not find better samples using the current search pattern, the pattern is reduced in size by exactly half, which leads to a grid coincident with 2^N resolution.

If not stated otherwise, the resolution parameter is set to $N = 24$ in this thesis. This leads to a precision of $2^{-24} = 6 \times 10^{-8}$ relative to the size of the design space, which is generally sufficient to meet the numerical accuracy demands of practical engineering problems. As an example, when GPS is used to identify the location of a defect on a structure with a length of 100 m using $N = 24$, an accuracy of $6 \mu\text{m}$ is achieved. This by far exceeds the practical requirements to such an identification, hence, the parameter N is not further considered in the following chapters.

2.3.4 Reuse of cached samples

As the design space is discretised using a grid and the search pattern is aligned on this grid, it is inevitable that points of the design space are sampled multiple times. To avoid the associated computational overhead, the result of every objective function evaluation $\mathbf{y} = f(\mathbf{x})$ is stored to a cache.

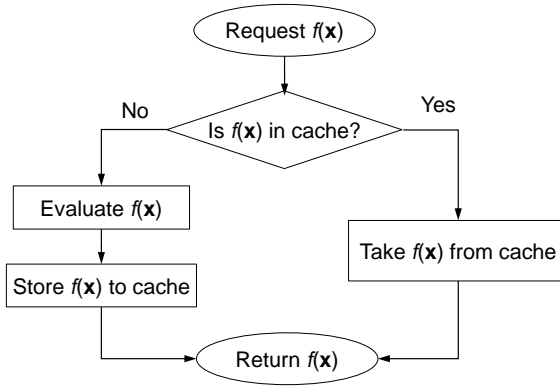


Figure 2.3: Flowchart of the Global Pattern Search caching scheme for objective function evaluations.

The logic of the cache is captured in Figure 2.3. This caching mechanism is invoked whenever the objective function $f(\mathbf{x})$ is requested. In case of a cache miss, the objective function is evaluated and the result is stored to the cache. In case of a cache hit, the result is taken from a matrix which stores the coordinates and the objective value of each sample in a separate row.

For the reuse of samples, it is very beneficial to employ a grid-based approach. This leads to exact matches when querying the sample cache. In contrast, when using an optimisation approach which generates its samples using pseudo-random numbers as the basis, the variability of the coordinates of cached samples has to be taken into account. Thus, random-based methods need to be designed, so that the statistical bias introduced by sample recycling does not negatively impact the convergence properties. As a result, sample reuse is rarely used in metaheuristic optimisers based on random numbers.

2.3.5 Verification using test functions

The sampling pattern of GPS is demonstrated in Figures 2.4 and 2.5 using the well-established Himmelblau test function [43]

$$f(\mathbf{x}) = (x_1^2 + x_2 - 11)^2 + (x_1 + x_2^2 - 7)^2. \quad (2.10)$$

In these plots, the topology of the test function is visualised using contour lines. The samples of the algorithm are colour-coded according to the iteration in

which they were generated. The Himmelblau function is a multi-modal function, i.e. it has multiple equally valued global minima.

The initial sample is situated in the centre of the design variable space and is therefore indicated using a dark blue colour. The samples from the last iterations of the optimisation runs are indicated using red colour and are located near a single global optimum or all of the the global optima of the Himmelblau test function, respectively. These examples also showcases the regular sampling grid and its refinement close to the optima of the test function.




Figure 2.4: Sampling pattern of GPS on the Himmelblau test function using $T=1$. Samples are colour-coded based on the iteration they were generated in.

For Figure 2.4, GPS is run with its parameter set to $T = 1$. This setting resembles the local Pattern Search algorithm. Consequently, the algorithm converges locally on one of the minima. Setting the parameter to $T = 20$ results in the algorithm converging towards all global minima simultaneously, as illustrated in Figure 2.5. This showcases that the GPS algorithm can be used to reliably solve multi-modal optimisation problems. This capability originates from a combination of the step width control and the ‘hall of fame’. Since the algorithm first explores the design variable space using a coarse grid, points close the global optima with large distances between them can be identified. The step width control then leads to local refinement and more and more points close to the global optima are evaluated. If one of the minima is actually lower than the others, eventually all T samples will end up in the vicinity of that global optimum. This occurs because the ‘hall of fame’ operates on all previously evaluated sampling points, which globally steers the sample generation towards areas with the lowest objective function values. However, in the Himmelblau function all minima have the same value, so that the best samples are located in all four minima and convergence occurs simultaneously at all locations. Of course, for this to happen, T needs to be high enough that at least one of the

‘hall of fame’ points stays close to each global optima at all times. Additionally, higher values of T lead to significantly increased computational cost.




Figure 2.5: Sampling pattern of GPS on the Himmelblau test function using $T=20$.

Figures 2.6 to 2.8 show the design variable history as well as the sampling patterns obtained by applying Algorithm 1 to various two-dimensional optimisation test functions. The corresponding objective value space is visualised using contour lines. Generally, a self-similar search pattern appears in the design variable space near the optima as the optimisation algorithm converges. In these plots, iterations are visualised using a colour bar. The parameter T was adjusted to the respective problem for robust convergence while maintaining a low number of objective function evaluations to illustrate best case performance. Higher values of T would lead to the same results, however the numerical cost would be higher. For most practical applications, reasonable parameter choices can be found in the range $5 < T < 50$. Section 2.3.6 discusses the optimal choice of T by means of numerical benchmarks.




Figure 2.6: Convergence on Camel6 function [80], $T=5$.

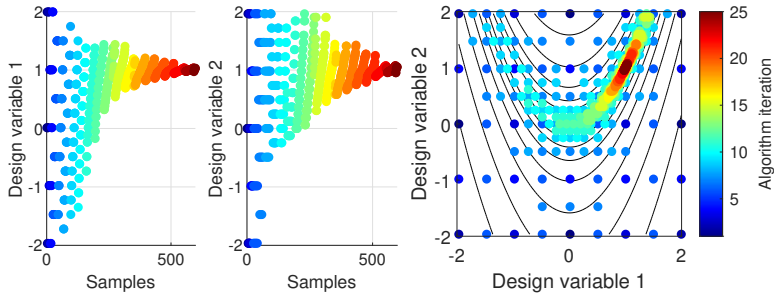


Figure 2.7: Convergence on Rosenbrock function [101], $T=20$.




Figure 2.8: Convergence on Schwefel function No. 128 [105], $T=20$.

The non-uniform scaling of the step width improves the performance of the algorithm on smooth problems. As can be seen in the sampling of the optimisation test functions, refinement on only one axis is often sufficient to improve the corresponding objective function value. The unbiased tracking of local optima leads to the discovery of all minima in the multi-modal test functions Camel6 [80] and Himmelblau [43].

The solutions to multi-modal problems are conducive to cluster analysis, e.g., by using an agglomerate hierarchical approach. This way, up to T distant solutions in the design variable space can be identified with only one optimisation run. This can also be useful in engineering optimisation tasks, feature many local minima. In that case, the near-optimal design solutions obtained using the optimisation algorithm may even contain a better solution than the actual global optimum. Such cases arise when there are other objectives or design aspects that are not considered by the objective function. Classical metaheuristic algorithms converge on only one minimum, restarts are thus required to iden-

tify all global minima. This issue is also addressed by multi-modal variants of metaheuristic algorithms [41], however, at the expense of added complexity.

The histories of the design variables show, that the algorithm performs as intended with an initial uniform exploration of the design space followed by a convergence close to local optima. Figure 2.8 shows that after 15 iterations, corresponding to about 250 samples, the algorithm constricts the sampling towards three local minima on the Schwefel test function. After several more iterations, however, the 20 ‘hall of fame’ samples all lie in the vicinity of the global minimum and the final convergence occurs at this point. Figures 2.6 and 2.7 show optimisation runs using the Camel6 and Rosenbrock test functions, which do not feature pronounced local minima. Here, the convergence towards the global minima follows immediately after the initial global exploration.

2.3.6 Benchmark using test functions

After making sure, that the proposed algorithm is able to converge to the desired global optima, it is compared to other algorithms from literature. The goal of this benchmark is to check whether the numerical efficiency is satisfactory. Further, the benchmark also yields data to judge upon the benefits and drawbacks of the proposed algorithm when compared to established global optimisation methods.

The benchmark is conducted using the engineering optimisation framework EngiO [10]. This framework is designed to provide an object-oriented interface between optimisation problems and algorithms, which makes it well-suited for benchmark tasks. The single-objective global optimisation algorithms Simulated Annealing [64], Particle Swarm Optimisation [59] and Genetic Algorithm [40] are used as the reference in this benchmark. Additionally, a pseudo-random sampling strategy is used as a worst-case baseline for the comparison. A benchmark regarding the multi-objective variant of GPS can be found in Section 2.4.4.

A parameter study is carried out to consider the influence of the hyper-parameters related to the number of function evaluations per iteration for the four considered global optimisation algorithms. The results of the parameter study can be used to determine the best-case performance and also to judge upon the sensitivity towards the numerical values of these hyper-parameters. The hyper-parameters considered for the algorithms are chosen according to Table 2.1. The hyper-parameters varied in the parameter study are denoted using intervals, while the fixed hyper-parameters are given as constant values. The fixed hyper-parameters in this benchmark are set to default values taken from literature, as they are mainly affect the fine-tuning of the algorithms and do not have a drastic impact on their convergence. In contrast, the hyper-parameters chosen for variation are the ones responsible for setting the balance

between global exploration and local convergence. Therefore, these parameters have to be set up such that the algorithm's optimisation strategy matches the complexity of the objective function. That way, the comparison to the proposed GPS algorithm is carried out on a similar basis.


Table 2.1: Optimisation hyper-parameter settings for benchmarks.

Algorithm	Hyper-parameter	Interval / Value
Random	-	-
Particle Swarm Optimisation	Number of particles	[2, 120]
	c_1	2.0
	c_2	2.0
	ω_1	0.9
	ω_2	0.4
Genetic Algorithm	Population size	[4, 240]
	Crossover percentage	0.8
	Crossover range	0.4
	Mutation range	0.2
	Selection type	Roulette wheel
Simulated Annealing	Number of particles	[4, 60]
	T_0	1
	α	0.95
Global Pattern Search	T	[1, 100]

The optimisation test functions and design variable intervals are summarised in Table 2.2. While many test functions operate on a fixed dimension, the functions chosen here can be used in N -dimensional settings. In the following, the cases $N = 2$ and $N = 10$ are investigated, i.e. the design variable spaces are comprised of two and ten design variables, respectively. To provide sufficient statistical significance, each test problem is run 100 times. For each of the algorithms, the maximum number of objective function evaluations is set to 2000, which is a number within a reasonable realm for many engineering optimisations and sufficient to allow for convergence on the considered test problems. Additionally, the boundary of the design variable space of individual optimisation runs for each algorithm is randomly shifted by 1% of its size for each run to obtain bias-free data. The median and 90% percentile values of the best objective function value achieved within each of these runs is then used for the evaluation. The six resulting benchmark cases are shown in Figures 2.9 to 2.14. In these figures, the best choice for the respective varied hyper-parameters coincides with the location of the minimum error value.

Table 2.2: Optimisation test function for benchmarks.

Test function	Interval	Formula
Sphere [66]	$[-10, 10]$	$f(\mathbf{x}) = \sum_{i=1}^N x_i^2$
Rosenbrock [101]	$[-2, 2]$	$f(\mathbf{x}) = \sum_{i=1}^{N-1} 100(x_{i+1} - x_i^2)^2 + (1 - x_i)^2$
Schwefel [105]	$[-500, 500]$	$f(\mathbf{x}) = \sum_{i=1}^N \sin \sqrt{ x_i }$

**Figure 2.9:** Benchmark results for Sphere function in 2 dimensions.

The optimisation test problems are set up so that the global optimum is shifted to the value zero, which facilitates a logarithmic axis scaling. Thereby, minute differences in the algorithm performance close to the minimum can be clearly distinguished. In all of the benchmark cases, the random search results are shown as constant over the parameter range, as it has no parameter. Algorithms with median and 90% percentile values higher than those provided by random search therefore point to a poor performance of the optimisation algorithm.

Figures 2.9 and 2.10 show the results for the most straightforward test function investigated in the benchmark. The Sphere test function is symmetrical, convex and has only one minimum, therefore the conditions are ideal for quick convergence. The two-dimensional case shows that Particle Swarm Optimisation and the Genetic Algorithm behave similarly with respect to the varying hyper-parameter. The achieved error first sinks to a minimum and rises again




Figure 2.10: Benchmark results for Sphere function in 10 dimensions.




Figure 2.11: Benchmark results for Rosenbrock function in 2 dimensions.

as the parameter is increased. In contrast, Simulated Annealing and Global Pattern Search start at the minimum and the error only increases from there. This behaviour arises because both of these algorithms essentially become local optimisation algorithms with their hyper-parameters set to the minimum value. This is amplified in the 10-dimensional case, in which these two algorithms excel, while Particle Swarm Optimisation and the Genetic Algorithm exhibit minima in the error function similar to the two-dimensional case. For all algorithms, the 90% percentiles closely follow the median values, which indicates that the algorithms perform robustly.

Figures 2.11 and 2.12 display the results obtained for the Rosenbrock test function previously depicted in Figure 2.7. For this function, all algorithms exhibit a pronounced minimal error value with respect to the hyper-parameter values. This may be attributed to the design of the Rosenbrock function, which makes it very hard to solve using local optimisation strategies. In the two-




Figure 2.12: Benchmark results for Rosenbrock function in 10 dimensions.




Figure 2.13: Benchmark results for Schwefel function in 2 dimensions.

dimensional case, the Genetic Algorithm shows relatively poor performance which is rarely better than random sampling. In the 10-dimensional case, the algorithms reach similar error values. Notably, Global Pattern Search exhibits a significant dip around $T = 20$.

The last benchmark case considers the Schwefel test function, which features multiple local minima as depicted in Figure 2.8. The results shown in Figure 2.13 indicate that all algorithms except Simulated Annealing show pronounced minima in the errors in the two-dimensional case. In fact, Simulated Annealing did not produce satisfying results at all for this test problem. Figure 2.13 also exemplifies the robustness of the Global Pattern Search algorithm, as the 90% percentiles stay close to the medians, while the 90% percentiles of all other algorithms are strongly elevated. This indicates that Global Pattern Search was able to converge to solutions close to the optimum in all of the 100 runs,




Figure 2.14: Benchmark results for Schwefel function No. 128 in 10 dimensions.

while the other algorithms struggled to discern between the global and local optima in many cases.

The 10-dimensional case depicted in Figure 2.14 shows that complex high-dimensional problems are problematic for the Global Pattern Search approach. Both the medians and the 90% percentiles are above the levels of Random Search and therefore the algorithm performed poorly. In contrast, the other algorithms exhibit a close grouping of the respective graphs. The relatively poor performance of Global Pattern Search can be attributed to a slow convergence speed under these circumstances. In case of the Schwefel function, the algorithm has to track multiple local minima, which are far apart in the design variable space. For each of the tracked minima, it then has to evaluate a high number of axis-aligned variations, considering the high number of design variables. The sample caching scheme, which reduces the number of samples significantly when the tracked global best samples are clumped together, fails to reduce the number of samples in this case. That way, convergence is reached only after a high number of objective function evaluations, and objective function values have not declined sufficiently when reaching the 2000 evaluations set as the limit in this benchmark. Chapter 7.2 contains a brief discussion on possible ways to improve the GPS algorithm in this regard.

2.4 Multi-Objective Global Pattern Search

In this section, the design principle of the GPS algorithm is extended to account for multi-objective problems. To distinguish it from the single-objective variant, the multi-objective algorithm is referred to as Multi-Objective Global Pattern Search (MOGPS). This extension is designed to retain the benefits of the grid-based deterministic approach as well as having only a single parameter. The

following subsections introduce the algorithm in more detail and point out the extensions with respect to the single-objective algorithm. Finally, the complete pseudo-code is given.

2.4.1 Non-dominated sorting and Pareto frontiers

The process of identifying the non-dominated set in a set of points is referred to as non-dominated sorting [29]. This involves the numerical sorting of the objective function values and the application of the Pareto dominance operator [67]. Due to the high computational cost involved, the optimal algorithmic implementation is still an active topic of research [102]. Non-dominated sorting is not only used to identify the final non-dominated set resulting from an optimisation run, but is also applied as a part of optimisation algorithms themselves. For instance, the Non-dominated Sorting Genetic Algorithm-II (NSGA-II) [28] derives its name from this feature. NSGA-II also includes the notion of secondary frontiers, which are obtained by repetitive application of non-dominated sorting. Figure 2.15 illustrates this concept and shows how a second and third frontier emerge in parallel to the primary frontier. Numerically, the second frontier is obtained by subtracting the first frontier from the whole set of points and applying non-dominated sorting to the remaining subset. This process can be repeated, providing a set of frontier points for each application of the sorting algorithm. Hence, after carrying out i sorting operations, the i -th frontier emerges. Each frontier can contain an arbitrary number of samples, depending on their coordinates in the objective value space.




Figure 2.15: Illustration of secondary frontiers resulting from repeated non-dominated sorting.

The i -th frontiers are usually employed in the sample generation scheme, whenever the first frontier does not contain enough samples to account for the minimum number required by the algorithm. For instance, evolutionary algorithms like NSGA-II [30] are designed to maintain a constant population size.

If the first frontier contains too few points to fill up the chosen population size, points from the second frontier are utilised as well. In approaches with a constant population size, the case may also arise, in which not all points on an i -th frontier can be included in the determination of the next sample generation. Thus, a mechanism of reducing the number of points on a frontier has to be employed. In case of NSGA-II, this is achieved by a crowding distance metric.

2.4.2 Multi-objective extension for Pattern Search

Taking only a subset of non-dominated samples into account may lead to a statistical bias and thus to a non-uniform coverage of the design variable space. To circumvent this issue, the Multi-Objective Global Pattern Search algorithm takes the frontiers into account as a whole. Due to the varying number of samples on the i -th frontiers, this leads to a varying number of points being tracked in each iteration of the algorithm. Hence, MOGPS does not simultaneously track a constant number of globally best solutions. Instead, it tracks at least T globally best solutions taken from i -th frontiers in the objective function space. This set of tracked samples from the objective space is referred to as the ‘hall of fame’, as previously introduced.

The updating scheme for the ‘hall of fame’ of MOGPS is shown schematically in Figure 2.16. In this plot, the set of points sampled in all previous iterations is referred to as \mathbf{Y} , with $|\mathbf{Y}|$ indicating the number of samples contained in it. This set is split using repetitive non-dominated sorting to yield the primary and several secondary frontiers, as indicated by division lines in Figure 2.16. Hence, the cumulative number of samples contained in all i -th frontiers is $|\mathbf{Y}|$. The parameter T is then used to decide how many frontiers are used to populate the ‘hall of fame’ vector \mathbf{h} . This is done by summing up the number of points contained in each of the frontiers and once exceeding T , the result is rounded up by a full frontier. In the example shown in Figure 2.16, the first, second and third frontiers are included in the ‘hall of fame’, since the first and second frontiers alone would include less than T samples. As this approach rounds up to a full frontier, the first frontier is always included in its entirety.

The design of the ‘hall of fame’ updating scheme comes with the benefit of not needing an algorithm for the reduction of points on a frontier, as is the case for evolutionary algorithms like NSGA-II. Further, the variable number of tracked globally best solutions represents a contrast to the single-objective GPS algorithm. Thereby, the most important consequence of the proposed scheme for tracking non-dominated solutions is that the whole primary frontier is always tracked. In late iterations of an optimisation run, this can lead to several thousand points being tracked simultaneously. When compared to tracking only a constant number T of globally best solutions, as is the case for evolutionary algorithms including NSGA-II, this has a positive effect: Since especially sam-




Figure 2.16: Illustration of the ‘hall of fame’ updating scheme based on non-dominated sorting of \mathbf{Y} and the parameter T . The resulting ‘hall of fame’ is comprised of at least T points, rounded up to full i -th frontiers.

ples near the primary frontier are used as the base vectors in every iteration, virtually all new samples will eventually be situated on the frontier. This often results in a relatively uniform and well-resolved sampling of the Pareto frontier. However, if T is chosen too small, the algorithm may entirely miss parts of the frontier and converge locally.

2.4.3 Details of the algorithm

The procedure for the ‘hall of fame’ update is shown as a pseudo-code in Algorithm 2. In the MOGPS algorithm, the updated ‘hall of fame’ indices $\hat{\mathbf{h}}$ are initialised with an empty set and then iteratively filled using the matrix of previously sampled points \mathbf{Y} . In this iterative update, a set \mathbf{e} is computed which excludes the indices to all points already belonging to $\hat{\mathbf{h}}$. Using the set \mathbf{e} , non-dominated sorting is applied and the resulting frontier is added to $\hat{\mathbf{h}}$. This process is repeated until at least T indices are stored in $\hat{\mathbf{h}}$ or all previously sampled points contained in \mathbf{Y} are exhausted. The latter case occurs in the first few iterations of the algorithm, while less than T points have been sampled, and $|\mathbf{Y}| \leq T$.

The complete MOGPS scheme is shown in Algorithm 3. As introduced for the single-objective GPS, the caching works with integer coordinates while the evaluation of the objective function takes place in \mathbb{R}^n . The ‘hall of fame’ \mathbf{h} comprises indices of the sample cache instead of directly containing the coordinates to facilitate the set operations needed for the ‘hall of fame’ update.

During run-time, the algorithm needs to store the coordinates of every sampled point as well as the corresponding objective function values. This is facili-

Algorithm 2 Update the ‘hall of fame’

```

 $\hat{\mathbf{h}} \leftarrow \emptyset$ 
while  $(|\hat{\mathbf{h}}| < T) \wedge (|\hat{\mathbf{h}}| < |\mathbf{Y}|)$  do
     $e \leftarrow \{j : j \in [1..|\mathbf{Y}|], j \notin \hat{\mathbf{h}}\}$  {Exclude hall of fame}
     $\hat{\mathbf{h}} \leftarrow \hat{\mathbf{h}} \cup \text{nonDominatedSort}(\mathbf{Y}_e)$  {Add non-dominated set}
end while

```

Algorithm 3 Multi-Objective Global Pattern Search (MOGPS)

```

 $w_i \leftarrow 2^{N-1}$  {Initialise step width vector}
 $\mathbf{S} \leftarrow (\mathbf{w}), \mathbf{h} \leftarrow (1)$  {Initialise sample cache and ‘hall of fame’}
loop
     $\mathbf{b}_k \leftarrow \mathbf{S}(h_k)$  {Get base coordinates using ‘hall of fame’}
    for  $k = 1$  to  $|\mathbf{h}|$  do
        Generate  $2n$  sampling coordinates  $\mathbf{s}_j$  for  $\mathbf{b}_k$ 
        Clamp sampling coordinates  $\mathbf{s}_j$  to  $[0, 2^N]$ 
        Deduplicate  $\mathbf{s}_j$  using cache, calculate  $\mathbf{x}_j$  from  $\mathbf{s}_j$ 
         $\mathbf{y}_j \leftarrow f(\mathbf{x}_j)$  {Sampling}
        Update cache  $\mathbf{S}$  and objective value matrix  $\mathbf{Y}$ 
    end for
     $\hat{\mathbf{h}} \leftarrow \text{updateHallOfFame}(\mathbf{Y})$ 
    if  $|\hat{\mathbf{h}}| \neq |\mathbf{h}|$  then
         $\mathbf{h} \leftarrow \hat{\mathbf{h}}$ 
        continue loop
    end if
    if every  $w_i$  is 1 then
        break loop
    end if
     $w_{\max} \leftarrow \frac{1}{2} w_{\max}$  {reduce largest step width}
end loop

```

tated by the matrices \mathbf{S} , containing the integer coordinates \mathbf{s}_j , and \mathbf{Y} , containing the objective function vectors \mathbf{y}_j , respectively. The sample coordinates are needed to deduplicate samples, while the objective function values are needed for non-dominated sorting.

When using the MOGPS algorithm on a single-objective problem, non-dominated sorting becomes equivalent to linear sorting and the ‘hall of fame’ update thus becomes identical to that of the single-objective GPS algorithm. This means that the single-objective algorithm is a special case of the multi-objective method. Therefore, the MOGPS can also be regarded as a multi-objective generalisation of GPS. Further, the number of samples generated in each iteration is even larger than in the single-objective case. Since the evaluations of samples in a given iteration of the algorithm are independent of each other, it is ideally suited for parallel computing.

