

Special Topics & Experimental Techniques in Structural Dynamics, Volume 5

Dario Di Maio
Matthew Allen
Matthew Stefanski
Chad Walber



Proceedings of the 44th IMAC, A Conference and Exposition on Structural Dynamics 2026



Conference Proceedings of the Society for Experimental Mechanics Series

Series Editor

Kristin B. Zimmerman
Society for Experimental Mechanics, Inc.,
Bethel, USA

The Conference Proceedings of the Society for Experimental Mechanics Series presents early findings and case studies from a wide range of fundamental and applied work across the broad range of fields that comprise Experimental Mechanics. Series volumes follow the principle tracks or focus topics featured in each of the Society's two annual conferences: IMAC, A Conference and Exposition on Structural Dynamics, and the Society's Annual Conference & Exposition and will address critical areas of interest to researchers and design engineers working in all areas of Structural Dynamics, Solid Mechanics and Materials Research.

Dario Di Maio · Matthew Allen · Matthew Stefanski · Chad Walber
Editors

Special Topics & Experimental Techniques in Structural Dynamics, 2026, Vol. 5

Proceedings of the 44th IMAC, A Conference and Exposition on
Structural Dynamics 2026



River Publishers

Published, sold and distributed by:

River Publishers
Broagervej 10
9260 Gistrup
Denmark

www.riverpublishers.com

ISBN 97887-438-1444-3 (Hardback)

ISBN 97887-438-1416-0 (eBook)

<https://doi.org/10.13052/97887-438-1444-3>

Conference Proceedings of the Society for Experimental Mechanics

An imprint of River Publishers

© The Society for Experimental Mechanics, Inc. 2026

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, or reproduction in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use. The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Preface

Special Topics & Experimental Techniques in Structural Dynamics represents one of ten volumes of technical papers presented at the 44th IMAC, A Conference and Exposition on Structural Dynamics, organized by the Society for Experimental Mechanics, and held January 19-22, 2026. The full proceedings also include volumes on Nonlinear Structures & Systems; Dynamics of Civil Structures; Model Validation and Uncertainty Quantification; Dynamic Substructuring & Transfer Path Analysis; Computer Vision & SLDV for Structural Dynamics; Dynamic Environments & Durability Testing; Topics in Modal Analysis & Parameter Identification in Structural Dynamics; Data Science in Engineering; and Structural Health Monitoring, Metamaterials, Metastructures & Additive Manufacturing and Dynamics of Biomedical and Bioinspired Systems.

Each collection presents early findings from experimental and computational investigations on an important area within Structural Dynamics. *Special Topics & Experimental Techniques in Structural Dynamics* represents papers highlighting new advances and enabling technologies for Experimental Techniques and Rotating Machinery.

The organizers would like to thank the authors, presenters, session organizers, and session chairs for their participation in this track.

Editors:

Dario DiMaio – University of Twente, Enschede, Netherlands;

Matthew Allen – Brigham Young University, Provo, Utah, USA;

Matthew Stefanski – U.S. Air Force Research Laboratory, Wright-Patterson AFB, OH, USA;

Chad Walber – Michigan Technological University, Houghton, MI, USA

Contents

1	Verifying the Operation and Health of Dynamic Strain Gages on Experimental Structures	1
	Matthew Stefanski and William Boles	
2	The Role of Open-Source Hardware for Research and Teaching	9
	Johannes Maierhofer, Oliver M. Zobel, and Daniel J. Rixen	
3	Vibration Measurement and Mitigation Development for Multi-Shaker Test Stand	15
	Maura C. Hart, Jackson A. Jones, Kyle Crapo, Madyson G. Hodson, Connor Crandall, Mason Rollans, Matthew S. Allen, and Michael Denison	
4	Dynamic Recognition of Damage States in Wind Turbines Using Cepstral Coefficients and ANOVA Analysis	29
	Costanza Speciale, Sina Shid-Moosavi, Eleonora Maria Tronci, and Silvia Milana	
5	Enhanced Bearing Fault Diagnosis Using Convolutional Neural Networks and Signal Preprocessing Techniques on the CWRU Dataset	39
	M. Conti, G. Battista, S. Pavoni, and M. Vanali	
6	Aeroacoustics Structure Interaction Experiment in a Rectangular Duct with a Flush Mounted Panel	49
	Daniele Fabbri, Carlo Rosso, Fabio Bruzzone, Francesco Avallone, Lucas A. Bonomo, and Julio A. Cordioli	
7	Machine Learning Surrogates for Rapid Hertzian Stress Estimation and Static Transmission Error in Gear Pre-Sizing	55
	Daniele Fabbri, Carlo Rosso, and Fabio Bruzzone	
8	Acoustic Excitation for Fluidic Particle Separation	63
	Blake Johnson, Bradon Timms, Jin Yue Trousil, Pavel Vakhlamov, Milo Prisbrey, John Greenhall, and Nicholas Lieven	
9	Extending Rattlesnake Vibration Controller to MIMO Sine Control and Other Significant Improvements	77
	Cody M. Langston and Dan Rohe	
10	Experimental Evaluation of Three-Dimensional Mode Shapes and Frame-Consistent Comparison with Numerical Models	85
	Umidjon Usmanov, Giuseppe Battiato, and Christian Maria Firrone	
11	Collapsing Distributed Fiber Sensors to Point Measurements to Boost Signal-to-noise using Custom-designed Fiber Coils	95
	Angel H. Chafino, Serena A. Vu, Reganne M. Watts, Steven R. Anton, Carly M. Donahue, Andrew K. Harter, Nathan T. Maier, and William A. Redford	
12	Reduced Order Structural Modeling for Control System Simulation	107
	Max Dommers, Matthew S. Pearson, and Adam C. Johnson	
13	Assessing Floquet Theory as an Alternative to Multi-Blade Coordinate Transformation in Wind Turbine Stability Analysis	119
	Jose J. Mora Amaro, Emmanuel Branlard, and Riccardo Riva	
14	Using Minimum Entropy Decomposition for Rolling Element Bearing Fault Diagnosis in a Permanent Magnet Motor	123
	Surendra (Suri) Ganeriwala	



Chapter 1

Verifying the Operation and Health of Dynamic Strain Gages on Experimental Structures

Matthew Stefanski and William Boles

Abstract Unlike Accelerometers, microphones, and most other sensors used to measure the mechanical dynamics of structures, dynamic strain gages lack a direct calibration method to a known force input. This paper provides practical guidance for engineers new to structural dynamic experiments on evaluating the integrity and proper response of dynamic strain gages. The techniques will include post-installation inspection, impedance measurements, and signal integrity checks that can be used to identify and mitigate potential sources of error. By implementing these methodologies, engineers can improve the reliability and accuracy of dynamic strain measurements, leading to more trustworthy experimental results.

Keywords Instrumentation · Sensors · Strain · Gage · Structure

The best way to ensure proper function of a thin film strain sensor is to ensure that the gage selection process and gage installation procedures were performed properly. However, the purpose of this paper is to discuss the best practices and procedures to ensure accurate gage functions AFTER the sensor has been installed.

When working with thin film strain gages, the quality control checks of **visual inspection** and process of gaging resistance check involve two primary measurements to ensure a proper installation: **insulation resistance** and the **shift in gage resistance**. These three checks are critical for verifying the quality of the installation and are usually performed with magnifiers and equipment such as a digital multimeter or Vishay Model 1300 Gage Installation Tester / GIT box.

Visual Inspection

The first step to ensure your newly installed strain sensor performs properly is to conduct a thorough **visual inspection**. Preferably under magnification due to the small size particulates voids that may exist, encapsulated under the strain sensor grid.

Trapped voids or foreign particulates within the bond layer of a thin film strain sensor can have significant negative effects on the sensor's performance and accuracy. These defects disrupt the uniform transfer of strain from the test specimen to the sensor's grid, leading to inaccurate and unreliable measurements.

Matthew Stefanski

Branch Chief Structural Validation, Air Force Research Laboratory, 2790 D St. WPAFB, OH 45433

e-mail: matthew.stefanski@us.af.mil

William Boles

Senior Instrumentation Engineer, Air Force Research Laboratory, 2790 D St. WPAFB, OH 45433

e-mail: william.boles@us.af.mil

Adverse Results of Voids and Particulates

The presence of these defects creates a poor mechanical connection between the strain sensor and the article under test, resulting in several adverse outcomes:




Fig. 1 Particulates in the Adhesive Bond Layer

1. **Reduced Sensitivity and Measurement Error:** The most direct consequence is a reduction in the sensor's sensitivity. Because the strain is not fully transferred, the change in resistance will be less than would be expected for the actual strain. The presence of these defects also introduces non-linearity and hysteresis into the measurements. The sensor may not return to its original zero point after being unloaded, causing a permanent zero shift.
2. **Non-Uniform Strain Transfer:** Voids and particulates prevent the strain sensor's grid from being perfectly bonded to the specimen. This creates localized areas where the strain is not accurately transmitted, leading to a distorted and uneven strain field across the sensor. This is a critical issue because the sensor's function relies on a one-to-one transfer of strain.

Premature Failure: These defects can act as stress concentrators, particularly at the edges of a void or around a particulate. This can lead to premature failure of the bond layer or the sensor itself, especially under cyclic or fatigue loading conditions. The presence of a void also makes the installation more susceptible to environmental factors like moisture and corrosive agents, which can lead to further degradation of the bond.

Cracks in the metallic grid of a thin film strain sensor have severe negative consequences, primarily causing inaccurate readings, signal instability, and eventual sensor failure. These cracks disrupt the intended electrical pathway, altering the fundamental relationship between strain and resistance. Visual inspection should be performed prior to mechanical cycle and periodically throughout test cycling, paying close attention to non-repeatable results at identical load point during the cycling.

Adverse Results of Cracks in the Grid

The primary purpose of a strain sensor's metallic grid is to provide a stable, uniform change in resistance with applied stress. When a crack occurs, this relationship is fundamentally broken.




Fig. 2 Voids/Bubbles in Adhesive Bond Layer

- **Non-Linearity and Hysteresis:** Cracks introduce an unpredictable change in the electrical path, making the sensor's output non-linear. The resistance change will no longer be proportional to the applied strain. This also causes hysteresis, where the output for a given strain value is different depending on whether the load is increasing or decreasing.

Electrical Effects and Complete Failure

Cracks can have various electrical effects that degrade the sensor's performance and can ultimately lead to sensor failure.

- **Increased Noise:** A crack acts as a tiny, unstable switch. As the crack opens and closes with minute movements or vibrations, it generates significant electrical noise in the sensor's signal, masking the actual strain measurement.
- **Open Circuit:** If the crack propagates completely through the metallic grid, it will create an **open circuit**. This means the electrical path is broken entirely, and the sensor will cease to provide any output. The resistance will become infinite, rendering the sensor completely useless.
- **Localized Heat Generation:** In cases where a crack significantly reduces the cross-sectional area of the grid, the current density increases at that point. This can cause localized heating, which further alters the resistance and can accelerate the propagation of the crack, leading to a faster failure.

Standard checks for installed thin film strain sensors primarily focus on **installed gage resistance** and **insulation resistance** to verify the quality of the installation before use. These checks are crucial for ensuring accurate and reliable data, as poor installations can lead to significant measurement errors.

Installed Gage Resistance or Shift Resistance: Gage shift in resistance is a crucial measurement for verifying the quality of a strain gage installation. An excessively high gage resistance during the installation of a thin film dynamic strain sensor can lead to several problems, primarily related to inaccurate measurements and poor signal quality. The main symptoms include a **low signal-to-noise ratio (SNR)** and **reduced sensitivity**. This check is taken to ensure the sensor hasn't been damaged during mounting or the installation process. It's important because a significant shift in the gage's resistance from its original, unstrained value can indicate damage from improper handling, excessive clamping force, or other installation issues.




Fig. 3 Cracking of the Sensor Grid

Reduced Sensitivity and Low Signal-to-Noise Ratio (SNR)

The symptoms of excessive gage resistance are directly related to the principles of Wheatstone bridge circuits, which are commonly used to measure strain. The output voltage of a Wheatstone bridge is proportional to the change in resistance of the strain gage. When the nominal resistance (R_g) of the gage is excessively high, it can negatively impact the bridge's performance.

The voltage output of a quarter-bridge circuit is given by the formula:

$$V_{out} = V_{in} \cdot \frac{GF \cdot \epsilon}{4}$$

Where:

- V_{out} is the output voltage
- V_{in} is the input voltage
- GF is the gage factor
- ϵ is the strain

While this formula doesn't explicitly show resistance, the noise in the circuit does. Thermal noise, which is a key contributor to the noise floor, is proportional to the square root of the resistance. The formula for thermal noise voltage (V_n) is:

$$V_n = \sqrt{4 \cdot k \cdot T \cdot R \cdot \Delta f}$$

Where:

- k is the Boltzmann constant
- T is the temperature in Kelvin
- R is the resistance
- Δf is the bandwidth

As the **gauge resistance** (R) increases, the **thermal noise voltage** (V_n) also increases. This leads to a higher noise floor.

A higher noise floor, combined with the relatively small signal generated by the strain, results in a **low signal-to-noise ratio (SNR)**. A low SNR makes it difficult to accurately measure small changes in resistance caused by strain, leading to noisy and unreliable data. The sensor's ability to detect small strains is therefore significantly reduced, effectively making it less sensitive. This can be problematic in applications where precise measurements of low-amplitude vibrations or deformations are required.

Other Symptoms

In addition to the fundamental issues of sensitivity and SNR, excessive gauge resistance can also manifest in other symptoms:

- **Difficulties with Data Acquisition:** High resistance can be a challenge for some data acquisition (DAQ) systems, which are optimized for common strain gage resistances (e.g., 120 Ω or 350 Ω). It may require a DAQ system with a higher input impedance to prevent loading effects that would further diminish the signal.
- **Drift and Instability:** High resistance can sometimes be a symptom of a poor connection or a damaged sensor element, which can cause the resistance to drift over time or with changes in temperature. This leads to an unstable zero point and inaccurate long-term measurements.

How to check: To measure the shift in resistance, you compare the **installed resistance** of the strain gage to its **nominal resistance**. This is typically done with a specialized instrument like the Vishay Micro-Measurements Model 1300 Gage Installation Tester, which is designed to measure this deviation with high precision but can be accomplished using a digital multimeter capable of precise resistive measurement.

- **Nominal Resistance:** This is the value specified by the manufacturer, such as 120 Ω or 350 Ω on the sensor packaging.
- **Installed Resistance:** This is the resistance of the gage after it has been bonded to the test specimen and the lead wires have been attached.

The measurement is often displayed as a **percentage deviation** from the nominal value. A resistance shift greater than the recommended limits, found in Vishay Application Note VMM-15, entitled "Quality Control of Strain Gage Installations", indicates a potential problem with the gage installation.

Recommended Values and What They Indicate

Vishay Micro-Measurements documents provide specific guidelines for acceptable resistance shifts, which can vary depending on the type of adhesive used:

1. **Room-Temperature-Curing Adhesives:** For these adhesives, the resistance shift should not normally exceed 0.5% of the nominal gage resistance. A shift greater than this limit usually indicates damage to the gage from improper handling or clamping during the installation process.
2. **Elevated-Temperature-Curing Adhesives:** With these adhesives, a larger shift is expected because of the difference in thermal expansion between the strain gage and the test specimen. This creates thermal stress in the gage as it returns to room temperature after curing. While the magnitude of this shift can vary, it should never exceed 2% and should be uniform among identical installations to within about 0.5%.

By performing this simple test, you can identify installation flaws like a damaged grid, poor solder connections, or lead wire damage before you begin your critical strain measurements.

Insulation Resistance (Leakage to Ground) :

This check measures the electrical resistance between the strain gage's grid and the test specimen. It's a fundamental test to ensure the gage is properly insulated from the surface it's bonded to.

Bad insulation resistance, or leakage to ground, in a thin film dynamic strain sensor creates a stray current path that can significantly interfere with measurements. This short circuit from the sensor's active element to the grounded test structure introduces a variable resistance in parallel with the strain gage, causing several critical symptoms.



Fig. 4 Vishay Model 1300 Gage Installation Tester

A low insulation resistance can cause extraneous signals to be introduced into the measurement, leading to inaccurate reading. It's often caused by trapped foreign matter, moisture, residual soldering flux, or damage to the gage's backing material.

Key Symptoms

- **Drift and Instability:** A primary symptom is a measurement baseline that **drifts** and is **unstable**. The leakage path's resistance isn't constant; it changes with temperature, humidity, and even the applied strain itself. This causes the output signal to fluctuate erratically, making it impossible to establish a reliable zero point for the measurement. You'll see the strain wander without any real load being applied.
- **Reduced Sensitivity and Low Output:** The leakage resistance creates a parallel path for current to bypass the strain gauge. This effectively lowers the overall resistance of the sensor circuit, which in turn reduces the change in resistance (ΔR) for a given strain. As a result, the output voltage from the Wheatstone bridge circuit is lower than expected, making the sensor less sensitive to strain.
- **Increased Noise:** Leakage to ground can introduce noise into the signal. The stray current path acts as an antenna, picking up electrical noise from the surrounding environment. This noise is superimposed on the small strain signal, leading to a poor **signal-to-noise ratio (SNR)**. This symptom is particularly problematic for dynamic measurements, where the signal is already small and easily obscured by noise.
- **Incorrect Strain Readings:** The most direct consequence is a false strain reading. The change in resistance from the leakage current is indistinguishable from the change in resistance caused by actual strain. This can lead to erroneous

tensile or compressive strain values, even if the structure is at rest. The error is unpredictable and can change with environmental conditions, making it impossible to correct in post-processing.

Causes and Troubleshooting

The primary cause of bad insulation resistance is damage to the thin film's dielectric or insulation layer. This can happen during installation due to:

- **Moisture or Contamination:** Trapped moisture, flux residue from soldering, or other contaminants under the sensor can create a conductive path. This is a common issue and is often a reason for the unstable, fluctuating readings.
- **Mechanical Damage:** A scratch, crack, or pinhole in the sensor's protective coating or substrate can expose the active grid to the grounded surface.
- **High Temperatures:** In high-temperature applications, the insulation resistance of the ceramic or polymer substrate naturally decreases, increasing the risk of leakage.

To test for this problem, you can perform an **insulation resistance test** (often called a “megger” test) using a specialized multimeter or insulation tester. You would apply a low DC voltage between the strain gauge circuit and the ground of the test structure and measure the resistance. A healthy installation should have an insulation resistance of **over 20,000 MΩ** (megaohms), although a reading at 10,000 MΩ or greater is generally considered the minimum acceptable value. Unstable readings or values under 10,000 MΩ indicate a problem.



Chapter 2

The Role of Open-Source Hardware for Research and Teaching

Johannes Maierhofer, Oliver M. Zobel, and Daniel J. Rixen

Abstract Open-source hardware and software are playing an increasingly important role in science and education. Open-source software for data analysis is already well-established, with Python Packages such as NumPy, SciPy and newly established machine learning packages such as scikit-learn, TensorFlow, and PyTorch.

In the field of structural dynamics Python Packages such as pyFRF, pyFBS, pyUFF, SDynPy, and pyIDI get more and more attention. However, the earlier stages of measurement – data acquisition and signal conditioning – remain heavily dependent on proprietary systems.

This poses challenges in both research and teaching. For students, commercial measurement systems often act as black boxes. The complexity and feature overload of professional software can obscure the core principles. Here, open-source acquisition tools provide a powerful alternative: by focusing on essential steps and exposing every detail of the processing chain, they enable a hands-on understanding of the underlying physics and signal processing.

From a research perspective, open-source hardware and software unlock new possibilities. For example, integrating live quality metrics directly into the measurement workflow enables immediate feedback. For example, channel consistency checks or impact repeatability for Virtual Point Transformations can be integrated into custom workflows. This prevents errors from propagating and reduces the need for costly repetitions.

However, building custom acquisition tools on top of proprietary platforms is not only technically challenging but also cost-prohibitive, especially in academic settings. A more inclusive solution is to develop open-source acquisition systems based on vendor-independent hardware. OASIS – the Open Acquisition System for IEPE Sensors – is one such initiative. It aims to create a flexible, transparent, and extensible platform suitable for both advanced research applications and introductory teaching labs.

This contribution outlines the motivation behind OASIS, highlights the benefits of open-source hardware in education and research, and invites the community to collaborate, share experiences, and help shape its future.

Keywords Testing equipment · Experimental dynamics · Open-source · Software · Hardware

Johannes Maierhofer

Chair of Applied Mechanics, TUM School of Engineering and Design,
Technical University of Munich, Boltzmannstr. 15, 85748 Garching, Germany
e-mail: j.maierhofer@tum.de

Oliver M. Zobel

Chair of Applied Mechanics, TUM School of Engineering and Design,
Technical University of Munich, Boltzmannstr. 15, 85748 Garching, Germany
e-mail: oliver.zobel@tum.de

Daniel J. Rixen

Chair of Applied Mechanics, TUM School of Engineering and Design,
Technical University of Munich, Boltzmannstr. 15, 85748 Garching, Germany
e-mail: rixen@tum.de

The Success of Open-Source Software

Historical Review

The history of open-source software begins with the early days of computing, when source code was freely distributed with personal computers [1]. Unix, initially released for free by AT&T, became increasingly restricted after version 7 in 1979, as the company recognized its commercial value [2]. A major turning point followed in 1983, when licensing changes led researchers and universities to abandon Unix in favor of the emerging BSD (Berkeley Software Distribution) [2]. In the same year, a court decision established that software could be copyrighted [3], prompting BSD to remove all AT&T code [2].