2.4.4 Verification using test functions

In this section, the MOGPS algorithm is numerically compared to the benchmark algorithm NSGA-II [28] by means of analytical optimisation test functions. The objective of the comparison is to verify that the numerical performance of the MOGPS method is appropriate for practical engineering optimisation tasks. The intention is to discuss characteristics of the considered algorithms on the basis of different kinds of optimisation problems rather than proving superior numerical properties. NSGA-II is chosen as the benchmark algorithm because it is arguably the most successful derivative-free multi-objective optimiser to date. Despite its age of 20 years, NSGA-II is still among the most efficient methods [20].

In the following, the performance and convergence properties of the MOGPS algorithm are discussed using the Poloni test function [93]. Due to the prevalence of metaheuristic approaches, most test problems are not designed with grid-based approaches in mind. Since grid-based approaches generate sample points on the boundaries and along the axes of the design variable space in early stages of the optimisation run, they are very likely to converge towards optima in these regions. In essence, this means that grid-based approaches have an unfair advantage when applied to optimisation problems with optimal regions on the boundaries or axes. The Poloni test function featured in this section is chosen for detailed analysis, because its optimal solutions are situated in part directly on boundaries and in part arbitrarily in the design variable space. This helps to reduce possible bias effects introduced by the test function itself when considering the grid-based MOGPS algorithm. The definition of the Poloni test

function is given by Poloni et al. [93]

$$\mathbf{f}(\mathbf{x}) = \left(\begin{array}{c} [1 + (A_1 - B_1(x_1, x_2))^2 + (A_2 - B_2(x_1, x_2))^2] \\ (x_1 + 3)^2 + (x_2 + 1)^2 \end{array} \right) \quad (2.11)$$

where

$$\begin{aligned} A_1 &= 0.5 \sin(1) - 2 \cos(1) + \sin(2) - 1.5 \cos(2) \\ A_2 &= 1.5 \sin(1) - \cos(1) + 2 \sin(2) - 0.5 \cos(2) \\ B_1(x_1, x_2) &= 0.5 \sin(x_1) - 2 \cos(x_1) + \sin(x_2) - 1.5 \cos(x_2) \\ B_2(x_1, x_2) &= 1.5 \sin(x_1) - \cos(x_1) + 2 \sin(x_2) - 0.5 \cos(x_2) \\ \text{s.t. } &[-\pi \quad -\pi]^T \leq \mathbf{x} \leq [\pi \quad \pi]^T. \end{aligned}$$

The results of the benchmark runs are assessed on the basis of numerical quality metrics. These metrics are chosen for the comparison of the numerical performance with respect to the total number of samples as well as the volume of the dominated objective function space. This approach is motivated by engineering optimisations with numerically expensive objective functions like finite element calculations. In such optimisation problems, the computing time is proportional to the number of calculated samples and optimisers thus have to be assessed accordingly. To this end, two quality metrics are employed: The hypervolume metric, which measures the dominated objective function space, and the yield ratio, which measures the ratio of non-dominated samples generated by the algorithm.

The hypervolume metric is used extensively in literature [13, 100, 133] and evaluates the dominated hypervolume relative to a fixed reference point

$$\text{HV} \equiv \Lambda \left(\bigcup_{\mathbf{y} \in \mathbf{Y}} \{ \mathbf{y}' \mid \mathbf{y} \prec \mathbf{y}' \prec \mathbf{y}_{\text{ref}} \} \right), \quad (2.12)$$

where Λ denotes the Lebesgue measure, \mathbf{Y} is the set of evaluated objective function vectors and \mathbf{y}_{ref} is a reference point. The point \mathbf{y}_{ref} has to be dominated by all points on the Pareto frontier, which essentially means that its components must be larger than those of every point on the frontier. Equation 2.12 therefore describes the space enclosed between the non-dominated points generated by the algorithm and the fixed reference point \mathbf{y}_{ref} . An introduction of the Pareto dominance concept can be found in Section 2.1.1. Figure 2.17 shows an example of the hypervolume metric when applied to a set of points in two dimensions. Because the reference point has to be dominated by all points, a higher hypervolume metric indicates a better approximation of the Pareto frontier. The maximum value depends on the geometry of the Pareto frontier associated with

the optimisation problem as well as on the chosen reference point. Since the hypervolume metric is based on the dominated space, it is influenced by the Pareto-optimality of individual points as well as the resolution of the frontier.

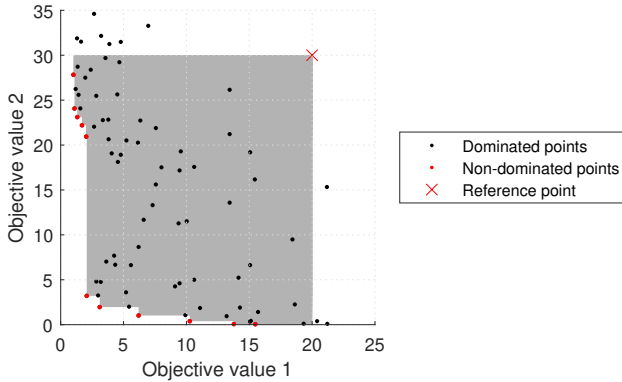


Figure 2.17: Illustrated example for the hypervolume metric. Shaded area corresponds to the hypervolume considered in Equation 2.12.

Further, the yield ratio is proposed as a measure of the efficiency in finding non-dominated points

$$\text{YR} \equiv \frac{\text{Number of non-dominated samples}}{\text{Number of total samples}} = \frac{|\{\hat{\mathbf{y}} \in \mathbf{Y} \mid \nexists \mathbf{y} \in \mathbf{Y} : \mathbf{y} \preceq \hat{\mathbf{y}}\}|}{|\mathbf{Y}|}, \quad (2.13)$$

where $|\mathbf{Y}|$ denotes the cardinality of the set \mathbf{Y} . This concept is similar to other cardinality-based metrics, including the comparative metrics Ratio of Non-dominated Individuals [63, 100, 111] as well as the Set Coverage Metric [134]. A similarly motivated intrinsic metric is the Overall Non-dominated Vector Generation Ratio [119]. Yield ratios are defined in the interval $[0, 1]$, where low values indicate poor performance in terms of finding non-dominated points. Contrarily, the value 1 resembles an ideal solution to the optimisation problem, in which every sample is also a non-dominated solution. Thus, a high yield ratio indicates that the algorithm efficiently uses samples to improve the resolution of the non-dominated set. A ‘perfect’ algorithm would therefore maximise the hypervolume metric while achieving a yield ratio of 1.

Figures 2.18 and 2.19 show results obtained with MOGPS for the Poloni test function. The algorithm is parameterised using $T = 16$ and terminated after 500 objective function evaluations. These values were chosen as they are close to the ones used in practical optimisation runs in Chapters 3 and 5. The

figures display all samples generated in the optimisation run as well as the final non-dominated set frontier. The frontier is highlighted using a colour gradient based on the first objective value to provide a mapping between the objective and the design variable space.




Figure 2.18: Objective value space obtained using MOGPS, showing all samples (black dots) and the final non-dominated set using a colour bar with respect to objective value 1.

The sampling pattern displayed in Figure 2.19 reflects the grid-based and self-similar search pattern of the MOGPS algorithm. The plot demonstrates a good balance between global sampling and a strong convergence near the Pareto-optimal areas. Regarding the objective value space shown in Figure 2.18, the concentration of sampling points near the Pareto frontier is evident.




Figure 2.19: Design variable space and sampling pattern for the Poloni test function obtained using MOGPS. Colour bar same as Figure 2.18.

In order to compare the performance and convergence properties of the MOGPS algorithm to NSGA-II, the dependencies of the yield ratio and the hypervolume metric on the number of objective function evaluations are evaluated for different parameter sets for the value T and the population size ‘pop’, respectively. The results are shown in Figure 2.20. For this comparison, the benchmark algorithm NSGA-II is run with the mutation probability $p_{\text{mut}} = n^{-1}$ and the crossover probability $p_{\text{cross}} = 0.9$, as suggested by the original authors [28]. According to the same literature, the simulated binary crossover distribution parameters are set to $\eta_c = 10$ and $\eta_m = 10$. To account for the metaheuristic nature of NSGA-II, the results are averaged over 10 optimisation runs.

Figure 2.20 a) shows the trend of the hypervolume metric over the number of objective function evaluations for both algorithms and different settings. The hypervolume metric is plotted using logarithmic scaling and is normalised using the maximum hypervolume achieved in all runs and by both algorithms. The metric was evaluated with the reference point (20|30) in the objective value space, leading to $\max(\text{HV}) \approx 536.09$ for the best run overall. Using this value, the individual runs and algorithms can be compared in terms of a common metric. Figure 2.20 depicts horizontal trends of the hypervolume metric for $T = 1$ and for $\text{pop} = 8$ indicate local convergence. This occurs because parts of the Pareto frontier are not discovered and the algorithms thus concentrate on only one part of the frontier. Higher settings for T and pop lead to a higher global coverage of the design variable space, which can result in the recovery of more separate parts of the frontier, but at the same time slows down the convergence. Except for the two aforementioned parameter settings, both

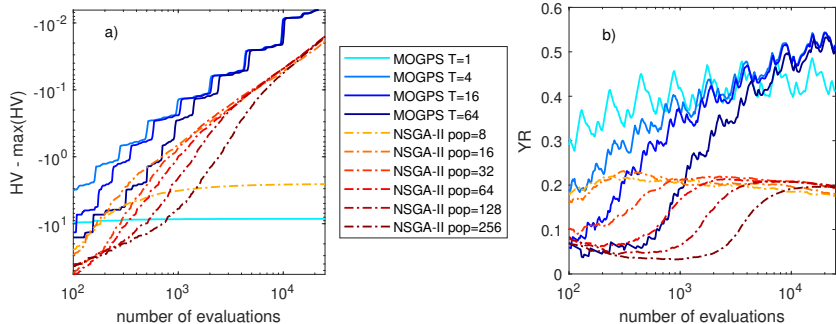


Figure 2.20: Evolution of a) the hypervolume metric and b) the yield ratio over the number of objective function evaluations on the Poloni test function for different parameter sets for MOGPS and NSGA-II.

algorithms continually converge towards the Pareto frontier, which is evident from the continually increasing hypervolume metric.

Figure 2.20 b) shows the evolution of the yield ratio metric over the course of the optimisation runs. The cases $T = 1$ and $\text{pop} = 8$ exhibit a high yield ratio metric, however they have been identified as locally convergent using Figure 2.20 a) and can therefore be disregarded. In general, the MOGPS algorithm is able to consistently improve the yield ratio over the course of an optimisation run. Low settings of T lead to high yield ratios early in the optimisation run. As the algorithm becomes increasingly more local for lower T , this behaviour is expected. The reason for the convergence of the yield ratios of the runs with $T > 1$ is that the number of non-dominated samples increases for high numbers of objective function evaluations. As the ‘hall of fame’ is always populated with the non-dominated samples, the included globally best solutions become identical for high numbers of objective function evaluations, as discussed in Section 2.4.2. In contrast, the NSGA-II runs exhibit peaks in the yield ratio, where the corresponding number of objective function evaluations depend on the chosen population sizes: Lower population sizes are suitable for optimisation runs with few objective function evaluations, whereas higher settings lead to a more global coverage of the design variable space. The yield ratios of NSGA-II eventually converge once the whole population is located in Pareto-optimal areas of the design variable space. In this example, the yield ratio achieved using MOGPS is significantly better than could be achieved using NSGA-II. This can be attributed to the deterministic nature of the algorithm, which places new samples directly adjacent to the Pareto-optimal regions instead of stochastically sampling across larger regions of the design variable space.

There is, however, a downside inherent to grid-based multi-objective approaches in general: For highly non-linear problems, the grid-based sampling in the design variable space can lead to a non-uniform resolution of the resulting non-dominated set, since the non-linearity distorts the sampling grid. An example for this behaviour is shown in Figure 2.21. This plot focuses on the lower part of the Pareto frontier shown in Figure 2.18. The Pareto frontier is shown in a cyan colour and the non-dominated points obtained with a run of the MOGPS algorithms are shown in magenta. In this plot, the non-dominated samples are positioned on a stepped pattern. This behaviour is the result of the grid-based approach, where the samples are uniformly distributed in the design variable space. Due to the non-linear nature of the Poloni test function, the grid-like sampling of the design variable space is distorted when it is transformed to the objective value space. This results in non-uniform distributions, which also vary along the Pareto frontier. While not an issue for the practical applications discussed in this thesis, there may be other research fields, where this behaviour is not desirable.

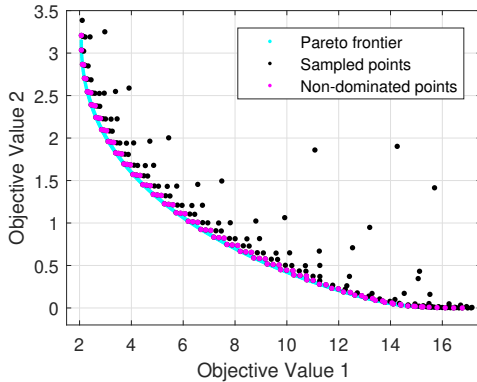


Figure 2.21: Part of the objective value space of the Poloni test function showing the distribution of samples obtained with MOGPS using $T = 16$ and 500 evaluations.

A practical mitigation of the non-uniform coverage issue is not easily achievable for MOGPS, since a regularisation of the design variable space would be necessary, which requires prior knowledge of the underlying problem. Deb [29] also discussed the biasing problems associated with non-uniform objective functions. His solution to this issue was the design of an explicit crowding distance metric for NSGA-II. This results in an unbiased and undistorted distribution of the non-dominated points, as depicted in Figure 2.22.

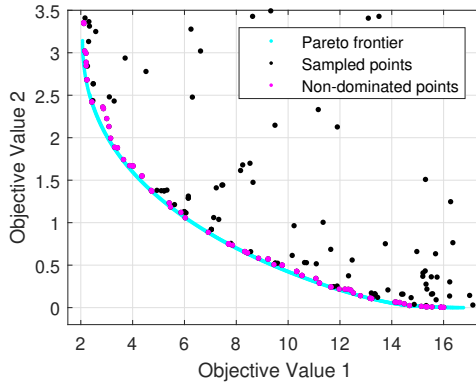


Figure 2.22: Part of the objective value space of the Poloni test function showing the distribution of samples obtained with NSGA-II using $\text{pop} = 32$ and 500 evaluations.

Due to the positioning of the Pareto optimal areas, the Poloni test function is ideally suited to additionally discuss the issue of local convergence. The trend of the hypervolume metric shown in Figure 2.20 indicates local convergence for $T = 1$. The same behaviour is observed when NSGA-II is run with the insufficiently small population size $\text{pop} = 8$. Hence, a setting of $T = 1$ is used to investigate the local convergence phenomenon for MOGSPS. Figure 2.23 displays the resulting objective value space of the Poloni function. In this case, the algorithm fails to discover the upper part of the Pareto frontier, cf. Figure 2.18.

The local convergence behaviour is caused by a failure to track sampling points close to the position of the upper section of the Pareto frontier in the first few iterations. Figure 2.24 shows the corresponding design variable space. Due to the initially small size of the ‘hall of fame’, the algorithm converges rapidly without much global exploration. This, in turn, means that the algorithm entirely misses the part of the Pareto frontier close to $\mathbf{x} = [1 \ 2]^T$ and thus converges locally. In contrast, by choosing a sufficiently high value for T , global convergence is achieved, as demonstrated in Figure 2.19. As is the case for the population size parameter in NSGA, some numerical experimentation or experience is required in order to determine the minimum setting for T , which leads to global convergence.

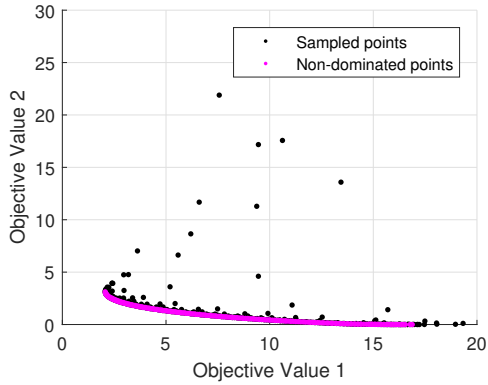


Figure 2.23: Objective value space of the Poloni test function obtained with MOGPS using $T = 1$ and 500 evaluations.

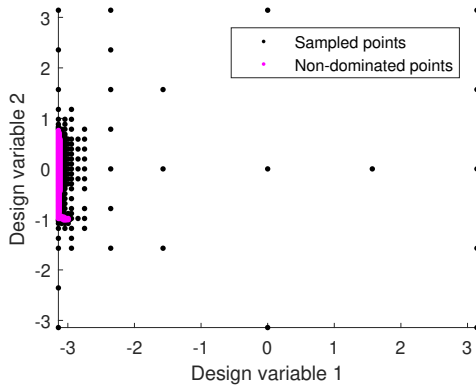


Figure 2.24: Design variable space of the Poloni test function obtained with MOGPS using $T = 1$ and 500 evaluations.

Chapter 3

Damage localisation using model updating

Using the optimisation algorithm described previously, a model-based damage localisation method is presented in this chapter. The finite element method is used for the mechanical modelling. The dynamic properties of the finite element model can thereby directly be compared to the properties of the measured structure. By exploiting this direct comparability, an optimisation problem is formulated to update the model to reflect the measured structure, therefore identifying structural damage.

This chapter contains parts of research published by the author of this thesis as well as submitted articles [44, 45, 126].

3.1 Fundamentals of model updating

3.1.1 Finite element model updating

Finite Element (FE) model updating was first proposed in the 1990s, notable early contributors were Friswell and Mottershead [37] as well as Link [71]. While the basic idea of modifying the FE model to match the observed behaviour of the structure is followed by all approaches, there is a multitude of possibilities to achieve this goal. These possibilities mainly arise because the parameterisation of the model and the objective function of the optimisation problem are independent of each other. The following sections present some of these possibilities in detail. Hence, almost any combination of parameterisation and objective function formulation yield a valid model updating strategy. Of course, only some of these combinations are advisable. The research in this field therefore mainly deals with finding robust, accurate and numerically efficient formulations.

In the model updating framework, there are at least four separate states and associated data sets of the monitored structure involved. They are generated by combinations of measured and simulated systems, as well as of reference and analysis states. The reference and analysis states are also referred to as healthy and damaged, respectively. In Table 3.1, the states are denoted using M (measured), S (simulated), 0 (reference) and 1 (analysis). For damage localisation, the states M0, M1 as well as S0 are constant and only S1 is variable and subject to model updating.

Table 3.1: States of the structure considered in model updating.

	Measurement	Simulation
Reference (healthy)	M0	S0
Analysis (damaged)	M1	S1

3.1.2 Parameterisations

The parameterisation of the FE model is an important aspect of model updating. While the goal of parameterisation is always to make the mechanical properties of the structure dependent on a number of parameters, this can be achieved in different ways. The most straightforward option is to assign one parameter for each individual element. As FE models often have hundreds or thousands of elements, this approach leads to a large amount of parameters. To reduce the number of parameters, a strategy commonly used is the assignment of one design variable to a group of elements supposedly having similar mechanical properties, which are then referred to as substructures or super-elements [62, 70]. Following this approach, the super-elements are assigned using various criteria. They can be uniformly distributed along the structure, or manually assigned to regions, where the probability of an emerging defect is known to be high. Further, super-elements can be assigned based on experience, prior knowledge or in accordance with results of preprocessing methods. The more super-elements are employed, the higher the number of parameters becomes. A different approach, which is motivated by a smooth variation of the properties along the structure, uses damage distribution functions. The distribution functions can take various forms, for example quadratic polynomials [113]. The damage distribution functions further lead to a formulation with comparatively few parameters.

In this thesis, a parameterisation based on damage distribution functions is chosen, due to the advantages of this formulation for numerical optimisation. A detailed discussion is given in Section. The reasoning and details behind this choice are given in Section 3.2.1.

3.1.3 Objective functions

Besides the parameterisation, the formulation of the objective function is another important factor in the process of model updating. The objective function numerically compares the experimental results with the results of the FE analysis. For this reason, it has to be sensitive even to small changes in the structural behaviour [53, 62].

The objective function is formulated on the basis of the available measurement data from the structure monitored. Using this data, metrics are extracted, which are used to calculate a difference between the model and the measurement. In the context of model updating, the difference between measurement and model is also referred to as the error. By minimising the error, the closest possible match is achieved, thereby fitting the model to the data. Several metrics have been proposed as the basis for the error, such as the modal assurance criterion [21], squared frequency differences [104] as well as transmissibility functions [75]. Another possibility is the formulation of compound metrics, e.g. a weighted sum of the mode shape and eigenfrequency errors [106].

The choice of metrics is dependent on the context and the goals of the monitoring task at hand. In this thesis, structures are assumed to be loaded by ambient excitation sources. Under these circumstances, approaches based on the comparison of time-domain signals obtained under transient loading are usually not applicable, since the ambient excitation forces are unknown. A natural way to extract the dynamic properties of a structure with ambient excitation is the operational modal analysis, which is introduced in more detail in Section 3.1.4. Hence, modal analysis is the method of choice to analyse the vibration data in a way that is directly comparable to finite element model calculations. Error metrics for model updating are, therefore, often based on modal parameters, i.e. eigenfrequencies, mode shapes, and possibly damping characteristics.

A popular approach for the objective function is based on mode shapes, since these dynamic features can be obtained experimentally in a relatively high quality [15]. Even though the accuracy achievable for the extraction of mode shapes is lower than for eigenfrequencies, the information content is considerable, due to their vectorial nature. To evaluate the correlation between the relevant mode shapes, the Modal Assurance Criterion (MAC) [2] is calculated. This criterion determines the degree of similarity between two mode shape vectors Φ_{M1} and Φ_{S1}

$$\text{MAC}_i(\mathbf{x}) = \frac{(\Phi_{M1,i}^T \Phi_{S1,i}(\mathbf{x}))^2}{(\Phi_{M1,i}^T \Phi_{M1,i}) (\Phi_{S1,i}^T(\mathbf{x}) \Phi_{S1,i}(\mathbf{x}))}, \quad (3.1)$$

where i denotes the mode number and the indices M and S denote measured and simulated quantities, respectively. The MAC returns a value of one, if the compared mode shapes are linearly dependent and a value of zero, if they

are linearly independent. For point-symmetrical structures, the MAC should be replaced with the S2MAC [27], which takes into account rotations of the mode shapes along the symmetry axis. A common way to formulate the error between numerical and experimental data is to sum up individual MAC values over multiple modes [71]. Considering N_{modes} eigenfrequencies, the objective function is formulated as

$$\varepsilon_{\text{MAC}}(\mathbf{x}) = \sum_{i=1}^{N_{\text{modes}}} 1 - \text{MAC}_i(\mathbf{x}), \quad (3.2)$$

where $\varepsilon_{\text{MAC}}(\mathbf{x})$ is the error metric based on the MAC. The differences in individual modes are obtained by the term $(1 - \text{MAC})$ and are accumulated over all considered modes. A drawback of this approach is that the simulation model has to be calibrated to closely match the actual measurement data, which may be hard to achieve in practice.

Alternatively, the mode shapes vectors can be compared directly

$$\varepsilon_{\Phi}(\mathbf{x}) = \sum_{k=1}^{N_{\text{modes}}} |(\Phi_{\text{S1},k}(\mathbf{x}) - \Phi_{\text{S0},k}) - (\Phi_{\text{M1},k} - \Phi_{\text{M0},k})|^2, \quad (3.3)$$

where ε_{Φ} is the mode shape error based on the measured mode shapes Φ_{M} and the simulated mode shapes Φ_{S} . The subscript $(\cdot)_0$ refers to the undamaged state, while $(\cdot)_1$ refers to the damaged one. The design variables only influence the simulation results for the damaged case, with all other terms of Equation 3.3 remaining constant during the optimisation run. While Equation 3.1 accepts arbitrarily scaled mode shapes vectors $\Phi_{(\cdot)}$, Equation 3.3 requires the mode shapes to be normalised. The normalisation is necessary due to the subtraction of mode shape vectors, which is similar to the enhanced COMAC metric [49]. The objective function according to Equation 3.3 is formulated in a way that cancels out much of the modelling error by subtracting the reference state from the analysis state. This is achieved, because the metric compares the relative change due to the damage between the model and measurement data. Such a formulation can be advantageous, as the difference between a damaged state and the reference state is often smaller than the difference between the simulation model and the actual structure. By cancelling out the initial mismatch, only the sensitivity towards damage is taken into account, which usually matches the actual structural behaviour significantly better than the absolute values of the modal parameters. As a result, the application of a relative metric often allows for omitting a model calibration, while still achieving a reasonable sensitivity. This can lead to a more practical approach overall, as demonstrated in Section 5.2.