Richard Stallman, concerned about the ethical implications of closed software, began promoting the term *Free Software* at MIT's AI Lab [1]. In 1983, he launched the GNU project, aiming to develop a fully open operating system. The motivation was partly practical—for example, a printer that repeatedly jammed but could not be fixed due to proprietary restrictions [4]. Importantly, *free* referred to freedom rather than zero cost [2]. The release of the GNU Compiler Collection (GCC) in 1986 provided a robust foundation for free software, and it remains widely used today.

From a Personal Project to Industrial Backbone

A look into the history of open-source projects must include the Linux project. What began in 1991 as Linus Torvalds' personal hobby—a small Unix-like kernel written for educational purposes—grew into one of the most influential open-source successes. Released under a free license, the Linux kernel developed into the operating system powering servers, supercomputers, and mobile devices worldwide. Its success also enabled new business models: companies such as Red Hat built sustainable enterprises not by selling licenses, but by providing long-term support and services [5].

On the one hand, this leads to the conclusion: open software systems can combine stability for industry with freedom for users [6]. On the other hand, the question arises: Can this be true also for open hardware systems? In this publication, we want to focus our view towards scientific measurement hardware. This is closely related to the way research (and teaching) is programmed at the time. Currently, Python is common practice in many fields.

Python in Science

During the last 10 years, Python became the standard language for data analysis. NumPy established fast array operations [7], and SciPy made advanced algorithms widely available [8]. These tools are community-driven yet form the core of daily scientific work. In vibration analysis, `pyFBS` is an example of how even advanced methods can be shared openly and reused across groups [9]. Openness here means that methods are easier to verify and teaching is closer to real research practice.

First Steps in Hardware

Compared to software, most measurement hardware is still closed and expensive. Open designs are now emerging that make instruments more accessible and adaptable [10]. They allow students to see what happens inside a system and give researchers more control over their setups. At the same time, hardware brings new responsibilities, such as ensuring safe use and compliance with regulation. The OASIS project follows this path, using Creative Commons licensing for hardware and MIT for software, and shows how open hardware can enter both teaching labs and research groups [11].

Open-Source Hardware Defined

Open-source hardware (OSH) extends the ideas of free and open software into the physical domain. The Open Source Hardware Association (OSHW)¹ describes it as hardware whose design files are publicly available, so that anyone can study, modify, distribute, make, and sell the design or derived versions [10]. This requires that the essential artifacts of a project – schematics, layout files, bills of materials (BoM), and firmware – are openly documented.

Licensing is central to this concept. For software, permissive licenses such as MIT or Apache are widely used [7, 8]. For hardware, CERN has developed the Open Hardware License (OHL)² in variants that mirror software traditions: a permissive

¹<https://oshwa.org/>

²<https://cern-ohl.web.cern.ch/home>




Fig. 1 Overview of open-source hardware and the required components: the hardware itself, licensing, and documentation

version (OHL-P), a weakly reciprocal version (OHL-W), and a strongly reciprocal version (OHL-S). These provide legal clarity for sharing and reusing hardware designs. In practice, projects often combine a hardware license with a separate software license. For example, the OASIS project³ uses Creative Commons Attribution 4.0 International for hardware and the MIT license for its software components.

Standards for documentation quality are another pillar. DIN SPEC 3105 [12, 13] defines how open hardware should be documented so that others can reproduce and verify it. Certification programs such as the OSHWA certification build on these ideas and give users a simple signal that a project meets basic openness criteria. Together, licenses and standards establish the framework that allows communities to build sustainable and legally robust hardware ecosystems.

Sustainable Data and File Formats

For sustainable open-source hardware projects, the choice of data formats is as important as licensing and governance.

A first example comes from source code itself: storing code as plain UTF-8 text has become an obvious standard. Practically every editor can open the files without affecting their meaning. This universality made open-source software collaboration possible at a global scale.

For hardware design files the situation is more complex. CAD and ECAD models are central artifacts, yet commercial tools use closed formats that are rarely interoperable. Projects such as KiCAD⁴ and FreeCAD⁵ provide open formats, but these are not directly compatible with proprietary systems. This tension makes the decision for open formats a critical sustainability factor: once adopted, a community must ensure long-term support and backward compatibility. A format must be widely supported while not limiting the abilities of the underlying tool.

A second example concerns experimental data. Here, reproducibility requires reliable and long-term storage. HDF5⁶ has established itself as a format that combines efficiency with wide software support. It allows large measurement datasets to be stored together with metadata and can be accessed from many programming environments. This makes it a robust choice for open scientific hardware projects where data must remain usable for years.

³<https://gitlab.com/oasis-acquisition>

⁴<https://www.kicad.org/>

⁵<https://www.freecad.org/>

⁶<https://www.hdfgroup.org/solutions/hdf5/>

Opportunities and Challenges

Open-source hardware invites new ways of building and teaching with instruments. It reduces entry barriers and makes methods visible. At the same time, it shifts effort toward documentation, validation, and governance. This section summarizes gains and pitfalls, then treats certification in depth.

Opportunities

- **Transparency.** Designs, firmware, and data paths are inspectable and auditable, which helps teaching and method verification [10, 14].
- **Adaptation.** Features can be added at the measurement stage (e.g., online quality checks) without waiting for vendor roadmaps [10].
- **Access.** Costs fall for entry setups; as an example scenario, an 8-channel EU build can be in the range of 500–1000 EUR total, whereas proprietary systems are often priced in the range of up to 1000 EUR per channel and include recurring contracts for software (illustrative only).

Challenges

- **Upkeep.** Communities must maintain designs, toolchains, and releases; this is a real cost [6].
- **Validation.** Without vendor warranties, labs must prove fitness for use.
- **Governance.** Clear versioning, issue handling, and contribution rules are needed for trust [10].

Certification and Compliance

Open status does not change regulatory obligations. Compliance attaches to the product and its intended use, not to the license.

CE (EU). For most laboratory electronics, the manufacturer issues a Declaration of Conformity after demonstrating compliance with the relevant directives and harmonised standards. Typical elements include: a technical file (schematics, layout, BoM, risk assessment), electromagnetic compatibility (EMC) and safety testing where applicable, user instructions, and product labeling. A Notified Body is only required in defined cases. Good practice for open projects: treat the repository as the living technical file; tag releases that correspond to tested hardware; archive test reports with the release.

UL/NRTL (US). In US workplaces, OSHA expects electrical equipment to be evaluated by a Nationally Recognized Testing Laboratory to the applicable standard. Listing is separate from CE and from openness. If open hardware is to be deployed in such environments, plan the enclosure, mains interface, wiring, and markings so that a later listing is feasible. Using listed external power supplies and observing creepage/clearance rules reduces rework.

Software and updates. Law increasingly treats firmware and updates as part of the product lifecycle. Open projects should define an update policy, keep change logs, and document security fixes. Release artefacts (bitstreams, binaries) should be reproducible from tagged sources.

Documentation quality. Standards and community norms help here. DIN SPEC 3105 [12, 13] stresses complete documentation for reproducibility; educational deployments benefit from that discipline [14, 10]. License choices should be explicit: hardware under CC BY 4.0 or CERN–OHL variants; software under permissive terms where appropriate [7, 8].

Quality Assurance and Data Trust

Calibration against traceable references is essential for credibility. Record calibration factors, uncertainties, and environmental conditions with the data. Maintain BoM/SBoM, firmware hashes, and versioned schematics so that past measurements remain interpretable. Prefer durable data formats (e.g., HDF5) so that large datasets and metadata travel together.

Teaching and Institutional Policies

Universities differ in lab rules; however, recurring elements exist: risk assessment, supervision, equipment checks, and written SOPs. Open hardware supports learning because students can see every step; it remains the instructor's duty to bound voltages, enclose mains circuits, and document safe operation [14].

Illustrative Example: Vibration Analysis

This section demonstrates how **pyFBS** [9] (software) and **OASIS-UROS** [11] (hardware and firmware) can serve as concrete examples that highlight both the potential and the limitations of open-source tools in vibration analysis.

pyFBS

The pyFBS project (University of Ljubljana / TUM) is a Python package for Frequency-Based Substructuring (FBS), Transfer Path Analysis, and modal identification methods, including multi-reference and modal tests [9]. It implements advanced methods such as the Virtual Point Transformation (VPT), System Equivalent Model Mixing (SEMM), Frequency Response Function (FRF) synthesis, and provides a 3D display of sensors, excitations, and modal shapes [9]. The package includes both numerical and experimental datasets ranging from simple beam-like structures to more complex assemblies, which allow users to practice realistic workflows. Licensed under MIT, pyFBS is free to use, extend, and integrate into teaching and research with minimal hurdles.

The main strength of pyFBS lies in its ability to provide researchers and students with access to advanced substructuring methods without requiring expensive software. Its open nature ensures reproducibility, as datasets and workflows are publicly available and the methods and transformations can be inspected in detail. Although the software is accessible to everyone, generating meaningful results still requires considerable expert knowledge and practical experience. In commercial solutions, the high cost can act as a natural filter, ensuring that mainly established experts use the tools with the required attention. In contrast, open-source software invites broader use, which makes the role of expertise even more central. Experimental accuracy, for example, continues to depend on high-quality data, especially in high-frequency or low-noise scenarios. Methods such as expanded VPT remain sensitive to factors like sensor placement and alignment, and overall performance is directly linked to data quality.

OASIS-UROS

OASIS-UROS is an open acquisition system for IEPE sensors (Integrated Electronics Piezoelectric), designed primarily for academic research and teaching. Building on earlier OASIS versions [15, 16], it incorporates improved design choices. Its architecture is based on an ESP32 microcontroller, combined with an 8-channel AD7606C-18 ADC, SD-card caching, adjustable voltage ranges, and an upgraded IEPE front-end. The system is capable of simultaneously acquiring eight channels with sampling rates of up to approximately 36 kHz across all channels, supporting oversampling and benefiting from the improved front-end design. Validation against commercial hardware in modal analysis has shown that in certain test cases, OASIS-UROS achieves performance close to commercial systems, though not fully identical. All hardware, firmware, and software are openly published and reproducible, with the project listed in OSHA's certified register⁷ under CC-BY-4.0 for hardware and MIT for software.

The strengths of OASIS-UROS lie in its cost-effectiveness and its pedagogical value. It enables students to study and experiment with all layers of the acquisition chain, from the analog front-end to the firmware and sampling chain, making it highly suitable for teaching and exploratory research. Its performance is sufficient for many vibration tasks, such as modal analysis and FRF estimation, particularly when extreme precision or very high bandwidths are not required. At the same time, limitations remain. Compared to high-end commercial systems, OASIS-UROS cannot match the lowest possible noise floors, the highest sampling rates, or the guarantees of certified safety and EMC compliance. Assembly, calibration, and user expertise can influence the results, and practical factors such as enclosures, input protection, power supplies, and the measurement environment remain critical for achieving a reliable data acquisition system.

⁷<https://certification.osha.org/de000150.html>

Lessons from pyFBS and OASIS-UROS

The comparison between pyFBS and OASIS-UROS illustrates broader lessons for open-source vibration analysis. Software alone is not sufficient: high-quality data acquisition hardware is essential for reliable results. The analog front end must be carefully validated for noise, gain, sampling rate, and synchronization, as otherwise software methods can yield misleading conclusions. Open licensing – MIT for software, CC-BY for hardware – and transparent publication of materials are key enablers for inspection, derivation, reuse, and building trust. At the same time, performance trade-offs are unavoidable. As requirements for channel count, sampling rates, or formal certification increase, so do cost and complexity. Consequently, open-source hardware is particularly compelling for teaching, prototyping, and non-mission-critical research, or as a supplement to commercial equipment, but does not fully replace high-end certified systems in all contexts.

Conclusion

Openness in hardware and software is not simply better or worse than proprietary systems. It shifts the balance of effort: away from procurement and vendor lock-in, and toward quality assurance, documentation, and certification. This redistribution can be highly valuable in teaching, where transparency deepens understanding, and in exploratory research, where flexibility and adaptability outweigh the need for certified performance.

The examples of pyFBS and OASIS-UROS show that advanced methods and practical hardware can both be shared under open licenses without losing relevance to the community. For long-term success, however, projects must take the next steps: align documentation with DIN SPEC 3105, define clear certification pathways (CE, UL/NRTL), and encourage reproducible community practice. By doing so, open-source hardware can establish itself as a trustworthy complement to commercial systems in research and education.

References

1. Maracke, C. “Free and Open Source Software and FRAND-based patent licenses”. In: *The Journal of World Intellectual Property* 22.3-4 (2019), pp. 78–102. ISSN: 1747-1796. DOI: [10.1111/jwip.12114](https://doi.org/10.1111/jwip.12114).
2. Weber, S. *The Success of Open Source*. Harvard University Press, 2005. ISBN: 9780674018587.
3. Nussbaum, J. L. “Apple Computer, Inc. v. Franklin Computer Corporation Puts the Byte Back into Copyright Protection for Computer Programs”. In: *Golden Gate University Law Review* 14.2 (1984). URL: <https://digitalcommons.law.ggu.edu/cgi/viewcontent.cgi?article=1344&context=gulrev>.
4. “Transcript of Richard M. Stallman’s speech “Free Software: Freedom and Cooperation””. In: (2001). URL: <https://www.gnu.org/philosophy/rms-nyu-2001-transcript.txt>.
5. Young, R. “How Red Hat Software Stumbled Across a New Economic Model”. In: *First Monday* (1999). DOI: [10.3998/3336451.0004.304](https://doi.org/10.3998/3336451.0004.304).
6. Duparc, E. “Archetypes of open-source business models”. In: *Electronic Markets* 32 (2022), pp. 609–623. DOI: [10.1007/s12525-022-00557-9](https://doi.org/10.1007/s12525-022-00557-9).
7. Harris, C. R., Millman, K. J., Walt, S. J. van der, Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., Kerkwijk, M. H. van, Brett, M., Wilson, A., Wilhelm, K. F., Hickey, H., Weckesser, S. J., Wilson, N. M., Millman, E., and al., et. “Array programming with NumPy”. In: *Nature* 585 (2020), pp. 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2).
8. Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., Walt, S. J. van der, Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., Mulbregt, P. van, and SciPy 1.0 Contributors, the. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17.3 (2020), pp. 261–272. DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
9. Bregar, T., Mahmoudi, A. E., Kodrič, M., Ocepek, D., Trainotti, F., Pogačar, M., Göldeli, M., Čepon, G., Boltežr, M., and Rixen, D. J. “pyFBS: A Python package for Frequency Based Substructuring”. In: *Journal of Open Source Software* 7.69 (2022), p. 3399. DOI: [10.21105/joss.03399](https://doi.org/10.21105/joss.03399).
10. Oellermann, M., Jolles, J. W., Ortiz, D., Seabra, R., Wenzel, T., Wilson, H., and Tanner, R. L. “Open Hardware in Science: The Benefits of Open Electronics”. In: *Integrative and Comparative Biology* 62.4 (May 2022), pp. 1061–1075. ISSN: 15577023. DOI: [10.1093/ich/iac043](https://doi.org/10.1093/ich/iac043).
11. Zobel, O. M., Maierhofer, J., Köstler, A., and Rixen, D. J. “OASIS-UROS: Open acquisition system for IEPE sensors upgraded, refined, and overhauled software”. In: *HardwareX* 23 (2025), e00650. ISSN: 2468-0672. DOI: <https://doi.org/10.1016/j.ohx.2025.e00650>.
12. *DIN SPEC 3105-1:2020-07, Open Source Hardware – Part 1: Requirements for technical documentation*. DOI: [10.31030/3173063](https://doi.org/10.31030/3173063).
13. *DIN SPEC 3105-2:2020-07, Open Source Hardware – Part 2: Community-based assessment*. DOI: [10.31030/3173062](https://doi.org/10.31030/3173062).
14. Heradio, R., Chacon, J., Vargas, H., Galan, D., Saenz, J., De La Torre, L., and Dormido, S. “Open-Source Hardware in Education: A Systematic Mapping Study”. In: *IEEE Access* 6 (2018), pp. 72094–72103. ISSN: 2169-3536. DOI: [10.1109/access.2018.2881929](https://doi.org/10.1109/access.2018.2881929).
15. Zobel, O. M., Maierhofer, J., Köstler, A., and Rixen, D. J. “OASIS: Bridging the Open-Source Gap Between Testing and Data Processing”. In: June 2025. ISBN: 978-961-7187-17-5.
16. Zobel, O. M. and Maierhofer Johannes and Rixen, D. J. “OASIS: Open Acquisition System for IEPE Sensors: For Academic Research and Teaching Purposes”. In: *Sensors & Instrumentation and Aircraft/Aerospace Testing Techniques, Volume 8*. Ed. by Walber, C., Stefanski, M., and Seidlitz, S. Springer Nature Switzerland, 2024, pp. 79–86. ISBN: 978-3-031-34938-6. DOI: [10.1007/978-3-031-34938-6_9](https://doi.org/10.1007/978-3-031-34938-6_9).



Chapter 3

Vibration Measurement and Mitigation Development for Multi-Shaker Test Stand

Maura C. Hart, Jackson A. Jones, Kyle Crapo, Madyson G. Hodson, Connor Crandall, Mason Rollans, Matthew S. Allen, and Michael Denison

Abstract Multiple Input Multiple Output (MIMO) vibration testing is used in many industries to simulate the environmental conditions that components will encounter, allowing for the assessment of their durability. A modular frame mounted with multiple shakers has been developed by Sandia National Laboratories that can produce complex vibration environments at a fraction of the cost and setup time currently required by typical shaker systems. MIMO random vibration testing aims to replicate the complex vibrations of a Device Under Test (DUT) by precisely controlling the drive signal sent to each shaker. In some test conditions, this setup experiences a sharp rise in drive voltage, which can make the test impossible to run with available shakers or amplifiers. The vibration environment within the frame itself has not been studied, and may be the source of these test failures. To better understand the dynamics of the test frame, a roving hammer modal test was conducted to identify the vibratory modes of the frame within the frequency band normally used for random vibration testing (20-2000 Hz). A single-shaker configuration was used to excite a suspended test article and the frame's response was measured. Two vibration treatments—viscoelastic isolation and particle damping—were explored. Their damping effect on antiresonance amplitude and drive voltage to the test article was quantified. It was found that the particle damping method was consistently effective at damping all frequencies, while viscoelastic damping had varied results depending on frequency and configuration. In the system studied here, the particle damping treatment decreased the RMS shaker voltage by a factor of 2 across the full test range and up to a factor of 20 around the antiresonances.

Keywords Dynamic testing environment · NEWT · Passive damping · Particle damping · Viscoelastic isolator

Introduction

Dynamic environment testing is essential for the structural qualification testing of everything from aerospace components [1] to consumer products [2]. Tests in industry are dominated by single-axis tests on large shakers, yet there is increasing

Maura C. Hart · Jackson A. Jones · Kyle Crapo · Madyson G. Hodson · Connor Crandall · Mason Rollans
Undergraduate Students, Brigham Young University, Department of Mechanical Engineering
e-mail: hartmc11@byu.edu; jaj268@byu.edu; kcrapo3@byu.edu; mgd35@byu.edu; ccran29@byu.edu; mrollans@byu.edu

Matthew S. Allen
Professor, Brigham Young University, Department of Mechanical Engineering, matt.allen@byu.edu
e-mail: matt.allen@byu.edu

Michael Denison*
Senior Members of Technical Staff, Sandia National Laboratories
e-mail: mdeniso@sandia.gov

*This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <https://www.energy.gov/downloads/doe-public-access-plan>.

interest in replacing these with multi-axis tests [3, 4]. Cross et al. recently proposed a new and robust multi-shaker system which they named the Natural Excitation Within a Test frame (NEWT) [5]. The NEWT frame is built using standard 3x3 in. (76x76 mm) 80/20 beams (i.e. see <http://8020.net>) and designed with a high degree of modularity, so one can use it to position a large number of shakers around a test article and reproduce the article's dynamic response in all axes simultaneously and up to levels comparable to those achieved on large shaker systems [5]. A NEWT frame is relatively inexpensive (i.e. compared to commercial, multi-axis test systems) and, because it is modular, it is easy to replace a shaker that becomes damaged in the course of testing.

Although the NEWT frame has been shown to be capable of replicating realistic dynamic environments, vibrations within the NEWT frame itself seem to limit the performance of the system. It has been observed that the shakers sometimes draw increasing amounts of electrical power and yet are still unable to reach the desired test conditions. The problem can sometimes be remedied by rearranging the shakers or tightening connections to the frame, or other adjustments. While the authors are not aware of a work that has studied this in detail, these problems are thought to occur due to antiresonances between particular shakers and the control accelerometer(s). An antiresonance in one of the shaker-control accelerometer transfer functions would cause the voltage to the shaker to increase significantly over a narrow range of frequencies. The aim of this study is to recreate a NEWT test frame, document a simplified baseline vibration environment in a Single Input Single Output (SISO) test, and develop passive vibration mitigation solutions for the recreated frame that reduce 1.) the vibration in the frame and 2.) the power input to the shakers when achieving a basic vibration environment in a component. This solution can then be applied in the future to a multi-shaker, Multiple Input Multiple Output (MIMO) testing frame, such as that in [5].