Another common method to compute the difference between numerical and experimental data is based on the eigenfrequencies [71]. A least squares error metric considering N_{modes} eigenfrequencies is

$$\varepsilon_f(\mathbf{x}) = \sum_{k=1}^{N_{\text{modes}}} \left(\frac{f_{S1,k}(\mathbf{x}) - f_{S0,k}}{f_{S0,k}} - \frac{f_{M1,k} - f_{M0,k}}{f_{M0,k}} \right)^2, \quad (3.4)$$

where ε_f is the eigenfrequency error. The eigenfrequencies extracted from measurement and the simulated eigenfrequencies are denoted as f_M and f_S , respectively. The notation for the reference and analysis state is the same as in Equation 3.3. In Equation 3.4, the eigenfrequencies are normalised to weigh the contribution of all modes equally. This normalisation prevents higher frequencies from dominating the metric and means that changes in lower frequencies will also affect the objective function. This frequency normalisation is similar to the formulation of $\varepsilon_\Phi(\mathbf{x})$, where the normalised mode shapes are used. Since this formulation employs the same difference scheme as Equation 3.3, it also allows for the usage of uncalibrated models.

In previous works on model updating, the mode shape and eigenfrequency error metrics are frequently weighted and summed up to yield a compound metric [104, 106]. This is expressed by an objective function according to

$$\varepsilon(\mathbf{x}) = \alpha_f \varepsilon_f(\mathbf{x}) + \alpha_\Phi \varepsilon_\Phi(\mathbf{x}), \quad (3.5)$$

where α_f and α_Φ are the weighting factor for the frequency and mode shape metric respectively. However, the weighting factors which yield the best model updating result are unknown prior to the optimisation run. A common way to solve this issue is to weigh all metrics equally [104], but this assumption may lead to sub-optimal identification results.

The application of multi-objective optimisation can remove this shortcoming [83], since both error metrics can be solved for simultaneously. In fact, the Pareto frontier resulting from the multi-objective formulation contains the optimal solutions for all conceivable weighting factor combinations. This approach thus provides additional information about the uncertainty of the damage analysis. In contrast, such information is not available in a single-objective optimisation formulation, where only a single solution results from the optimisation.

3.1.4 Modal analysis

The model updating approach presented in this thesis is based on modal parameters which result from modal analyses of FE models and based on measurement data. Generally, the equation of motion for a structure with multiple degrees of freedom is expressed as

$$\mathbf{M}\ddot{\mathbf{u}}(t) + \mathbf{B}\dot{\mathbf{u}}(t) + \mathbf{K}\mathbf{u}(t) = \mathbf{p}(t), \quad (3.6)$$

where \mathbf{M} , \mathbf{B} and \mathbf{K} denote the mass, damping and stiffness matrices, respectively. Further, $\mathbf{u}(t)$ is the time-dependent displacement vector, $\mathbf{p}(t)$ is an exciting force vector and t expresses time [87]. The derivatives with respect to time lead to a differential equation system capable of damped oscillation.

Since most civil engineering constructions can not be directly modelled using multiple degree of freedom systems, they are spatially discretised. For the purpose of this thesis, FE modelling is employed to achieve this [12]. By also discretising time, Equation 3.6 can be solved numerically to yield the transient dynamic response of a structure. By disregarding the damping and under the absence of an exciting force, the structure will oscillate in its eigenfrequencies. This oscillation takes the form of a harmonic vibration, such that Equation 3.6 can be simplified to

$$(\mathbf{K} - \omega_{0i}^2 \mathbf{M})\mathbf{u}_{0i} = \mathbf{0}, \quad (3.7)$$

where ω_{0i} is the angular frequency of the vibration and \mathbf{u}_{0i} contains the amplitudes of the shape of the deformation. Since this is a generalised eigenvalue problem, ω_{0i} is referred to as the eigenfrequency and \mathbf{u}_{0i} is referred to as the eigenvector or alternatively as the mode shape vector. The modal parameters of an FE model can thus be directly computed by solving the generalised eigenvalue problem stated in Equation 3.7.

In addition to the modal parameters of the mechanical model, the mode shapes and eigenfrequencies from measurement data are necessary to perform model updating using objective functions based on modal parameters. As the dynamic exciting forces are usually unknown, operational modal analysis is necessary. To this end, many data-driven methods to extract modal parameters have been proposed, with notable methods being the Stochastic Subspace Identification (SSI) [118] and Frequency-Domain Decomposition (FDD) [18]. For the purpose of model updating, the results of any modal analysis method can be used. Because the structures investigated in this thesis exhibit well-separated modes [17], the relatively simple FDD method is used, which yields adequate results for these structures.

Vibration data can be captured using a multitude of sensor types, including displacement gauges, geophones, accelerometers or strain gauges. Of these sensor types, accelerometers are the simplest ones to handle since they can be employed independently of their orientation relative to earth's gravity field, do not require a fixed reference system and are easily removable. Hence, the simulations as well as the experiments considered in this thesis were carried out using accelerometer time series data.

3.2 Damage localisation method

The damage localisation using model updating presented in this thesis is achieved using a multi-objective formulation based on both eigenfrequency and mode shape error metrics. The finite element model is parameterised using damage distribution functions. This novel combination enables a balance between accuracy, robustness and computing times, which can not be readily achieved using a single-objective formulation. The details of this parameterisation and the objective function are given in the next sections.

3.2.1 Parameterisation of damage distribution

The FE model is parameterised using the damage distribution functions defined by Bruns et al. [22]. These distribution functions are intended for use with slender beam-like structures. The approach features only three parameters, namely the damage position, intensity and extent. To obtain a high spatial resolution using the super-element approach, a lot more parameters would be necessary. The low number of parameters is a major advantage, since the numerical complexity of optimisation problems grows exponentially with the number of variables. This rapid growth also came to be known as the ‘curse of dimensionality’ [9]. In addition to the computational costs, the problem can become ill-posed when too many parameters are used [106]. Ill-posed problems can have several local minima in addition to the actual solution which do not represent physically meaningful solutions. This can lead to the ‘wrong’ local minimum being picked as the solution, which in turn leads to identification results which do not match reality. In the context of model updating, the optimisation problem is usually uniquely solvable when the number of variables is in the low to medium single-digit range.

The parameterisation using damage distribution functions controls the structural stiffness and the associated parameters are expressed as

$$\mathbf{x} = [\mu \quad D \quad \sigma]^T, \quad (3.8)$$

where μ is the centre of the damage, D is the damage intensity and σ describes the spatial extent of the damage distribution. These parameters are summarised in the design variable vector \mathbf{x} , which is solved for by the numerical optimiser.

The damage distribution function is used to scale the initial stiffness EI^0 of each element. This is accomplished using a scaling factor θ_i for each element i

$$EI_i^\theta = \theta_i \cdot EI_i^0. \quad (3.9)$$

The values of θ_i are calculated using a cumulative distribution function $F(s_i|\mu, \sigma)$ which operates on the node positions s_i along the structure. Additionally, the cumulative distribution function is clamped within the interval

$0 \leq s_i \leq L$, which is written as $F(s_i|\mu, \sigma, 0, L)$. The parameter L denotes the length of the structure. Following this approach, the stiffness scaling factors θ_i are defined as

$$\theta_i = 1 - DL \frac{F(s_{i+1}|\mu, \sigma, 0, L) - F(s_i|\mu, \sigma, 0, L)}{s_{i+1} - s_i}. \quad (3.10)$$

Since the value range of the stiffness scaling factors is not restricted to positive values by Equation 3.10, it is possible that negative θ_i values may arise for low values of σ . However, negative stiffness values would lead to meaningless results. To avoid this issue, all models with negative stiffness values are rejected prior to FE calculation. Since this approach creates a discontinuity in the objective function, a constraint is added in order to soften this discontinuity. Therefore, the minimum stiffness scaling factor is used to formulate an inequality constraint, which restricts values below 15% of the original stiffness. The constraint is mathematically expressed as

$$\min_i (\theta_i) > 0.15. \quad (3.11)$$

For the relatively mild damage cases regarded in this thesis, Equation 3.11 does not significantly influence the outcome of the optimisation runs. It rather only becomes relevant in the initial iterations of the optimisation algorithm, when the sampling pattern is not yet focused close to the global optimum. However, the constraint helps to guide the optimisation algorithm away from unreasonable areas in the design variable space and, thus, aids in increasing the convergence speed.

3.2.2 Objective function formulation

The model updating formulation chosen in this thesis has two objectives. The eigenfrequency error as well as the mode shape error are minimised according to

$$\begin{aligned} & \text{minimise} \begin{pmatrix} \varepsilon_f \\ \varepsilon_\Phi \end{pmatrix} \\ & \text{s.t. } [0 \quad 0 \quad 0]^T \leq \mathbf{x} \leq [L \quad D_{\max} \quad \sigma_{\max}]^T \\ & \text{s.t. } \min_i (\theta_i) > 0.15, \end{aligned} \quad (3.12)$$

where \mathbf{x} is the design variable vector, L is the length of the structure, D_{\max} is the maximum damage intensity and σ_{\max} is the maximum damage extent to be considered. The inequality constraint given in Equation 3.11 is also part of this formulation.

The multi-objective approach leads to a more complex numerical problem compared to a single-objective problem. However, there are also several advantages associated with employing multiple objectives. First, a weighting of the objectives prior to the optimisation run is not necessary, as detailed in Equation 3.5. This lack of weighting removes the chance of misidentifying the damage due to improperly chosen weighting factors. Second, multi-objective optimisations are more robust, since the tendency of convergence towards local minima is damped by the presence of secondary objectives which help to guide the algorithm. Third, there are multiple solutions to a multi-objective problem, which convey an impression of the certainty of damage identification.

3.3 Verification using numerical simulation

For the numerical verification of the proposed optimisation algorithm and model updating method, a cantilever beam structure is chosen to represent a typical engineering problem. Mechanical parameters of this structure are given in Table 3.2. The beam is discretised using 36 one-dimensional finite elements which obey Bernoulli linear beam theory. Nine virtual acceleration sensors are placed uniformly along the structure, starting at the tip. A schematic illustration of the considered finite element model is shown in Figure 3.1. The dynamic excitation consists of pink noise which is applied uniformly along the structure. Additionally, Gaussian-distributed white noise is added to the time series of the virtual sensors to simulate measurement noise. The resulting amplitude spectrum is shown in Figure 3.2 using a singular value spectrum plot [98]. The peaks at the eigenfrequencies are well-separated due to the one-dimensional nature of the simulated beam.

Table 3.2: Properties of simulated cantilever beam

Property	Value
Elastic Modulus	200 GPa
Length	1 m
Height	5 mm
Width	50 mm
Density	7850 kg/m ³
Damping Ratio	0.15 %

Damage is simulated by reducing the element stiffness in predetermined regions. The objective of the model updating therefore is to identify the damage in terms of severity and location. Since the same numerical model is used to simulate the damaged structure and for the model updating, it would in princi-




Figure 3.1: Illustration of simulated cantilever beam structure, finite element discretisation and sensor positions.




Figure 3.2: Singular value spectrum of undamaged simulated cantilever beam showing the peaks belonging to the first five vibration modes.

ple be possible to perfectly solve such a problem. However, in practice there are modelling, measurement and modal identification errors which make obtaining a valid solution for a model updating problem considerably harder due to the added uncertainty. In this regard, some measures were taken to complicate the task and rule out the possibility of a perfect solution: First, the time domain simulation contains significant amounts of noise from the excitation and added measurement noise. Second, a relatively short duration of the time series data is chosen to increase the uncertainty in the modal analysis. Third, the modal parameters for the ‘measured’ reference and analysis states are identified from time domain simulations using operational modal analysis, thus introducing uncertainty in the modal parameters. As a last measure, the distribution of the actual damage intentionally does not match the distribution function used for the model updating. The effect of the intentionally added noise is also reflected in the singular value spectrum shown in Figure 3.2. Naturally, the quality of the results would improve relative to the ones presented here, if less noise, longer analysis times and a damage distribution function matching the actual damage were used instead.

Twelve scenarios were simulated, which consist of the reference state as well as regions of reduced stiffness with varying damage severity. The damage leads

to a decrease in the eigenfrequencies of the structure as well as to alterations in the bending mode shapes. The simulated scenarios are summarised in Table 3.3. The first three scenarios function as control samples, in which realisations of the reference state are simulated using an undamaged model. The remaining nine scenarios feature damage in three positions along the beam with increasing severity. The centres of the damage distribution are located at a quarter, half and three quarters of the beam length, respectively. Damage is applied by uniformly reducing the stiffness of four neighbouring elements close to the damage centre. The smallest damage features a 3 % stiffness reduction and is designed to be quite hard to identify. The 10 % damage level represents a structural defect which should be above the threshold of monitoring systems, while the 30 % damage level is severe enough to judge upon the accuracy of the localisation and quantification of the model updating method.

Table 3.3: Damage scenarios and numbering.

		Damaged elements		
		8 ... 11	17 ... 20	26 ... 29
Relative stiffness reduction	0 %	1	2	3
	3 %	4	5	6
	10 %	7	8	9
	30 %	10	11	12

The influence of the damage scenarios on the eigenfrequencies as well as their absolute values are shown in Figure 3.3. The frequency deviation due to the smallest damage is much less than 1 %, while the largest considered damage leads to a deviation of approximately 5 %. The eigenfrequencies exhibit a relatively low sensitivity towards damage which has to be compensated for by very accurate modal parameter extraction techniques.

An illustration of the three most severe damage scenarios is shown in Figure 3.4, which also depicts the first five bending mode shapes. In these mode shapes, the local effect of damage can be observed based on the mode shapes. Deviations between the reference case and the damage scenarios are most pronounced at locations with a high mode shape curvature close to the damage. The figure also makes clear that mode shapes by themselves are not very sensitive towards structural damage. Even though the stiffness reduction of 30 % is quite severe, the changes in the mode shapes are barely visible. Hence, monitoring based on modal parameters is dependent on highly accurate modal identification methods in order to produce significant damage identification results.

For the verification, the modal parameters are extracted from time-domain simulations with a duration of one minute each. The beam model is excited dynamically using a pink noise force spectrum uniformly distributed along the




Figure 3.3: Illustration of relative frequency deviation of the considered damage scenarios for the first five bending modes. Note, that the damage scenarios are sorted according to damage position.

beam. The discrete time domain simulation is carried out using a sampling frequency of 8000 Hz and an output sampling frequency of 4000 Hz. The relatively high data rate is necessary to avoid aliasing effects. The enhanced frequency domain decomposition method [19, 120] is used to identify modal parameters from the time series data. For model updating, the identification results obtained for the reference state represent the measured healthy state M0, as defined in Table 3.1. The identification results of the twelve damage scenarios are considered as the analysis state M1.

To verify the robustness of the model updating scheme towards mismatched damage distributions, a Gaussian damage distribution is chosen for the model updating procedure. This distribution therefore intentionally differs from the uniform distribution used to simulate the damage scenarios. The damage distribution is parameterised according to the formulation introduced in Section 3.2.1 using the centre of damage μ , the damage intensity D and the damage extent σ . The domain of the optimisation is restricted to a hypercube which encompasses the length of the structure for the damage location, a maximum damage level of 0.3 and a maximum damage extent of 1 m. These values were chosen such that the considered damage scenarios are situated inside the domain.

The MOGPS algorithm is run with its parameter set to $T = 50$ and considering 5000 objective function evaluations for each of the damage scenarios. Figure 3.5 shows the evolution of the design variables for an exemplary optimisation run. This figure demonstrates, that the sampling pattern does not fully converge towards a single value for each design variable, which would appear as a straight line in the plot. Instead, a broader band forms, in which the sampling pattern still remains discernible. This band corresponds to points near the Pareto-optimal areas of the design space. The damage parameter for the

depicted scenario can be calculated analytically as

$$D_{\text{actual}} = \frac{(1 - \theta) L_D}{L} \approx 0.03, \quad (3.13)$$

where θ is the stiffness scaling factor, L_D is the length of the damage and L is the length of the beam. In Equation 3.13, the L and L_D can also be expressed in terms of number of finite elements, as the element lengths are uniform. The damage location for the depicted scenario is $\mu_{\text{actual}} = 0.5$ m and was also accurately identified by the model updating scheme. The damage extent for all the scenarios is approximately $\sigma_{\text{actual}} \approx 0.05$ m. The objective space for the same damage scenario is shown in Figure 3.6. It reveals, that an unsymmetrical continuous Pareto frontier has formed which includes large variations of the mode shape error and small variations concerning the eigenfrequency error. It exemplifies that predetermining weighting factors in order to obtain a single-objective model updating problem would be a hard task in practice due to the severe non-linearity of the frontier.

The non-dominated solutions for all damage scenarios can further be visualised in terms of the stiffness distribution along the beam, as shown in Figure 3.7. This plot also allows for the visual confirmation of the damage quantification and localisation. Here, the impact of the statistical variance created by the simulated noisy measurement data becomes apparent, as the model updating method struggles to identify the smaller damage scenarios.

The top row of Figure 3.7 shows the control cases with no damage and the identified distributions reflect that. The second row contains the results for the smallest damage. While the method yields solutions different from the control case, the damage distributions were not recovered very well. For scenarios 4 and 6, the solutions which emphasise the mode shape error fit the actual damage better, while in scenario 5, solutions with a higher weight on the eigenfrequency error lead to better results. The mediocre performance for these scenarios is expected, as the short simulation time leads to a low spectral resolution for the operational modal analysis and the small damage leads to very small changes in the dynamic behaviour. Even though damage localisation is not successful in most of the previously discussed cases, the identification is possible in all cases. In this instance, damage identification refers to model updating results which point to damage distributions that are significantly different from those obtained in the reference cases. The third row shows the results for the medium damage cases and at least one distribution found by the multi-objective approach matches the actual damage well in each case. However, the results also showcase the downside of this approach, as many identified distributions point to incorrect damage locations. Notably, solutions with more weight towards the eigenfrequency error lead to a localisation mirrored about the centre point

of the structure in scenario nine. This may be attributed to the lower spatial information contained in the eigenfrequencies as opposed to mode shapes.

The results for damage scenario 11, which is showcased in Figures 3.5 and 3.6, can be seen in the centre of the last row. While the other two scenarios in the last row are identified with a tight damage distribution, a higher scatter is apparent for scenario 11 due to the large extent of the Pareto frontier for this particular scenario. Here, solutions which emphasise the eigenfrequency error more than the mode shape error lead to a less accurate localisation. The relatively high scatter in this scenario can be attributed to nodes in the mode shapes close to the damage, as shown in Figure 3.4. Since there is no mode shape curvature at the nodes of vibration modes, damage occurring at the nodes does not impact the respective mode shapes or eigenfrequencies. This essentially conceals the damage from detection using the affected modes and hence reduces the information available to the model updating method.

Minor deficiencies aside, these findings confirm that the model updating method as well as the optimisation algorithm can fulfil their respective tasks in a verification setting. The localisation and quantification are adequately accurate for higher damage severity levels. Additionally, the results confirm that an optimal a-priori weighting between eigenfrequency errors and mode shape errors is not possible. The multi-objective approach also highlights the high uncertainty in the identification of some damage scenarios. Consequently, the multi-objective approach is able to provide significantly more information from a single optimisation run than would be possible using single-objective optimisation. In addition, the single-objective results are also recovered by the multi-objective approach in form of the end-points of the Pareto frontier. While conceptually, it would be preferable to obtain a single solution, the results show that monitoring methods need to deal with the vagueness inherently contained in the underlying measurement data. Of course, the results presented in this section only give a rough insight into the performance of the proposed method, as only a single statistical realisation was examined for each damage scenario. Further, just one type of structure with a single sensor layout was shown here, so further research is necessary to generalise the presented findings. For instance, investigations regarding the effect of employing different damage distribution functions were carried out by Wolniak et al. [126]. A validation of the presented multi-objective approach, using measurement data obtained from an experimental structure, is presented in Chapter 5.




Figure 3.4: Illustration of stiffness distribution and first five bending mode shapes for damage scenarios 10 to 12.




Figure 3.5: Evolution of design variables during the optimisation run for damage scenario 11.

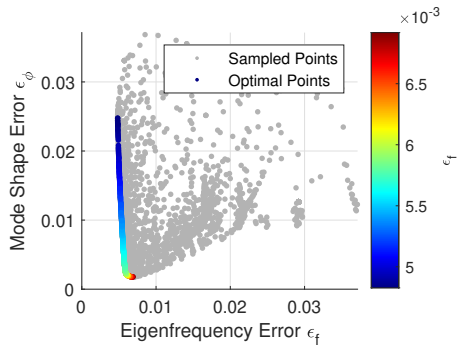


Figure 3.6: Objective value space and non-dominated (optimal) points for damage scenario 11.

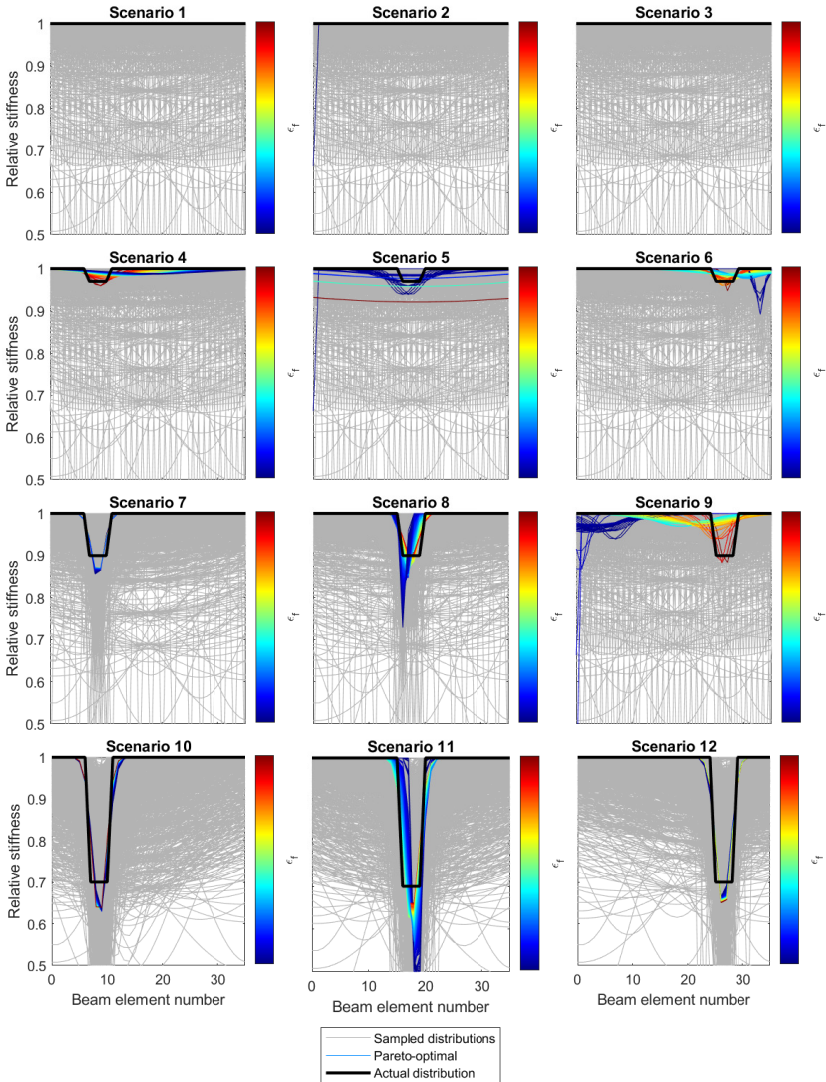


Figure 3.7: Stiffness distribution of damage scenarios and identification results of model updating. The order is according to Table 3.3.

Chapter 4

Damage localisation using impulse response identification

Damage localisation using data-driven vibration measurement data has been an active research topic for a long time and many methods have been developed to tackle the problem. Data-driven damage identification based on output-only methods allows for a quick real-time assessment of the monitored structure. The results can be used to decide whether more computationally intensive methods, such as model updating, should be employed to evaluate potential damage. Hence, data-driven monitoring is a crucial part of practical damage localisation frameworks.

In this thesis, a data-driven damage localisation scheme based on multiple-input Finite Impulse Response (FIR) filters is proposed. To the knowledge of the author, this approach has not yet been covered in the context of structural health monitoring and may have conceptual advantages over more traditional approaches. Since a relatively high filter model order is required, novel techniques to increase the numerical efficiency are proposed as well. The following sections introduce the identification of FIR filters, describe the resulting damage localisation method and present the verification results.

This chapter contains parts of research published by the author of this thesis as well as preprints [46, 47].

4.1 Theory of impulse response filters

Damage localisation in an output-only setting and without a physics-based model is generally achieved by detecting changes in the characteristics of the dynamical system. This is accomplished by comparing the reference behaviour,

which is identified from a reference measurement data, to the observed behaviour when monitoring the structure. Based on this comparison, damage-sensitive features are extracted, which are used to identify potential damage and its location.