The effect of vibration within a similar frame was recently studied by Kramer et al. [6], who constructed a two-shaker frame and evaluated its performance in two configurations: with viscoelastic strips added to increase its damping, and with bars and cross braces added to increase its stiffness. They showed that both treatments changed the mode shapes of the frame, although any increase in damping was not quantified, nor was the change in stiffness. On the other hand, they found that all frames had control difficulty leading to over-test (i.e. resonances where the shaker voltage could not be made low enough to avoid exceeding the acceleration environment) at exactly the same frequencies. They also showed that all frames had nearly the same inverse FRFs, i.e. voltage over acceleration, indicating that the treatments did not alter the anti-resonances significantly. This is peculiar, because vibration absorber (or tuned-mass damper) theory [7] shows that the depth of an antiresonance decreases as damping is increased.

This study builds on [6] in several ways. First, the frame studied here is constructed from stiffer 3 in. (76 mm) 80/20 tubes, so it is similar to that of Cross et al. [5]. Specifically, a NEWT frame was constructed from 14 extruded 80/20 aluminum bars, sized 3x3x30 in (76.2x76.2x762 mm) and 3x3x36 in (76.2x76.2x914 mm), total dimensions of 36x36x36 in (914x914x914 mm), as shown in Figure 4a, and is the subject of this study. Second, we quantify the effect of two different damping treatments on the frame. Particle damping, the use of freely moving particles confined in a cavity, has been studied as a passive damping solution [8, 9, 10], and is implemented here. A roving hammer modal test is used to quantify the effect of the particle damping treatment on a single beam in the test frame, and on the frame as a whole, and found to increase the damping by an order of magnitude or more at some frequencies. Another common damping solution is to add viscoelastic material between components, typically in a way that minimizes the transmission of vibrations from one part to another. In our study this is done by inserting viscoelastic pads between the beams comprising the frame or beneath the shaker.

For both damping treatments, we construct a closed-loop control problem in which antiresonances are present and quantify the reduction in the shaker drive voltage required to reach the environment of interest on the Device Under Test's (DUT). In these tests, the vibration in the test frame was measured with an accelerometer recording along the axis of the shaker. The voltage and current to the shaker were also measured to quantify the effect of antiresonances on the voltage required to reach the environment.

Note that, while this study focuses solely on passive damping solutions, other ideas were considered. We also considered implementing a tuned mass damper. A tuned mass damper is passive, but is typically tuned to address a resonance at a specific frequency, such as the natural frequency of a building [11]. We attempted to implement two such techniques, but found that it was difficult to tune them for the anti-resonance frequencies of interest. Indeed, considerable experimentation was required to even identify the anti-resonances that contributed most to difficulties in our test system. Additionally, all of the tuned mass damper designs that we attempted were found to add several other modes to the system, potential creating new antiresonances and exacerbating the problem rather than addressing it. Hence, in the end we decided to focus on solutions that increased the damping in the frame over a wide (20-2000 Hz) range of frequencies.

The following section details the damping methods used and the tests that were performed. Results of those tests are presented in Sec. , followed by conclusions in Sec. .

Methods

To evaluate the performance of damping methods on a simple system, we initially performed roving hammer modal testing on a single, suspended 80/20 beam. We tested the beam both with and without the central cavity filled with silica, in order to document the baseline effectiveness of particle damping on the dynamics of the beam (Figure 1). Roving hammer testing was then performed on a completed NEWT frame with and without particle based damping treatments applied to compare the effect of the damping on the modes existing within a complex test frame. Another test we used to quantify the damping effect on the frame was a random noise vibration test on a DUT with both a damped and undamped test frame.




Fig. 1 Picture of experimental setup for roving hammer tests

Single Beam Roving Hammer Testing

A 3x3x36 in. (76.2x76.2x914.4 mm) 80/20 beam was suspended, without particle damping treatment, by bungee cables to approximate free-free boundary conditions and isolate the beam modes from any existing modes of the frame. A modal hammer test was used to measure Frequency Response Functions (FRFs) of the beam at 10 points along the top and side faces of the beam as shown in Figure 2. Each point was struck in the x and z directions, with corners also being struck in the y direction, to obtain 24 FRFs.




Fig. 2 Top view sketch of hammer impact points

The natural frequencies, damping ratios and mode shapes of the beam were obtained from the measured FRFs using the Algorithm of Mode Isolation (AMI) [12]. By isolating the modes of the beam and fitting modal model to the FRF, damping ratios were calculated for each configuration. To test the damped system, the hollow core of the beam was filled with approximately 1.7 L of 4075 white silica sand and end caps were placed at the end of each beam, as shown in Figure 3, to prevent leakage. This volume of sand fills the beam fully without packing. The modal hammer test was repeated using the same parameters and a new FRF of the damped case was produced.


Full Frame Roving Hammer Testing

To examine how the damping affected modal behavior within a complex frame, a modal hammer test using the same testing hardware as the single beam test was performed, with impact points being selected at the corners and midpoints of select




Fig. 3 Picture of beam with plastic end cap, One cap was used on each end and the cavity was filled with silica

beams to simplify the analysis as shown by the small squares in Figure 4b. This test was performed first on an undamped frame and then on the same frame with all the beam cores filled with approximately 1.7 L (1.4 L in the shorter 30 in beams) of white silica sand and sealed with end caps to prevent leaking. For both tests, FRFs were generated and damping ratios calculated for each identified mode using the same methods as in the single beam tests.



(a) Completed NEWT frame



(b) Test Geometry in Siemens TestLab®

Fig. 4 Experimental setup for hammer tests of the complete frame

Random Noise Vibration Testing

Using the same test frame from our full-frame roving hammer tests, we bolted a single Modal Shop Shaker Model 2075E to a crossbeam and oriented it parallel to the ground. The shaker was connected to an acrylic stinger with an aluminum collar as shown in Figure 5. The stinger was then fastened to a second aluminum collar which was attached to the DUT using super glue. We attached a control accelerometer opposite the stinger connection point on the DUT. To monitor frame behavior during the test, we also attached a second accelerometer to the test frame along the axis of actuation on the beam on which the shaker was mounted. To monitor the power going to the shaker, we recorded the voltage output of our control system in the software and recorded the current using a current probe attached to the shaker power cable.

For our controller we used a DataPhysics Abacus and the Random Noise package from the SignalStar Matrix software. The Data Physics system automatically computes the drive voltage needed by the shaker to match a given environment at the control accelerometer over the frequency range of interest (20-2000 Hz). The software measures and stores the Power Spectral Density (PSD) and FRF of the control accelerometer, of the accelerometer on the base of the frame, and of the voltage sent to the shaker. Current PSDs were also recorded in some tests, and found to have a similar appearance to the voltage PSDs, so this paper focuses on the voltages.





Fig. 5 Labeled image of the initial test configuration for random vibration testing

During initial tests, we used both a 1 lb mass and a 10 lb dumbbell as DUTs. We hung the DUT from Kevlar thread, which we then suspended from a crossbeam using bungee cables to approximate free-free boundary conditions and minimize energy being transmitted through the DUT into the test frame. With this test setup, no antiresonances were identified by visual inspection of PSDs or FRFs generated by the SignalStar software (see Appendix A, Figure 18). We hypothesize that it may not be possible to obtain anti-resonances in a single-shaker system with a rigid DUT, as the dynamic coupling of multiple shakers through the DUT and frame cause the undesirable antiresonances. As only one shaker was available, we simulated this multi-shaker coupling by using the single shaker to drive one of the cross beams of the frame. We established a direct connection to a horizontal cross beam of the frame using dental cement between the DUT beam and the stinger collar (Figure 6). The dental cement replaced Loctite super glue because the super glue failed and caused separation from the beam during testing. With the new setup, we were able to visually identify several frequencies at which antiresonance occurred, as shown in Figure 11, which is described subsequently.



(a) Shaker connection to DUT crossbeam



(b) Control accelerometer and shaker

Fig. 6 Photos of the Random Noise experimental setup

Using this configuration, we ran random noise tests with a maximum voltage output of 3 V to the amplifier to prevent damage to the shaker, a reference acceleration of $5 \times 10^{-8} g^2/\text{Hz}$, and a frequency range of 20-2000 Hz. These tests were repeated across several undamped cases before adding particle damping or viscoelastic isolation. For the particle damped




Fig. 7 Image of viscoelastic isolator installation with labels

tests, all beams comprising the frame were filled as previously mentioned. For viscoelastic isolation, 70 duro Sorbothane, 0.1 in (2.54mm) thick pads were placed between the shaker and the test frame as shown in Figure 7.

Results

Roving Hammer Tests

Each peak in the FRF plots is a resonance within the test frame. Figure 8 shows the composite (or average) of the measured frequency responses from both the single-beam and full-frame tests. In both cases the system with silica sand (plotted as a red line) has wider, less sharp resonances, indicating an increase in damping. The single beam shows only about five major resonances, so this effect is clear for that system. In contrast, for the undamped frame at least 136 resonances were identified below 2kHz, while only about 25 could be extracted once the silica was added to the frame. Additionally, the silica is seen to be most effective above about 300 Hz; modes below that frequency do not show a strong increase in damping.

The damping ratios for both systems were estimated from the measured FRFs using AMI, and are shown in Figure 9. The 80/20 beam without any treatment had damping ratios around 0.001, and these increased to 0.01-0.1 when the beam was






Fig. 8 Composite frequency responses from roving hammer tests. Red lines show the result for the system damped using silica powder and blue lines show the undamped system

filled with silica, an increase by a factor of 10-100. The damping in the full frame in the undamped configuration was a bit higher, between 0.001-0.005 for the higher frequency modes, and as high as 0.03 for the lower modes, as seen in Figure 9b. This increase in damping is presumed to come from the bolted joints in the assembly. The silica increases the damping of the higher frequency modes to about 0.05-0.08, an increase of a factor of 10-80. Granted, the damping was large enough in the high frequency modes that it was hard to extract reliably. In contrast, the damping ratios of the modes below about 300 Hz are similar to those of the bare frame. The silica seems to add relatively little damping at those frequencies. Those frequencies are below the first bending mode of a free-free 80/20 beam.




(a) Damping Ratios of Single Beam



(b) Damping Ratios of Full Frame

Fig. 9 Damping ratios extracted from roving hammer test on each system

The mode shapes of the full frame were examined to confirm the type of motion that was present at each frequency. A few sample mode shapes are shown in Figure 10. The modes below 100 Hz seemed to be dominated by the stiffness of the joints, with the beams moving rigidly. At about 100 Hz the vertical (Z-direction) beams begin to show first-mode type deformation (e.g. see Mode 5 in Figure 10b), and by 500 Hz the cross beams can be seen to bend as well (see Mode 41 in Figure 10c). Higher modes are expected to correlate to higher order bending of the various beams that comprise the frame, although the measurement grid used wasn't sufficient to visualize those modes.



(a) Mode 3 at 51.0 Hz

(b) Mode 5 at 97.7 Hz

(c) Mode 41 at 535.6 Hz

Fig. 10 Deformation Shapes of Selected Modes. Dashed black lines show the undeformed frame, while blue lines show the deformation pattern of the mode in question

Random Vibration

Frequency Response Function and Power Spectral Density

One of the most useful sets of data from the random vibration shaker tests is the FRF, a ratio of DUT acceleration to the shaker voltage input. This information shows how much voltage is needed to simulate a test environment of a certain amplitude and

frequency. These plots show at what frequencies resonances and antiresonances occur as indicated by peaks and valleys respectively. At antiresonances, the shaker requires a large amount of voltage to achieve a given amplitude of vibration in the DUT, which potentially compromises controllability of the DUT. Significant antiresonances were visually identified at around 580, 616.25, and 1357.5 Hz from Figure 11. The severity of the antiresonance and its effect on controllability can be gauged by the depth of the valley in the acceleration FRF.




Fig. 11 Control accelerometer FRF with the undamped frame. Antiresonance bands used for the RMS values are highlighted in red

Each antiresonance indicates a frequency in which the vibration at the control accelerometer is minimal for a given shaker voltage. Per the well known theory [7], the energy from the shaker is driving vibration in other parts of the structure, which vibrate in a way that cancels the applied force, allowing the control accelerometer to remain stationary. Antiresonances generally emerge due to a linear combination of the modes near the antiresonance frequency. To obtain an understanding of how the frame might move in each antiresonance, some of the modes that were identified from the data in Figure 8b are shown in Figure 12. The control accelerometer is located near the center of the beam in the front of the frame (near $(x, y, z) = (16, 0, 15)$ in those figures). For the first two antiresonances, it seems that the lateral bars (along the x -axis) vibrate, as do the top corners of the frame, to cancel the motion at the control accelerometer. The antiresonance at 1357.5 Hz, which is the strongest in Figure 11, seems to involve motion of the top of the frame and motion of the legs where they rest on the carpet.

Another important data set from the random vibration shaker tests is the PSD of the voltage sent to the shaker. Because the DUT is controlled to a constant amplitude across the frequency range, the voltage PSD looks very similar to the reciprocal




Fig. 12 Deformation shapes of the modes identified from the data in Figure 8b nearest to each antiresonance in Figure 11

of the acceleration FRF. This means that antiresonance points are characterized by peaks instead of valleys. The PSD can be integrated over a frequency band to find a Root Mean Squared value (RMS) of the voltage over that range of frequency. In our analysis, we compared the RMS voltages from a broadband range of 20-2000 Hz as well as narrower bands at 550-650 Hz and 1300-1400 Hz. These two narrower bands focus on identified antiresonances.

Undamped System Sensitivity

While performing the random vibration shaker tests on the undamped frame, it was noted that the dynamics of the frame and DUT are very sensitive to setup and are challenging to recreate. For example, the torque to which the frame bolts are tightened significantly alters the system dynamics. Figure 13 shows the results from an undamped shaker test with the frame joints tightened to 11 ft-lb as well as an undamped test with 8 ft-lb bolted joints. The change in bolt torque caused both a frequency shift and amplitude changes to minima and maxima in the plots. The resonance and antiresonance frequencies typically shifted upwards as the bolt torque increased, which makes sense considering that higher torques likely increase the stiffness of the joints. The difference in amplitude of the antiresonances is quantified by the RMS values of those frequency bands, which are shown in Table 1.




Fig. 13 Comparison of random noise results when using an undamped frame with variable bolt torque at joints

Table 1 Comparison of voltage RMS values for an undamped frame at different frame bolt torques

	Voltage RMS Over Antiresonance Bands (V^2)		
	550-650 Hz	1300-1400 Hz	20-2000 Hz
11 ft-lb	7.556×10^{-7}	8.583×10^{-4}	1.097×10^{-3}
8 ft-lb	8.415×10^{-8}	4.437×10^{-4}	6.369×10^{-4}

Even while holding bolt torque in the frame constant, removing beams from the frame and then reattaching them in the same location significantly changed the amplitude of resonances and antiresonances. Figure 14 shows FRFs and PSDs of multiple undamped tests. Between each of these runs, the frame was partially disassembled and then reassembled with identical geometry and bolt tightness. While the general profile of these plots remains constant, there are differences in the amplitudes of the minima and maxima. This is an especially important result for large valleys in the FRF where frame antiresonances occur as these require significant increases in the shaker voltage and hence could make the test impossible to run using the given shaker.

Viscoelastic Isolators

As discussed above, the antiresonances were found to move when the system was disassembled and reassembled. As a result, we had difficulties justifying which undamped test the viscoelastic material was comparable to since we had to loosen bolt torques and change other boundary conditions in order to add it to the frame. We repeated viscoelastic and undamped tests




Fig. 14 Comparison of results across several random noise tests with no change to frame parameters

in an attempt to get a better idea of what the factor of improvement would be. During this testing, we also tested a few configurations of the viscoelastic material to determine which was the most effective. The viscoelastic material was tested in two places: 1.) between the shaker and linear slider and 2.) between the shaker beam and the rest of the frame. Additionally, several different bolt torques were used on the bolts attaching the viscoelastic to the frame. Two rounds of testing with viscoelastic are shown below. In the first, the bolts for the frame were tightened to 11 ft-lb everywhere except where the viscoelastic was placed. The bolts used to tighten the viscoelastic were only tightened to 8 ft-lb, as this seemed to give the largest improvement.

As shown in Figure 15, there were two major bands where antiresonance problems were found. Looking at the FRF, the antiresonances occurred at 580, 616.25, and 1357.5 Hz. By visual inspection of the FRF, we can see that the viscoelastic material helped the antiresonances by reducing the depth of the valley in the graph. This can be quantified by finding the RMS value from the PSD, as shown in the three rightmost columns in Table 2. Three RMS values were obtained from each of the PSDs in Figure 15b. The RMS value for the 550-650 Hz band gives the average change over the first two antiresonances, while the 1300-1400 Hz band focuses on the antiresonance at 1357.5 Hz. Additionally, the RMS value is given for the entire frequency range (20-2000 Hz) as the last column in Table 2. The viscoelastic material reduced the RMS voltage at both antiresonances, although the factor of improvement was larger for the higher frequency antiresonance. It can also be seen in Figure 15a that the viscoelastic treatment shifts and damps several modes of the system in the 100-500 Hz region, and seems to generally increase the damping. This frequency band does not produce significant antiresonances in this setup, so this doesn't affect the RMS shaker voltage required. Nevertheless, this does suggest that the sorbothane is improving the dynamics of the system.

Table 2 Comparison of voltage RMS values for a frame with and without Sorbothane padding with frame bolts tightened to 8 and 11 ft-lb

Voltage RMS Over Antiresonance Bands (V^2)						
	8 ft-lb (see Figure 15)			11 ft-lb (see Figure 16)		
	550-650 Hz	1300-1400 Hz	20-2000 Hz	550-650 Hz	1300-1400 Hz	20-2000 Hz
No Treatment	8.415×10^{-7}	4.437×10^{-4}	6.368×10^{-4}	7.556×10^{-7}	8.583×10^{-4}	1.097×10^{-3}
Sorbothane	6.730×10^{-7}	2.018×10^{-4}	5.094×10^{-4}	4.393×10^{-7}	1.571×10^{-4}	3.834×10^{-4}
Improvement	1.252	2.199	1.250	1.720	5.463	2.861

Another set of tests was performed and the results are shown in Figure 16. For these tests, the frame was disassembled and reassembled with all of the bolt torques at 8 ft-lb.

This second set of results demonstrates a problem encountered with viscoelastic material. Though the viscoelastic material still helped at the antiresonance problem bands, the factor of improvement was less than previous tests. The numerical results are shown in the left three columns in table 2. With the lower preload on the bolts, Figure 16a shows that the 650-1300 Hz frequency band has a similar character to that observed in Figure 15a, however, this FRF has a new antiresonance




Fig. 15 Comparison of results across random noise tests with and without Sorbothane pad with 11 ft-lb preload on the frame bolts




Fig. 16 Comparison of results across secondary random noise tests with and without Sorbothane pad with 8 ft-lb preload on the frame bolts

at 1126.3 Hz, which would increase the shaker voltage at this frequency and likely causes issues for controllability. Visually, the modes around 100-500 Hz in Figure 16a were not damped as much as seen in Figure 15a. Additionally, Sorbothane moved many of the resonances and antiresonances making the tests less predictable and harder to compare. We were unable to conclude why this antiresonance appeared in this test and not in the others; further research would be needed to understand why this antiresonance appeared.

Silica Particle Damping

Some of the most interesting and promising results came from testing the silica particle damping-treatment. Figure 17a shows that silica adds significant damping to the FRF between the control accelerometer and the shaker voltage. Many of the extrema are completely smoothed, especially above 700 Hz. Hence, the voltage PSD in Figure 17b is also smoother. Visually, the shaker voltage doesn't have significant peaks at any of the identified antiresonances. Quantitatively, this is reflected in the RMS values over the bands of interest as shown in Table 3. Around the lower antiresonance, the silica decreased the voltage RMS by a factor of 8.67. At the higher antiresonance, the factor of improvement is 20.51. Note that the voltage RMS over the whole frequency range improves by only a factor of 2.65. This is because the RMS shaker voltage over the entire frequency band is dominated by a peak at 60 Hz. If the RMS was instead computed over the 100-2000 Hz band, the RMS voltage decreased by a factor of 4.8 instead of 2.65.




Fig. 17 Comparison of results across random noise tests with and without silica particle damping

Table 3 Comparison of voltage RMS values for a frame with and without silica particle damping

	Voltage RMS Over Antiresonance Bands (V^2)		
	550-650 Hz	1300-1400 Hz	20-2000 Hz
No Treatment	8.415×10^{-7}	4.437×10^{-4}	6.368×10^{-4}
Silica	9.711×10^{-8}	2.1632×10^{-5}	2.399×10^{-4}
Improvement	8.67	20.51	2.65

While this treatment performs well for the antiresonances that are of primary interest, and in general at mid to high frequencies, there remain significant dynamics below 700 Hz. The same was seen in our roving hammer test results (e.g. see Figure 8b) where it was noted that particle damping was less effective below 500 Hz. If a test setup has antiresonances below 500 Hz, one would imagine that it is typically possible to reconfigure the frame to move them up to higher frequencies. Hence, these results suggest that silica particle damping can be an effective damping treatment to counter test frame antiresonances.

Conclusions

This study explored passive damping solutions to mitigate vibration in a frame similar to the NEWT frame [5], which was recently proposed by Sandia National Laboratories to facilitate multi-shaker testing. It was found that using particle damping in the frame, by filling the hollow center of the beam with silica sand, significantly increased the damping in the frame, particularly above 500 Hz. This lowered the amplification of most of the frame's resonances and decreased the power draw of the shaker at the antiresonances. The silica damping treatment is ready to undergo testing in a full-scale MIMO frame with many shakers.