4.1.1 Data-driven damage localisation

A central idea in data-driven damage localisation is pinpointing sensor positions close to structural changes, which may indicate damage. Depending on the way that the underlying structural dynamic identification is formulated, the extracted damage-sensitive features take various forms. This section describes some notable damage localisation methods and damage-sensitive features, that were proposed in the past decades.

As demonstrated in Chapter 3, accurate localisation results can be obtained using modal parameters and finite element model updating. However, in a purely data-driven context, modal parameters are much harder to interpret because no information on the modal dynamics of the structure is available beyond the observations contained in the data. Nonetheless, damage localisation based on modal parameters can be accomplished e.g. by using the mode shape curvatures [86]. This method exploits the local change in stiffness of structures due to damage and the sensitivity of mode shapes towards this effect. While for beam-like structures with a uniform stiffness distribution, this works reasonably well, the mode shape curvature does not perform particularly well for non-uniform structures. This is due to the correlation between the stiffness distribution of the underlying structure and the resulting curvature changes. Another downside of this method is that the mode shape curvatures are computed for each eigenmode separately and it is often unclear which combination of modes provides the best localisation accuracy.

While many authors have proposed damage detection methods based on parametric identification [117, 131], the localisation based on parametric models is a considerably harder endeavour. A method based on the power of difference processes obtained using state estimators was proposed by Lenzen and Vollmering [68] and was extended by Wernitz et al. [123]. This method is based on the stochastic subspace identification of system models and a projection method to enhance the sensitivity. While control theory-based parametric approaches are powerful tools, they are often mathematically complicated and some have many parameters which have to be tuned appropriately to achieve satisfying results.

Several authors successfully applied localisation approaches, which use transmissibility functions to describe the dynamic behaviour on a local level [75, 121]. This approach, however, has drawbacks regarding the numerical stability, as discussed by Chesné and Deraemaeker [25]. Due this practical limitation in the frequency-domain identification of multiple-input transmissibility functions,

usually only the single-input topology is used. In the latter case, the transmissibility from one sensor to another sensor is identified for several pairs of sensors. Additionally, transmissibility functions are very sensitive to changes in the location of excitation sources. These limitations associated with the multiple-input frequency-domain identification, however, can be reduced using time-domain identification [46].

In this thesis, the modal parameter approach is used as part of the presented model updating method. Thus, to supplement the relatively complex model updating process with a less computationally complex method, a non-parametric time-domain multiple-input approach is taken for the data-driven method introduced in this thesis.

4.1.2 Transmissibility functions and impulse response filters

Impulse response filters can capture the transmission behaviour of linear time-invariant systems. In mechanical structures, they can be used to describe the vibration transmission between sensor positions. Transmissibility functions in a discrete time setting provide a mapping from a transient signal $x(z)$ to another transient signal $\hat{y}(z)$, where z denotes the time-discrete Laplace symbol [108]. This can be expressed graphically as a signal flow, where the transmissibility function $T(z)$ acts as a digital filter.

$$x(z) \longrightarrow \boxed{T(z)} \longrightarrow \hat{y}(z)$$

The transmissibility function corresponding to the same filter is

$$T(z) = \frac{\hat{y}(z)}{x(z)}, \quad (4.1)$$

where $T(z)$ is the transmissibility function, $x(z)$ is the spectrum of the measured signal and $\hat{y}(z)$ is the spectrum of the estimated signal. The transmissibility can in practice be computed from the frequency response of the system [121].

In contrast, the finite impulse response filter corresponding to the transmissibility function shown above can be expressed as [96]

$$\hat{y}[i] = \sum_{j=0}^M b[j] x[i - j], \quad (4.2)$$

where M is the filter order and the coefficient vector \mathbf{b} is comprised of $M + 1$ elements. The indices i and j denote the time step and the filter coefficients, respectively. This formulation implies that the finite impulse response filter ignores all samples outside the range $[i - M, i]$. The finiteness of FIR filters is

due to the fact that they do not consider input samples which lie more than M samples in the past. The impulse response is thus truncated in time. With the formulation according to Equation 4.2, causality of the filter is also implied, which means that changes at the output \hat{y} must not temporally precede changes in the input x .

Impulse response filters and transmissibility functions express the same concept, i.e. they provide a mathematical description of the transmission path between two measurement points. The difference is that the former is formulated in time domain and the latter in frequency domain.

4.1.3 System and control theory considerations

In this section, the FIR filter topology is analysed from a system and control theory standpoint. In detail, the properties finiteness, controllability, observability and invertability are regarded, as these constitute the defining characteristics for any given linear filter model.

In terms of the discrete version of the Laplace transformation, also known as the z -transformation, the structure of a single-input and single-output finite impulse response filter can be expressed as [108]

$$G(z) = \frac{b[0] + b[1]z^{-1} + b[2]z^{-2} + \dots + b[M]z^{-M}}{1}, \quad (4.3)$$

where $G(z)$ is the transfer function and $b[0..M]$ are the coefficients of the FIR filter. The denominator of the transfer function is unity, which indicates that the filter has no poles and is thus unconditionally stable. Due to the unconditional stability, for any finite input, the filter response will always be finite. This is important for practical dynamic systems, since the predicted amplitude must be bounded for any conceivable input signal.

To examine the controllability and observability, the state-space model of a single-input/single-output system is formulated

$$\begin{aligned} \mathbf{x}_{ss}[i+1] &= \mathbf{A}_{ss}\mathbf{x}_{ss}[i] + \mathbf{b}_{ss}x[i] \\ y[i] &= \mathbf{c}_{ss}^T\mathbf{x}_{ss}[i] + d_{ss}x[i], \end{aligned} \quad (4.4)$$

where \mathbf{A}_{ss} is the $(M \times M)$ -dimensional state matrix, \mathbf{b}_{ss} is the input vector, \mathbf{c}_{ss}^T is the output vector and d_{ss} is the feedthrough factor. The input time series x and the output time series y are incorporated into the system using the time step i . In order to represent a FIR filter, the parameters are set to

$$\begin{aligned} A_{ss,kl} &= \delta_{k-1,l} & b_{ss,k} &= \delta_{k,1} \\ c_{ss,l}^T &= b[l+1] & d_{ss} &= b[0], \end{aligned} \quad (4.5)$$

where the Kronecker delta function δ is used to describe the contents of \mathbf{A}_{ss} and \mathbf{b}_{ss} . In this formulation, the state space matrix \mathbf{A}_{ss} has ones on the lower secondary diagonal and is zero otherwise. This means that the entire internal state is shifted in each time step. The input vector \mathbf{b}_{ss} feeds the most recent input $x[i]$ to the beginning of this delay line. The output vector \mathbf{c}_{ss}^T contains the coefficients $b[j]$ of the FIR filter, starting from the second coefficient. Since the first coefficient $b[0]$ of the filter needs to be applied without any delay, it is contained in the feedthrough factor d_{ss} . The controllability matrix following from this state space representation is $\mathbf{C} = [\mathbf{b}_{ss} \quad \mathbf{A}_{ss}\mathbf{b}_{ss} \quad \cdots \quad \mathbf{A}_{ss}^M\mathbf{b}_{ss}] = \mathbf{I}$, meaning that the output is unconditionally controllable. The observability matrix $\mathbf{O} = [\mathbf{c}_{ss} \quad \mathbf{A}_{ss}^T\mathbf{c}_{ss} \quad \cdots \quad (\mathbf{A}_{ss}^M)^T\mathbf{c}_{ss}]^T$ only has full rank if all filter coefficients are non-zero, in which case the system is observable as well.

Since there are no restrictions imposed on the filter coefficients contained in \mathbf{c}_{ss} , the zeros of Equation 4.3 are usually not bounded to the unit circle. This means that an inverse filter obtained by flipping the numerator and denominator of the transfer function is generally unstable. However, it is possible to obtain a stable least-squares optimal inverse finite impulse filter using Equation 4.9 [94].

4.2 Damage localisation method

To localise structural damage, the reference structure is dynamically excited and finite impulse filters are derived using acceleration data. In this process, data from multiple adjacent sensors is used as an input to derive the impulse response to the neighbouring sensor positions. The derived filters are applied to obtain an estimation of the transient response for the reference as well as analysis states. Residual signal energies between measured and predicted data are calculated, which increase locally when structural damage occurs, enabling the localisation.

The filter identification and the computation of the residuals are discussed in the following sections.

4.2.1 Single-input impulse response identification

An FIR filter is identified by finding a vector of filter coefficients \mathbf{b} , which transforms the input signal \mathbf{x} , so that it approximates the measured output signal \mathbf{y} as good as possible. This time domain identification enables capturing transient dynamics as well as the stationary response. This means that the residual signal power between the filter response and the measured output signal has to be minimised. Filter identification also requires that the time series \mathbf{x} and \mathbf{y} have to be recorded using a fixed sampling rate and be synchronised

to prevent temporal drift over time. The unconstrained linear least-squares optimisation problem yielding \mathbf{b} can be expressed as [96]

$$\underset{\mathbf{b}}{\text{minimise}} \sum_{i=M+1}^m \left(y[i] - \sum_{j=0}^M b[j] x[i-j] \right)^2, \quad (4.6)$$

where m denotes the number of measured time steps \mathbf{x} and \mathbf{y} . In contrast to Equation 4.2, both the signal \mathbf{x} and the signal \mathbf{y} are measured quantities. In order to make this problem tractable using linear algebra, Equation 4.2 is rewritten using a matrix-vector multiplication [96]

$$\mathbf{T} \cdot \mathbf{b} = \mathbf{y}, \quad (4.7)$$

where \mathbf{T} is an $(m - M - 1) \times (M + 1)$ Toeplitz matrix containing data from m measurement time steps. The matrix contains the samples of the filter input $x[i]$ shifted to the respective temporal positions as defined in Equation 4.2, so that

$$T_{ij} = x[i - j]. \quad (4.8)$$

Equation 4.7 assumes the form of an over-determined system of equations, when the number of samples m is larger than twice the order of the FIR M . To attenuate the influence of measurement noise and signal contamination, it is beneficial to derive the coefficients \mathbf{b} with $m \gg M$, so that an average solution is achieved. In practical applications, this usually means several minutes of measurement data are used, which can easily amount to several hundred thousand data samples.

In this thesis, only causal filter responses are considered, i.e., no temporally preceding samples are used as inputs. This is done in an effort to make the method capable of real-time monitoring. However, the method is not restricted to that and by adjusting the indices that establish the shift matrix, it is possible to also consider non-causal filters [108].

Using the Moore-Penrose pseudoinverse, a least-squares optimal solution to Equation 4.7 is obtained, which is known as the Affine Projection Algorithm [85]

$$\mathbf{b} = \left(\mathbf{T}^\top \mathbf{T} \right)^{-1} \mathbf{T}^\top \mathbf{y}. \quad (4.9)$$

Equation 4.9, however, usually does not yield a usable filter in practice, since measurement noise can deteriorate the quality of the identification. This can be mitigated in part by providing large amounts of measurement data from diverse excitation states. However, the high-frequency response usually contains

artefacts caused by overfitting, since the signal-to-noise ratio worsens in the high-frequency domain due to low excitation forces.

A method to reduce the effects of overfitting is the regularisation of the system of equations. This can be achieved by employing the Tikhonov regularisation [115], which results in a regularised form of the Affine Projection Algorithm [32]

$$\mathbf{b} = (\mathbf{T}^\top \mathbf{T} + \lambda \mathbf{I})^{-1} \mathbf{T}^\top \mathbf{y}, \quad (4.10)$$

where $\lambda \mathbf{I}$ is the identity matrix multiplied by a regularisation factor. Low values of λ lead to solutions close to those obtained without any regularisation applied. High values lead to very smooth spectra, while underestimating the vibration level significantly. The regularisation therefore serves to minimise the energy contained in the filter coefficients. As a result, the value of the regularisation parameter is a compromise between accuracy and smoothness and has to be tuned to specific measurement setups to obtain good identification results.

The smoothing effect of the regularisation parameter defined in Equation 4.10 is dependent on the filter order and the signal amplitude of the input. Hence, it is normalised to achieve a comparable effect across a wide range of scenarios. A normalised regularisation parameter is proposed

$$\lambda = \lambda_0 \frac{\|\mathbf{T}^\top \mathbf{T}\|_F}{M + 1}, \quad (4.11)$$

where $\|\cdot\|_F$ indicates the Frobenius norm and $\lambda_0 \in \mathbb{R}_{\geq 0}$ is the normalised smoothing factor. The Frobenius norm is linked to the energy of the input signal and its usage as the basis for the regularisation parameter leads to a smoothing effect proportional to the signal amplitude. It is defined as

$$\|\mathbf{A}\|_F := \sqrt{\sum_i \sum_j |A_{ij}|^2}, \quad (4.12)$$

where A_{ij} are the elements of a matrix \mathbf{A} .


The smoothing is thus also invariant to the number of samples considered in the shift matrix \mathbf{T} . To eliminate the dependency on the filter order as well, Equation 4.11 includes a division by $M + 1$.

4.2.2 Expansion to the multiple-input/multiple-output case

Even though vibration estimation can be achieved using a single sensor time series as an input, it can be drastically improved by considering multiple input sensors. This is the case, since the vibrations recorded at additional sensor positions can contribute additional information about the vibration state of

the considered structure. Hence, the filter identification method is extended to account for several input time series. The basic idea for multiple-input finite impulse response filters was presented by Powell et al. [94]. An advanced method for data-driven identification of such filters was introduced by Chen et al. [24] and is incorporated in the MATLAB system identification toolbox [72]. As these previous approaches are numerically expensive and thus very time-consuming for high model orders, special consideration is given to the numerical performance of the identification procedure presented in this thesis.

The multiple-input/multiple-output finite impulse response estimation can be graphically modelled as follows



where the 1, 2 and 3 measurement directions of triaxial sensors are considered as an example. The matrix containing the filter coefficients \mathbf{b} is generally unsymmetrical, thus identifications have to be carried out individually for each filter in the matrix. For the sake of convenience, a triaxial case is assumed in this derivation, however, the formulation can be readily extended to account for any number of input and output channels. For example, measurement data from two triaxial sensors could be used as the inputs $x_{1...6}$, and data from only one triaxial sensor could be used as the outputs $\hat{y}_{1...3}$. Consequently, the matrix \mathbf{b} would assume a rectangular shape. Since the system of equations is overdetermined anyway, no special treatment is required for such cases.

Similar to Equation 4.7, there is also an equivalent matrix expression that is conducive to the solution by a pseudoinverse

$$\mathbf{T} \cdot \mathbf{b} = \begin{bmatrix} \mathbf{T}_1 & \mathbf{T}_2 & \mathbf{T}_3 \end{bmatrix} \cdot \begin{bmatrix} \mathbf{b}_{11} & \mathbf{b}_{12} & \mathbf{b}_{13} \\ \mathbf{b}_{21} & \mathbf{b}_{22} & \mathbf{b}_{23} \\ \mathbf{b}_{31} & \mathbf{b}_{32} & \mathbf{b}_{33} \end{bmatrix} = \mathbf{Y}, \quad (4.13)$$

where \mathbf{T}_1 denotes the shift matrix associated with the first input channel, \mathbf{T}_2 denotes the second channel and so forth. The time series of the measured outputs are denoted using the matrix \mathbf{Y} , which has a size of $(m - M - 1) \times n_y$, where n_y denotes the number of output channels. The coefficients of the impulse responses \mathbf{b}_{11} through \mathbf{b}_{33} can therefore be determined simultaneously, according to Equation 4.10. Since the pseudoinverse itself is not dependent on the output vectors \mathbf{y} , it only needs to be computed once, independently of the number of output channels. In the case of one triaxial input and one triaxial output, the filter coefficients are obtained using

$$\begin{bmatrix} \mathbf{b}_{11} & \mathbf{b}_{12} & \mathbf{b}_{13} \\ \mathbf{b}_{21} & \mathbf{b}_{22} & \mathbf{b}_{23} \\ \mathbf{b}_{31} & \mathbf{b}_{32} & \mathbf{b}_{33} \end{bmatrix} = \left(\begin{bmatrix} \mathbf{T}_1^\top \\ \mathbf{T}_2^\top \\ \mathbf{T}_3^\top \end{bmatrix} \cdot [\mathbf{T}_1 \quad \mathbf{T}_2 \quad \mathbf{T}_3] + \lambda \mathbf{I} \right)^{-1} \cdot \begin{bmatrix} \mathbf{T}_1^\top \\ \mathbf{T}_2^\top \\ \mathbf{T}_3^\top \end{bmatrix} \cdot \mathbf{Y}. \quad (4.14)$$

Equation 4.14 thus yields the nine filter coefficient vectors describing the vibration transfer from every input channel to every output channel.

The regularisation parameter as per Equation 4.11 also applies to the multiple-input case. Since the matrix \mathbf{T} has more elements in the multiple-input case compared to the single-input case, the numerical value of its matrix norm increases. Consequently, the value of the regularisation parameter λ according to Equation 4.11 increases as well. Thus, the smoothing effect of the regularisation remains constant.

In the following sections, the equations are derived without loss of generality using the single-input/ single-output topology to simplify the notation.

4.2.3 Interpolated finite impulse response filters

It is possible to solve Equation 4.14 efficiently by exploiting the Toeplitz structure of the matrix \mathbf{T} to evaluate the expression $\mathbf{T}^\top \mathbf{T}$. However, the computer memory required to store the square matrix $\mathbf{T}^\top \mathbf{T}$ is in many cases prohibitively large. Therefore, the numerical filter identification scheme outlined above can be further improved by considering a filter topology with less coefficients. Sparse finite impulse response filters, also known as ‘tap delay’ filters [108], can significantly reduce the numerical complexity of FIR filters while maintaining a high fidelity. Instead of identifying $M + 1$ filter coefficients for the full model order, a smaller number of $N + 1$ coefficients is used. The non-zero coefficients are referred to as filter ‘taps’ while the remaining coefficients are usually set to zero.

The achievable sparsity depends on the high frequency damping of the system, where strongly damped systems are conducive to high reduction ratios. For band-stop filter design, a reduction of non-zero coefficients exceeding 50% can be achieved in practice [55].

To achieve the best possible sparse filter quality, it is advisable to place the filter taps where the amplitude of the dense filter coefficients is highest. Generally, in mechanical systems with viscous damping, high frequency components decay faster than low frequency components. This means that most of the energy and high-frequency oscillation is contained in the beginning of the impulse response. The filter taps should thus be placed densely at the beginning of the impulse response, in order to be able to capture the high frequency content. The taps in the tail of the filter should be placed at larger distances, since low-

frequency signals dominate in this section [55]. For this purpose, a quadratic function to place the filter taps is proposed

$$t_k = k - 1 + (M - N) \left(\frac{k - 1}{N} \right)^2, \quad (4.15)$$

where $k \geq 1$ denotes the filter tap index and t_k is the corresponding tap position. Equation 4.15 forces a tap spacing greater than one, i.e., above the sampling period. This prevents the coefficients from being overdetermined in the identification process. Further, the tap position for the index $k = N + 1$ coincides with the filter coefficient $M + 1$, placing the last tap at the end of the filter.

The tap positions t_k resulting from Equation 4.15 are real-valued. For a direct application of sparse filters, they would have to be rounded to the next integer. However, the use of sparse filters leads to a poor spectral signal quality. This can be overcome by employing Interpolated Finite Impulse Response (IFIR) [84]. IFIR are based on interpolation functions, which are assigned to each filter tap. These interpolation functions act as window functions, and are usually of the Bartlett type [77]. A more detailed discussion of these functions is provided in the latter part of this section.

The interpolation functions can be formulated using the real-valued tap positions t_k derived using Equation 4.15. The use of interpolation functions results in a bandwidth limitation and thus improves the spectral quality. The interpolation functions are combined in the $M \times N$ interpolation matrix \mathbf{H} . By applying this matrix to the sparse filter coefficients, the dense filter coefficients are recovered

$$\mathbf{b} = \mathbf{H}\tilde{\mathbf{b}}, \quad (4.16)$$

where $\tilde{\mathbf{b}}$ are the sparse coefficients. Further, the interpolation coefficients are constrained to have a unit sum for each sparse coefficient k

$$\sum_{j=-\infty}^{\infty} H_{jk} \stackrel{!}{=} 1. \quad (4.17)$$

Equation 4.17 leads to a normalisation of the signal content in each of the sparse coefficients and is beneficial for regularisation. Even if the filter coefficients are only defined in the interval $[0, M]$, the interpolation functions may extend beyond this interval. The sum in Equation 4.17 is thus defined over the interval $[-\infty, \infty]$ to achieve a consistent formulation.

Equation 4.7 can be rewritten using Equation 4.16

$$\mathbf{T}\mathbf{H}\tilde{\mathbf{b}} = \mathbf{y}, \quad (4.18)$$

again representing an overdetermined system of equations. The order of multiplications in Equation 4.18 is modified

$$\tilde{\mathbf{T}} = \mathbf{T}\mathbf{H} \implies \tilde{X}_{ik} = \sum_{j=0}^M H_{jk} x[i - j], \quad (4.19)$$

where the shift matrix \mathbf{T} is eliminated and replaced by a convolution of the input with the interpolation matrix to yield the sparse shift matrix $\tilde{\mathbf{T}}$. A similar scheme was put forward for adaptive IFIR filters by Wu et al [129]. Equation 4.10 can be used to determine the sparse filter coefficients by substituting \mathbf{T} for $\tilde{\mathbf{T}}$ and \mathbf{b} for $\tilde{\mathbf{b}}$

$$\tilde{\mathbf{b}} = \left(\tilde{\mathbf{T}}^T \tilde{\mathbf{T}} + \lambda \mathbf{I} \right)^{-1} \tilde{\mathbf{T}}^T \mathbf{y}. \quad (4.20)$$

The regularisation parameter λ defined in Equation 4.11 also applies for interpolated filters, since the signal energy content in \mathbf{T} and $\tilde{\mathbf{T}}$ is roughly the same. This is the case, since Equation 4.17 forces the interpolation filters to have a unit gain in the passband. The dense filter coefficients \mathbf{b} are recovered from the sparse filter coefficients $\tilde{\mathbf{b}}$ using Equation 4.16, and can subsequently be used for the vibration prognosis.

Triangular interpolation functions, also known as Bartlett windows, are often used for computational resource efficiency in real-time processing [77]. In this work, Gaussian window functions are proposed, which allow for smoother interpolations than Bartlett windows while sacrificing some numerical efficiency.

The centres of the Gaussian windows are placed at the real-valued tap positions, such that $\mu_k = t_k$. A parametrisation of the window using

$$\sigma_k = \frac{1}{2} \begin{cases} t_{k+1} - t_k & \text{if } k = 1 \\ t_k - t_{k-1} & \text{if } k = N + 1 \\ \frac{t_{k+1} - t_{k-1}}{2} & \text{otherwise} \end{cases} \quad (4.21)$$

is proposed, where σ_k is the standard deviation of the Gaussian interpolation function associated with the k -th filter tap. This way, neighbouring tap positions have a distance of approximately two standard deviations. The different cases arise due to the first and last tap position, which do not have neighbours to both sides. Since the area under the curve of a Gaussian distribution is always one, Equation 4.17 is readily fulfilled. An illustration of the interpolation matrix is shown in Figure 4.1, which also depicts the overlap between the windows resulting from Equation 4.21. The interpolation functions are clipped outside the interval $[0, M]$, however, this only affects the first and last function and, thus, has negligible influence on the results. Figure 4.2 displays a single interpolation function taken from this matrix.




Figure 4.1: Plot of the interpolation coefficient matrix with $N = 10$ and $M = 40$. The individual interpolation functions are each highlighted with a different colour. The coefficients are connected with solid lines.




Figure 4.2: Interpolation coefficients of interpolation function $k = 8$ from Figure 4.1 shown using stem plot. Centre of the function at $t_k = 21.7$ indicated using a dashed line.

Another benefit of Gaussian window functions is the theoretically infinite side lobe suppression in the frequency domain [103]. In practice, the window length is limited to $M + 1$, but side lobe suppression down to the numerical noise floor is still achievable. This is indicated by Figure 4.3, which is obtained from the interpolation function shown in Figure 4.2, where the flat section of the graph close to -250 dB is caused by the numerical round-off error. To reduce the number of non-zero elements in \mathbf{H} and, thus, minimise the multiplications required to compute Equation 4.19, a truncated Gaussian function is used. This truncation is carried out by setting the filter coefficients with values below 1×10^{-5} to zero. This truncation is tuned to obtain a side lobe suppression greater than -100 dB. The spectrum resulting from truncation is superimposed in Figure 4.3.




Figure 4.3: Frequency response spectrum for the Gaussian interpolation function shown in Figure 4.2. A truncation of the interpolation function leads to the emergence of side lobes.

4.2.4 Computational performance

The evaluation of equation 4.20 is computationally intensive, due to the matrix multiplication $\tilde{\mathbf{T}}^T \tilde{\mathbf{T}}$. The computational complexity is of the order $\mathcal{O}(mN^2)$, where m is the number of measurement data samples and N is the number of filter taps. This is valid under the assumption that $m \gg M$, i.e., that the number of considered measurement data samples is much larger than the filter order. The numerical complexity implies that the number of filter taps N should be as low as possible, since it has a quadratic impact on the computing time.

Further, the shift matrix $\tilde{\mathbf{T}}$ takes up more memory than typically available on a desktop computer, even for small numbers of measurement samples. This slows down the computation of $\tilde{\mathbf{T}}^T \tilde{\mathbf{T}}$, because the speed of this matrix multiplication is limited by the memory bandwidth. To reduce the size of the shift matrix, Equation 4.19 is computed in a batch operation for small chunks of measurement data. The terms $\tilde{\mathbf{T}}^T \tilde{\mathbf{T}}$ as well as $\tilde{\mathbf{T}}^T \mathbf{y}$ can thus be summed up iteratively to minimise the computing time. Additionally, splitting the measurement data into batches enables parallel computation, since the batches can be processed independently. Thus, the batch computation scheme also enables an efficient computation on many-core computers.

Due to the relatively low dimension of $\tilde{\mathbf{T}}^T \tilde{\mathbf{T}}$, the matrix inversion in Equation 4.20 has an insignificant impact on the computing time, if long time series of measurement data are employed. Hence, no special consideration of performance aspects is required for the matrix inversion.

4.2.5 Damage-sensitive residuals

Using the multiple-input filter topology, a vibration estimation can be calculated considering K input signals

$$\hat{y}[i] = \sum_{k=1}^K \sum_{j=0}^M b_k[j] x_k[i-j], \quad (4.22)$$

where k is the index corresponding to the input signals. For slender structures, it is straightforward to use adjacent sensors for the estimation of a measurement position in between. Figure 4.4 illustrates this concept for a simple beam structure using a biaxial sensor setup.

The impulse responses identified using Equation 4.13 contain the structural dynamics on a local level and for the full spectrum. When structural changes due to damage, environmental or operating conditions occur, the vibration prognosis will then no longer reflect the actual structural behaviour. Residuals are used to detect such changes with respect to the reference state. The residual ε is therefore defined as the difference between the estimated vibration signal and




Figure 4.4: Sketch of beam structure with three measurement locations. The time series $y[i]$ is estimated using the coefficient vectors \mathbf{b}_1 through \mathbf{b}_4 and the signals x_1 through x_4 obtained at adjacent sensor locations. Signals measured in the x -direction and y -direction are marked red and blue, respectively.

the measured one

$$\varepsilon[i] = y[i] - \hat{y}[i]. \quad (4.23)$$

The average residual power P is calculated by accumulating the residuals for the whole time series as a root mean square value to obtain a damage-sensitive feature

$$P = \sqrt{\frac{1}{|\varepsilon|} \sum_i (\varepsilon[i])^2}, \quad (4.24)$$

where $|\varepsilon|$ refers to the number of samples – or cardinality – of ε . For long data series with a zero mean value, this metric yields similar results to the statistical standard deviation.

Using the residuals of the reference state and the damaged cases, a relative metric can also be constructed. The comparison between the two states is achieved by dividing the analysis state power by the reference power

$$\tilde{P} = \frac{P_{\text{ana}}}{P_{\text{ref}}} - 1, \quad (4.25)$$

where the subscripts ‘ana’ and ‘ref’ refer to the analysis and reference states, respectively. Relative metric calculations are conducted individually for each sensor position. Subsequently, potential damage sites are localised by determining the relative residual power with the maximum value. The metric according to Equation 4.25 can lead to negative values for \tilde{P} when the analysis data set contains less signal power compared to the reference data set. However, since the residuals can not be directly interpreted in a mechanical sense, this is not deemed to be a major issue. It rather shows that damage quantification is not easily possible using residual methods.




Figure 4.5: Exemplary spectrum of ‘measured’, estimated and residual signal of reference case for the fifth sensor.

4.3 Verification using numerical simulation

To verify the FIR-based damage localisation scheme, the simulation model and damage cases introduced in Section 3.3 are reused. The multiple-input filters representing the reference state are identified using the exact same one-minute time series data sets as were employed for the model updating approach, in an effort to compare both methods on the same basis. These filters are employed to compute the residuals for the reference and analysis data sets. The filters are set up so that the two neighbouring sensor positions are used as the input for the prognosis, similar to the scheme depicted in Figure 4.4. For the outermost sensors, the closest two sensors were chosen as the inputs. For the filter identification, the parameters $M = 2000$, $N = 550$ and $\lambda_0 = 0.03$ were used, based on a model order selection method [47]. Considering the sampling rate of 4000 Hz, this means that the impulse response filters have a time-span of 0.5 s.

Figure 4.5 shows an exemplary spectrum of an acceleration sensor on the simulated beam as well as the spectrum estimated by the FIR filters. Additionally, the residual spectrum is given, which is computed from the difference of the actual signal and the estimation time series using Equation 4.23. The spectrum of the estimation precisely follows the actual signal, which means that the two signals are close to identical. Consequently, the residual spectrum is dominated by the noise floor. This finding indicates a very good match between the identified filters and the observed system.

Since the residuals at all sensor positions have to be regarded for damage localisation, all the spectra of all sensors should be examined at once for a better overview. Hence, Figure 4.6 shows the spectra of all sensor positions and the respective residuals. The spectra of the estimations themselves are left out to reduce the amount of redundant information. A slight correlation between the residuals and the actual signal can be observed in some channels, which




Figure 4.6: Exemplary spectrum of ‘measured’ and residual signals for reference case and all sensors.




Figure 4.7: Exemplary spectrum of ‘measured’ and residual signals for all sensors and damage scenario 11, according to Table 4.1.

means that the residual spectra rise in frequency ranges where the ‘measured’ spectrum has peaks. This indicates sub-optimally identified filters or excitation signals which could not be picked up at the sensors which were used as the filter inputs. However, overall the residuals mostly exhibit a flat frequency response, which points to an adequate identification quality.

The relatively flat residual spectra observed in the reference case are contrasted by the damaged case as displayed in Figure 4.7. In this plot, the residual spectrum has significant sharp peaks rising from the noise floor. These peaks are exactly correlated with the frequencies of the peaks observed in the ‘measured’ spectrum. This, in turn, means that the estimation does not match well with measurement and a system change, e.g. due to damage, is likely. For residual-based methods, the sensitivity towards damage depends on the ratio between the amplitude of the peaks in the damaged case and the noise floor in the reference case. This sensitivity depends on many factors, such as the model order, the signal-to-noise ratio in the measurement data, modal damping or the density of the sensor network.

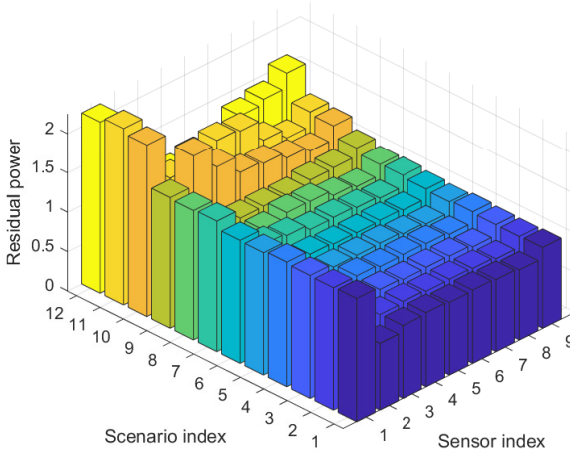


Figure 4.8: Residual power for all considered scenarios. Sensor indices indicated according to Figure 3.1.

The rise of the residual spectrum due to damage is also reflected by the residual power, which is calculated using Equation 4.24. Figure 4.8 shows this metric at all sensor locations and for all 12 damage scenarios as defined in Table 4.1. It shows that damage increases the residual power at all sensor locations and in each damage case. This rise in the power is especially recognisable in the damage scenarios 10 through 12. For reference, the damage cases defined in Section 3.3 are reproduced in Table 4.1.

Table 4.1: Damage scenarios and numbering.

		Damaged elements		
		8 ... 11	17 ... 20	26 ... 29
Relative Stiffness Reduction	0 %	1	2	3
	3 %	4	5	6
	10 %	7	8	9
	30 %	10	11	12

The general idea behind residual-based damage localisation is that the location of the maximum residual power should match the actual position of the

damage. A localisation of the damage directly from the residual power shown in Figure 4.8 is, however, not possible in this numerical example. Figure 4.9 shows that the relative residual power metric, i.e. Equation 4.25, improves the quality of the damage-sensitive feature substantially and leads to the desired results.

The first three scenarios depicted in Figure 4.9 are realisations of the reference state and show a distribution of the relative residual power without prominent features. Further, the amplitude of the relative power in these scenarios reaches 0.03 at maximum. The diagrams for the other scenarios have a distinctly different characteristic. In the damaged cases, the relative power is significantly higher than in the reference case. In the scenarios 4 through 6, which belong to the 3% stiffness reduction, the maximum power doubles to about 0.06. The larger damage introduced in the scenarios 7 through 12 leads to more pronounced elevations of the residual power levels. Aside from the absolute amplitude, the distributions of the relative power in the damaged cases appear similar to each other. This relates to the damage positions, which are the same for each level of damage severity. The actual damage positions for the damage scenarios are indicated in the diagrams using dashed ellipses. As desired for damage localisation, the maxima of the relative residual power roughly coincides with the position of the damage. While localisation can be verified in this case, the residual power at sensor locations far away from the damage is still significant. This can be attributed to the changes in global dynamics which accompany local damage. An ideal residual-based method would yield zero at all sensor positions except the one close to the damage location. The presented method therefore has much room for improvement in this regard.

The data-driven localisation method based on multiple-input FIR filters is very sensitive toward damage and can detect even the smallest damage case which features only 3% stiffness reduction. However, the localisation accuracy can not exceed the resolution of the sensor network. This means that when a high spatial resolution is desired, a large effort has to be undertaken for the experimental instrumentation. While a data-driven approach can be employed without knowledge of the mechanical properties of the structure monitored, the vibration data still reflects its structural dynamics in terms of the modal parameters. Hence, sensor placement can also have an decisive effect on the quality of the results, even though mode shapes are not employed explicitly in such an approach. Additionally, this numerical example shows that residuals by themselves are not capable of damage quantification, as the relative residual power amplitudes have no straightforward relation to the stiffness reduction introduced into the structure. A thorough comparison of the model updating approach and the data-driven method presented in this section can be found in Chapter 7.

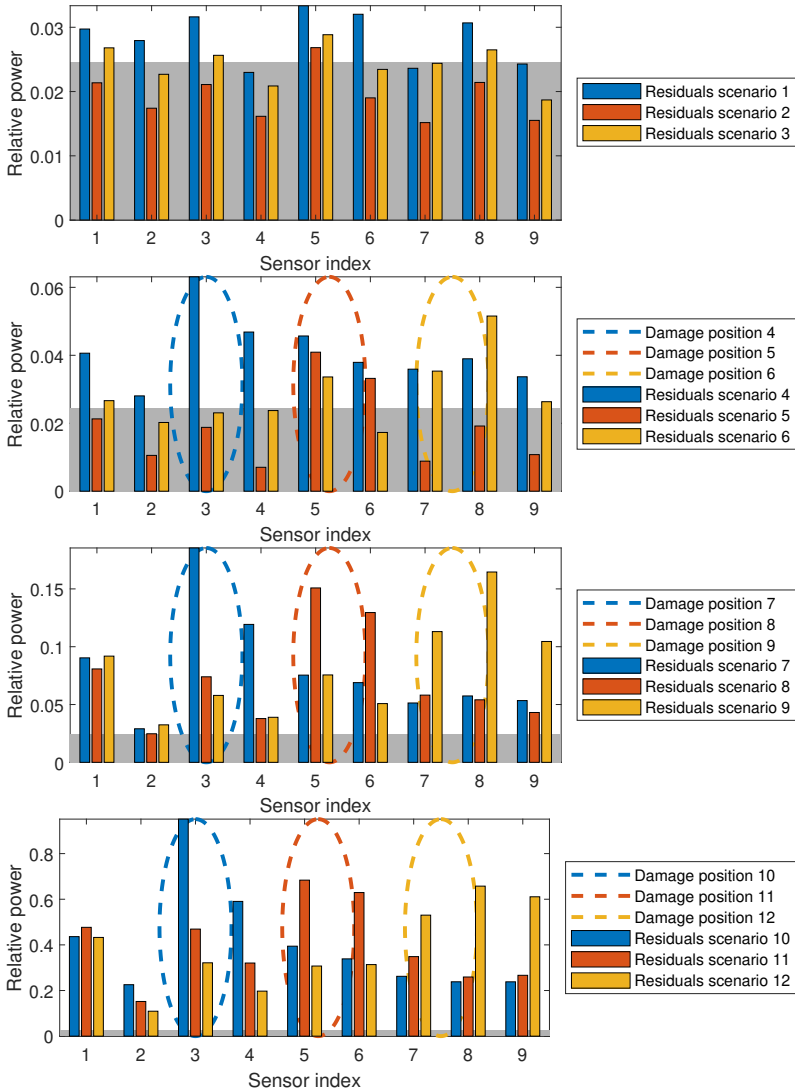


Figure 4.9: Relative residual power and damage locations. Scenarios with equal damage severity are combined into one diagram, respectively. Grey area denotes average power level of reference state, note the different axis scaling of the relative power metric.

Chapter 5

Experimental validation

The two damage localisation methods presented in the preceding chapters are validated using an experimental structure. The experimental setup, damage mechanisms, data sets and the results are discussed in the following sections. The validation is performed exemplarily using a single damage position to show that the methods are able to localise damage considering real-world measurement data. The first investigated damage scenario is a symmetrical damage case and the second scenario considers an unsymmetrical damage case.

This chapter contains parts of research published by the author of this thesis as well as submitted articles [44, 46, 52, 124].

5.1 Girder-mast structure

The experimental structure is a girder-mast, which is situated in an outdoor test facility. It consists of three equal sections which are bolted together. The segments have a length of 3 m, resulting in a total height of 9 m. Each segment has three legs and consists of seven bracing levels as well as short connecting sections at the ends. The legs are manufactured from steel tubing with an outer diameter of 30 mm and the bracings are solid steel rods with an diameter of 10 mm. The bay height is 400 mm. The structure is placed on a massive steel plate which is rigidly connected to a concrete foundation, providing a clamped boundary condition. A photograph of the structure is shown in Figure 5.1.

The measurement setup consists of nine measurement levels, which are each instrumented with biaxial Integrated Electronics Piezo Electric (IEPE) accelerometers in the horizontal plane. The locations of the measurement levels as well as the damage location are indicated in Figure 5.2. The sensors are mounted to one of the three legs using clamps and the cabling is run down on



Figure 5.1: Girder-mast structure located at the outdoor facility. Acceleration sensors are arranged in a biaxial setup along one of the three legs.




Figure 5.2: Schematic figure of measurement levels (ML), damage location as well as definition of global coordinate system.

that same leg of the structure. A temperature sensor to measure the material temperature is fitted as well. Besides the dynamic excitation by environmental conditions, no further sources of vibration are present close to the mast, which prevents the contamination of the measurement data. Wind excitation is thus the main contributor to the structural acceleration. Detailed information regarding the considered girder-mast structure was published by Wernitz et al. [124]. This publication also contains a link to an open-access archive containing the measurement data.

5.1.1 Measurement campaign

On several levels of the structure, a reversible damage mechanism enables the disconnection of the bracings, which changes the structural stiffness locally. The damage position used in the proceeding discussion is illustrated in Figure 5.2.

A photograph of the damage mechanism is depicted in Figure 5.3 and shows threaded connecting rods inserted into the bracings. The damage is activated by unscrewing and removing the rods in one or multiple bracings. By removing the rods, the structural stiffness of the respective bracing effectively becomes zero which consequently leads to a different dynamic response of the structure. For the purpose of the examined measurement campaign, either all braces or a single brace were removed from the respective bay. This means that the considered damage is either symmetrical or unsymmetrical about the axis of the structure. In addition, the removal of a single brace has less impact on the structural stiffness than the removal of all three braces. The threaded rods were completely removed from the structure, which technically also leads to a mass difference. However, the impact of the mass deficiency is negligible in comparison to the effect of the stiffness change.



Figure 5.3: Closeup of reversible damage mechanism incorporated into the structure.

For the validation, two time series are used, with one representing the state when all bracings were intact. In the second time series, the damage is activated. To obtain a reasonable baseline for comparison, the temperature and excitation level of the two time series were selected to be close to each other. In particular, both data sets were recorded at a temperature of $11\text{ }^{\circ}\text{C}$ and a wind speed of 5 m/s for the symmetrical damage. The data sets for the unsymmetrical damage were each recorded a temperature of $10\text{ }^{\circ}\text{C}$ and a wind speed of 5 m/s . The measurement time series were digitised using a sampling rate of 1652 Hz and each have a length of 600 s . The symmetrical damage is discussed in Sections 5.2 and 5.3, whereas the unsymmetrical damage is discussed in section 5.5.

5.1.2 Impact of structural damage

To obtain an initial impression of the damage impact on the structural dynamics, it is useful to regard the power spectral density of the vibration data. Since the structure is unconstrained at the top end, the maximum vibration amplitude due to wind excitation occurs at the top. Therefore, the spectra of the sensors at the top of the structure are depicted in Figure 5.4, considering the symmetrical damage scenario.




Figure 5.4: Spectra of reference and damaged states at the top measurement level for symmetrical damage. Graphs show x -direction (top) and y -direction (bottom).

From Figure 5.4, two main conclusions can be drawn: First, the power spectral density graphs of the two data sets exhibit a vertical offset. This points to an elevated vibration amplitude of the damaged time series compared to the reference time series. This is not only caused by the damage but mainly due a differing wind load which leads to a higher excitation level. Second, a horizontal shift to the left can be recognised in the spectra. This indicates a decrease of the eigenfrequencies of the structure, caused by the damage. These observations hold for measurement data obtained for both the x -direction and the y -direction.

5.2 Model-based localisation of symmetrical damage

The finite element model updating method is validated using the measurement data obtained from the girder-mast structure for the symmetrical damage scenario. Operational modal analysis is used to obtain eigenfrequencies and eigenmodes of the structure. A finite element model of the structure is parameterised using a Gaussian damage distribution function. The model updating problem is then solved using the Multi-Objective Global Pattern Search algorithm. More details are given in the following sections.

5.2.1 Modal analysis

The modal parameters of the structure are identified by applying the frequency domain decomposition method [19] on the measured 10 min data sets. The first three bending modes as well as the first torsion mode are considered for the model updating process. These modes are chosen due to a high spectral power and signal to noise ratio. Visualisations of the mode shapes and the corresponding eigenfrequencies are shown in Figure 5.5 for the undamaged and the damaged conditions, respectively.




Figure 5.5: Mode shapes and eigenfrequencies identified from measurement data for the girder-mast structure shown in Figures 5.1 and 5.2. Undamaged state indicated using dashed blue lines, damaged state indicated using solid red lines.

The damage location at the bottom of the structure clearly influences the modal behaviour on a local scale, as the mode shapes show significant deviations close to the damage position. Regarding the eigenfrequencies, the damage has quite significant impact, leading to deviations exceeding 10% in most of the

considered modes. Considering the first bending mode, the damage only has a minor impact.

5.2.2 Finite element model

The model updating is performed using a finite element model of the girder-mast structure, in which the legs and braces are modelled using first-order beam elements. The cross-sections are assigned according to the measured dimensions of the structure. As the structure is made of construction steel, a Young's modulus $E = 210$ GPa and a Poisson's ratio of $\nu = 0.3$ are used as the mechanical material properties. The density is specified as $\rho = 7850$ kg/m³. The modal analysis of the model is carried out using the finite element solver Abaqus. Figure 5.6 shows the resulting mode shapes and eigenfrequencies, which closely resemble the results obtained from the measurement data of the undamaged structure shown in Figure 5.5.




Figure 5.6: Finite element results for mode shapes and eigenfrequencies of the undamaged state.

While the mode shapes of the model fit to the measured data, the eigenfrequencies exhibit systematic discrepancies of approximately 15%. The most likely explanation for the observed deviation is that the elastic modulus or wall

thickness of the steel structure are slightly lower than assumed in the finite element model. This prohibits direct comparisons between the model and the measured data. However, the relative metrics defined in Section 3.1.3 are specifically designed to compensate for such errors. Hence, a meticulous calibration of the model can be omitted, which means that it can be directly employed for damage localisation without any modifications. While this saves time and effort, it also leads to a more stringent procedure compared to an approach involving model calibration. In particular, the mechanical model and the measurement data are kept separate from each other before the application of the error metric. Therefore, the mixing of these domains happens at the last possible stage of the procedure.

5.2.3 Parameterisation

The finite element model is parameterised using the damage distribution function formulation introduced in Section 3.2.1. The stiffness scaling factors θ_i used to scale the elastic modulus of the respective bays prior to the modal analysis of the model are

$$\theta_i = 1 - DL \frac{F(s_{i+1}|\mu, \sigma, 0 \text{ m}, 9 \text{ m}) - F(s_i|\mu, \sigma, 0 \text{ m}, 9 \text{ m})}{s_{i+1} - s_i}. \quad (5.1)$$

where i denotes the index of the bay, s_i is the bay's height above ground and F is a truncated Gaussian cumulative distribution function. Equation 5.1 therefore links the damage distribution parameters to the stiffness of each element in the model. These parameters are contained in the vector of unknowns

$$\mathbf{x} = [\mu \quad D \quad \sigma]^T, \quad (5.2)$$

where μ describes the damage location, D represents the damage intensity and σ parameterises the spatial damage extent. Even though the actual damage is introduced only to the diagonal braces, the stiffness scaling is applied to the legs as well in each bay of the model. This is done to obtain a more generalised approach, in which it is not necessary to know in advance, whether the damage occurred in a brace or in a leg. By applying the stiffness scaling to the whole bay, the model updating method is sensitive to damage both in the braces and the legs.

This leads to the formulation of the bounded and constrained multi-objective optimisation problem

$$\begin{aligned} & \text{minimise} \begin{pmatrix} \varepsilon_f \\ \varepsilon_\Phi \end{pmatrix} \\ & \text{s.t. } [0 \quad 0 \quad 0]^T \leq \mathbf{x} \leq [9 \text{ m} \quad 0.3 \quad 2 \text{ m}]^T \\ & \text{s.t. } \min_i (\theta_i) > 0.15. \end{aligned} \quad (5.3)$$

Equation 5.3 restricts the domain of the optimisation to a hypercube which encompasses the length of the structure for the damage location, a maximum damage level of 0.3 and a maximum damage extent of 2 m. These values were chosen to reflect realistic damage scenarios and such that the actual damage intensity and extent are contained within this domain. The objective functions for the multi-objective optimisation are based on the relative eigenfrequency and mode shape errors, as defined in Equations 3.3 and 3.4. The relative formulations of these metrics provide a compensation for systematic deviations between the actual mechanical structure and the finite element model. An inequality constraint is applied to the minimum value of the stiffness scaling factors θ_i . This prevents the optimisation from searching for solutions which belong to excessively damaged structures and also prevents the occurrence of negative stiffness values. The constraint is enforced using the exterior linear penalty method [26]. Since the model updating is formulated in a multi-objective setting, the resulting penalty term is applied equally to both objective function values.

5.2.4 Damage identification and localisation

The MOGPS and NSGA-II algorithms are run with 2000 objective function evaluations using $T = 16$ and a population size of 64, respectively. These values were chosen according to the results obtained in the test function benchmarks discussed in Section 2.4.4. The resulting non-dominated solutions are shown in Figure 5.7. In this plot, the NSGA-II results are shown as black circles, while the MOGPS results are colour-coded according to ε_f . Both algorithms converge to similar solutions, with MOGPS achieving a higher resolution along the Pareto frontier, which confirms the findings made in Section 2.4.4. In practice, however, the results of both algorithms are sufficient for the purpose of damage localisation.

Figures 5.8 and 5.9 show the design variable space resulting from the two optimisation runs. Since the design space is three-dimensional, two projections of the design variable space are displayed. The colour-coding in these figures is consistent with the one used in Figure 5.7. These plots show that all Pareto-optimal solutions are concentrated in a region close to $\mu = 1$ m. The highest scatter is along the σ axis which means that the geometrical extent of the damage is identified with a relatively high uncertainty. Interestingly, the scatter is associated with solutions which belong to low eigenfrequency errors, while the solution which favour minimising the mode shape error are more concentrated. This may be attributed to the higher amount of spatial information contained in mode shape vectors as opposed to the scalar eigenfrequencies.