This study also explored other damping treatments. In particular, viscoelastic (Sorbothane) pads placed at the interface between the shaker and frame had varying degrees of success, depending on the configuration. They were generally found to increase the damping in the system, but they also moved resonances and anti-resonances, sometimes increasing the voltage needed by the shaker.

The antiresonances that were found in this study were not easy to predict, and they were found to change noticeably if the system was disassembled and reassembled. Hence, future studies should pay close attention to the configuration; damping is only expected to reduce the shaker voltage requirements when it reduces the severity of an antiresonance. While not elaborated here, several other configurations were studied as part of this work, and the acceleration to voltage FRF for one of these is shown in Figure 18; it does not exhibit any antiresonances and hence the damping treatment didn't change the shaker voltage for that setup.

References

1. Rudder, F. and Plumblee, H.E. "Sonic Fatigue Design Guide for Military Aircraft". Technical report, Flight Dynamics Laboratory, Wright-Patterson AFB, Ohio (1974)ifnextchar.gobble.
2. Foley, J. "Normal and abnormal dynamic environments encountered in truck transportation". *The Shock and Vibration Bulletin*, 39:31 (1969)ifnextchar.gobble.
3. Daborn, P.M., Roberts, C., Ewins, D.J., and Ind, P.R. "Next-Generation Random Vibration Tests". In Allemang, R., editor, *Topics in Modal Analysis II, Volume 8*, Conference Proceedings of the Society for Experimental Mechanics Series, pages 397–410, Cham (2014)ifnextchar.gobble. Springer International Publishing.
4. Schoenherr, T.F. "Derivation of Six Degree of Freedom Shaker Inputs Using Sub-structuring Techniques". In Mains, M. and Dilworth, B.J., editors, *Topics in Modal Analysis & Testing, Volume 9*, pages 5–14. Springer International Publishing, Cham (2019)ifnextchar.gobble. Series Title: Conference Proceedings of the Society for Experimental Mechanics Series.
5. Cross, K., Melendez, G., and Zwink, B. "Natural excitation within a test-frame: Experimental results". Proceedings of IMAC - International Modal Analysis Conference. IMAC (2023)ifnextchar.gobble.
6. Kramer, H., Schultze, J., Danforth, S., and Mann, B. "Dynamic analysis of a modular test stand for multi-axis vibration testing". Proceedings of IMAC - International Modal Analysis Conference. Unpublished Manuscript (2025)ifnextchar.gobble.
7. Ginsberg, J.H. *Mechanical and Structural Vibrations: Theory and Applications*. John Wiley and Sons, New York, 1st edition (2001)ifnextchar.gobble.
8. Lu, Z., Wang, Z., Masri, S.F., and Lu, X. "Particle impact dampers: Past, present, and future". *Structural Control and Health Monitoring*, 25 (2017)ifnextchar.gobble.
9. Lu, Z., Wang, Z., Masri, S.F., and Lu, X. "Particle impact dampers: Past, present, and future". *Structural Control and Health Monitoring*, 25(1):e2058 (2018)ifnextchar.gobble. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/stc.2058>.
10. Hollkamp, J.J. and Gordon, R.W. "Experiments with particle damping". volume 3327 of *Proceedings of SPIE - The International Society for Optical Engineering*, pages 2–12. SPIE (1998)ifnextchar.gobble.
11. Zhou, K. and Li, Q.S. "Vibration mitigation performance of active tuned mass damper in a super high-rise building during multiple tropical storms". *Engineering Structures*, 269 (2022)ifnextchar.gobble.
12. Allen, M.S. and Ginsberg, J.H. "A Global, Single-Input-Multi-Output (SIMO) Implementation of The Algorithm of Mode Isolation and Applications to Analytical and Experimental Data". *Mechanical Systems and Signal Processing*, 20:1090–1111 (2006)ifnextchar.gobble.

Appendix A




Fig. 18 Undamped FRF between DUT acceleration and shaker voltage, with 10 lb mass used as the DUT, demonstrating the lack of antiresonances



Chapter 4

Dynamic Recognition of Damage States in Wind Turbines Using Cepstral Coefficients and ANOVA Analysis

Costanza Speciale, Sina Shid-Moosavi, Eleonora Maria Tronci, and Silvia Milana

Abstract Traditional damage detection methods for dynamic systems often use regression or classification algorithms to distinguish between damaged and intact states, but they typically require case-specific parameter tuning. This need for manual adjustment limits their generalizability across different structures and varying operational conditions, reducing diagnostic accuracy. For complex systems like wind turbines, varying conditions, such as wind speed, temperature, and load demands, can complicate the choice of optimal parameters for reliable damage-sensitive features, leading to false positives and inconsistent diagnostics. This work leverages a semi-supervised damage detection algorithm based on a strategy combining Linear Discriminant Analysis (LDA) and Probabilistic Linear Discriminant Analysis (PLDA) with cepstral coefficients for damage identification and classification in wind turbines. While LDA and PLDA are traditionally supervised, this approach integrates them into an unsupervised framework that dynamically recognizes and categorizes emerging damage states as they appear, rather than relying on predefined training phases. ANOVA-based sensitivity analysis is incorporated to evaluate how user-defined parameters impact the performance of the classification outcomes, guiding an adaptive tuning strategy that ensures consistent detection across varied operational conditions. The proposed damage detection algorithm is validated using experimental data from the Aventa AV-7 wind turbine, assessing its ability to recognize and classify multiple damage states under varied operational conditions. ANOVA-based sensitivity analysis identifies key parameters for optimal performance, demonstrating the algorithm's adaptability and robustness. This validation confirms the framework's effectiveness in achieving consistent and reliable diagnostics across real-world environmental variability.

Keywords Semi-supervised damage detection · Cepstral coefficients · ANOVA · Sensitivity analysis · Wind turbine diagnostics

Introduction

Conventional approaches to damage detection in dynamic systems often employ regression or classification techniques to distinguish between healthy and damaged states. While effective in controlled cases, these methods usually rely on manually tuned parameters optimized for specific scenarios, which limits their ability to generalize across different structures and varying operational conditions. In the context of wind turbines, where external factors such as wind speed, temperature, and load demand fluctuate continuously, parameter sensitivity poses a significant challenge. In practice, this dependence can lead to unreliable extraction of damage-sensitive features, resulting in false alarms and inconsistent diagnostic performance.

A key challenge in Structural Health Monitoring (SHM) for wind turbines is the scarcity of labeled data representing damaged states. Realistic failure cases are rare, difficult to reproduce, and often poorly documented, making supervised

Costanza Speciale · Silvia Milana

Department of Mechanical and Aerospace Engineering, Sapienza University of Rome, Rome, Italy

e-mail: costanza.speciale@uniroma1.it; silvia.milana@uniroma1.it

Sina Shid-Moosavi

Department of Civil and Environmental Engineering, Northeastern University, Boston, MA, USA

e-mail: shidmoosavi.s@northeastern.edu

Eleonora Maria Tronci

Department of Civil and Urban Engineering, Center for Urban Science and Progress, New York University, New York, NY, USA

e-mail: emt377@nyu.edu

learning approaches impractical in many scenarios [1, 2]. As a result, unsupervised methods such as novelty detection have been widely applied, where a baseline model of the healthy state is used to flag significant deviations as potential indicators of damage [3, 4, 5, 6]. While effective for distinguishing healthy from damaged conditions, these approaches are generally limited to binary classification and often lack mechanisms to adapt when new data becomes available. For dynamic systems such as wind turbines, this creates a critical need for more flexible methods capable of evolving with changing operational and environmental conditions.

As described in the previous chapter, features such as cepstral coefficients can be extracted to characterize the dynamic response of wind turbines. Building on this, the present chapter focuses on a semi-supervised damage detection framework that integrates cepstral features with a hybrid classification strategy combining Linear Discriminant Analysis (LDA) for feature compression and Probabilistic Linear Discriminant Analysis (PLDA) for adaptive classification, enabling the system to learn from healthy operational data and to detect emerging damage scenarios without requiring labeled fault data. Although LDA and PLDA are traditionally applied in supervised learning contexts, they are embedded here within an unsupervised structure that enables adaptive identification and categorization of emerging damage states without reliance on fixed training datasets.

One of the main limitations of existing methods is the static nature of parameter selection, which typically requires offline sensitivity analysis. To address this issue, an exploratory Reinforcement Learning (RL)-based strategy is proposed to automate and optimize the tuning of parameters in cepstral-based diagnostic pipelines. In this proposed framework, an RL agent would interact with the classification system to adjust parameter settings dynamically, guided by a reward function designed to balance detection accuracy, false alarm rates, and robustness under variable operating conditions. Although not yet fully developed, this concept highlights the potential of dynamic optimization to enable continuous learning from new data in the field, thereby reducing reliance on predefined thresholds and improving adaptability to real-world variability.

The effectiveness of the proposed framework is evaluated using experimental data from the Aventa AV-7 wind turbine. Comparisons will be made between static parameter configurations and those adaptively tuned by the RL agent under diverse operational and damage scenarios. This chapter thereby addresses the limitations identified in conventional methods, advancing toward a more reliable and adaptive solution for structural health monitoring in wind farm applications.

Dataset Description

This research utilizes data from the Aventa AV-7 [7], a 7kW wind turbine operated by ETH Zurich (Figure 1), which has been extensively monitored to evaluate SHM methods as the turbine approaches the end of its design life. Since 2020, sensors have been deployed throughout the turbine to facilitate system identification, modal analysis, and failure detection studies, with data collected under four specific operational scenarios: normal operation, pitch system failure, aerodynamic imbalance, and rotor icing. For each condition, datasets include time series measurements recorded over specified dates between 2022 and 2023.



Fig. 1 Aventa AV-7 [7]

The Aventa AV-7 turbine operates with a variable-speed, variable-pitch control, beginning energy generation at a 2 m/s wind speed and ceasing at 14 m/s for a rated power of 7kW. It has a three-blade rotor with a 12.9 m diameter, positioned 18 m above ground on a concrete and steel tower. A belt-driven generator and a frequency converter enable variable-speed operation.

The turbine is outfitted with eleven accelerometers distributed along the tower, nacelle, main bearing, and generator, as well as two full-bridge strain gauges at the tower base for measuring axial and bending strains. All data is recorded at 200Hz, while SCADA (Supervisory Control and Data Acquisition) data, such as wind speed, rotor speed, and power output, is recorded at 7Hz, with temperature and humidity recorded at 1Hz. The accelerometers are paired and located at specific tower sections to capture the x - and y -axis accelerations aligned with wind direction, allowing comprehensive monitoring of the turbine's structural response.

The SCADA system provides insights into power, wind speed, rotor speed, and yaw angle, which are critical for analyzing turbine performance, especially under "normal operation" conditions. The power curve, derived from these SCADA readings, illustrates the relationship between wind speed and power output, with distinct operational regions (cut-in, rated, and cut-off speeds) [8].

The dataset further captures three specific fault conditions, providing valuable data for condition monitoring methods.

In February 2022, a failure occurred in the flexible coupling of the pitch drive system, which is critical for adjusting blade angles to optimize performance under changing wind conditions. SCADA data captured deviations in rotor speed, power output, and yaw position compared to normal operation. Before the complete failure, the pitch system became stuck, reducing the turbine's output to 1-2 kW while the rotor speed dropped to around 40 rpm. An emergency stop was triggered to prevent further damage, as indicated by sharp changes in SCADA measurements at approximately 270 seconds. Acceleration data from sensors across the tower and nacelle show a significant structural response to the failure, highlighting its impact on the turbine's dynamics.

In December 2022, rotor icing caused a near-complete halt in turbine operations, reducing rotor speed to almost zero and stopping power generation. This highlights the importance of early detection in cold climates where ice accumulation can significantly reduce aerodynamic efficiency and increase vibration.

Then, data collected after roughness tape was applied to one blade in December 2022 revealed imbalances that affected rotor speed and induced mechanical stress on key turbine components, underscoring the need to identify imbalances early to prevent failures.

This research focuses on two of these damage conditions: pitch drive failure and aerodynamic imbalance. The first corresponds to a sudden and critical rupture in the turbine's system, and is particularly challenging for SHM due to the need for immediate detection and intervention to prevent further damage. By concentrating on this sudden event, the study aims to improve SHM techniques' ability to quickly identify and respond to rapid, unforeseen failures, which are particularly difficult to manage.

The second, in contrast, typically develops gradually. It can be induced, for instance, by surface degradation or roughness accumulation on the blades. Although less abrupt than a pitch drive failure, it is still important to detect it promptly, as it increases loads and vibrations throughout the drivetrain and tower, potentially leading to more severe failures over time. Moreover, it also negatively impacts power production.

The dataset therefore provides a key resource for developing more effective methods to address these two conditions, selected as representative examples of the broader categories of sudden and gradual failures that typically characterize the operational life of wind turbines.

Semi-Supervised Damage Detection Framework

This section provides an overview of the framework employed for damage detection. After cepstral coefficients are extracted, they are used as damage-sensitive features within a statistical pattern recognition framework that integrates LDA and PLDA. The goal is to reduce the dependence on manually tuned parameters and enable adaptive identification of damage states under varying operational conditions.

LDA [9] is a supervised machine learning method commonly applied for dimensionality reduction and feature extraction. Its aim is to project high-dimensional data into a lower-dimensional latent space that maximizes class separability. This is achieved by identifying projection directions that maximize the distance between class means (between-class scatter) while minimizing variance within each class (within-class scatter) [10].

Let the training dataset consist of N samples $\{\mathbf{x}_n\}_{n=1}^N$, where each observation $\mathbf{x}_n \in \mathbb{R}^{D \times 1}$ belongs to one of K known classes. The algorithm begins by computing the class-wise means $\mathbf{m}_k$, the global mean $\mathbf{m}$, the within-class scatter matrix $\mathbf{S}_W$, and the between-class scatter matrix $\mathbf{S}_B$. The optimal projection matrix is obtained by solving the generalized

eigenvalue problem:

$$\mathbf{S}_B \mathbf{w} = \lambda \mathbf{S}_W \mathbf{w} \quad (1)$$

The eigenvectors $\{\mathbf{w}_i\}$ define the directions of maximum separability, and their eigenvalues $\{\lambda_i\}$ quantify the discriminative strength of each direction. The g leading eigenvectors are selected to construct the LDA projection matrix, which maps the original features into a reduced space optimized for classification.

A summary of the full LDA procedure is provided in Algorithm 1.

Algorithm 1 Linear Discriminant Analysis

Require: Training data $\{\mathbf{x}_i, y_i\}_{i=1}^N$ with K classes

- 1: Compute class means: $\mathbf{m}_k = \frac{1}{n_k} \sum_{\mathbf{x}_i \in C_k} \mathbf{x}_i$
 - 2: Compute global mean: $\mathbf{m} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i$
 - 3: Compute within-class scatter: $\mathbf{S}_W = \sum_{k=1}^K \sum_{\mathbf{x}_i \in C_k} (\mathbf{x}_i - \mathbf{m}_k)(\mathbf{x}_i - \mathbf{m}_k)^T$
 - 4: Compute between-class scatter: $\mathbf{S}_B = \sum_{k=1}^K n_k (\mathbf{m}_k - \mathbf{m})(\mathbf{m}_k - \mathbf{m})^T$
 - 5: Solve generalized eigenvalue problem: $\mathbf{S}_B \mathbf{w} = \lambda \mathbf{S}_W \mathbf{w}$
 - 6: Select top g eigenvectors: $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_g]$
 - 7: **return** LDA projection matrix $\mathbf{W}$
-

The first step in the framework is training the classification model using data representing the various operational classes of the undamaged state. In this study, 70% of the undamaged data were randomly selected for training, while the remaining 30% were reserved for testing. Figure 2 illustrates the overall process.




Fig. 2 Training phase flowchart

Classification is performed using a PLDA model [11], a probabilistic extension of LDA that provides robustness in complex classification scenarios, including the recognition of previously unseen classes. Unlike conventional LDA, which classifies based on distance to centroids in the latent space, PLDA models both samples and class centers as Gaussian realizations. This allows likelihood-based evaluation of class membership, enabling the model to identify new patterns beyond the training set. Such flexibility is especially important for SHM, where operational variability and emerging damage states are expected.

The PLDA framework assumes that each observation $\mathbf{x}$ and its class center $\mathbf{z}$ are generated from latent variables $\mathbf{u}$ and $\mathbf{v}$ through linear mappings $\mathbf{x} = \mathbf{m} + \mathbf{A}\mathbf{u}$ and $\mathbf{z} = \mathbf{m} + \mathbf{A}\mathbf{v}$, with $\mathbf{u} \sim \mathcal{N}(\mathbf{v}, \mathbf{I})$ and $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Psi})$. Model parameters $\{\mathbf{m}, \mathbf{A}, \mathbf{\Psi}\}$ are estimated from the latent representations derived via LDA, as summarized in Algorithm 2.

Algorithm 2 Probabilistic Linear Discriminant Analysis Training

Require: Projected training data $\{\mathbf{u}_i\}$ from LDA

- 1: Estimate global mean $\mathbf{m}$ and projection matrix $\mathbf{A}$
 - 2: Assume latent class means $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Psi})$
 - 3: **for** each training class k **do**
 - 4: Compute PLDA scores: $\log P(\mathbf{u}_p | \mathbf{u}_1^k, \dots, \mathbf{u}_{n_k}^k)$
 - 5: Fit a GEV distribution to scores
 - 6: Set threshold at 95th percentile
 - 7: **end for**
 - 8: **return** Parameters $\{\mathbf{m}, \mathbf{A}, \mathbf{\Psi}\}$ and score thresholds
-

Once trained, the model computes a PLDA score for each new observation, corresponding to the log-likelihood of class membership. Thresholds are established by fitting score distributions with a Generalized Extreme Value (GEV) model [12] and setting cutoffs at GEV_{Th} , ensuring robustness against variability and facilitating detection of previously unseen conditions.

The testing phase evaluates both undamaged (remaining 30%) and damaged data. Each data point, corresponding to a set of cepstral coefficients from one signal frame, is processed sequentially, making the method suitable for real-time SHM.




Fig. 3 Testing phase flowchart

For each frame, cepstral coefficients are projected into the latent space defined by LDA, ensuring dimensionality reduction while preserving discriminative features. The projected vector is evaluated using the trained PLDA model, which computes likelihood scores. Classification proceeds as follows: (i) if a point matches one class above its threshold, it is assigned to that class and thresholds are updated once a counter N_{UP} is reached; (ii) if a point fails to meet any threshold, it is used to define a new class once counter N_{NC} is reached; (iii) if a point exceeds thresholds for multiple classes, it is left unassigned.

The choice of N_{UP} and N_{NC} is critical, as they govern the adaptability of the framework. To guide parameter selection, two quality indices are defined. These indices provide a quantitative measure of the framework's performance:

$$I_{BIN} = \frac{\text{number of correctly classified data}}{\text{total predictions}} \quad (2)$$

$$I_{MULT} = \frac{\text{number of correctly classified data}}{\text{total predictions}} \cdot \frac{1}{1 + \text{number of missed or extra classes}} \quad (3)$$

The binary index I_{BIN} evaluates the accuracy of distinguishing between damaged and undamaged states, while the multi-class index I_{MULT} extends this evaluation to multiple damage scenarios. The latter penalizes cases in which additional spurious classes are created or true classes are not identified, thereby encouraging both accuracy and stability in the classification results.

Sensitivity Analysis

After presenting the framework, it was applied to the experimental data from the Aventa AV-7 turbine to evaluate wind turbine damage detection. As anticipated, for the training phase only undamaged data from normal turbine operation were used. These data were partitioned into four classes according to wind speed and rotor speed values, each representing a distinct operating region. For the testing phase, instead, damaged data were also included. Specifically, four additional classes were defined: three derived from measurements collected on the day of the pitch failure, corresponding to different damaged states (incipient failure, stuck pitch drive, and complete failure), and one representing the aerodynamic imbalance condition. Within this framework, several parameters were identified as influential: N_{CC} , the number of cepstral coefficients used as input features; N_{LDA} , the number of LDA components retained; GEV_{Th} , the percentile threshold derived from the GEV distribution; and the counters N_{UP} and N_{NC} , which regulate class updating and new-class creation, respectively, and can be further tuned during RL optimization.

In previous work [13], it was observed that the two input parameters N_{NC} and N_{UP} significantly affect the algorithm's performance. An analysis was carried out by performing multiple simulations with different combinations of these parameters. As shown in Figure 4, the mean and standard deviation of the two performance indices, I_{BIN} and I_{MULT} , vary substantially across the different cases.




Fig. 4 Mean and standard deviation (red lines) of the binary and multi-class performance indices [13]

To determine which parameters are most critical and to investigate their potential interactions, a factorial Analysis of Variance (ANOVA) was performed. The analysis was conducted using both the binary index I_{BIN} and the multi-class index I_{MULT} as performance metrics. A full factorial design was employed to span the parameter space, testing all possible combinations of factor levels. The factor levels are summarized in Table 1. The total number of experimental runs is equal to the product of the level counts, yielding 5,040 unique parameter combinations.