When comparing Figures 5.8 and 5.9, the different approaches to sample generation between deterministic optimisation algorithms and those based on random numbers becomes apparent. The MOGPS samples follow its axis-aligned

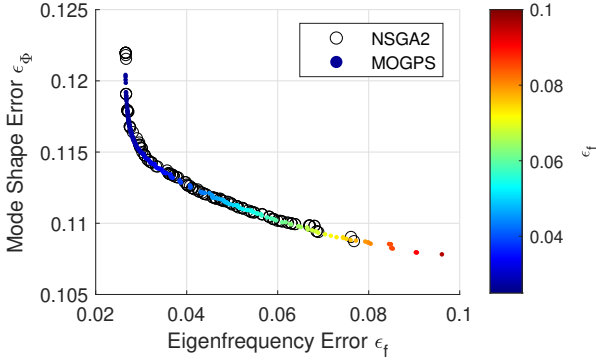


Figure 5.7: Pareto frontier of the multi-objective model updating problem.

pattern and concentrate towards the optimal regions. The shape of the optimal region is clearly discernible and uniformly surrounded by non-optimal samples. On the contrary, the NSGA-II results are dominated by scattering along the axes emanating from the optimal regions. Further, the outlines of the region containing solutions with low eigenfrequency errors are hardly recognisable in the NSGA-II samples, which correlates with the low number of points on the Pareto frontier shown in Figure 5.7.

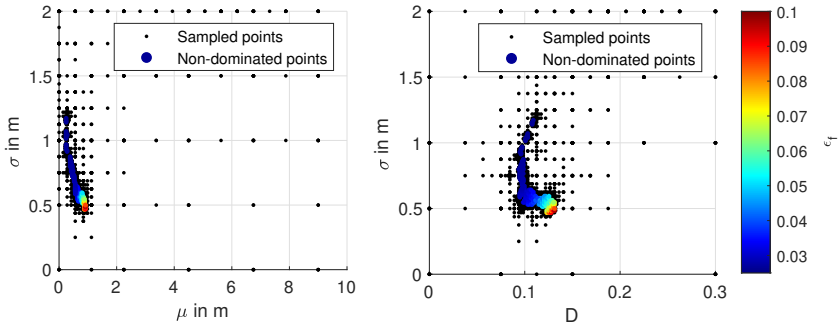


Figure 5.8: Design variable space for the multi-objective model updating using MOGPS.

Another way to analyse the optimisation results is to regard the sampling history. Therefore, Figure 5.10 displays the progression of samples during the optimisation runs of the two algorithms. For both algorithms, a convergence

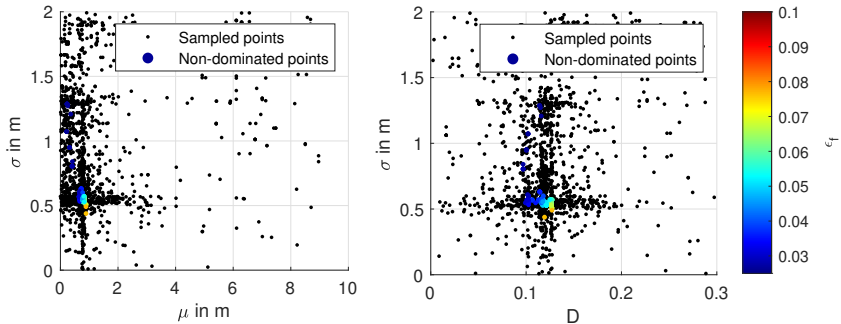


Figure 5.9: Design variable space for the multi-objective model updating using NSGA-II.

towards the optimal regions can be observed over the course of the runs, which is indicated by a denser sampling as the number of objective function evaluations increases. Contrary to the test functions shown for single-objective optimisation in Section 2.3.5, the sampling pattern does not narrow down to a line but instead to a broader interval. This is the case since the optimal region of the multi-objective model updating formulation is not point-like but rather a volumetric shape, which is shown using two orthogonal projections in Figure 5.8.

The non-dominated points identified by the algorithms can also be illustrated by the stiffness scaling factors associated with the corresponding design variables. Figure 5.11 shows the distributions with the same colour coding based on the eigenfrequency error as was used in Figures 5.7 to 5.9. In Figure 5.11, the distributions belonging to the lowest eigenfrequency error ε_f (i.e., blue colour) indicate a damage at the root of the structure. The distributions with the highest ε_f (i.e., red colour) point to damage at a height of approximately 1 m. The most accurate damage localisation is produced by solutions belonging to $\varepsilon_f \approx 0.05$ (i.e., cyan colour), which place the centre of the distribution in the bay where the damage actually occurred as indicated in Figure 5.11 by dashed lines.

The use of multi-objective optimisation in this verification enables a comprehensive insight into the most probable damage distributions of the structure. The assessment of the stiffness distribution functions in Figure 5.11 shows that no matter how the eigenfrequency and mode shape errors are weighted, the damage is always identified close to the actual damage location. However, the extremes of the Pareto frontier, i.e. the solutions belonging to single-objective optimisation using either only mode shapes or only eigenfrequencies, are located relatively far apart in the design variable space. This means that the

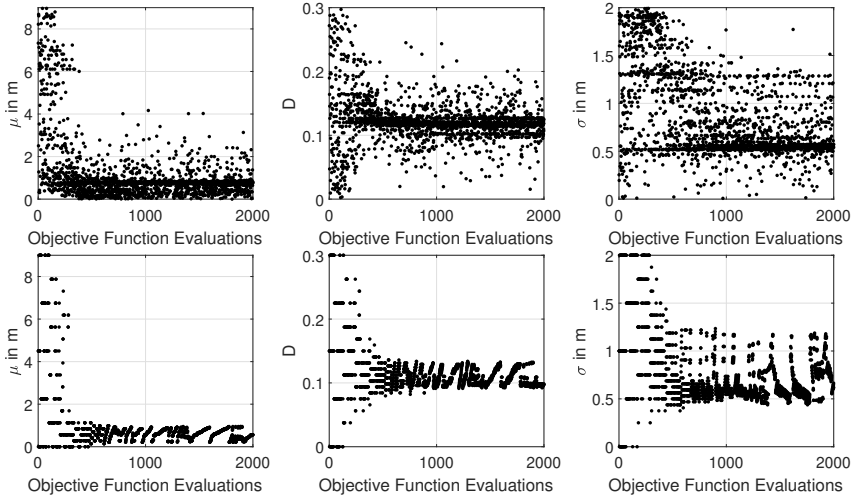


Figure 5.10: History of design variables for NSGA-II (top panel) and MOGPS (bottom panel).

multi-objective model updating approach is able to significantly increase the confidence level in terms of damage localisation compared to single-objective formulations. While NSGA-II is able to produce useful localisation results, the MOGPS algorithm exhibits favourable performance in this practical example as it is able to provide a significantly higher resolution of the Pareto frontier.

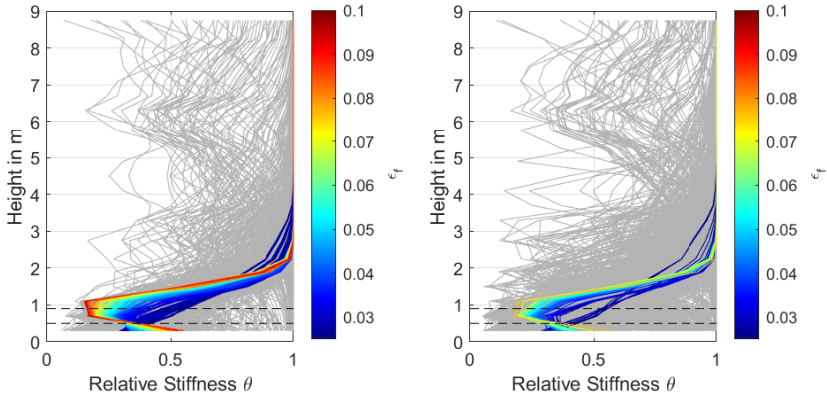


Figure 5.11: Stiffness scaling factors computed of the Pareto-optimal results for MOGPS (left panel) and NSGA-II (right panel). Position of experimentally damaged bay indicated with dashed lines.

5.3 Data-driven localisation of symmetrical damage

Using the same measurement data as before, the damage localisation method based on multiple-input finite impulse response filters is applied. In a first step, the coefficients of the finite impulse responses for the reference state are identified and discussed. In a second step, the residual power metric is applied to localise the damage.

5.3.1 Identification of impulse response filters

The scheme shown in Figure 4.4 is used to identify the filters using the validation measurement data set. The tower has a biaxial sensor setup with two sensors at each measurement level, hence, two estimation time series are calculated for each measurement level. Since the neighbouring measurement levels contain two sensors each, these are regarded as the inputs for the filters. Therefore, the 14 sensors in measurement levels ML2 through ML8, as defined in Figure 5.2, each have four contributing filters from the neighbouring levels. The filter inputs for the measurement levels ML1 and ML9 are set up to use the four closest sensors from the adjacent measurement levels. Using Equation 4.13, the filter coefficients are determined. In total, 72 finite impulse responses of the order $M = 4000$ were identified to obtain the vibration estimation for the local dynamics of the structure. Further, the parameters $N = 550$ for Equation 4.15




Figure 5.12: First 400 coefficients of the finite impulse responses contributing to the prognosis of sensor $5y$. These correspond to sensors from measurement levels ML4 and ML6 in both x and y -direction.




Figure 5.13: Spectra of the four finite impulse response filters for the adjacent sensors contributing to the acceleration signal prognosis for sensor $5y$.

and $\lambda_0 = 0.01$ for Equation 4.11 were used in this instance. An automatic model order selection method for these parameters was proposed by the author of this thesis [47]. As an example, the finite impulse response coefficients contributing to the prognosis of sensor $5y$ are shown in Figure 5.12.

The spectra corresponding to the filters are depicted in Figure 5.13. This figure shows that even though the estimated sensor is oriented in the y -direction, there are significant contributions of both the x and y -directions. This can be attributed to vibration modes, in which the different directions of motion are coupled, e.g. because the vibration occurs diagonally. Hence, the choice to include the orthogonal measurement direction in the filter identification process was reasonable.

5.3.2 Damage identification and localisation

To localise the structural damage, the residuals for the different damage states are calculated. The damage was introduced to weaken both the x and y -directions of the structure, therefore, residuals were expected to be roughly the same in both directions for this damage case.

Since the girder-mast structure is dynamically excited by ambient wind, the data exhibits varying signal amplitudes. To obtain useful results with the residual power method introduced in Section 4.2.5, the residuals need to be normalised to the signal amplitudes encountered in the respective data sets [38]. Therefore, Equation 4.23 is modified

$$\varepsilon[i] = \frac{y[i] - \hat{y}[i]}{\sqrt{\frac{1}{|y|} \sum_j (y[j])^2}}, \quad (5.4)$$

where the residuals are divided by the root mean square value associated with the vibration data. The power spectral densities of the reference and damaged case are shown in Figures 5.14 and 5.15. These plots also show the spectrum of the residuals. A clear difference between the reference and damaged cases can be observed in the frequency range between 0 Hz and 200 Hz, which is displayed separately in the graphs. While in the reference state, the residuals are close to the noise floor, they are strongly correlated with the measured signals in the damaged state. This suggests that the identified filters provide an adequate estimation of the structural dynamics up to 200 Hz and that the damage sensitivity is sufficient. In the range above 200 Hz, the residuals become more congruent with the measured spectra as the frequency increases. This means that the filters fail to capture the dynamic response in the high frequency range. When compared to the residual spectra shown in Section 4.3, which are obtained from a simulation, the filters perform a lot worse when applied to real measurement data. This is expected, since simulations usually fail to capture all of the environmental variability acting on an outdoor experimental structure. Additionally, the assumption of linear time invariance may be broken to some extent, as friction and aerodynamic interactions may contribute nonlinear effects.

To obtain an understanding of the stability of the residual power metric, two separate time series for the reference state are analysed. This is achieved by splitting the 10 min interval into two pieces of 5 min and using only the first 5 min for identification, while calculating residuals for both 5 min sets. The results of this procedure are shown in Figure 5.16. For the graphical presentation of the residuals, the two measurement directions are evaluated separately. The results show, that there is some variability of the residual power. However, this is to




Figure 5.14: Measured spectrum and residual spectrum showing all 18 sensors for the reference state.




Figure 5.15: Measured spectrum and residual spectrum showing all 18 sensors for the damaged state.

be expected due to the stochastic wind excitation. Further, the variability is small compared to the absolute residual power.

The residuals are not zero in the reference state because of effects, which are not adequately captured by the finite impulse response filters. These effects include inaccuracies in the filter identification, measurement noise, non-linear dynamics, etc. Residuals for the upper measurement levels of the structure are elevated, because the vibration amplitude in the considered girder-mast structure is highest at the locations farthest away from the foundation.

The damaged state was also analysed using the residual power metric, as presented in Figure 5.17. The measurement levels close to the damage location exhibit the highest residuals. However, judging only from this result, it would appear that more damage occurred in the y -direction, since the residual power in this direction is significantly higher than for the x -direction. This may also be exacerbated by the eccentric sensor layout, which captures translations as

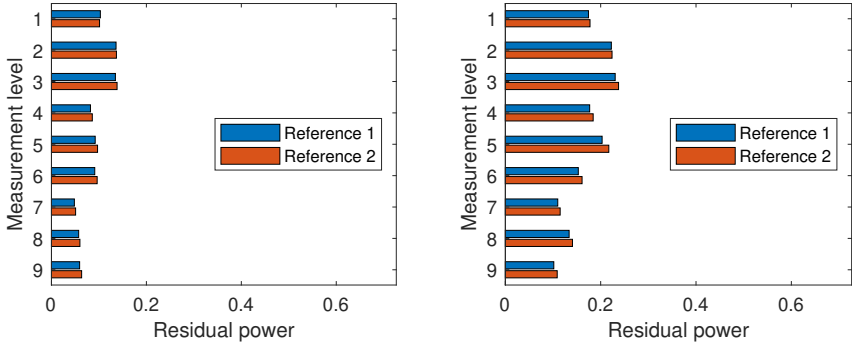


Figure 5.16: Comparison of the two reference states before introducing any damage. Residual power is shown for x -direction (left panel) and y -direction (right panel).

well as torsional motions in the y -direction, while capturing solely translations in the x -direction. To alleviate these issues, the relative residual power, as introduced in Equation 4.25, is used to further process the residuals. Figure 5.18 therefore displays the numerical results obtained using this metric. A significant improvement over the result shown in Figure 5.17 is achieved in terms of damage sensitivity. This way, Figure 5.18 demonstrates that a roughly equal damage in both directions of the structure as well as the approximate damage position can be clearly recovered from the residuals. While the relative residual is highest close to the actual damage at measurement level 9 in the x -direction, it is off by one measurement level in the y -direction.

As is typical for damage-sensitive features based on residuals, the relative power metric reacts to the structural damage at all positions. This ‘leakage’ is associated with changes in the global behaviour of the structure, which are overlaid with the local changes. Thus, a perfectly localised rise of the residual power at the position of the damage is not achieved. Rather, a gradual increase of the residuals towards the damage location is observed.

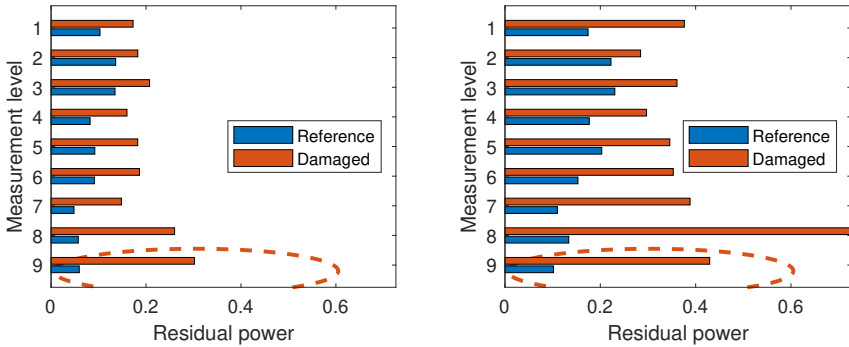


Figure 5.17: Comparison of the reference state and the damage state. Residual power is shown for x -direction (left panel) and y -direction (right panel). The true damage location is marked using a dashed ellipse.




Figure 5.18: Damage localisation results obtained using the relative residual power metric. Values are shown for x -direction (left panel) and y -direction (right panel). The true damage location is marked using a dashed ellipse.

5.4 Comparison of localisation results

The parameterisation of the model updating problem is designed under the assumption of a singular defect area. In contrast, the residuals of the filter-based estimation are not constrained to the detection of damage at a single location. Further, the model updating parameterisation allows for damage localisation on a continuous scale, as the damage location μ is a real number. Conversely, the data-driven method is limited in its spatial resolution by the sensor positions. Additionally, the model-based approach allows for the quantification of the damage in terms of mechanical stiffness, which the data-driven approach can not provide. Hence, the conceptual differences between these two approaches make it hard to compare the localisation results using a uniform metric.

Qualitatively, both methods were able to detect the introduced structural damage and narrow the damage position down to a region of approximately 1 m. Even though the localisation certainty is numerically expressed differently in the two approaches, the individual distributions resulting from the multi-objective optimisation and the relative residual power metrics point to a similar area of the structure. The damage distribution function used to parameterise the model updating problem leads to an exclusion of damage in the top part of the structure. This is due to the assumption of a single damage in the distribution function. Since the residuals are not constrained to find single damage location, they rise at all locations of the structure due to the damage. Thus, damage at the top of the structure can not be excluded by the data-driven method.

All in all, both methods deliver validation results which are in line with the findings of the verification. The localisation quality and accuracy are sufficient to clearly identify the damage introduced into the structure. Therefore, the validation with a symmetrical damage is concluded with satisfactory results.

5.5 Unsymmetrical damage scenario

Another pair of data sets was evaluated for damage at the same position but with a different severity. In this scenario, only a single brace was removed, leading to an unsymmetrical damage scenario. As a result, an unsymmetrical change of the vibration modes occurs, which leads to an overall unsymmetrical vibration behaviour of the structure.

Due to the unsymmetrical structural dynamic behaviour, the model updating method can not be applied in the same manner as in case of a symmetrical damage. An issue preventing the direct application of the formulation presented in Section 5.2.3 is shown in Figure 5.19. The figure illustrates a simulation result for the first torsional mode as well as the result of operational modal analysis based on measurement data. In the simulation, the damaged brace was com-

pletely removed from the model to closely resemble the actual damage scenario. Contrary to the symmetrical damage scenario depicted in Figure 5.5, the torsional mode shape does not bulge outwards close to the damage location, but rather inwards. With no further information given, this would point to stiffening of the structure, rather than structural damage. The reason for this behaviour is the sensor placement on a single leg of the structure, in which the sensors are situated on the opposite side of the damaged brace. As the neutral axis of the structure shifts away from the damaged brace, it moves closer to the leg equipped with sensors, which means that this leg experiences less vibration amplitude compared to the reference state. Due to this apparent stiffening effect, it is not possible to apply the radially symmetrical damage assumption stated in Equation 5.1, and, thus, a different model updating formulation needs to be used in the unsymmetrical scenario.

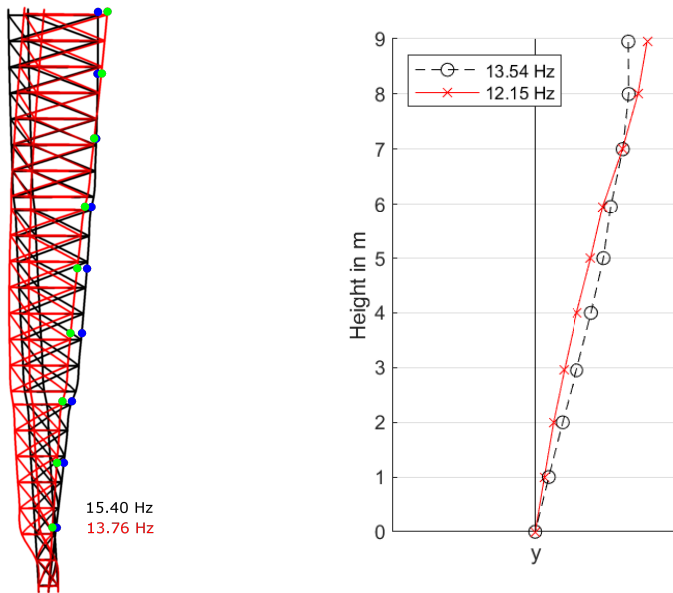


Figure 5.19: Comparison of first torsion mode for reference state and unsymmetrical damage close to the bottom. Simulation results shown in the left panel and experimental results shown in the right panel. Reference mode shape shown in black and damaged mode shape shown in red. Sensor amplitudes in the simulation result shown in blue and green, respectively.

In addition to the apparent stiffening effect shown in Figure 5.19, several more issues occur in the unsymmetrical damage scenario, which can not be readily addressed using the theory laid out in this thesis. First, the assumption of a radially homogeneous damage distribution along the structure is broken due to the asymmetry. While discrepancies between the damage distribution function and the actual damage are admissible to a small extent, the asymmetrical damage fundamentally changes the sensitivity of torsional modes. To tackle this issue, an alternative parameterisation of the finite element model would need to be employed. This alternative parameterisation would need to be designed to allow for individual stiffness alterations at each brace. Second, the operational modal analysis needs to be direction-sensitive to capture the asymmetry in the dynamic behaviour. This can be challenging for symmetrical structures like the girder mast investigated in this thesis. In such structures, the bending modes occur in orthogonal pairs at almost the same frequency [27]. Due to their spectral proximity, the closely spaced modes are hard to separate from each other, and therefore a clear identification is hard to achieve. To overcome this issue, a modal analysis technique more sophisticated than the basic Frequency-Domain Decomposition (FDD) would need to be employed. Third, in order to apply the mode shape error stated in Equation 3.3, the mode spaces of both the operational modal analysis domain and the finite element modal analysis domain need to be aligned with each other. This is not generally the case, since the bending mode shape pairs are usually not spatially orientated in the same direction in these domains. Additionally, in this specific structure, some torsional modes are spectrally close to bending modes, which can lead to a mixing of the modal subspace of the torsion and bending modes. This issue may be solved using rotations in the modal subspace in order to align the simulation mode space with the experimental one. The aforementioned challenges may eventually be overcome by employing advanced techniques not directly related to the model updating routine itself, which are therefore outside of the scope of this thesis.

The asymmetric damage scenario also poses challenges to the data-driven filter approach. The spectra for the reference and the damaged state are shown in Figures 5.20 and 5.21. While in the reference state, the residuals mostly stay close to the noise floor in the frequencies below 200 Hz, the damaged state exhibits a notable increase in the residual spectrum in Figure 5.21. From this observation, it can be concluded that damage detection is successful in this scenario. An interesting difference between the unsymmetrical damage scenario and the symmetrical one shown in Figure 5.15 is that the residuals only rise at some particular eigenfrequencies and not at all of them.

However, as the residual power metric shown in Figure 5.22 indicates, the residuals do not increase most at the damage location. In fact, the strongest




Figure 5.20: Measured spectrum and residual spectrum showing all 18 sensors for the reference state.




Figure 5.21: Measured spectrum and residual spectrum showing all 18 sensors for the damaged state.

increase of the metric is situated at the top of the structure, opposite to the damage location. When applying the relative residual power metric, as depicted in Figure 5.23, it becomes clear that the damage caused a major change in the structural dynamic behaviour of the structure. While the increase in the relative residual power is significant close the damage location, it is multiple times higher at the top of the structure. Similar to the findings made using modal analysis, these unexpected results may be caused by a combination of sensor placement and geometry. This may also explain why the residuals rise only at some of the eigenfrequencies and not at all of them. Hence, the results are coherent with the ones obtained by Wernitz [122], who found that damage detection was possible for the removal of a single brace, but localisation mostly failed.