Table 1 Factor levels and number of levels in the full factorial design

Factor (Parameter)	Levels	No. of Levels
N_{CC} (Number of cepstral coefficients)	6, 10, 15, 20, 30, 50, 75, 100	8
N_{LDA} (Number of LDA components)	1, 2, 3, 4, 5	5
GEV_{Th} (Percentile threshold)	0.85, 0.90, 0.95	3
N_{UP} (Update counter)	5, 10, 15, 20, 25, 30	6
N_{NC} (New-class counter)	10, 15, 20, 25, 30, 35, 40	7

The ANOVA was conducted using the `pg.anova` function from the `Pinguin` Python package [14, 15]. Effect sizes were quantified as partial eta-squared (η_p^2), which represents the proportion of variance explained by a factor after accounting for other sources of variation. For a one-way ANOVA, the decomposition of sums of squares is defined as:

$$\begin{aligned} SS_{\text{total}} &= SS_{\text{effect}} + SS_{\text{error}}, \\ SS_{\text{total}} &= \sum_i \sum_j (Y_{ij} - \bar{Y})^2, \quad SS_{\text{effect}} = \sum_i n_i (\bar{Y}_i - \bar{Y})^2, \\ SS_{\text{error}} &= \sum_i \sum_j (Y_{ij} - \bar{Y}_i)^2 \end{aligned} \quad (4)$$

with the F-statistic computed as:

$$F^* = \frac{MS_{\text{effect}}}{MS_{\text{error}}} = \frac{SS_{\text{effect}}/(r-1)}{SS_{\text{error}}/(n_t-r)}, \quad (5)$$

where r is the number of groups and n_t is the total sample size. Partial eta-squared is then given by:

$$\eta_p^2 = \frac{SS_{\text{effect}}}{SS_{\text{effect}} + SS_{\text{error}}}. \quad (6)$$

The η_p^2 values serve as sensitivity indices, providing a quantitative ranking of parameter contributions to variability in the output. This methodology follows best practices for reporting effect sizes in experimental design [16, 17]. Results for both I_{BIN} and I_{MULT} are presented in Figure 5. The analysis shows that N_{LDA} is the most influential parameter, followed by GEV_{Th} and N_{CC} . Importantly, part of the sensitivity of all parameters arises from interaction effects with other parameters, underscoring the need to account for parameter correlation rather than evaluating each factor in isolation. For instance, the effect of N_{CC} is strongly tied to the number of LDA components retained, as the discriminatory power of additional cepstral coefficients can only be realized if sufficient LDA dimensions are available. Similarly, the influence of GEV_{Th} depends on the interplay with both N_{UP} and N_{NC} , since the threshold for defining new classes and updating existing ones directly shapes how probabilistic scores are interpreted in practice.




Fig. 5 Sensitivity analysis results for framework parameters based on binary and multi-class performance indices

To further explore these relationships, pairwise interaction effects were visualized using sensitivity matrices, shown in Figure 6. These matrices display η_p^2 values for both main effects and two-way interactions across all parameter pairs. The binary index highlights strong interaction effects between N_{LDA} and both N_{CC} and GEV_{Th} , indicating that the influence of cepstral resolution and threshold settings depends heavily on how many LDA dimensions are retained. In the multi-class case, the same pairs also show moderate interaction effects, though GEV_{Th} emerges as more influential overall than N_{CC} .

These findings reinforce that parameter interactions, not just main effects, play a decisive role in shaping performance. Therefore parameter optimization for damage detection cannot be approached as a purely one-dimensional tuning exercise. Instead, parameter sets must be tuned collectively, as the performance of the framework emerges from the balance and



Fig. 6 Sensitivity matrices showing partial η_p^2 for main and pairwise interaction effects across all parameters on binary and multi-class performance indices

interaction among them. This interdependence further motivates the adoption of adaptive optimization strategies, such as reinforcement learning, which can dynamically adjust multiple parameters simultaneously in response to evolving operating conditions and data characteristics. Such an approach offers the potential to avoid suboptimal configurations that might result from static or independent parameter tuning, thereby enhancing the robustness and reliability of the detection system.

Conclusions

This work presented a semi-supervised framework for damage detection in wind turbines that integrates cepstral coefficients with LDA and PLDA. The proposed approach addresses key limitations of conventional supervised and unsupervised strategies by enabling dynamic recognition of damage states without relying on large labeled datasets or static parameter settings. Through application to experimental data from the Aventa AV-7 wind turbine, the framework demonstrated its ability to classify different failure modes, capturing the variability of real-world operating conditions.

A sensitivity analysis based on ANOVA quantified both single-parameter and interaction effects. In isolation, performance is dominated by the number of LDA components (N_{LDA}), followed by the GEV percentile threshold (GEV_{Th}) and the number of cepstral coefficients (N_{CC}), with the update and new-class counters (N_{UP} , N_{NC}) exerting smaller but non-negligible influence. At the same time, several interaction terms are significant—most notably $N_{LDA} \times N_{CC}$ and $N_{LDA} \times GEV_{Th}$ (and, to a lesser extent, $GEV_{Th} \times N_{UP}/N_{NC}$)—highlighting that multi-parameter tuning is necessary and providing clear directions for adaptive optimization.

Overall, the presented methodology advances structural health monitoring practices by combining interpretable feature extraction with adaptive classification and systematic parameter evaluation. Future work will focus on integrating an automated tuning strategy for online parameter adjustment, extending validation to larger datasets, and additional failure types.

Acknowledgments The authors would like to sincerely thank ETH Zurich and its collaborators for providing the “Aventa AV-7” dataset, which was invaluable for our study.

References

1. Farrar, C.R. and Worden, K. *Structural health monitoring: a machine learning perspective*. John Wiley & Sons (2012)
2. Regan, T., Beale, C., and Inalpolat, M. “Wind turbine blade damage detection using supervised machine learning algorithms”. *Journal of Vibration and Acoustics*, 139(6):061010 (2017)

3. Eltoumy, K., Gomaa, M., and Liang, X. “Unsupervised learning methods for data-driven vibration-based structural health monitoring: A review”. *Sensors*, 23(6) (2023)
4. Castellani, F., Garibaldi, L., Daga, A.P., Astolfi, D., and Natili, F. “Diagnosis of faulty wind turbine bearings using tower vibration measurements”. *Energies*, 13(6) (2020)
5. Movsessian, A., García Cava, D., and Tcherniak, D. “An artificial neural network methodology for damage detection: Demonstration on an operating wind turbine blade”. *Mechanical Systems and Signal Processing*, 159:107766 (2021)
6. Feijóo, M.d.C., Zambrano, Y., Vidal, Y., and Tutivén, C. “Unsupervised damage detection for offshore jacket wind turbine foundations based on an autoencoder neural network”. *Sensors*, 21(10) (2021)
7. Chatzi, E., Abdallah, I., Hofsäß, M., Bischoff, O., Barber, S., and Marykovskiy, Y. “Aventa AV-7 ETH Zurich Research Wind Turbine SCADA and high frequency Structural Health Monitoring (SHM) data” (2023)
8. Bilendo, F., Meyer, A., Badihi, H., Lu, N., Cambron, P., and Jiang, B. “Applications and modeling techniques of wind turbine power curve for wind farms—a review”. *Energies*, 16(1):180 (2022)
9. Xanthopoulos, P., Pardalos, P. M., and Trafalis, T. B. *Robust Data Mining*. Springer (2013)
10. Tharwat, A., Gaber, T., Ibrahim, A., and Hassanien, A.E. “Linear discriminant analysis: A detailed tutorial”. *Ai Communications*, 30:169–190, (2017)
11. Ioffe, S. “Probabilistic linear discriminant analysis”. In Leonardis, A., Bischof, H., and Pinz, A., editors, *Computer Vision – ECCV 2006*, pages 531–542, Berlin, Heidelberg (2006) Springer Berlin Heidelberg.
12. Kotz, S. and Nadarajah, S. *Extreme value distributions: theory and applications*. World Scientific (2000)
13. Speciale, C., Tronci, E.M., Betti, R., and Milana, S. “Cepstral coefficients for wind turbine damage detection: Recognizing unknown classes in real-time”. Available at SSRN 5173243.
14. Vallat, R. “Pingouin: statistics in python”. *Journal of Open Source Software*, 3(31):1026 (2018)
15. Team, P.D. “pingouin.anova documentation”. <https://pingouin-stats.org/build/html/generated/pingouin.anova.html> Accessed September 5, 2025.
16. Olejnik, S. and Algina, J. “Generalized eta and omega squared statistics: Measures of effect size for some common research designs”. *Psychological Methods*, 8(4):434–447 (2003)
17. Lakens, D. “Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and anovas”. *Frontiers in Psychology*, 4:863 (2013)



Chapter 5

Enhanced Bearing Fault Diagnosis Using Convolutional Neural Networks and Signal Preprocessing Techniques on the CWRU Dataset

M. Conti, G. Battista, S. Pavoni, and M. Vanali

Abstract Convolutional Neural Networks (CNN) are increasingly being used in the diagnostics of industrial machines, rotors, gears and bearings. Originally adopted for image recognition, they have been appreciated for their ability to extract important features from data. Both model-based and data-driven methods are extensively present in the literature for bearing diagnostics. In this paper, the symbiosis of the two to enhance the classification of faults is proposed. As deep learning method, LeNet-5, which is a kind of CNN architecture, is used and improved in this paper to obtain better accuracy results. Raw data provided by accelerometers has also been pre-processed with different mathematical models aimed at extracting the content linked to the fault from the signal. Envelope Analysis is the main technique employed for this target and can be paired with preprocessing models as Autoregressive Filter or Cepstrum Pre-Whitening. Another solution has been the use of a Median Filter paired with differentiation of the second order of the acceleration signal, to highlight impulses due to damage. Case Western Reserve University dataset has been taken as a reference to analyze the performance of the improved LeNet-5 network in classifying bearing faults in different conditions, such as speed and fault size, with different preprocessing. Two different fault sizes have been compared in order to evaluate the net ability to recognize the same fault at different sizes.

Keywords Convolutional Neural Networks · bearings · diagnostic · data preprocessing.

Introduction

Bearings are crucial components within rotary machines. They perform a dual function of supporting and enabling mutual rolling between the rotating parts. They consist of three objects: the inner race, the outer race and the rolling bodies, which can be balls or rollers. Each of the two races can be subjected to surface wear that manifests itself in small cracks. A third type of damage is called brinelling, which occurs when excessive loads on the bearing cause plastic deformation of the contact area between the rolling element and either the outer or inner race. When one of these faults happens, pulses are generated. These specific pulses can be detected in the vibration at particular frequencies that depend on the geometric characteristics of the bearing and the speed of the shaft on which it is mounted.

In recent years, in addition to model-based methods, data-driven machine learning and even more advanced deep learning methods have gained popularity to extract features from increasingly complex signals and map these features to bearing fault types. The employment of Neural Networks (NN) has been proven to be particularly useful in classifying the type of damage [1]. One of the most promising techniques is the Convolutional Neural Networks (CNN), which has shown strong effectiveness in this field. This type of NN is a class of deep learning based on convolutional layers, fully connected layers and a classifier [2]. Input given to CNN can be raw vibration signals or a pre-processed one.

In this paper, the use of various preprocessing strategies for improved classification of bearing faults is investigated. The techniques that, according to the literature, have shown the most promising results in clearly identifying the characteristic frequencies of bearing damage have been investigated. These techniques exploit either single processing approaches or the

M. Conti · G. Battista · S. Pavoni · M. Vanali

Università di Parma – Parco Area Delle Scienze, 181/A – 43124 Parma

e-mail: marcello.vanali@unipr.it; gianmarco.battista@unipr.it; stefano.pavoni@unipr.it; marcello.vanali@unipr.it

Table 1 Characteristics bearing fault frequencies for each velocity employed in the investigated dataset

	Fault Frequency [Hz]			
	1725 rpm	1750 rpm	1772 rpm	1797 rpm
BPFO	103.36	104.56	105.87	107.36
BPFI	156.13	157.94	159.93	162.19
BSF	135.91	137.48	139.21	141.16

combination of multiple methods, such as the envelope analysis [3], the Auto Regressive (AR) model [4], the Cepstrum Pre-Whitening (CPW) [5] and the Median Filter (MF) [6]. For comparison reasons, the same CNN network model is adopted, inspired by the LeNet-5 architecture with a few improvements [7]. The data employed in this work is the Case Western Reserve University (CWRU) dataset [8], which has been extensively treated in Bearing Diagnostics [9]. To better understand the improvement of the considered preprocessing techniques, LeNet5 models were initially trained using simple input data, such as the raw signal in the time domain and the power spectrum in the frequency domain. Then, the same strategy has been employed by providing pre-processed data as input, using common methods to highlight the characteristic components of bearing damage in the spectrum.

Results are compared in terms of accuracy in classifying the bearing faults. The different types of preprocessing are briefly explained in the following sections.

Experimental Data and Theory Background

When a fault occurs in a bearing, it generates a series of impulsive excitations caused by the interaction between the rolling elements and the localized defect. These excitations propagate through the structure and become evident in the vibration response. These distinctive behaviours manifest in the vibration signal at specific frequencies, which are determined by the bearing's geometry and the rotational speed of the shaft.

The characteristics frequencies associated to the bearings components are the following

$$BPFO = \frac{nf_r}{2} \left(1 - \frac{d}{D} \cos \theta \right) \quad (1)$$

$$BPFI = \frac{nf_r}{2} \left(1 + \frac{d}{D} \cos \theta \right) \quad (2)$$

$$BSF = \frac{Df_r}{2d} \left(1 - \left(\frac{d}{D} \cos \theta \right)^2 \right) \quad (3)$$

where $BPFO$ is Ball Pass Frequency Outer, $BPFI$ the same for inner race and BSF Ball Spin Frequency, n is the number of rolling bodies, f_r is the shaft frequency, d is the diameter of the rolling body, D is the primitive diameter of the bearing and θ is the contact angle between race and rolling body [10].

The experimental data are taken from the CWRU, which is an open-source dataset commonly used as a reference for developing new bearing faults detection techniques. The dataset is sampled at 48 kHz with different types of bearing damage (inner race, outer race faults and brinelling) at multiple damage levels (7 mils, 14 mils and 21 mils), operating at four various speeds (about 1725, 1750, 1772 and 1797 rpm). The characteristic fault frequencies for each case are shown in Table 1.

In the following section, the data processing techniques employed in this work to enhance the bearing fault diagnosis using a LeNet5 architecture are introduced.

LeNet5 Input Data

The experimental data used in this work are taken from the CWRU bearing dataset. It consists of the measurements of two accelerometers placed on the fan end and drive end bearing of the electrical motor benchmark. The vibration signals are sampled at a frequency of 48 kHz and each recording lasts for some seconds. Since the duration of the measurements is relatively short (5/10 seconds), in order to obtain a sufficiently large dataset to reliably train the networks, a compromise must be reached between the temporal length of the signal in the time domain and the frequency range in the spectrum. Different dimensions of the input data were tested to evaluate the best possible combination, resulting in the following outcome:

- **Time domain:** Signal segments of 1024 samples, corresponding to approximately 0.02 seconds, were selected as the input data length. This duration was considered sufficient to capture in the raw signal the information related to the characteristic bearing frequencies, which in each case are greater than 50 Hz (0.02^{-1} s).
- **Frequency domain:** Power spectra were computed using 1-second signal segments, providing a frequency resolution of 1 Hz. The first 1024 elements of the spectrum (covering the frequency range 0-1023 Hz) were then employed as input data, as this range is sufficient to capture the bearing fault characteristics.

Several studies have proven that the damaged signal is not perfectly periodic but cyclo-stationary. This is due to the fact that the rolling element, upon encountering the crack, tends to slip, changing the temporality of the impulses [11]. For this reason, methods were sought in the literature to isolate the deterministic component from the signal, i.e. linked to the normal periodic operation of the machine. Below are listed and explained the procedures employed to obtain the input data of the neural networks. All the methods proposed, except for the raw signal and the simple power spectrum, aim to better identify the characteristic component of the damage.

Raw Data

Segments of the signal were extracted in the time domain by selecting 1024 consecutive samples, with the initial point of each segment chosen randomly. This approach allows for a diverse representation of the signal characteristics while ensuring that each segment contains sufficient information to capture the dynamics and features of the bearing vibrations.

Power Spectrum

The input dataset for the Power Spectrum (PS) was obtained by segmenting the signal into 1-second sections, with the initial point of each segment chosen randomly. For each section, the Fast Fourier Transform (FFT) was computed and multiplied by its complex conjugate as follows

$$S_{xx}(f) = X(f) \cdot X^*(f) \quad (4)$$

where $X(f)$ is the FFT of the segment and $X^*(f)$ its complex conjugate.

Then, the first 1024 elements of the resulting PS were selected as input data.

Simple Envelope Analysis

Envelope analysis aims to demodulate the signal to extract the components related to the fault [12]. Mathematically, demodulation is performed by evaluating the absolute value of the discrete-time analytic signal, exploiting the Hilbert transform of the signal as follows

$$x_e(t) = |x(t) + j\tilde{x}(t)| \quad (5)$$

where $x(t)$ is the raw signal and $\tilde{x}(t)$ is its Hilbert transform [13].

Then, the PS of $x_e(t)$ is evaluated as in Section 3.2.

Autoregressive (AR) Filter

The AR model, as described in [4], predicts the next value in a time series by using a series of previous observations. Given a signal in the time domain $x(n)$ with $n = 1, 2, \dots, N$, the predicted value $x_p(n)$ at the n -th time step is calculated by the following equation

$$x_p(n) = - \sum_{k=1}^p a(k)x(n-k) \quad (6)$$

where $x_p(n)$ is a weighted sum of its p previous values in the sequence, called order of the model, $a(k)$ is the weighting coefficient, and $x(n-k)$ is the n -kth data.

The deterministic signal $x_p(n)$ could be seen also in a matrix form as $x_p = -Xa$. The matrix X is an $N \times p$ matrix based on the past series values of x . The coefficient $a(k)$ can be obtained by solving the Yule-Walker equation [14]

$$\begin{bmatrix} r_{xx}(0) & r_{xx}(-1) & \dots & r_{xx}(-p+1) \\ r_{xx}(1) & r_{xx}(0) & \dots & r_{xx}(-p+2) \\ \dots & \dots & \dots & \dots \\ r_{xx}(p-1) & r_{xx}(p-2) & \dots & r_{xx}(0) \end{bmatrix} \begin{bmatrix} a(1) \\ a(2) \\ \dots \\ a(p) \end{bmatrix} = - \begin{bmatrix} r_{xx}(1) \\ r_{xx}(2) \\ \dots \\ r_{xx}(p) \end{bmatrix} \quad (7)$$

where $r_{xx}(k)$ is the autocorrelation function of $x(n)$, which is defined by

$$r_{xx}(k) = \frac{1}{N} \sum_{n=1}^N x(n)x(n-k), \quad k = 0, \dots, p-1 \quad (8)$$

The residual signal is estimated using a p order of the AR model. There are several criteria for choosing the optimal solution of order p . To select the proper p order, the Kurtosis is an excellent parameter to indicate the presence of pulses within the signal that are often associated with damage, such as in the case of bearings [15]. In this paper, an order $p = 10$ is selected because the Kurtosis of the residual signal begins to settle.

Once the deterministic signal is estimated by this process, the residual signal is given by subtracting $x_p(n)$ to $x(n)$.




Fig. 1 Scheme of the AR model processing

The resulting residual signal was finally processed with the envelope analysis and the PS using the method presented in Section 3.2. We decided not to use spectral kurtosis to identify a bandpass filter in which the impulse content, and therefore the kurtosis, was maximum because it did not lead to better results.

Cepstrum Pre-Whitening (CPW)

In rotary machines, there are elements such as shafts, gears meshing, and motors that generate periodic but non-sinusoidal signals that are given by the sum of families of multiple harmonics of a fundamental frequency. Using the Cepstrum Transform (CT), it is possible to group these families through a single peak at a given quefrency. The CPW [5] is very easy to implement and very useful to automate bearing diagnostics since it does not require any parameters. It is based on the real cepstrum as follows

$$C(T) = \mathcal{F}^{-1} \{ \log [|\mathcal{F}(x(t))|] \} \quad (9)$$

where $x(t)$ is the raw signal, $\mathcal{F}$ and $\mathcal{F}^{-1}$ indicate the Fourier Transform and its inverse operation and T is the new domain of the CT that is called Quefrency. The CPW exploits the CT characteristics to separate the cyclostationary component of the signal from the periodic one, highlighting the components represented by the bearing faults. The process consists in setting a zero value for the whole real cepstrum. Then the signal is transformed back in the frequency domain and is recombined with the phase of the original signal and inverse transformed to the time domain. Making this operation the bearing signal that is not periodic but cyclo-stationary of the second order will not be filtered and will remain as a residual signal. The residual signal $x_{cpw}(t)$ obtained with CPW can be implemented by the following equation

$$x_{cpw}(t) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}(x(t))}{|\mathcal{F}(x(t))|} \right\} \quad (10)$$

where $x(t)$ is the raw signal, $\mathcal{F}$ and $\mathcal{F}^{-1}$ indicate the Fourier transform and its inverse operation. Finally, the resulting signal $x_{cpw}(t)$ was processed with the envelope analysis and the PS using the method presented in Section 3.2.

Median Filter (MF)

The MF processing is an alternative method to the EA. A kind of envelope-squared signal is obtained in two steps. Firstly, the acceleration signal is differentiated: this process enhances the contribution of jerk forces and lifts the high-frequency content of the signal, where impacts caused by sudden jumps due to bearing faults are characterized [16]. The selection of the optimum differentiation order is performed by means of the examination of the maximum Kurtosis in the differentiated signal. In this work, the acceleration signal was processed with a second order derivative, obtaining the so-called Jounce $\left(\frac{m}{s^4}\right)$.