As pointed out in this section, while the overall concept of the damage localisation approaches discussed in this thesis is sound, the application to particular

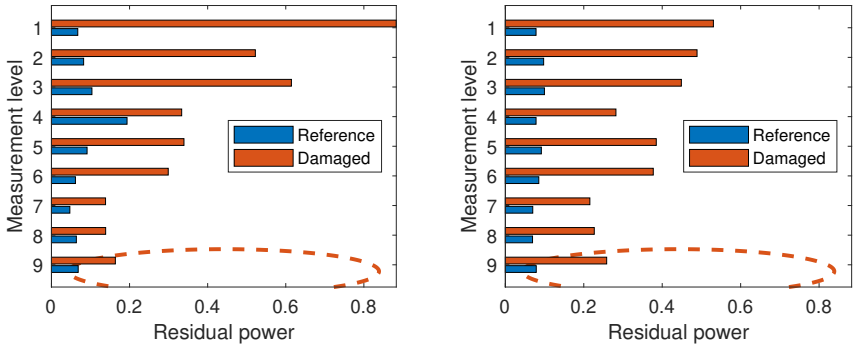


Figure 5.22: Comparison of the reference state and the damage state. Residual power is shown for x -direction (left panel) and y -direction (right panel). The true damage location is marked using a dashed ellipse.

structures and damage scenarios can be very challenging. As such, future improvements are needed to overcome these problems.




Figure 5.23: Damage localisation results obtained using the relative residual power metric. Values are shown for x -direction (left panel) and y -direction (right panel). The true damage location is marked using a dashed ellipse.

5.6 Comparison of numerical performance

The model-based and the data-driven localisation method follow vastly different concepts, so the numerical complexity also varies greatly.

To conduct localisation on a given set of data using the data-driven method, the measurement has to be processed using the identified estimation filter model. The filter model itself has to be identified only once for a single reference state, while it can be applied to an unlimited number of analysis states. While the identification of a 10 min data set takes several computing minutes, the application of the filters to obtain residuals can be computed in seconds. This means that vast amounts of measurement data can be processed quickly using this approach.

In contrast, the damage localisation using finite element model updating takes several hours of computation, which is dominated by the time allocated for the finite element solver. In principle, this amount of computing time has to be dedicated to each analysis data set. Significant reductions in computing time of the finite element could in principle be achieved by reducing the spatial resolution of the mesh, but this would negatively affect the obtainable localisation precision. However, a significant advantage of the Global Pattern Search optimisation algorithm is its deterministic nature. Since it samples at the exact same points in the design variable space when solving an optimisation problem for a second time, this can be used to cut down the computational cost. Since the modal parameters of a monitored structure usually do not vary significantly in consecutive data sets, a vast part of the optimisation trajectory during model updating is coincident with previous runs. By memorising previous results of

the finite element model evaluations, many modal analysis calculations could thus be omitted by following this approach.

Still, even a single evaluation of the finite element model takes approximately as long as the application of the finite impulse filters to a complete 10 min data set. As a conclusion, the data-driven approach based on finite impulse response filters is well-suited for continuous monitoring systems. Conversely, the finite element model updating method should only be applied to a limited number of selected data sets due to its high demand on computing power.

Chapter 6

Benefits and limitations

6.1 Numerical optimisation

As the ‘no free lunch’ theorems of optimisation [127] state, there can not be an optimisation algorithm, which performs well on every optimisation problem. That said, in certain applications such as model updating, the objective functions are mostly smooth. Thus, it is possible to tailor algorithms such that they exhibit robust and quick convergence under these circumstances. Considering the Global Pattern Search approach presented in this thesis, very good results were achieved when applied to the multi-objective model updating formulation. A prototype version of the MOGPS algorithm was also successfully applied to practical engineering problems by Berger et al. [11] as well as by Haldar et al. [42].

The presented algorithms however also have conceptual weaknesses. Due to the deterministic sample generation scheme, the method lacks robustness when applied to very noisy objective functions. Further, the sampling pattern formulated using Cartesian coordinates performs increasingly worse the higher the number of design variables of the problem becomes. A possible way to fix this issue is discussed in Section 7.2. Another weakness is the reliance on Pareto sorting, which is a complex and time-consuming process especially for high numbers of objectives and many objective function evaluations.

6.2 Damage localisation based on model updating

There are many ways to formulate the model updating problem on the basis of modal parameters. The formulation chosen here was especially guided by the strengths of the optimisation algorithm proposed in this thesis. Hence, a

low-dimensional parametrisation was used, which is conducive to the solution with the Global Pattern Search algorithm. The low dimensionality also leads to unique solutions and enhances the numerical stability. However, this formulation is only able to detect a single defect, while it will not yield proper results for multiple-defect cases. Additionally, a Gaussian shape is assumed for the distribution of the stiffness reduction, which does not always represent the actual damage distribution accurately. To tackle these issues, additional knowledge about the most probable defects would have to be incorporated into the method. A general weakness of the proposed model updating approach prevented the application of this approach to the validation scenario with an unsymmetrical damage: Model updating can only be applied if the damage distribution parameterisation is able to approximate the actual damage and if the vibration modes of the measurement and the simulation can be exactly matched to each other.

The multi-objective formulation chosen to include both mode shapes and eigenfrequencies has both positive and negative implications. On the positive side, the robustness of the damage localisation is increased compared to a single-objective formulation. This is the case, since all possible combinations of weighting ratios between mode shapes and eigenfrequencies are employed to find the best-fitting damage scenarios. This greatly increases the chance, that the actual damage distribution is among these solutions. On the negative side, the multi-objective scheme does not result in a single damage identification, but in multiple identifications. Thus, a sensible post-processing scheme needs to be employed to reduce the number of solutions and yield a manageable damage metric. Further, the verification using numerical simulation examples shows, that the sensitivity towards damage may increase for cases with small damage, when compared to a single-objective formulation.

Uncertainties exist in the measurement data, identification of modal parameters and the finite element model. However, in this thesis these uncertainties were ignored and the model updating problem was treated as a deterministic one. While proper results could be achieved overall, some identifications yielded erroneous damage distributions. Hence, uncertainty quantification of the modal parameters and a probabilistic model updating approach would enable damage localisation with a higher statistical confidence.

6.3 Data-driven damage localisation

An advantage of the localisation based on finite impulse response identification is its simplicity and the low number of parameters that have to be adjusted to obtain acceptable results. Energy residuals can be computed in real-time even on low-end computers, because the necessary signal convolution can be imple-

mented very efficiently using vector processing. This enables online monitoring and damage localisation using existing hardware.

However, data-driven output-only methods are only capable of giving a first indication about the damage location and their usefulness is highly dependent on the density and positioning of the sensor network. The validation scenario with an unsymmetrical damage showed that while the damage could be detected, the localisation failed, partly due to the complexity of the structure and the sensor setup. Moreover, the interpretation of energy residuals is not trivial for real structures, due to the ‘leakage’ of changes in the global behaviour to residuals at positions far away from the damage location. This problem is common to most localisation techniques based on residual energies. To some extent, sophisticated classifiers can overcome this limitation and increase the selectivity. Additionally, damage quantification based on residual values is not trivial, due to the nonlinear relation between the residuals and mechanical properties of the monitored structure. Farrar and Worden [36] even go as far as ruling out the possibility of quantification in an unsupervised output-only setting. Thus, model-based techniques would be well-suited to accurately pinpoint and also quantify damages indicated by data-driven methods during online monitoring.

6.4 Environmental and operating conditions

The methods and numerical examples discussed in this thesis are not applied to data sets containing variable environmental and operating conditions such as temperature, wind speed, actuator positions etc. However, they are designed with varying conditions in mind.

Due to the mathematical structure of finite impulse response filters, the mixing of several identified systems is achievable by linear interpolation. Even extrapolations are possible to a limited extent with basic linear methods. As a result, filter identifications can be carried out for a range of environmental and operating conditions and can later be interpolated to match the conditions encountered in the analysis time frame. Such interpolations can even be applied on a sample-by-sample basis, given that the changes in the environmental conditions are slow compared to the eigenfrequencies of the observed structure.

Regarding the model updating scheme, many environmental and operating conditions can be implemented directly into the physically-based finite element model. That way, the model can represent diverse influences such as softening and residual stresses due to temperature, stiffening due to body forces or varying boundary conditions. As these effects need to be modelled explicitly, this approach is evidently more labour-intensive than the interpolation scheme possible for the finite impulse response filters. However, the relative error metrics

used to formulate the objective function may help in reducing the modelling fidelity necessary to achieve satisfactory sensitivity towards damage.

Another aspect omitted in this thesis is the statistical post-processing of the identification results. Statistical analysis is especially required when dealing with varying conditions and large data sets acquired in long-term monitoring campaigns. Considering the data-driven method presented in this thesis, a detection threshold has to be derived from the reference data, which needs to be high enough to suppress false positives but also low enough to enable detection of minor damage cases. Similarly, thresholds also have to be applied to the results of the model updating approach. For this purpose, the objective values may also be considered in order to judge the quality of particular identification runs.

Chapter 7

Conclusion

7.1 Summary

This thesis presented a novel approach to damage localisation based on model updating using multi-objective optimisation and an innovative approach based on multiple-input finite impulse response filters. Both of the approaches contain elements which represent advancements over the state of research. These include a novel deterministic multi-objective optimisation algorithm, a multi-objective model updating approach using relative error metrics, as well as the residual-based damage localisation using multiple-input FIR filters. Additionally, a comparison between the two approaches was conducted using the same numerical and experimental data, and the advantages and disadvantages were elaborated. On the one hand, model updating allows for a relatively precise and rugged damage identification using modal parameters, which can readily be physically interpreted. To yield meaningful results, an appropriate numerical model of the structure has to be created and validated. On the other hand, the impulse response filters can be identified purely using measurement data, making it easily applicable. However, the localisation accuracy and robustness using the data-driven approach are worse than what is achievable with model updating.

In summary, a multi-stage approach seems to be advisable for practical monitoring applications, where a data-driven method is used to detect and coarsely locate possible damages and a method based on a mechanical model is used to pinpoint and quantify them afterwards. That way, efficient filtering techniques can be used for real-time analysis and the numerically expensive mechanical model has to be employed only when necessitated by indications of the filtering technique.

7.2 Future research topics

Regarding the Global Pattern Search optimisation algorithm, changes to the sampling pattern to increase the performance for high-dimensional problems should be explored. Several authors [5, 8, 116] have proposed generalised sampling patterns, which should be investigated in this regard. Another interesting area is the development of enhancements to tackle reliability-based and robust optimisation problems. Here, metamodel-based approaches similar to the concept of Efficient Global Optimisation [58] may be an interesting approach. Such algorithms can be used in robust model updating schemes which take into account the identification uncertainty of modal parameters and, therefore, enable quantifying the confidence of damage identifications.

Due to the grid-based sampling strategy of Global Pattern Search, the model updating scheme introduced in this thesis may even be suitable for efficient long-term monitoring. In this setting, the model updating has to be carried out for a large number of different sets of modal parameters extracted from continuous measurement data. As the monitored structure does not significantly change its structural dynamics over time, the algorithm samples along similar trajectories in the design space in each of the optimisation runs. The grid-based approach means that the points along these trajectories are exactly the same in each run, so that the results can be reused for subsequent runs. Hence, almost no new evaluations of the finite element model are necessary after processing a large number of data sets, which drastically reduces the numerical effort. A similar approach was already investigated for efficient α -level optimisation by Hübler and the author [51]. With regards to the validation scenario with an unsymmetrical damage, further research into dealing with the modal subspace of closely spaced modes in model updating is needed.

While taking into account changing environmental and operating conditions is relatively easy for the model updating method due to the underlying physically-based material model, this is more challenging for the filter-based method. A straightforward way to resolve this issue is to employ a piecewise linear interpolation of multiple filter identifications in the reference state [6, 132]. Due to the structure of finite impulse response filters, interpolation of the filter outputs is equivalent to interpolation of the filter coefficients themselves, which makes this approach also appealing from a numerical performance standpoint. To this end, further experiments focusing on distributed dynamic excitation due to wind interaction as well as other environmental conditions are going to be carried out. Investigations on the sensitivity to changing locations and types of excitation sources are also planned.

Bibliography

- [1] Nizar Faisal Alkayem, Maosen Cao, Yufeng Zhang, Mahmoud Bayat, and Zhongqing Su. Structural damage detection using finite element model updating with evolutionary algorithms: A survey. *Neural Computing and Applications*, 30(2):389–411, 2018. doi:10.1007/s00521-017-3284-1.
- [2] Randall J Allemang and David L Brown. A correlation coefficient for modal vector analysis. In *Proceedings of the 1st international modal analysis conference*, volume 1, pages 110–116. SEM Orlando, 1982.
- [3] Randall J Allemang and David L Brown. Experimental modal analysis and dynamic component synthesis. Volume 1. Summary of technical work. Technical report, University of Cincinnati, Structural Dynamics Research Laboratory, 1987.
- [4] Siu-Kui Au. Uncertainty law in ambient modal identification – Part II: Implication and field verification. *Mechanical Systems and Signal Processing*, 48(1):34 – 48, 2014. doi:10.1016/j.ymssp.2013.07.017.
- [5] Charles Audet and J. E. Dennis. Analysis of generalized pattern searches. *SIAM Journal on Optimization*, 13(3):889–903, 2002. doi:10.1137/S1052623400378742.
- [6] G. Iulia Bara, Jamal Daafouz, Frédéric Kratz, and José Ragot. Parameter-dependent state observer design for affine LPV systems. *International Journal of Control*, 74(16):1601–1611, 2001. doi:10.1080/00207170110086953.
- [7] Robert Barthorpe. *On Model- and Data-based Approaches to Structural Health Monitoring*. PhD thesis, University of Sheffield, Sheffield, UK, 2011.

- [8] Paul Belitz and Thomas Bewley. New horizons in sphere-packing theory, Part ii: Lattice-based derivative-free optimization via global surrogates. *Journal of Global Optimization*, 56(1):61–91, 2013. doi:10.1007/s10898-012-9866-7.
- [9] Richard Bellman. *Dynamic programming*. Princeton University Press, 1957.
- [10] Ricarda Berger, Marlene Bruns, Andreas Ehrmann, Ayan Haldar, Jan Häfele, Benedikt Hofmeister, Clemens Hübler, and Raimund Rolfes. EngiO – Object-oriented framework for engineering optimization. *Advances in Engineering Software*, 153:102959, 2021. doi:10.1016/j.advengsoft.2020.102959.
- [11] Ricarda Berger, Benedikt Hofmeister, Cristian G. Gebhardt, and Raimund Rolfes. A two-objective design optimisation approach for blending repairs of damaged compressor blisks. *Aerospace Science and Technology*, 105:106022, 2020. doi:10.1016/j.ast.2020.106022.
- [12] Josef Betten. *Finite Elemente für Ingenieure 1: Grundlagen, Matrixmethoden, Elastisches Kontinuum*, volume 83. Springer-Verlag, 2013.
- [13] N. Beume, C. M. Fonseca, M. Lopez-Ibanez, L. Paquete, and J. Vahrenhold. On the complexity of computing the hypervolume indicator. *IEEE Transactions on Evolutionary Computation*, 13(5):1075–1082, 2009. doi:10.1109/TEVC.2009.2015575.
- [14] Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [15] Marc Boeswald, Dennis Goege, Ulrich Fuellekrug, and Yves Govers. A review of experimental modal analysis methods with respect to their applicability to test data of large aircraft structures. In *Proc. of the International Conference on Noise and Vibration Engineering ISMA*, pages 2461–2482, 2006.
- [16] Paul T Boggs and Jon W Tolle. Sequential quadratic programming. *Acta numerica*, 4:1–51, 1995. doi:10.1017/S0962492900002518.
- [17] Rune Brincker and Carlos Ventura. *Introduction to operational modal analysis*. John Wiley & Sons, 2015.

-
- [18] Rune Brincker, Lingmi Zhang, and Palle Andersen. Modal identification from ambient responses using frequency domain decomposition. *Proceedings of the International Modal Analysis Conference - IMAC*, 1, 01 2000.
- [19] Rune Brincker, Lingmi Zhang, and Palle Andersen. Modal identification of output-only systems using frequency domain decomposition. *Smart Materials and Structures*, 10(3):441–445, 2001. doi:10.1088/0964-1726/10/3/303.
- [20] Dimo Brockhoff, Thanh-Do Tran, and Nikolaus Hansen. Benchmarking numerical multiobjective optimizers revisited. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation, GECCO '15*, pages 639–646. ACM, 2015. doi:10.1145/2739480.2754777.
- [21] Marlene Bruns, Benedikt Hofmeister, Tanja Griebmann, and Raimund Rolfes. Comparative study of parameterizations for damage localization with finite element model updating. In *Proceedings of the 29th European Safety and Reliability Conference (ESREL)*, pages 1125–1132, Singapore, 2019. Research Publishing Services. doi:10.3850/978-981-11-2724-3_0713-cd.
- [22] Marlene Bruns, Benedikt Hofmeister, Clemens Hübler, and Raimund Rolfes. Damage Localization Via Model Updating Using a Damage Distribution Function. In *Structural Health Monitoring 2019*, pages 1–10, Lancaster, PA, 2019. DEStech Publications, Inc. doi:10.12783/shm2019/32202.
- [23] Bastien Chapuis. *Introduction to Structural Health Monitoring*. Springer International Publishing, Cham, 2018. doi:10.1007/978-3-319-69233-3_1.
- [24] Tianshi Chen, Henrik Ohlsson, and Lennart Ljung. On the estimation of transfer functions, regularizations and gaussian processes—revisited. *Automatica*, 48(8):1525–1535, 2012. doi:10.1016/j.automatica.2012.05.026.
- [25] Simon Chesné and Arnaud Deraemaeker. Damage localization using transmissibility functions: A critical review. *Mechanical Systems and Signal Processing*, 38:569–584, 2013. doi:10.1016/j.ymsp.2013.01.020.
- [26] Carlos A Coello Coello. Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: A survey of the state of the art. *Computer Methods in Applied Mechanics and Engineering*, 191(11):1245 – 1287, 2002. doi:10.1016/S0045-7825(01)00323-1.

- [27] W. D'AMBROGIO and A. FREGOLENT. Higher-order mac for the correlation of close and multiple modes. *Mechanical Systems and Signal Processing*, 17(3):599–610, 2003. doi:10.1006/mssp.2002.1468.
- [28] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. doi:10.1109/4235.996017.
- [29] Kalyanmoy Deb. *Multiobjective optimization using evolutionary algorithms*. Wiley-Interscience series in systems and optimization. Wiley, Chichester, 2002.
- [30] Kalyanmoy Deb. *Multi-objective Optimisation Using Evolutionary Algorithms: An Introduction*, chapter 1, pages 3–34. Springer London, London, 2011. doi:10.1007/978-0-85729-652-8_1.
- [31] Gaurav Dhiman and Vijay Kumar. Spotted hyena optimizer: A novel bio-inspired based metaheuristic technique for engineering applications. *Advances in Engineering Software*, 114:48–70, 2017.
- [32] Paulo SR Diniz. *Adaptive filtering*. Springer, 1997.
- [33] E. Dolan, R. Lewis, and V. Torczon. On the local convergence of pattern search. *SIAM Journal on Optimization*, 14(2):567–583, 2003.
- [34] Zhongdong Duan, Guirong Yan, Jinping Ou, and B.F. Spencer. Damage localization in ambient vibration by constructing proportional flexibility matrix. *Journal of Sound and Vibration*, 284(1):455 – 466, 2005. doi:10.1016/j.jsv.2004.06.046.
- [35] Yu.G. Evtushenko and M.A. Posypkin. A deterministic algorithm for global multi-objective optimization. *Optimization Methods and Software*, 29(5):1005–1019, 2014. doi:10.1080/10556788.2013.854357.
- [36] Charles R Farrar and Keith Worden. *Structural health monitoring: A machine learning perspective*. John Wiley & Sons, 2012.
- [37] M. I. Friswell and J. E. Mottershead. *Finite Element Model Updating in Structural Dynamics*, volume 38 of *Solid Mechanics and its Applications*. Springer Netherlands, 1995. doi:10.1007/978-94-015-8508-8.
- [38] MICHAEL L. Fugate, Hoon Sohn, and CHARLES R. Farrar. Vibration-based damage detection using statistical process control. *Mechanical Systems and Signal Processing*, 15(4):707–721, 2001. doi:10.1006/mssp.2000.1323.

-
- [39] Zong Woo Geem, Joong Hoon Kim, and G. V. Loganathan. A New Heuristic Optimization Algorithm: Harmony Search. *SIMULATION*, 76(2):60–68, 2016. doi:10.1177/003754970107600201.
- [40] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1st edition, 1989.
- [41] David E Goldberg, Jon Richardson, et al. Genetic algorithms with sharing for multimodal function optimization. In *Genetic algorithms and their applications: Proceedings of the Second International Conference on Genetic Algorithms*, pages 41–49. Hillsdale, NJ: Lawrence Erlbaum, 1987.
- [42] A. Haldar, E. Jansen, B. Hofmeister, M. Bruns, and R. Rolfes. Analysis of novel morphing trailing edge flap actuated by multistable laminates. *AIAA Journal*, 58(7):3149–3158, 2020. doi:10.2514/1.J058870.
- [43] David M. Himmelblau. *Applied Nonlinear Programming*. McGraw-Hill, 1972.
- [44] Benedikt Hofmeister, Marlene Bruns, Clemens Hübler, and R. Rolfes. Multi-objective global pattern search: Effective numerical optimisation in structural dynamics. *Institutional Repository of Leibniz Universität Hannover*, 2021. doi:10.15488/11024.
- [45] Benedikt Hofmeister, Marlene Bruns, and Raimund Rolfes. Finite element model updating using deterministic optimisation: A global pattern search approach. *Engineering Structures*, 195:373–381, 2019. doi:10.1016/j.engstruct.2019.05.047.
- [46] Benedikt Hofmeister, Clemens Jonscher, Clemens Hübler, and Raimund Rolfes. Damage localisation by residual energy from multiple-input finite impulse response prognosis. In *European Workshop on Structural Health Monitoring*, pages 711–719. Springer International Publishing, 2021. doi:10.1007/978-3-030-64908-1_66.
- [47] Benedikt Hofmeister, Stefan Wernitz, Tanja Griebmann, Clemens Hübler, and Raimund Rolfes. Data-driven vibration prognosis using multiple-input finite impulse response filters and application to railway-induced vibration of timber buildings. Preprint, 2022. doi:10.15488/12855.
- [48] Robert Hooke and T. A. Jeeves. "direct search" solution of numerical and statistical problems. *J. ACM*, 8(2):212–229, 1961. doi:10.1145/321062.321069.

- [49] David L Hunt. Application of an enhanced coordinate modal assurance criterion. In *10th International modal analysis conference*, volume 1, pages 66–71, 1992.
- [50] Moritz W. Häckell and Raimund Rolfes. Monitoring a 5MW offshore wind energy converter – Condition parameters and triangulation based extraction of modal parameters. *Mechanical Systems and Signal Processing*, 40(1):322–343, 2013. doi:10.1016/j.ymsp.2013.04.004.
- [51] Clemens Hübler and Benedikt Hofmeister. Efficient and user-friendly α -level optimisation for application-orientated fuzzy structural analyses. *Engineering Structures*, 247, 2021. doi:10.1016/j.engstruct.2021.113172.
- [52] Clemens Hübler, Benedikt Hofmeister, Stefan Wernitz, and Raimund Rolfes. Validierung von daten- und modellbasierten Methoden zur Schadenslokalisierung. *Bautechnik*, 2022. doi:10.1002/bate.202200015.
- [53] Mahmoud M. Jahjough. *A modified adaptive harmony search algorithm approach on structural identification and damage detection*. PhD thesis, Gottfried Wilhelm Leibniz Universität Hannover, Hannover, 2016.
- [54] Mahmoud M Jahjough and Udo Nackenhorst. A modified harmony search approach on structural identification and damage detection of wind turbine supporting structures. *Journal of Vibroengineering*, 18(1), 2016. URL: <https://www.jvejournal.com/article/16694>.
- [55] A. Jiang, H. K. Kwan, Y. Zhu, X. Liu, N. Xu, and X. Yao. Peak-error-constrained sparse FIR filter design using iterative L1 optimization. In *2016 24th European Signal Processing Conference (EUSIPCO)*, pages 180–184, 2016. doi:10.1109/EUSIPCO.2016.7760234.
- [56] Seung-Seop Jin, Soojin Cho, Hyung-Jo Jung, Jong-Jae Lee, and Chung-Bang Yun. A new multi-objective approach to finite element model updating. *Journal of Sound and Vibration*, 333(11):2323–2338, 2014. doi:10.1016/j.jsv.2014.01.015.
- [57] D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181, 1993. doi:10.1007/BF00941892.
- [58] Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998. doi:10.1023/A:1008306431147.