Once the best differentiation order is identified, the signal is squared and filtered with a median filter. That represents a better alternative to low-pass filtration and can be compared to an envelope process with robust smoothing. The median filter length (order N) is selected in order to have the first 10 harmonics of the characteristic fault frequencies [17]. Therefore, considering a signal with a sampling frequency f_s , the order of the filter N is obtained through the following equation

$$N = \frac{f_s}{10f_c} \quad (11)$$

where f_c is the characteristic frequency of the bearing fault.

Finally, the resulting filtered signal was processed with the envelope analysis and the PS using the method presented in Section 3.2.

The functions employed for the CNN training are summarized in Figure 2.




Fig. 2 Representation of the signals employed for the network's training. In blue, the 1024 samples selected as input data for each function

LeNet5 Architecture and Training

The LeNet-5 model is a classic CNN model proposed by [18]. It is widely used in several fields due to its high accuracy in the classification process. Usually, its architecture consists of an input represented by an image converted in grayscale, a series of convolutional layers followed by pooling and fully connected layers, which finally end in a classification output vector. In Figure 3 an example of a LeNet-5 model is shown.

The dimensions of the input images in 2D are obtained from the 1D signal from segments of 1024 samples transformed into greyscale images of 32x32 pixels, where each sample is normalised according to the equation

$$p'_i = \frac{p_i - p_{min}}{p_{max} - p_{min}} \cdot 255 \quad (12)$$




Fig. 3 Generic LeNet5 architecture

where p_i is the single value of the function to normalize and p_{min} and p_{max} are respectively its minimum and maximum. In Figure 4n, an example of the .png files employed as input data for the LeNet5.




Fig. 4 Images obtained through the preprocessing techniques employed for the networks as input

An improved LeNet-5 is employed, starting from 32 x 32 images with three convolutional layers of 32 kernels 3x3 followed by a ReLu layer and a batch normalization layer to improve the stability and effectiveness of the net and three pooling layers 2x2, two fully connected layers (100 and 4 neurons), a softmax and a classification layer. The architecture of the improved LeNet-5 is shown in Figure 5.

To compare the performance of LeNet-5 models with different processing methods, the same optimization procedure was applied. The input image size was fixed at 32 x 32 pixels. The resulting 2D dataset was split into 70% for training and




Fig. 5 Improved LeNet5 architecture

30% for testing. The training set was first fed into the LeNet5 network to train the model. Then, the testing set was used to evaluate the model's performance by analyzing the results.

The models were trained and validated using data from the medium wear state (14 mils), while the test sets included the worst damage level (21 mils). The training and validation sets consisted of 2000 observations (1400 training and 600 validation), covering all available shaft velocities. The test set was composed of 400 observations for each case (100 for each velocity). An exception was made for the healthy state and the outer race condition. In fact, the healthy condition presents only two velocities; therefore, the observations considered in the test set were only 200. Regarding the race case, damage was induced in three different positions, resulting in a total of 1200 observations for that case. The number of test observations employed for each condition and shaft velocity is reported in Table 2.

Table 2 Number of observation employed for each case to validate the improved LeNet5s accuracy

	1725 rpm	1750 rpm	1772 rpm	1797 rpm
Brinelling	100	100	100	100
Inner Race	100	100	100	100
New	0	100	100	0
Outer Race	300	300	300	300

Results are presented in the following section.

Results

The six improved LeNet5 models were trained and validated using the different input data mentioned previously:

- Raw Data
- Power Spectrum
- Simple Envelope Analysis + Power Spectrum
- AR Processing + Envelope Analysis + Power Spectrum
- Cepstrum Pre-Whitening + Envelope Analysis + Power Spectrum
- Median Filter + Power Spectrum

For each model and processing, ten different networks have been trained and tested to evaluate the overall efficiency.

Results are reported in terms of the average of the classifications among the ten networks trained for each processing strategy in Figure 6.



Fig. 6 Confusion matrices obtained through the validation of the trained LeNet5s for each input preprocessing strategy

Results show that all the networks trained can detect the healthy condition properly, with an accuracy of almost 100%. However, some distinctions have to be deepened:

- Raw Data training is unbalanced in the classification, where Brinelling and Outer Race are often confused.
- Power Spectrum performs poorly overall. The damage classes are heavily confused.
- Envelope is slightly better compared to the previous methods. However, the Outer Race condition is still misclassified.
- Auto Regressive improves the classification of the Outer Race.
- Cepstrum Pre-Whitening provides more balanced results. However, Outer Race is still heavily confused, even with the healthy state, which could be problematic in a diagnostic strategy.
- The Median Filter gives the best overall results. Brinelling and Inner Race are classified with high accuracy, while Outer Race achieves a relatively good outcome. Misclassifications are lower and more balanced.

Concluding, among the six preprocessing approaches, the Median Filter achieves the best performance, ensuring high accuracy across all three fault classes and minimizing confusion between damaged and healthy bearings. CPW comes second, with acceptable but less balanced results. Power Spectrum and Envelope are the weakest methods, with a strong bias toward misclassifying cases as OR.

Conclusion

Adopting the improved LeNet5 convolutional neural network makes it possible to recognize, with good accuracy, the class of damage of the bearing, also with different entities of fault. The preprocessing of the signal increases the performance of the model, as shown in Figure 6. In particular, several methods (and more traditional), such as Raw Data, Power Spectrum, envelope analysis and the AR model, tended to misclassify specific classes (notably Outer Race). While the Median Filter led to a superior discrimination, reducing the misclassification rate and enhancing overall robustness.

The significance of these outcomes lies in the method's potential applicability for early fault detection in bearings. By minimizing the incorrect labelling of damaged cases as healthy, the Median Filter approach enhances diagnostic reliability.

References

1. B. Li, M.-Y. Chow, Y. Tipsuwan, and J. C. Hung, "Neural-network-based motor rolling bearing fault diagnosis," *IEEE Transactions on Industrial Electronics*, vol. 47, no. 5, pp. 1060–1069, 2000, doi: 10.1109/41.873214.
2. V. Pandhare, J. Singh, and J. Lee, "Convolutional Neural Network Based Rolling-Element Bearing Fault Diagnosis for Naturally Occurring and Progressing Defects Using Time-Frequency Domain Features," in *2019 Prognostics and System Health Management Conference (PHM-Paris)*, IEEE, May 2019, pp. 320–326. doi: 10.1109/PHM-Paris.2019.00061.
3. S. A. McInerny and Y. Dai, "Basic vibration signal processing for bearing fault detection," *IEEE Transactions on Education*, vol. 46, no. 1, pp. 149–156, Feb. 2003, doi: 10.1109/TE.2002.808234.
4. S. Kim, D. An, and J.-H. Choi, "Diagnostics 101: A Tutorial for Fault Diagnostics of Rolling Element Bearing Using Envelope Analysis in MATLAB," *Applied Sciences*, vol. 10, no. 20, p. 7302, Oct. 2020, doi: 10.3390/app10207302.
5. P. Borghesani, P. Pennacchi, R. B. Randall, N. Sawalhi, and R. Ricci, "Application of cepstrum pre-whitening for the diagnosis of bearing faults under variable speed conditions," *Mech Syst Signal Process*, vol. 36, no. 2, pp. 370–384, Apr. 2013, doi: 10.1016/j.ymssp.2012.11.001.
6. W. A. Sethares, A. J. Milne, S. Tiedje, A. Prechtel, and J. Plamondon, "Spectral Tools for Dynamic Tonality and Audio Morphing," *Computer Music Journal*, vol. 33, no. 2, pp. 71–84, Jun. 2009, doi: 10.1162/comj.2009.33.2.71.
7. S. Li, G. Xie, W. Ji, X. Hei, and W. Chen, "Fault Diagnosis of Rolling Bearing Based on Improved LeNet-5 CNN," in *2020 IEEE 9th Data Driven Control and Learning Systems Conference (DDCLS)*, IEEE, Nov. 2020, pp. 117–122. doi: 10.1109/DDCLS49620.2020.9275100.
8. "https://engineering.case.edu/bearingdatacenter."
9. Y. Jin, L. Hou, and Y. Chen, "A new rotating machinery fault diagnosis method based on the Time Series Transformer," Aug. 2021, [Online]. Available: <http://arxiv.org/abs/2108.12562>
10. R. B. Randall and J. Antoni, "Rolling element bearing diagnostics—A tutorial," *Mech Syst Signal Process*, vol. 25, no. 2, pp. 485–520, Feb. 2011, doi: 10.1016/j.ymssp.2010.07.017.
11. R. B. RANDALL, J. ANTONI, and S. CHOBASAARD, "THE RELATIONSHIP BETWEEN SPECTRAL CORRELATION AND ENVELOPE ANALYSIS IN THE DIAGNOSTICS OF BEARING FAULTS AND OTHER CYCLOSTATIONARY MACHINE SIGNALS," *Mech Syst Signal Process*, vol. 15, no. 5, pp. 945–962, Sep. 2001, doi: 10.1006/mssp.2001.1415.
12. D. Hochmann and E. Bechhoefer, "Envelope bearing analysis: theory and practice," in *2005 IEEE Aerospace Conference*, IEEE, 2005, pp. 3658–3666. doi: 10.1109/AERO.2005.1559671.

13. H. Luo, X. Fang, and B. Ertas, "Hilbert Transform and Its Engineering Applications," *AIAA Journal*, vol. 47, no. 4, pp. 923–932, Apr. 2009, doi: 10.2514/1.37649.
14. R. B. Randall, *Vibration-based Condition Monitoring*. Wiley, 2011. doi: 10.1002/9780470977668.
15. N. Sawalhi, R. B. Randall, and H. Endo, "The enhancement of fault detection and diagnosis in rolling element bearings using minimum entropy deconvolution combined with spectral kurtosis," *Mech Syst Signal Process*, vol. 21, no. 6, pp. 2616–2633, Aug. 2007, doi: 10.1016/j.ymssp.2006.12.002.
16. N. Sawalhi, "Rolling element bearings localized fault diagnosis using signal differencing and median filtration," *Journal of Vibroengineering*, vol. 20, no. 3, pp. 1322–1339, May 2018, doi: 10.21595/jve.2017.18254.
17. J. B. Bednar, "Applications of median filtering to deconvolution, pulse estimation, and statistical editing of seismic data," *GEOPHYSICS*, vol. 48, no. 12, pp. 1598–1610, Dec. 1983, doi: 10.1190/1.1441442.
18. Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998, doi: 10.1109/5.726791.



Chapter 6

Aeroacoustics Structure Interaction Experiment in a Rectangular Duct with a Flush Mounted Panel

Daniele Fabbri, Carlo Rosso, Fabio Bruzzone, Francesco Avallone,
Lucas A. Bonomo, and Julio A. Cordioli

Abstract This work presents a preliminary study motivated by the need to design and conduct controlled experiments on the aeroacoustics response of a flush mounted flexible panel embedded in an otherwise rigid rectangular duct. The overarching goal is to understand how structural flexibility influences the generation and propagation of aerodynamic sound through fluid structure interaction. While the full experimental campaign will involve the characterization of different panel materials and thicknesses under grazing flow, this contribution focuses on outlining the proposed test rig and reporting preliminary numerical simulations that guide the design process. The simulations presented here highlight the expected frequency shifts with respect to the in vacuo structural modes and emphasize the differences between purely acoustic structure interaction and aeroacoustics structure interaction cases. These results provide an initial indication of the coupling regimes that will be most relevant for subsequent experiments. Ultimately, the study aims to establish a controlled benchmark test case and generate reference datasets for the validation of multiphysics numerical approaches in aeroacoustics fluid structure interaction.

Keywords Multiphysics · Aeroacoustics · Structure · ASI · FEA

Introduction

The need of couple different physics in the same framework was well described by *Korobkin et al.* in [1] describe the multiphysics coupling in hydroelasticity and emphasize a key principle that also applies to aeroacoustics: “to make significant progress ... we need stronger interaction and better collaboration between computational, experimental and theoretical researchers.” In the same spirit, the purpose of this paper is to describe the experimental set up of an aeroacoustics structure interaction case describing also some preliminary numerical results. The theoretical basis of aeroacoustics originates from *Lighthill's* acoustic analogy [2], which models aerodynamic sound generation as turbulence induced quadrupole sources in a compressible fluid. While effective for free space turbulence such as jets, *Lighthill's* formulation excludes the effects of solid boundaries. This limitation was addressed by *Curle* [3], and later by *Ffowcs Williams* and *Hawkings* [4], who introduced dipole and monopole source terms to account for stationary and moving surfaces. These extensions enabled the modeling of sound generated by flow structure interactions, including panel flutter and traveling wave instabilities, where wall motion directly contributes to acoustic radiation. Although nonlinear effects and complex geometries lie beyond the reach of such analytical models, their contribution to the foundations of fluid–structure interaction is significant. This multiphysics approach can be crucial to find strategies useful in the noise reduction as the comfort is increasingly important in engineering design, as it directly affect workers' performance and well being. In the same way, appropriate illumination plays a central role in ensuring human comfort [5, 6]

Early analytical studies of aeroacoustics structure interaction relied on simplified flow representations to clarify the underlying mechanisms. *Zolotarev* [7] analyzed the dynamics of a thin elastic plate in a fluid channel using potential flow

Daniele Fabbri · Carlo Rosso · Fabio Bruzzone · Francesco Avallone Department of Mechanical Engineering
Politecnico di Torino, TO 10129, Italy
e-mail: daniele.fabbri@polito.it; carlo.rosso@polito.it; fabio.bruzzone@polito.it; francesco.avallone@polito.it

Lucas A. Bonomo · Julio A. Cordioli
Federal University of Santa Catarina, 88040-900, Florianópolis, Brazil
e-mail: lucas.bonomo@lva.ufsc.br; julio.cordioli@ufsc.br

theory, enforcing continuity of velocity and pressure at the interface. His formulation showed that even a light fluid can noticeably alter plate frequencies and, in the case of near coincident acoustic and structural resonances, produce mode splitting into two coupled branches. He also identified critical flow velocities associated with flutter instability, where no real valued natural frequencies exist. In a complementary approach, *Tang et al.* [8] investigated the interaction of a single vortex with a flexible wall, modeling the vortex as a singularity in an irrotational background flow. Their results highlighted that wall vibration primarily radiates sound through a monopole mechanism linked to volumetric flux, with dipole effects playing a secondary role when the vortex passes very close to the surface. Together, these works establish both the global modal coupling and the local vortex excitation aspects of aeroacoustics structural phenomena, serving as useful theoretical references for more advanced Navier–Stokes-based models.

Huang et al. [9, 10, 11] carried out a series of analytical and experimental studies on acoustic–panel interaction in ducts without flow, demonstrating how panel length, tapering, and damping influence transmission loss and providing strategies for passive low-frequency noise control. Building on this, D’Elia [12] developed a compact, tunable liner concept that combines vibrating beams with micro slit cavities, validated through analytical impedance models and experiments. More recently, *Hessenthaler et al.* [13] proposed a comprehensive analytical framework for FSI with sixteen canonical cases intended as benchmarks for validating numerical schemes. The simplifications of analytical methods, however, limit their applicability, to capture nonlinearities and complex geometries, numerical approaches are required. Building on the analytical work of Huang et al. [9], Fan et al. [14] introduced a fully coupled time-domain numerical framework for aeroacoustics–structural interaction (AASI) in ducts and pipes. Their method combines a compressible Navier–Stokes solver based on the Conservation Element/Solution Element (CE/SE) scheme [15] with a one dimensional membrane model of the flexible panel, coupled through a predictor–corrector partitioned scheme. This approach captures acoustic, vortical, and entropy fluctuations without linearization, making it suitable for studying nonlinear panel–flow interaction in confined geometries. The model reproduces Huang’s frequency domain results in the absence of flow and then explores Mach dependent effects. Simulations identify nonlinear modal transitions, shifts in acoustic radiation from upstream to downstream dominance, and mode excitation at supersonic speeds induced by shock. Transmission loss exhibits four regimes, from rapid decay in the very low subsonic range to partial recovery in the low-supersonic regime. These results confirm that in tube and pipe systems, accurate prediction of aeroacoustics structural coupling requires fully time resolved numerical simulations to capture the regime dependent dynamics of panel vibration and wave radiation. Other numerical method which. Several studies have explored the interaction between turbulent flows and flexible panels, often through experimental campaigns or simplified numerical models [16, 17, 18]. However, their direct application to configurations inside ducts remains only partially clarified. As also noted by Lam in his thesis, current numerical methodologies can capture aeroacoustics structural interaction at low flow speeds but neglect convective terms at the wall boundary, making them unsuitable for higher flow regimes where curvature effects and instabilities such as divergence or Kelvin–Helmholtz waves may arise. Moreover, existing approaches often fail to account for large structural deformations, rely on two dimensional assumptions, and do not systematically address the influence of panel material properties, duct geometry, or multiple panel configurations commonly encountered in practice. These limitations underscore the need for further research into fully coupled, three dimensional turbulent flow panel interaction within duct systems, in order to better explain and predict their aeroacoustics response.

Methodology

The aeroacoustic structure–interaction experiment is conducted in the UFSC Grazing Flow Acoustic Impedance Facility [19, 20], composed of modular rectangular ducts with a cross section of $40 \times 100 \text{ mm}^2$. A 3 dimensional schematic of the test rig is shown in Fig 1. Quasi anechoic terminations at the inlet and outlet minimize acoustic reflections, while eight compression drivers are distributed both upstream and downstream of the test section to generate sound fields up to 150 dB propagating with and against the flow. Grazing flow is supplied by an external compressed air system, sustaining Mach numbers up to 0.7. Flow speed is controlled by differential pressure transmitter connected to a 2 mm diameter pitot tube at the inlet, calibrated according to ISO 3966:2008 [21], with temperature monitored by a transmitter.

In the present study, the conventional sample holder is replaced by a flush-mounted flexible panel, enabling direct investigation of aeroacoustic structure interaction (AeroASI) mechanisms under grazing flow. An array of eight equally spaced flush mounted pressure field microphones is located downstream the panel section to capture the acoustic field for modal decomposition of incident, reflected, and transmitted waves. Panel vibrations are measured with a scanning laser Doppler vibrometer, while structural parameters are characterized independently through bench-top modal testing. In addition to the acoustic measurements, the downstream flow field will also be characterized. To capture the velocity profile, hot wire prob and a Pitot tube will be considered, allowing direct comparison between mean and fluctuating flow quantities. These measurements will provide key data on boundary layer perturbations and on the possible development of shear layer instabilities




Fig. 1 USFC Impedance Tube Test rig

downstream of the panel, as well as some vortex acoustic coupling, thereby supporting a more complete understanding of the coupled aeroacoustics structure coupling. The configuration used is figured in 2




Fig. 2 Experimental scheme

In this work, the structure is considered the central object in the aeroacoustics fluid–structure interaction. The vibrating plate is not treated as a passive boundary, but rather as a dynamic element whose flexibility and material properties fundamentally govern the coupling mechanisms. To investigate this, three different structural materials will be studied for the plate. By systematically varying the elastic modulus, density, and damping characteristics, we will alter the flexibility of the structure and thus its ability to interact with the surrounding fluid and acoustic fields. This approach follows the consolidated understanding that the structural response dictates not only the modal behavior but also the efficiency of energy transfer between flow, vibration, and sound radiation.

Results - preliminar numerical case

A preliminary numerical investigation of the test rig was performed to quantify the influence of the mean flow on the aeroacoustic–structural interaction comparing it with an acoustic structure interaction (ASI) case. The computational domain consists of the central section of the duct, where the flexible plate is mounted, with perfectly matched layers (PMLs) applied at both inlet and outlet to avoid spurious reflections (Fig. 3). The plate, modeled as an aluminum structure with a thickness of 1 mm, is clamped at both ends and located at the center of the domain.

The analysis was carried out in two stages. First, the mean flow field was obtained by solving the steady turbulent flow using the SST model. A fully developed velocity profile was imposed at the inlet, while a zero pressure condition was

prescribed at the outlet. This velocity profile was subsequently mapped into the acoustic domain, which was represented by two fluid regions adjacent to the plate. A background flow acoustics coupling was employed to account for the convective effects of the mean flow. On the structural side, the flexible plate was coupled to the acoustic field through an aeroacoustic–structure boundary condition, ensuring continuity of pressure and normal velocity across the interface. To assess the role of mean flow, two configurations were compared: the full aeroacoustic–structural interaction with background flow, and a vibroacoustic case without mean flow (ASI). The acoustic excitation was introduced using a plane wave source imposed at the inlet, allowing for the study of the structural response under incident sound waves both with and without the influence of the mean flow.




Fig. 3 Numerical geometry

In the low frequency range of interest (20–1000 Hz), the simulations were conducted and the results showing the sound pressure level (SPL) over the plate 4, downstream of the duct 5, and the Y component displacement response of the plate 6. The uncoupled natural frequencies of the plate are indicated in the next figures by vertical dotted red lines. These lines mark the positions of the first three in vacuo frequencies, providing a direct reference for comparison with the coupled vibroacoustic and AeroASI results. For clarity, the first natural frequencies are also summarized in Table 1.

Table 1 First three natural frequencies of the plate

Mode	Natural Frequency [Hz]
1	124.9
2	344.3
3	674.7




Fig. 4 SPL over the plate




Fig. 5 SPL downstream

Examining the sound pressure levels, the ASI simulations display generally higher SPL values than the AeroASI results. This difference arises because in ASI the coupling is limited to the acoustic and structural domains, so the acoustic wave interacts directly with the plate. In contrast, in AeroASI the flow solution is first mapped onto the acoustic mesh, and the subsequent aeroacoustic–structure coupling leads to additional damping, which reduce the net SPL observed in the duct.

Moreover, in vibroacoustic case the SPL spectra show sharper peaks and deeper notches associated with the structural resonances, while when the background fluid is considered these features are smoother and shifted. The spikes are broadened by the interaction between the convected flow field, the scattered acoustic waves, and the vibrating plate. This highlights the fundamental difference between the two approaches: ASI emphasizes the direct acoustic–structure resonance, whereas AeroASI captures the more complex aeroacoustic feedback mechanisms that tend to dampen and spread out the resonance effects.

The comparison between the in vacuo natural frequencies of the plate and the coupled simulations highlights a systematic frequency shift. In the first mode, the coupled resonance appears slightly below the uncoupled structural frequency, which indicates that at low frequency the fluid mainly contributes through an added mass effect. Conversely, for the second and third modes the coupled resonances are shifted to higher values than the in vacuo frequencies. This trend is a well documented feature in fluid structure interaction problems [22, 9] and is mainly attributed to the additional stiffness introduced by the compressible fluid. Rather than acting as a pure added mass, the surrounding air modifies the effective impedance of the system, leading to a frequency dependent behavior: mass like at low frequency and stiffness like at higher frequencies.




Fig. 6 Plate displacement - Y component

Conclusions

The present work has compared two configurations for a flexible plate in a duct: the acoustic–structure interaction case, where only the coupling between the acoustic wave and the structural response is considered, and the aeroacoustic–structure interaction case, where the presence of a mean flow is included through a stationary fluid solution mapped onto the acoustic structural domain.

The results highlight that the coupled resonance frequencies differ from the in-vacuo values: the first mode shifts slightly lower due to fluid mass loading, while the higher modes shift upward as a consequence of the additional stiffness introduced by the compressible fluid. Moreover, the sound pressure level distributions clearly show that ASI predicts sharper peaks and higher overall SPL compared to AeroASI, where the mean flow introduces damping, smoothing the spectral response.

This study is intended as a preliminary phase of a broader process aimed at understanding the impact of structural flexibility on noise propagation and reduction methods. The numerical analyses presented here serve to identify the type of behavior that can be expected when experimental investigations will be carried out, and to build the necessary knowledge for modeling more complex scenarios, such as the interaction between a grazing flow and a flexible structure in terms of aeroacoustic noise.

Acknowledgments This publication is part of the project PNRR-NGEU which has received funding from the MUR – DM 117/2023

References

1. Korobkin, A., Părau, E.I., and Vanden-Broeck, J.M. “The mathematical challenges and modelling of hydroelasticity”. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1947):2803–2812 (2011)ifnextchar.gobble.
2. Lighthill, M.J. “On sound generated aerodynamically i. general theory”. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 211(1107):564–587 (1952)ifnextchar.gobble.
3. Curle, N. “The Influence of Solid Boundaries upon Aerodynamic Sound”. *Proceedings of the Royal Society of London Series A*, 231(1187):505–514 (1955)ifnextchar.gobble.
4. Williams, J.E.F. and Hawkins, D.L. “Sound Generation by Turbulence and Surfaces in Arbitrary Motion”. *Philosophical Transactions of the Royal Society of London Series A*, 264(1151):321–342 (1969)ifnextchar.gobble.
5. Carlin, A. and Verso, V.R.M.L. “Is lighting based on photopic requirements in assembly workstations also optimal for non-visual performance in a human-centric perspective? a case study in the automotive sector”. *International Journal of Mechanics and Control*, 26(01):185–212 (2025)ifnextchar.gobble.
6. Engel, M.S., Fiebig, A., Pfaffenbach, C., and Fels, J. “A review of the use of psychoacoustic indicators on soundscape studies”. *Current Pollution Reports*, 7(3):359–378 (2021)ifnextchar.gobble.
7. *Fluid-Structural-Acoustical Interactions of a Thin Plate in a Cannel With Flowing Fluid*, volume 4th International Symposium on Fluid-Structure Interactions, Aeroelasticity, Flow-Induced Vibration and Noise: Volume I of *ASME International Mechanical Engineering Congress and Exposition* (1997)ifnextchar.gobble.
8. Tang, S.K., Leung, R.C.K., So, R.M.C., and Lam, K.M. “Acoustic radiation by vortex induced flexible wall vibration”. *The Journal of the Acoustical Society of America*, 118(4):2182–2189 (2005)ifnextchar.gobble.
9. Huang, L. “A theoretical study of passive control of duct noise using panels of varying compliance”. *The Journal of the Acoustical Society of America*, 109(6):2805–2814 (2001)ifnextchar.gobble.
10. Huang, L. “A theoretical study of duct noise control by flexible panels”. *The Journal of the Acoustical Society of America*, 106(4):1801–1809 (1999)ifnextchar.gobble.
11. Huang, L., Choy, Y.S., So, R.M., and Chong, T.L. “Experimental study of sound propagation in a flexible duct”. *The Journal of the Acoustical Society of America*, 108(2):624–631 (2000)ifnextchar.gobble.
12. PhD thesis.
13. Hessenthaler, A., Balmus, M., Röhrle, O., and Nordsletten, D. “A class of analytic solutions for verification and convergence analysis of linear and nonlinear fluid-structure interaction algorithms”. *Computer Methods in Applied Mechanics and Engineering*, 362:112841 (2020)ifnextchar.gobble.
14. Fan, K.H. et al. “Computational aeroacoustic-structural interaction in internal flow with ce/se method” (2018)ifnextchar.gobble.
15. Lam, G., Leung, R., Yu, K.F., and Tang, S.K. *Computational Duct Aeroacoustics Using CE/SE Method*.
16. Springer, M., Scheit, C., and Becker, S. “Fluid-structure-acoustic coupling for a flat plate”. *International Journal of Heat and Fluid Flow*, 66:249–257 (2017)ifnextchar.gobble.
17. Schafer, F., Muller, S., Uffinger, T., Becker, S., Grabinger, J., and Kaltenbacher, M. “Fluid-structure-acoustic interaction of the flow past a thin flexible structure”. *AIAA Journal*, 48(4):738–748 (2010)ifnextchar.gobble.
18. Schoder, S., Falk, S., Wurzing, A., Lodermeier, A., Becker, S., and Kniesburges, S. “A benchmark case for aeroacoustic simulations involving fluid-structure-acoustic interaction transferred from the process of human phonation”. *Acta Acustica*, 8 (2024)ifnextchar.gobble.
19. Quintino, N.T., Bonomo, L.A., Pereira, L.M., Cordioli, J.A., and Avallone, F. *Boundary Layer Effects on Experimental Impedance Education*.
20. Bonomo, L.A., Quintino, N.T., Cordioli, J.A., Avallone, F., Jones, M.G., Howerton, B.M., and Nark, D.M. *A Comparison of Impedance Education Test Rigs with Different Flow Profiles*.
21. “Measurement of fluid flow in closed conduits – velocity area method using pitot static tubes” (2008)ifnextchar.gobble. Standard.
22. Fahy, F. and Gardonio, P. *Sound and Structural Vibration: Radiation, Transmission and Response*. Academic Press. Elsevier/Academic (2007)ifnextchar.gobble.



Chapter 7

Machine Learning Surrogates for Rapid Hertzian Stress Estimation and Static Transmission Error in Gear Pre-Sizing

Daniele Fabbri, Carlo Rosso, and Fabio Bruzzone

Abstract The dimensioning and dynamic evaluation of spur and helical gears rely on different strategies, typically addressed separately in conventional design. In this work, we integrate both within a machine learning framework by developing two complementary surrogates, each based on a distinct dataset. The first dataset is generated analytically from ISO based Hertzian stress calculations, representing the traditional dimensioning route. These values are used to train regression models such as Support Vector Machines, Random Forests, Gradient Boosted Trees, and Feed Forward Neural Networks, providing a fast surrogate for assessing surface durability and pitting resistance. In contrast, the second dataset targets the prediction of the static transmission error (STE), a key excitation source in gear dynamics, and is constructed using a semi analytic approach. To capture the meshing evolution of STE, both Feed Forward and Sequence to Sequence GRU models are investigated. By unifying dimensioning through ISO based stresses with dynamic evaluation via STE prediction, this work establishes a foundation for comprehensive gear design surrogates that accelerate exploration and support multi objective optimization.

Keywords Machine Learning · Gears · Dynamic · Neural Networks · Static Transmission Error

Introduction

In the context of gear engineering the use of machine learning (ML) has been grown in the recent years for modeling, diagnosis, and optimization. Conventional approaches, including analytical models, finite element simulations (FEM), or experimental campaigns, are typically resource intensive and limited in their adaptability to the wide range of gear geometries and operating conditions. By contrast, surrogate models offer the ability to learn complex nonlinear relationships between inputs and outputs, supporting accurate predictions, early fault detection, and efficient exploration of the design space. In general, ML techniques are classified into two main groups: supervised and unsupervised learning.

Unsupervised learning can reveal hidden structures in unlabeled data through clustering or rule-based methods [1], with applications in anomaly detection [2]. However, gear engineering mainly relies on supervised regression, addressing both regression tasks (e.g., stress or noise prediction) and classification tasks (e.g., healthy vs. faulty gears) [3, 4]. Such models have been successfully applied to predict stresses, wear, and acoustic emissions, with performance typically assessed using *MAE*, *MSE*, *RMSE*, and R^2 [5].

In the present work, two complementary evaluation strategies are adopted. For single value predictions, such as contact pressure, a parameter of primary importance in gear design, *RMSE*, is used to quantify deviations from reference values, while R^2 assesses the overall explanatory capacity of the model. For sequence predictions, namely the Static Transmission Error (*STE*) [6] evaluated over the mesh cycle, the surrogate model also considers the agreement in curve shape between the

Daniele Fabbri · Carlo Rosso · Fabio Bruzzone
Department of Mechanical and Aerospace Engineering
Politecnico di Torino, 10129 Torino, Italy
e-mail: daniele.fabbri@polito.it; carlo.rosso@polito.it; fabio.bruzzone@polito.it

Carlo Rosso · Fabio Bruzzone
GeDy TrAss srl
10128 Torino, Italy
e-mail: carlo.rosso@polito.it; fabio.bruzzone@polito.it

ML output and the semi analytical reference solution. This dual perspective ensures that the surrogate models are not only numerically accurate but also capable of faithfully reproducing the physical trends of gear behavior.

Recent research has also addressed the use of machine learning for predicting both the design of the gears and the Static Transmission Error (STE), one of the main sources of dynamic excitation in gear systems [7]. Papalexis et al. [8] developed neural network surrogates trained on more than 10,000 finite element simulations of polymer gears, achieving a mean absolute percentage error below 0.5% in replicating full STE curves. Urbas et Al. [9] demonstrated the use of multiple regressors including Random Forest, AdaBoost, and NN to predict the shape factor Y_z of non involute polymer gears from 456 FEM simulations, achieving R^2 values above 0.9 and confirming the feasibility of real time stress estimation. On the experimental side, Hyoungh et al. [10] predicted gear whine noise using micro geometry and test bench data from 12 gear sets. After LASSO-based feature selection, ensemble models such as XGBoost outperformed OLS, improving R^2 by up to 63.5%. Haefner et Al [11] developed a Neural Network based surrogate model to predict tooth root stress in micro gears using measured shape deviations. The model trained on FEM outputs, resulted in a model capable of providing not only point predictions but also confidence intervals [12].

Computational methods are also widely applied in gear design optimization, where the objective is to identify geometric configurations that minimize stress, noise, or vibration. Evolutionary approaches such as Genetic Algorithms (GAs) have been employed to explore the design space [13], but they typically depend on costly simulations like FEM or Loaded Tooth Contact Analysis (LTCA) to evaluate candidate designs [14, 15], creating a significant computational bottleneck.

Overall, recent studies demonstrate that machine learning has become a practical and scalable tool across gear engineering, from stress prediction and fault diagnosis to noise estimation and design optimization. Surrogate models, in particular, enable fast approximations of computationally expensive analyses [16], accelerating iterative workflows within the broader framework of data driven model order reduction [17].

In this context, our work is based on the ISO 6336 analytical standard [18] as a reliable data source for training surrogate models, while also extending the framework to the prediction of the Static Transmission Error (STE). By addressing both Hertzian contact stress and STE, the proposed approach provides efficient tools for assessing surface durability, pitting resistance, and dynamic performance directly at the design stage. The objective of this work is to establish a foundational methodology for developing comprehensive gear design surrogate models that can be integrated into multi objective optimization frameworks, enabling the rapid identification of optimal and robust gear pair configurations.

In Section 2, the procedure used to generate the datasets is described. The first dataset is constructed entirely from analytical contact stress values computed using the ISO 6336 standard, covering a broad range of gear macro geometries and loading conditions. In parallel, a second dataset is built for the prediction of static transmission error (STE), using the same gear input parameters but adopting different learning strategies. For this task, both a Feed Forward Neural Network trained via backpropagation and a Sequence to Sequence (Seq2Seq) GRU based model are considered.

The overall machine learning workflow and results will be presented in section third, detailing the training and evaluation strategies adopted for various regression models, including Linear Regression, Ensemble Algorithms, and Neural Networks. All models were implemented using widely used machine learning libraries that are highly optimized for regression tasks, such as *scikit-learn* [19] and *Keras* [20] through *TensorFlow* [21].

Finally, the conclusions in Section summarize the key findings and outline future directions in the development of surrogate models for gear design.

Dataset preparation

The aim of this work is to develop surrogate models for spur gear pairs targeting two complementary outputs: the Hertzian contact stress, $\sigma_{H1,2}$, computed from the ISO 6336-2 formulation, and the STE calculated from *GeDy TrAss* software which employs a nonlinear and Non-Hertzian approach to determine the deformations of the teeth under load and their contact conditions [22, 23, 24]. The dataset spans a wide range of gear macro geometry and loading conditions (Table 1), and only safe configurations are retained. These data are modeled with Supervised Regression Methods, including Elastic Net, Ensemble Algorithms, and Neural Networks. For the STE dataset, the input features are different and include macro-geometric parameters such as torque T , transmission ratio τ , addendum h_a , module m , face width b , pinion teeth z_1 and helix angle β . Furthermore, micro-geometric modifications are introduced in the form of parabolic tip relief C_a from a start diameter of d_{Ca1} , as well as facewidth crowning C_β . The target output is a sequence of 20 discrete points $ste_1, \dots, ste_{20}$, representing the transmission error over a mesh cycle. Since STE is inherently sequential, the learning strategy relies on Sequence to Sequence GRU models. Together, these two datasets enable benchmarking of regression and sequence based machine learning approaches for durability (via stress prediction) and dynamic performance (via STE prediction) in gear design.

Hertzian contact stress - Dataset preparation

The stress dataset consists of about 5×10^6 samples with 15 initial features and one target label, the Hertzian stress σ_{H1} . After removing non-numeric entries and outliers, the dataset was analyzed to ensure suitability for regression. Skewness and kurtosis [25] were computed to assess distributional properties; strongly skewed features were normalized using either the $\log_{10}$ or *Yeo–Johnson* transformation [26], improving feature symmetry.

To mitigate multicollinearity [27, 28], Pearson correlation analysis [29] was performed both with the target and across features. Redundant variables with perfect or near perfect correlation were removed, considering also their physical meaning. The final dataset contained 14 features and one label, summarized in Table 1.

Table 1 Features and label - Hertzian stress dataset

Type	Symbol	Dataset label	Range
Target variable	σ_{H1}	sigma_H1	$14 \div 792 MPa$
Feature	T	Input_T	$10 \div 3000 Nm$
Feature	h_a	Input_ha	$1.25 \div 1.8$
Feature	h_f	Input_hf	$1 \div 1.6$
Feature	ρ_a	Input_rhoa	$0.1 \div 0.33$
Feature	τ	Input_tau,	$2 \div 6$
Feature	α_t	alphan	$20 \div 21.17$
Feature	b	b	$10 \div 70 mm$
Feature	β	beta	$0 \div 20^\circ$
Feature	d_{w1}	dw1	$15.28 \div 241.2 mm$
Feature	i	i	$25.19 \div 792.97 mm$
Feature	m_n	mn	$1 \div 6 mm$
Feature	x_1	x1	$-0.6 \div 0.9$
Feature	x_2	x2	$-0.6 \div 0.9$
Feature	z_1	z1	$17 \div 35$

Material: 34CrNiMo6, ISO code V,
Brinell Hardness 319HB, Hardened and Tempered

STE - Dataset preparation

The static transmission error (STE) dataset is smaller than the stress dataset, consisting of about 1350 samples. The input space includes 10 geometric and load features: torque T , transmission ratio τ , addendum h_a , module m , face width b , number of pinion teeth z_1 , helix angle β , start diameter of profile modification d_{Ca1} , amount of parabolic tip relief C_a , and amount of facewidth crowning C_β . Since the start diameter of the profile modification is dependent on the macro-geometry configuration of the gear currently analyzed, its values have been determined proportionally to the pitch diameter D_p of the pinion and gear respectively and to the addendum h_a in three discrete levels. To ensure a smooth exploration of the multidimensional design space in a limited number of simulations, the design space was sampled using Latin Hypercube Sampling (LHS) [30] on the target 1350 simulations.

The obtained STE dataset was then reformulated into an autoregressive *seq2seq* with Gated Recurrent Units (GRUs) learning task in order to capture the cyclic structure of the target over the gear meshing. Each gear configuration generated a sequence of 20 static transmission error values, complemented by 10 geometric and load parameters together (Table 2) with two synthetic cycle features, $\sin(\lambda)$ and $\cos(\lambda)$, where $\lambda \in [0, 1]$ as done in [8], they represent the normalized mesh position sampled at 20 equally spaced points. These two additional features provided the model with an explicit and continuous representation of the cycle progression, ensuring that the learning process was not limited to purely discrete indices. Before training, the target STE values were normalized using a *MinMax* transformation.

Methodology and results

The following methodology was applied consistently across all models and datasets to ensure a rigorous and comparable model development process. The first step was to perform a grid search [31] on a stratified subset of the training data to

Table 2 Features and label - STE dataset

Type	Symbol	Dataset label	Range
Target variables	$ste_{1...20}$	$ste_{1...20}$	$0 \div 2.61$
Feature	T	Input_T	$10 \div 1010Nm$
Feature	h_a	Input_ha	$1.25 \div 1.8$
Feature	τ	Input_tau,	$1.5 \div 6$
Feature	b	b	$10 \div 80mm$
Feature	β	beta	$0 \div 0.35^\circ$
Feature	m_n	m	$1 \div 5mm$
Feature	z_1	z1	$17 \div 30$
Feature	d_{Ca1}	dCa1	$D_p \div D_p + \frac{2}{3}h_a \cdot m$
Feature	C_a	Ca	$0 \div 6\mu m$
Feature	C_{beta}	Cbeta	$0 \div 6\mu m$

Material: 34CrNiMo6, ISO code V,
Brinell Hardness 319HB, Hardened and Tempered

identify appropriate hyperparameters. This preliminary stage was crucial for tuning key parameters such as regularization strength, tree depth, learning rate, and network architecture, depending on the algorithm. The optimal hyperparameter configuration was selected by analyzing training and validation loss curves, with particular attention to instability indicators such as oscillations or divergence. The models were then retrained with the best hyperparameters on the full training dataset to maximize the use of available data and improve generalization. Evaluation was conducted on a separate held out test set, where all predictive metrics were reported to validate each model's performance. In addition to accuracy, the training time of each algorithm was measured and reported, providing a broader comparison of efficiency versus predictive power. All computations were performed on a workstation equipped with an *i9-13900F* processor (2.00 GHz) and 64 GB of RAM. The models were trained using CPU implementations of the machine learning frameworks. GPU acceleration was not employed in this study but it is regarded as a promising avenue for future work to reduce training times, particularly for complex neural networks or larger datasets.

From the Hertzian stress prediction results, it was evident that neural networks are well suited for capturing the nonlinear relationships between geometry, loading conditions, and stress response. However, the prediction of static transmission error presents a different challenge, as it is inherently sequential and reflects the mesh cycle of a gear. Initially, we attempted to model the full 20 point STE sequence with a simple Feed Forward network producing all outputs simultaneously, but this approach failed to capture the dependencies of the outputs to the mesh cycle. For this reason, we adopted a sequence to sequence architecture with GRU units, which leverages the autoregressive structure of the data to reconstruct the shape of the output curve point by point. This approach demonstrated improved performance, even though the STE dataset was smaller than the stress dataset, highlighting the importance of matching the learning strategy to the physical nature of the target values.

Hertzian stress contact

A wide set of machine learning models was trained and evaluated on the original ISO 6336 stress dataset, which includes more than five million samples and 14 features. The results show a clear ranking in terms of predictive performance, and computational cost. The linear regression baseline with Elastic Net regularization demonstrated moderate accuracy, with a test RMSE above 70 MPa and $R^2 = 0.85$. While interpretable and efficient, its limited capacity to capture nonlinear effects restricted performance. Support Vector Regression with an RBF kernel achieved significantly better accuracy ($R^2 = 0.968$, RMSE = 32.5 MPa), but required subsampling (50,000) due to the high computational cost of kernel methods, providing a good balance between cost and performance, though scalability to the full dataset remains limited. Tree based ensemble methods offered an attractive compromise between predictive power and efficiency. Random Forest achieved strong results ($R^2 = 0.974$, RMSE = 30 MPa) with very fast training, highlighting its robustness and speed. Histogram-based Gradient Boosting outperformed all classical methods with an RMSE of 14.8 MPa and $R^2 = 0.994$, although it required longer training and showed signs of potential overfitting. Neural networks delivered the best overall accuracy. The scikit-learn MLPRegressor (3 layers: 200-100-50) achieved an RMSE of 3.43 MPa, converging stably with early stopping. A deeper Keras-based architecture (160-192-96 neurons with ℓ_2 regularization) confirmed this performance with a test RMSE of 4.63 MPa, though training times were much longer (almost 2900 seconds). Despite the higher computational burden, neural networks proved particularly effective in capturing the strong nonlinearities of the Hertzian stress formulation. Overall,

the results indicate that while ensemble methods like Random Forest and Gradient Boosting provide highly competitive accuracy at low computational cost, neural networks clearly stand out in terms of predictive precision. This confirms that deep learning approaches are especially suitable for surrogate modeling of ISO based Hertzian stress, where complex nonlinear relationships between gear macro geometry and stress response dominate. In Table 3 a comparison between different models is listed.