-
- [59] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95 – International Conference on Neural Networks*, volume 4, pages 1942–1948, 1995. doi:10.1109/ICNN.1995.488968.
- [60] Gaëtan Kerschen, Keith Worden, Alexander F. Vakakis, and Jean-Claude Golinval. Past, present and future of nonlinear system identification in structural dynamics. *Mechanical Systems and Signal Processing*, 20(3):505 – 592, 2006. doi:10.1016/j.ymsp.2005.04.008.
- [61] Hamed Haddad Khodaparast, John E. Mottershead, and Kenneth J. Badcock. Interval model updating with irreducible uncertainty using the kriging predictor. *Mechanical Systems and Signal Processing*, 25(4):1204–1226, 2011. doi:10.1016/j.ymsp.2010.10.009.
- [62] Gyeong-Ho Kim and Youn-Sik Park. An improved updating parameter selection method and finite element model update using multiobjective optimisation technique. *Mechanical Systems and Signal Processing*, 18(1):59–78, January 2004. doi:10.1016/S0888-3270(03)00042-6.
- [63] Mifa Kim, Tomoyuki Hiroyasu, Mitsunori Miki, and Shinya Watanabe. SPEA2+: Improving the performance of the strength pareto evolutionary algorithm 2. In *International Conference on Parallel Problem Solving from Nature*, pages 742–751. Springer, 2004. doi:10.1007/978-3-540-30217-9_75.
- [64] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science (New York, N.Y.)*, 220(4598):671–680, 1983. doi:10.1126/science.220.4598.671.
- [65] Masaru Kitahara, Sifeng Bi, Matteo Broggi, and Michael Beer. Non-parametric bayesian stochastic model updating with hybrid uncertainties. *Mechanical Systems and Signal Processing*, 163:108195, 2022. doi:10.1016/j.ymsp.2021.108195.
- [66] T. Krink, J.S. Vesterstrom, and J. Riget. Particle swarm optimisation with spatial particle extension. In *Proceedings of the 2002 Congress on Evolutionary Computation*, volume 2, pages 1474–1479, 2002. doi:10.1109/CEC.2002.1004460.
- [67] H. T. Kung, F. Luccio, and F. P. Preparata. On finding the maxima of a set of vectors. *J. ACM*, 22(4):469–476, 1975. doi:10.1145/321906.321910.

- [68] Armin Lenzen and Max Vollmering. An output-only damage identification method based on H_∞ theory and state projection estimation error (SP2E). *Structural Control and Health Monitoring*, 24:1–11, 2017. doi:10.1002/stc.2003.
- [69] R.I. Levin and N.A.J. Lieven. Dynamic finite element model updating using neural networks. *Journal of Sound and Vibration*, 210(5):593–607, 1998. doi:10.1006/jsvi.1997.1364.
- [70] R.I. Levin and N.A.J. Lieven. Dynamic finite element model updating using simulated annealing and genetic algorithms. *Mechanical Systems and Signal Processing*, 12(1):91–120, 1998. doi:10.1006/mssp.1996.0136.
- [71] Michael Link. Updating of analytical models—review of numerical procedures and application aspects. In *Proc., Structural Dynamics Forum SD2000*, pages 193–223. Research Studies Press, Baldock, UK, 1999.
- [72] Lennart Ljung. System identification toolbox reference. Reference, The MathWorks, Inc, Natick, MA, 2017.
- [73] J Maeck, M Abdel Wahab, B Peeters, G De Roeck, J De Visscher, W.P De Wilde, J.-M Ndambi, and J Vantomme. Damage identification in reinforced concrete structures by dynamic stiffness determination. *Engineering Structures*, 22(10):1339–1349, 2000. doi:10.1016/S0141-0296(99)00074-7.
- [74] N.M.M. Maia, J.M.M. Silva, and A.M.R. Ribeiro. The transmissibility concept in multi-degree-of-freedom systems. *Mechanical Systems and Signal Processing*, 15(1):129 – 137, 2001. doi:10.1006/mssp.2000.1356.
- [75] Graeme Manson, Keith Worden, and David Allman. Experimental validation of a structural health monitoring methodology: Part III. Damage location on an aircraft wing. *Journal of Sound and Vibration*, 259(2):365–385, 2003. doi:10.1006/jsvi.2002.5169.
- [76] R. T. Marler and J. S. Arora. Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26(6):369–395, 2004. doi:10.1007/s00158-003-0368-6.
- [77] A. Mehrnia and A. N. Willson. On optimal IFIR filter design. In *2004 IEEE International Symposium on Circuits and Systems (IEEE Cat. No.04CH37512)*, volume 3, pages III–133, 2004. doi:10.1109/ISCAS.2004.1328701.

- [78] Seyedali Mirjalili and Andrew Lewis. The whale optimization algorithm. *Advances in engineering software*, 95:51–67, 2016. doi:10.1016/j.advengsoft.2016.01.008.
- [79] Seyedali Mirjalili, Seyed Mohammad Mirjalili, and Andrew Lewis. Grey wolf optimizer. *Advances in engineering software*, 69:46–61, 2014. doi:10.1016/j.advengsoft.2013.12.007.
- [80] Marcin Molga and Czesław Smutnicki. Test functions for optimization needs. Technical report, Department of Automation, Mechatronics and Control Systems, Polytechnika Wroclawska, 2005. URL: <http://new.zsd.iiar.pwr.wroc.pl/files/docs/functions.pdf>.
- [81] H.L. Moore. Cours d'Économie Politique. By Vilfredo Pareto, Professeur à l'Université de Lausanne. *The Annals of the American Academy of Political and Social Science*, 9(3):128–131, 1897. doi:10.1177/000271629700900314.
- [82] J. E. Mottershead and M. I. Friswell. Model updating in structural dynamics: A survey. *Journal of Sound and Vibration*, 167(2):347–375, 1993. doi:10.1006/jsvi.1993.1340.
- [83] Javier Naranjo-Pérez, María Infantes, Javier Fernando Jiménez-Alonso, and Andrés Sáez. A collaborative machine learning-optimization algorithm to improve the finite element model updating of civil engineering structures. *Engineering Structures*, 225:111327, 2020. doi:10.1016/j.engstruct.2020.111327.
- [84] Y. Neuvo, Dong Cheng-Yu, and S. Mitra. Interpolated finite impulse response filters. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32(3):563–570, 1984. doi:10.1109/TASSP.1984.1164348.
- [85] Kazuhiko Ozeki and Tetsuo Umeda. An adaptive filtering algorithm using an orthogonal projection to an affine subspace and its properties. *Electronics and Communications in Japan (Part I: Communications)*, 67(5):19–27, 1984. doi:10.1002/ecja.4400670503.
- [86] A.K. Pandey, M. Biswas, and M.M. Samman. Damage detection from changes in curvature mode shapes. *Journal of Sound and Vibration*, 145(2):321–332, 1991. doi:10.1016/0022-460X(91)90595-B.
- [87] Mario Paz and Young Hoon Kim. *Structural dynamics*. Springer, 2019.

- [88] Bart Peeters and Guido De Roeck. Reference-based stochastic subspace identification for output-only modal analysis. *Mechanical Systems and Signal Processing*, 13(6):855–878, 1999. doi:10.1006/mssp.1999.1249.
- [89] Nikolai Penner. *Monitoring ambient angeregter baodynamischer Systeme durch mehrschichtige Perzeptren*. PhD thesis, Gottfried Wilhelm Leibniz Universität Hannover, Hannover, 2021.
- [90] Ricardo Perera and Antonio Ruiz. A multistage FE updating procedure for damage identification in large-scale structures based on multiobjective evolutionary optimization. *Mechanical Systems and Signal Processing*, 22(4):970–991, 2008. doi:10.1016/j.ymsp.2007.10.004.
- [91] Ricardo Perera, Antonio Ruiz, and Carlos Manzano. An evolutionary multiobjective framework for structural damage localization and quantification. *Engineering Structures*, 29(10):2540 – 2550, 2007. doi:10.1016/j.engstruct.2007.01.003.
- [92] E. Peter Carden and James M.W. Brownjohn. ARMA modelled time-series classification for structural health monitoring of civil infrastructure. *Mechanical Systems and Signal Processing*, 22(2):295–314, 2008. doi:10.1016/j.ymsp.2007.07.003.
- [93] Carlo Poloni, Andrea Giurgevich, Luka Onesti, and Valentino Pediroda. Hybridization of a multi-objective genetic algorithm, a neural network and a classical optimizer for a complex design problem in fluid dynamics. *Computer Methods in Applied Mechanics and Engineering*, 186(2-4):403–420, 2000. doi:10.1016/S0045-7825(99)00394-1.
- [94] R. E. Powell and W. Seering. Multichannel Structural Inverse Filtering. *Journal of Vibration, Acoustics, Stress, and Reliability in Design*, 106(1):22–28, 01 1984. doi:10.1115/1.3269147.
- [95] Ashok Prajapati, James Bechtel, and Subramaniam Ganesan. Condition based maintenance: A survey. *Journal of Quality in Maintenance Engineering*, 22:384–400, 2012. doi:10.1108/13552511211281552.
- [96] John G Proakis. *Digital signal processing: principles algorithms and applications*. Pearson Education India, 2001.
- [97] Shiqiang Qin, Juntao Kang, and Qiuping Wang. Operational modal analysis based on subspace algorithm with an improved stabilization diagram method. *Shock and Vibration*, 2016, 2016. doi:10.1155/2016/7598965.

-
- [98] Carlo Rainieri and Giovanni Fabbrocino. *Operational modal analysis of civil engineering structures*. Springer, 2014.
- [99] R Rao. Rao algorithms: Three metaphor-less simple algorithms for solving optimization problems. *International Journal of Industrial Engineering Computations*, 11(1):107–130, 2020. doi:10.5267/j.ijiec.2019.6.002.
- [100] N. Riquelme, C. Von Lüken, and B. Baran. Performance metrics in multi-objective optimization. In *2015 Latin American Computing Conference (CLEI)*, pages 1–11, 2015. doi:10.1109/CLEI.2015.7360024.
- [101] H.H. Rosenbrock. An automatic method for finding the greatest or least value of a function. *The Computer Journal*, 3(3):175–184, 1960. doi:10.1093/comjnl/3.3.175.
- [102] P. C. Roy, K. Deb, and M. M. Islam. An efficient nondominated sorting algorithm for large number of fronts. *IEEE Transactions on Cybernetics*, 49(3):859–869, 2019. doi:10.1109/TCYB.2017.2789158.
- [103] T. K. Roy and M. Morshed. Performance analysis of low pass FIR filters design using Kaiser, Gaussian and Tukey window function methods. In *2013 2nd International Conference on Advances in Electrical Engineering (ICAEE)*, pages 1–6, 2013. doi:10.1109/ICAEE.2013.6750294.
- [104] Karsten Schröder, Cristian Guillermo Gebhardt, and Raimund Rolfes. A two-step approach to damage localization at supporting structures of offshore wind turbines. *Structural Health Monitoring*, 17(5):1313–1330, 2018. doi:10.1177/1475921717741083.
- [105] Hans-Paul Schwefel. *Numerical optimization of computer models*. Wiley, Chichester, 1981.
- [106] Ellen Simoen, Guido De Roeck, and Geert Lombaert. Dealing with uncertainty in model updating for damage assessment: A review. *Mechanical Systems and Signal Processing*, 56-57:123–149, 2015. doi:10.1016/j.ymsp.2014.11.001.
- [107] Jesse D. Sipple and Masoud Sanayei. Finite element model updating using frequency response functions and numerical sensitivities. *Structural Control and Health Monitoring*, 21(5):784–802, 2014. doi:10.1002/stc.1601.
- [108] Julius Orion Smith. *Introduction to digital filters: with audio applications*, volume 2. Julius Smith, 2007.

- [109] Zhongqing Su and Lin Ye. *Identification of damage using Lamb waves: From fundamentals to applications*. Springer, 2009.
- [110] Kenneth Sörensen. Metaheuristics—the metaphor exposed. *International Transactions in Operational Research*, 22(1):3–18, 2015. doi:10.1111/itor.12001.
- [111] Kay Chen Tan, Tong Heng Lee, and Eik Fun Khor. Evolutionary algorithms for multi-objective optimization: Performance assessments and comparisons. *Artificial intelligence review*, 17(4):251–290, 2002. doi:10.1023/A:1015516501242.
- [112] Dmitri Tcherniak and Lasse L Mølgaard. Active vibration-based structural health monitoring system for wind turbine blade: Demonstration on an operating vestas v27 wind turbine. *Structural Health Monitoring*, 16(5):536–550, 2017. doi:10.1177/1475921717722725.
- [113] Anne Teughels, Guido De Roeck, and Johan A.K. Suykens. Global optimization by coupled local minimizers and its application to FE model updating. *Computers & Structures*, 81(24-25):2337–2351, 2003. doi:10.1016/S0045-7949(03)00313-4.
- [114] Anne Teughels, Johan Maeck, and Guido De Roeck. Damage assessment by FE model updating using damage functions. *Computers & Structures*, 80(25):1869–1879, 2002. doi:10.1016/S0045-7949(02)00217-1.
- [115] Andreï Nikolaevich Tikhonov, Aleksandr S Leonov, and Anatolij Grigorevic Yagola. *Nonlinear ill-posed problems*. Springer Netherlands, 1998.
- [116] Virginia Torczon. On the convergence of pattern search algorithms. *SIAM Journal on Optimization*, 7(1):1–25, 1997. doi:10.1137/S1052623493250780.
- [117] Stavroula Tsiapoki, Moritz W Häckell, Tanja Griebmann, and Raimund Rolfes. Damage and ice detection on wind turbine rotor blades using a three-tier modular structural health monitoring framework. *Structural Health Monitoring*, 17(5):1289–1312, 2018. doi:10.1177/1475921717732730.
- [118] P. van Overschee and B.L. de Moor. *Subspace Identification for Linear Systems: Theory — Implementation — Applications*. Springer US, 2012.
- [119] David Allen Van Veldhuizen. *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*. PhD thesis, Air Force Institute of Technology, Wright Patterson AFB, OH, USA, 1999.

-
- [120] Tong Wang, Lingmi Zhang, and Yukio Tamura. An operational modal analysis method in frequency and spatial domain. *Earthquake Engineering and Engineering Vibration*, 4:295–300, 2005. doi:10.1007/s11803-005-0012-0.
- [121] Wout Weijtjens, Gert De Sitter, Christof Devriendt, and Patrick Guillaume. Operational modal parameter estimation of mimo systems using transmissibility functions. *Automatica*, 50(2):559 – 564, 2014. doi:10.1016/j.automatica.2013.11.021.
- [122] Stefan Wernitz. *Damage Localization in Data-Driven Vibration-Based Structural Health Monitoring Using Linear Quadratic Estimation Theory*. PhD thesis, Gottfried Wilhelm Leibniz Universität Hannover, Hannover, 2022.
- [123] Stefan Wernitz, Eleni Chatzi, Benedikt Hofmeister, Marlene Wolniak, Wanzhou Shen, and Raimund Rolfes. On noise covariance estimation for kalman filter-based damage localization. *Mechanical Systems and Signal Processing*, 170:108808, 2022. doi:10.1016/j.ymsp.2022.108808.
- [124] Stefan Wernitz, Benedikt Hofmeister, Clemens Jonscher, Tanja Griebmann, and Raimund Rolfes. A new open-database benchmark structure for vibration-based structural health monitoring. *Structural Control and Health Monitoring*, 29(11):e3077, 2022. doi:10.1002/stc.3077.
- [125] Stefan Wernitz, Dorian Pache, Tanja Griebmann, and R. Rolfes. Damage localization with SP2E under changing conditions. In *Structural Health Monitoring 2019*, pages 1–9, 2019. doi:10.12783/shm2019/32505.
- [126] Marlene Wolniak, Benedikt Hofmeister, Clemens Jonscher, Matthias Fankhänel, Ansgar Loose, Clemens Hübler, and Raimund Rolfes. Validation of an FE model updating procedure for damage assessment using a modular laboratory experiment with a reversible damage mechanism. Preprint, 2022. doi:10.15488/12787.
- [127] David H Wolpert and William G Macready. No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997. doi:10.1109/4235.585893.
- [128] Stephen J. Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, Jun 2015. doi:10.1007/s10107-015-0892-3.
- [129] Cheng-Shing Wu and An-Yeu Wu. A novel cost-effective multi-path adaptive interpolated FIR (IFIR)-based echo canceller. In *2002 IEEE International Symposium on Circuits and Systems. Proceedings (Cat. No.*

- 02CH37353*), volume 5, pages V–V. IEEE, 2002. doi:10.1109/ISCAS.2002.1010738.
- [130] Xin-She Yang. *Nature-inspired metaheuristic algorithms*. Luniver Press, Frome, 2. Ed. edition, 2010.
- [131] Ruigen Yao and Shamim N. Pakzad. Autoregressive statistical pattern recognition algorithms for damage detection in civil structures. *Mechanical Systems and Signal Processing*, 31:355–368, 2012. doi:10.1016/j.ymsp.2012.02.014.
- [132] Qinghua Zhang and Lennart Ljung. From structurally independent local LTI models to LPV model. *Automatica*, 84:232–235, 2017. doi:10.1016/j.automatica.2017.06.006.
- [133] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003. doi:10.1109/TEVC.2003.810758.
- [134] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evol. Comput.*, 8(2):173–195, 2000. doi:10.1162/106365600568202.

Mitteilungen des Instituts für Statik und Dynamik der Leibniz Universität Hannover

1	R. Rolfes/ C. Hühne	Eröffnungskolloquium (Tagungsband)	2005
2	H. Rothert/ M. Kaliske/ L. Nasdala	Entwicklung von Materialmodellen zur Alterung von Elastomerwerkstoffen unter besonderer Berücksichtigung des Sauerstoffeinflusses (DFG-Abschlußbericht)	2005
3	L. Nasdala	Simulation von Materialinelastizitäten bei Nano-, Mikro- und Makrostrukturen – Stabilitätsprobleme, Schädigungs- und Alterungsprozesse bei Kohlenstoffnanoröhren und Elastomerwerkstoffen (Habilitationsschrift)	2005
4	C. Hühne	Robuster Entwurf beulgefährdeter, unversteifter Kreiszyllinderschalen aus Faserverbundwerkstoff (Dissertationsschrift)	2006
5	L. Nasdala/ K.-U. Schröder	Finite Element Applications in Structural Analysis (Skript zur Hörsaalübung)	2006
6		Klausuraufgabensammlung, 4. Auflage	2007
7	R. Rolfes/ W.-J. Ge- rasch/ D. Rotert	Vorlesung Tragwerksdynamik	2007
8	K.-H. Elmer/ K. Betke/ Th. Neumann	Standardverfahren zur Ermittlung und Bewertung der Belastung der Meeresumwelt durch die Schallmission von Offshore-Windenergieanlagen (Abschlussbericht zum BMU-Forschungsvorhaben)	2007
9	K.-U. Schröder	Zur nichtlinearen Berechnung von Stahlbeton- und Verbundbauteilen (Dissertationsschrift)	2007
10	G. Ernst	Multiscale Analysis of Textile Composites – Stiffness and Strength (Dissertationsschrift)	2008
11	G. Haake	Systemidentifikation mit Autoregressiven Modellen und Validierung numerischer Strukturmodelle bei Offshore-Windenergieanlagen (Dissertationsschrift)	2010
12	T. Griebmann	Dynamisches Tragverhalten von Stahlbetonbiegebalcken im Experiment und in der Simulation (Dissertationsschrift)	2011
13	S. Zerbst	Global Approach for Early Damage Detection on Rotor Blades of Wind Energy Converters (Dissertationsschrift)	2011
14	N. Wiczorek	Semiaktive Schwingungsdämpfung leichter Fußgängerbrückenkonstruktionen (Dissertationsschrift)	2011
15	B. Krieges- mann	Probabilistic Design of Thin-Walled Fiber Composite Structures (Dissertationsschrift)	2012
16	J. Reetz	Schadensdiagnose an Tragstrukturen von Windenergieanlagen mit der Multiparameter-Eigenwertproblem-Methode (Dissertationsschrift)	2012

17	R. Rolfes/ D. Rotert	Vorlesung Baustatik	2012
18	R. Rolfes/ D. Rotert	Vorlesung Stabtragwerke	2012
19	H. Krüger	Ein physikalisch basiertes Ermüdungsschädigungsmodell zur Degradationsberechnung von Faser-Kunststoff-Verbunden (Dissertationsschrift)	2012
20	S. Czichon	Multi scale Failure Analysis of Fibre Reinforced Polymers with production induced Porosity Defects (Dissertationsschrift)	2013
21	T. Pahn	Inverse Load Calculation for Offshore Wind Turbines (Dissertationsschrift)	2013
22	M. Fricke/ B. Neddermann/ A. Lübben/ J. Gabriel	Realistische Hydroschallszenarien auf der Basis von Prognosemodellen und Monitoring für den Bau von Offshore-Windparks in der deutschen Nordsee („HyproWind“) (Abschlussbericht zum BMU-Forschungsvorhaben)	2014
23	M. Vogler	Anisotropic Material Models for Fiber Reinforced Polymers (Dissertationsschrift)	2014
24	M. Fricke	Ein physikalisch basiertes Gesamtmodell für hydroakustische Immissionsprognosen bei Offshore-Pfahlrammungen (Dissertationsschrift)	2015
25	M. Häckell	A holistic evaluation concept for long-term structural health monitoring (Dissertationsschrift)	2015
26	J. Rustemeier	Optimierung von Blasenschleiern zur Minderung von Unterwasser-Rammschall (Dissertationsschrift)	2016
27	S. Hühne	A two-way loose coupling procedure for buckling and damage analysis of composite structures (Dissertationsschrift)	2016
28	A. Meurer	Filtering Geometric Imperfection Patterns for Analysis and Design of Composite Shell Structures (Dissertationsschrift)	2017
29	A. Dean	Material Modeling of Short Fiber Reinforced Polymeric Composites: Theory, Numerical Aspects, and Applications (Dissertationsschrift)	2017
30	M. Bishara	Compressive Failure of Polymer Composites Including Fiber Kinking and Interaction of Failure Mechanisms (Dissertationsschrift)	2017
31	K. Schröder	Advanced Model Updating Strategies for Structural Dynamic Systems (Dissertationsschrift)	2018
32	S. Tsiapoki	Transmissibility-Based Monitoring and Combination of Damage Detection Decisions within a Holistic Structural Health Monitoring Framework (Dissertationsschrift)	2018
33	S. R. Nabavi	Failure analysis of polycrystalline silicon-based photovoltaic modules considering the effects of residual stresses and mechanical loading (Dissertationsschrift)	2018

34	S. Scheffler	Ein neuer Modellierungsansatz zur systematischen numerischen Untersuchung des Versagensverhaltens von Verbindungen in FVK (Dissertationsschrift)	2018
35	J. Häfele	A numerically efficient and holistic approach to design optimization of offshore wind turbine jacket sub-structures (Dissertationsschrift)	2019
36	C. Hübler	Efficient probabilistic analysis of offshore wind turbines based on time-domain simulations (Dissertationsschrift)	2019
37	M. Akterskaia	Global-local progressive failure analysis of composite panels including skin-stringer debonding and intralaminar damage (Dissertationsschrift)	2019
38	C. Gebhardt	Robust computational procedures for the nonlinear dynamic analysis of beam and shell structures (Habilitationsschrift)	2020
39	A. Haldar	Multistable morphing structures using variable stiffness laminates (Dissertationsschrift)	2020
40	R. Unger	Multi-scale constitutive modelling of nanoparticle/epoxy nanocomposites: Molecular simulation-based methods and experimental validation (Dissertationsschrift)	2020
41	J. Fankhänel	A Multi-Scale Framework for Nanocomposites including Interphase and Agglomeration Effects (Dissertationsschrift)	2020
42	N. Penner	Monitoring ambient angeregter baodynamischer Systeme durch mehrschichtige Perzeptren (Dissertationsschrift)	2021
43	C. Gerendt	A finite element-based continuum damage model for mechanical joints in fiber metal laminates under static and fatigue loading: Theory and experimental validation (Dissertationsschrift)	2022
44	M. Brod	Damage prediction of unidirectional fiber composites under cyclic loading with different amplitudes (Dissertationsschrift)	2022
45	R. Berger	Multi-Objective Structural Optimization of Repairs of Blisk Blades (Dissertationsschrift)	2022
46	G. Balokas	Metamodel-based uncertainty quantification for the mechanical behavior of braided composites (Dissertationsschrift)	2022
47	S. Wernitz	Damage Localization in Data-Driven Vibration-Based Structural Health Monitoring Using Linear Quadratic Estimation Theory (Dissertationsschrift)	2022
48	B. Daum	On the computational analysis of microbuckling via mesoscale approaches (Habilitationsschrift)	2022
49	N. Safdar	A Stochastic Failure Investigation of Composites under Combined Compression-Shear Loads (Dissertationsschrift)	2022
50	S. van den Broek	Tailoring Structures Using Stochastic Variations of Structural Parameters (Dissertationsschrift)	2023

- 51 B. Hofmeister Vibration-based damage localisation: Impulse response identification and model updating methods (Dissertationsschrift) 2023