Table 3 Performance comparison of surrogate models on the original ISO 6336 stress dataset. Subset sizes are indicated where full dataset training was computationally prohibitive.

Model	Test RMSE [MPa]	Test R^2	Training Time [s]
Elastic Net (SGD)	70.29	0.852	38
Support Vector Regression (RBF, 50k samples)	32.47	0.968	106
Random Forest (50k samples)	29.62	0.974	5
Histogram Gradient Boosting (100k samples)	14.82	0.994	46
Neural Network – MLPRegressor (200-100-50)	3.43	0.99	560
Neural Network – Keras (160-192-96)	4.63	0.98	2880

STE

In this part of the work, the objective is to train a surrogate model to predict the Static Transmission Error (STE) curve of a gear pair over a full meshing cycle. The methodology follows a sequence based learning approach, where each curve is treated as a temporal sequence of 20 discretized steps representing different positions along the mesh cycle. The dataset includes both geometric and operational parameters, two additional columns are included to encode the phase of each sample in the cycle in a smooth and periodic way. The target variable is the STE value at each of these 20 positions. Additionally, a column stores the previous value in the sequence to support autoregressive learning.

Once the data is organized, it is split into training and test sets using a standard 80/20 ratio. The architecture used to model this problem is a Sequence to Sequence (Seq2Seq) model with Gated Recurrent Units (GRUs) [32], The idea is to use GRU rather than other methods as LSTM (Long Short Term Memory) because the problem has a short length sequences (20 steps of STE) where the data is not deeply hierarchical in time and GRUs can achieve very similar performance to LSTMs, often with faster convergence and less risk of overfitting[33, 34]. This approach seems to be a novelty in the context of gear dynamics, as the other works used a traditional approach in which the neural networks outcome 20 values in the output layers.

The encoder takes as input the full sequence of geometric and meshing parameters and compresses it into a latent representation of dimension 256, which serves as the internal state passed to the decoder. The decoder then takes the previous STE values (in a sequence of 20 points) and learns to reconstruct the future STE waveform step by step, conditioned on the encoded representation. The final dense layer is wrapped in a TimeDistributed layer, ensuring that a single scalar STE value is output for each of the 20 time steps. The loss function used is the Huber loss, the model is trained over 100 epochs with a batch size of 64, and both training and validation performance are tracked. After training, an example prediction is generated using the trained model. A correction is then applied to the first point of the predicted curve.

Finally, the predicted STE curve and the real curve are plotted together for visual comparison. This provides qualitative insight into the model's accuracy and smoothness in reproducing the physical behavior of gear meshing transmission error as shown in Figures 1 and 2. The result shows a close alignment between the predicted and actual curves, demonstrating the Seq2Seq model's effectiveness in capturing the dynamic structure of the problem. This is confirmed also by the Peak-to-Peak Transmission Error (PPTE) and by the metrics of the model summarized in Table 4

Table 4 Performance of the STE prediction using Seq2Seq GRUs .

Model	Test RMSE [MPa]	Test R^2	PPTE Error
Seq2Seq GRUs	0.00456	0.934	0.86%




Fig. 1 Result STE curve true vs predicted values- First curve




Fig. 2 Result STE curve true vs predicted values- Second curve

Conclusions

This paper should be intended as a starting point for applying different machine learning techniques to gear design, highlighting both dimensioning and dynamic behavior analysis. The section on Hertzian stress prediction demonstrates the effectiveness of using neural networks to capture complex nonlinear relationships, even with moderately deep architectures. In the second part, a novel approach for the regression of the Static Transmission Error (STE) is reconstructed using a specialized type of Recurrent Neural Network: the Seq2Seq GRU. The promising results confirm that traditional feedforward neural networks would struggle to retain temporal information about the meshing cycle and the evolution of the error over time. In contrast, the sequence to sequence approach succeeds in accurately reproducing the overall shape and trend of the STE curve, even when trained on a relatively small dataset.

Acknowledgments This publication is part of the project PNRR-NGEU which has received funding from the MUR – DM 117/2023

References

1. Aggarwal, C.C. *Data Classification: Algorithms and Applications*. Chapman & Hall/CRC, 1st edition (2014)ifnextchar.gobble.
2. Chandola, V., Banerjee, A., and Kumar, V. “Anomaly detection: A survey”. *ACM Computing Surveys (CSUR)*, 41(3):1–58 (2009)ifnextchar.gobble.
3. Bansal, S., Sahoo, S., Tiwari, R., and Bordoloi, D. “Multiclass fault diagnosis in gears using support vector machine algorithms based on frequency domain data”. *Measurement*, 46(9):3469–3481 (2013)ifnextchar.gobble.

4. LI, X., LI, J., QU, Y., and HE, D. “Semi-supervised gear fault diagnosis using raw vibration signal based on deep learning”. *Chinese Journal of Aeronautics*, 33(2):418–426 (2020)ifnextchar.gobble.
5. Botchkarev, A. “A new typology design of performance metrics to measure errors in machine learning regression algorithms”. *Interdisciplinary Journal of Information, Knowledge, and Management*, 14:045–076 (2019)ifnextchar.gobble.
6. Czákó, A., Řehák, K., Prokop, A., Rekem, J., Láštík, D., and Trochta, M. “Static transmission error measurement of various gear-shaft systems by finite element analysis”. *Journal of Measurements in Engineering*, 12(1):183–198 (2023)ifnextchar.gobble.
7. Bruzzone, F. and Rosso, C. “Sources of excitation and models for cylindrical gear dynamics: A review”. *Machines*, 8(3) (2020)ifnextchar.gobble.
8. Papalexis, C., Sakaridis, E., Terpos, K., Kalligeros, C., Tsolakis, A., and Spitas, V. “Neural network surrogates for finite element models in loaded tooth contact analysis of polymeric gears”. *Mechanism and Machine Theory*, 214:106127 (2025)ifnextchar.gobble.
9. Urbas, U., Zorko, D., and Vukašinović, N. “Machine learning based nominal root stress calculation model for gears with a progressive curved path of contact”. *Mechanism and Machine Theory*, 165:104430 (2021)ifnextchar.gobble.
10. Sun Hyoung, L. and Park, K.P. “Development of a prediction model for the gear whine noise of transmission using machine learning”. *International Journal of Precision Engineering and Manufacturing*, 24 (2023)ifnextchar.gobble.
11. Haefner, B., Biehler, M., Wagner, R., and Lanza, G. “Meta-model based on artificial neural networks for tooth root stress analysis of micro-gears”. *Procedia CIRP*, 75:155–160 (2018)ifnextchar.gobble.
12. for Guides in Metrology, J.C. “Jcgm 101: Evaluation of measurement data - supplement 1 to the “guide to the expression of uncertainty in measurement” - propagation of distributions using a monte carlo method”. Technical report, JCGM (2008)ifnextchar.gobble.
13. Cerrada, M., Zurita, G., Cabrera, D., Sánchez, R.V., Artés, M., and Li, C. “Fault diagnosis in spur gears based on genetic algorithm and random forest”. *Mechanical Systems and Signal Processing*, 66-67:812–828 (2016)ifnextchar.gobble.
14. Artoni, A. “A methodology for simulation-based, multiobjective gear design optimization”. *Mechanism and Machine Theory*, 133:95–111 (2019)ifnextchar.gobble.
15. Bonori, G., Barbieri, M., and Pellicano, F. “Optimum profile modifications of spur gears by means of genetic algorithms”. *Journal of Sound and Vibration*, 313(3-5):603–616 (2008)ifnextchar.gobble.
16. Forrester, A.I.J. and Keane, A.J. “Recent advances in surrogate-based optimization”. *Progress in Aerospace Sciences*, 45(1-3):50–79 (2009)ifnextchar.gobble.
17. Hartmann, D., Herz, M., and Wever, U. “Model order reduction a key technology for digital twins”. In *Digital Twin: The Vision and the Future*, pages 167–179. Springer, Cham (2018)ifnextchar.gobble.
18. “Calculation of load capacity of spur and helical gears”. ISO 6336 (2006)ifnextchar.gobble.
19. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Van der Plas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. “Scikit-learn: Machine learning in python”. *CoRR*, abs/1201.0490 (2012)ifnextchar.gobble.
20. Chollet, F. et al. “keras” (2015)ifnextchar.gobble.
21. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mane, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viegas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. “Tensorflow: Large-scale machine learning on heterogeneous distributed systems” (2016)ifnextchar.gobble.
22. Bruzzone, F., Maggi, T., Marcellini, C., Rosso, C., and Delprete, C. “Proposal of a novel approach for 3d tooth contact analysis and calculation of the static transmission error in loaded gears”. *Procedia Structural Integrity*, 24:178–189 (2019)ifnextchar.gobble.
23. Bruzzone, F., Maggi, T., Marcellini, C., and Rosso, C. “Gear teeth deflection model for spur gears: Proposal of a 3d nonlinear and non-hertzian approach”. *Machines*, 9(10) (2021)ifnextchar.gobble.
24. Rosso, C. and Bruzzone, F. “On the influence of the actual load sharing factor in increasing the power density in gearboxes”. In Petuya, V., Quaglia, G., Parikyan, T., and Carbone, G., editors, *Proceedings of I4SDG Workshop 2023*, pages 420–429, Cham (2023)ifnextchar.gobble. Springer Nature Switzerland.
25. Cain, M.K., Zhang, Z., and Yuan, K.H. “Univariate and multivariate skewness and kurtosis for measuring nonnormality: Prevalence, influence and estimation”. *Behavior Research Methods*, 49(5):1716–1735 (2017)ifnextchar.gobble.
26. Yeo, I.K. and Johnson, R.A. “A new family of power transformations to improve normality or symmetry”. *Biometrika*, 87(4):954–959 (2000)ifnextchar.gobble.
27. Shrestha, N. “Detecting multicollinearity in regression analysis”. *American Journal of Applied Mathematics and Statistics*, 8(2):39–42 (2020)ifnextchar.gobble.
28. Kim, J.H. “Multicollinearity and misleading statistical results”. *Korean Journal of Anesthesiology*, 72(6):558–569 (2019)ifnextchar.gobble.
29. Pearson, K. “Note on regression and inheritance in the case of two parents”. *Proceedings of the Royal Society of London*, 58:240–242 (1895)ifnextchar.gobble.
30. Iman, R.L., Davenport, J.M., and Zeigler, D.K. “Latin hypercube sampling (program user’s guide). [lhc, in fortran]”. Technical report, Sandia Labs., Albuquerque, NM (USA) (1979)ifnextchar.gobble.
31. Liu, J., Liang, G., Siegmund, K.D., and Lewinger, J.P. “Data integration by multi-tuning parameter elastic net regression”. *BMC Bioinformatics*, 19(1):369 (2018)ifnextchar.gobble.
32. Sutskever, I., Vinyals, O., and Le, Q.V. “Sequence to sequence learning with neural networks” (2014)ifnextchar.gobble.
33. Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. “Empirical evaluation of gated recurrent neural networks on sequence modeling” (2014)ifnextchar.gobble.
34. Balti, H., Abbes, A.B., and Farah, I.R. “A bi-gru-based encoder–decoder framework for multivariate time series forecasting”. *Soft Computing*, 28(9):6775–6786 (2024)ifnextchar.gobble.

Chapter 8

Acoustic Excitation for Fluidic Particle Separation



Blake Johnson, Bradon Timms, Jin Yue Trousil, Pavel Vakhlamov, Milo Prisbrey,
John Greenhall, and Nicholas Lieven

Abstract During diving operations, eddies generated near sediment beds cause the sediment to mix with water, reducing visibility for divers. Good visibility is critical for applications such as municipal maintenance or coastal construction, where obscured vision can lead to costly delays; naval operations, where turbid water can hamper navigation; and cave diving, where reduced visibility is deadly if the diver cannot locate the cave exit. In this work, we use Ultrasound Separation (UltraSep) as a means of clarifying silt-laden water to improve underwater visibility via a handheld device. The approach uses ultrasonic standing waves that drive particles to coalesce at the standing wave nodes and flocculate, increasing visibility through the rest of the fluid. To determine the effect of UltraSep on fluid clarification, we model the ultrasound, fluid, and hydrodynamic physics in a UltraSep device. Then, we experimentally measured the improvement in visibility and investigated its dependence on excitation power and fluid flow rate. The results from the experimentation with UltraSep not only prove its clarification ability, but also demonstrate that this ability follows a trend—which suggests that this process can be optimized to maximize visibility improvement. The simulations revealed that the acoustic resonance frequency that the chamber should be operated at depends on the initial particle concentration and also changes as the particles move to the nodes and the fluid becomes inhomogeneous. These results suggest the need for implementation of a control system to dynamically adjust the frequency of a device while it operates.

Keywords Ultrasound · separation · filtration · navigation

Blake Johnson
Department of Mechanical and Aerospace Engineering
University of California San Diego, San Diego, CA 92093
e-mail: bjohnson@lanl.gov

Bradon Timms
Department of Mechanical Engineering
Purdue University, West Lafayette, IN 47907
e-mail: timms@lanl.gov

Jin Yue Trousil
Department of Mathematics and Computer Science
Whitworth University, Spokane, WA 99251
e-mail: bjohnson@lanl.gov

Pavel Vakhlamov · Milo Prisbrey · John Greenhall
Los Alamos National Laboratory
Los Alamos New Mexico USA, 87545
e-mail: pvakhlamov@lanl.gov; mprisbrey@lanl.gov; jgreenhall@lanl.gov

Nicholas Lieven
Imperial College London
London SW7 2AZ, UK
e-mail: n.lieven@imperial.ac.uk

Introduction

Water bed surfaces are often comprised of small particulates, commonly known as silt, that can be easily disturbed and mixed into the surrounding water, reducing visibility through the water (increasing turbidity). This phenomenon can impede or even introduce life-threatening hazards for underwater work such as underwater optical imaging [1], construction [2], inspection [3], and cave diving. For instance, in 2014, the Diver Alert Network reported that visibility was a contributing factor in 31% of diving incidents [4]. There are few existing measures to mitigate this problem. The Divers Alert Network suggests using underwater lights, but cedes that silt-laden water can drastically reduce visibility due to light scattering from the particles. Divers may also use physical lines to help navigate without vision or use strobe lights to mark landmarks. However, these methods do not address the cause of the poor visibility, which is the silt suspended in the water[5]. Low visibility can also increase the time divers take to complete their tasks or potentially force them to surface, both of which lead to increases in cost, required time, and material waste. For object detection, divers may employ handheld tools with magnetic [6] and sonar capabilities [7] to assist them in reduced visibility conditions. However, sonar tools are limited by the strength of the unit and scattering and magnetic tools do not give divers an awareness of their environment and spatial constraints. Thus, there is a critical need for a method of clarifying silt-laden underwater environments by removing particulate. Most existing technologies that can remove particulate or silt rely on filters/membranes or hydrocyclones[8]. However, hydrocyclone technology requires high power input causing them to be large and heavy, and filter/membrane systems saturate quickly and require disposable replacements parts.

In this work, we present a method for improving underwater visibility by separating the silt from the water via Ultra-sonic Separation (UltraSep). UltraSep utilizes ultrasonic waves to generate Acoustic Radiation Force (ARF) to separate and concentrate particles suspended in a fluid. UltraSep has been adopted in a range of applications including, manipulating and separating individual cells or particles[9], separating cells for flow cytometry[10], focusing powders in laser powder bed fusion additive manufacturing[11], arranging user-specified patterns of microparticles to functionalize structure fabricated via stereolithography[12], and water clarification to remove contaminants[13].

However, acoustic separation techniques are primarily used for biomedical applications currently, as their ability to sort biological samples of relatively uniform size and material from fluids in confined, small (μm -scale) environments non-destructively perfectly suits the needs of the biomedical community[9]. As a result, most research has been focused in this realm with micro-fluidic devices[14]. In contrast, for open underwater activities, the fluid is in an unconstrained body and it can contain particles with a broad range of sizes and material properties. Additionally, existing work does not account for the change in acoustic resonance before steady state as particles move within the bulk medium and create inhomogeneities.

In this work, we demonstrate that UltraSep can be used to improve visibility in silt-laden water, specifically identifying the impact of flow rate and amplitude on the clarification process. In this method, the device pumps water from a larger container through a small acoustic chamber, using the acoustic radiation force from the ultrasonic standing wave to group silt particles so that they coalesce, thus reducing turbidity. We address a real-world application of this technology in the form of a hand-held chamber that can be passed through the water at the will of the diver to clear a region of silt. In addition to the exploring of the the clarification process, we also will simulate the acoustic resonance stability in COMSOL Multiphysics® [15] as the particle concentration changes.

Methods

There are multiple physical phenomena involved in the UltraSep process. The primary acoustic radiation force (ARF), which imparts momentum from the incident wave as it scatters from the particles; secondary radiation force, wherein the scattered wave from each particle scatters again when it interacts with neighboring particles[9]; and viscous drag force, which resists particle motion with respect to the local fluid motion, govern particle movement. In this paper, we excite a flow cell chamber at one of its resonant frequencies to generate a standing ultrasound wave in the cell. Standing waves comprise pressure nodes, where the pressure is zero, and antinodes, where the pressure oscillates periodically between $\pm P_0$, the amplitude of the standing wave. In the standing wave, particles are driven to either the nodes or antinodes based on their density ρ_p and sound speed c_p , in relation to the bulk fluid properties (ρ_f, c_f) [16].

For a particle of radius R in a standing wave with wavelength $\lambda \gg R$, the equation of the time-averaged primary ARF is given as[17]

$$\vec{F} = -\nabla U, \quad (1)$$

where U is the Gor'kov Acoustic Radiation Potential (ARP), given as

$$U = 2\pi R^3 \left[f_{\text{comp}} \frac{1}{3\rho c^2} \overline{p_{in}^2} - f_{\rho} \frac{\rho}{2} \overline{v_{in}^2} \right]. \quad (2)$$

Here, $\overline{p_{in}^2}$ and $\overline{v_{in}^2}$ are the mean-square fluctuations of the wave pressure and fluid velocity at the particle position. The force amplitude due to pressure and velocity is determined by the coefficients f_{comp} and f_{ρ} , which are the compressibility and density contrast factors, respectively, and are given as

$$f_{\text{comp}} = 1 - \rho c^2 / \rho_s c_s^2, \text{ and} \quad (3)$$

$$f_{\rho} = 2(\rho_s - \rho) / (2\rho_s + \rho). \quad (4)$$

We note that in the case of silt particles in water we find $f_{\text{comp}} > 0$ and $f_{\rho} > 0$, which means that the particles will be driven to the nodes of a standing wave, where $\overline{p_{in}^2} = 0$ and $\overline{v_{in}^2}$ is maximum. For the remainder of this manuscript, we will focus on this scenario, where particles are driven to the nodes.

In addition to ARF, which drives the particles to the nodes, there is drag force $\vec{F}_d$ that resists motion due to ARF and drives the particles with the fluid flow. For small particles at low Reynolds numbers, the drag force is given using Stokes' drag equation[9]

$$\vec{F}_d = -6\mu R\pi \vec{v}. \quad (5)$$

The Stokes drag is influenced by the fluid viscosity μ , the particle radius r_p , and relative velocity between particle and flow v .

In this work, we represent silt separation for an open-channel device both experimentally and with numerical simulations (COMSOL Multiphysics®). The experimental setup is shown in Figure 1.




Fig. 1 Depiction of Experimental SetUp

We have represented the motion of the UltraSep silt-separation device through open water by instead pumping water in a tank through the UltraSep device. A tank containing a silt water mixture is separated into two sides. A pump circulates the water from the right side of the tank to the left. The resulting pressure difference forces silt-laden water through the UltraSep device, back to the right side of the tank. From here on, the language in the paper will orient itself to the UltraSep device—calling the “left side” the “inlet” and “right side” the “outlet”. We drive the $17 \times 17 \times 32$ mm borosilicate glass UltraSep device with a piezoelectric ultrasonic transducer (PZT-5H, center frequency 2 MHz) to generate vertical standing waves across the device cross section. We investigate the effect of the driving voltage, applied through a power amplifier, on the separation performance. A digital camera is bolted to a fixed position in front of the tank to video the separation process and act as a turbidity sensor.

During each separation experiment, we first turn on the pump to control the flow rate, and agitate both the inlet and outlet sides of the tank to homogenize the suspension of silt. We begin recording video of the flow and then excite the UltraSep transducer. The camera takes a video with a total duration 5 minutes while the system operates to capture the initial flow without separation and steady-state behavior once the silt separation process begins.

To quantify the performance of the UltraSep process, we analyze the water turbidity from the recorded video frames. Two spacial regions of the video frames were inspected: the “inlet region” just before the flow enters the UltraSep device, and the “outlet region” just after the flow exits the UltraSep device, as indicated in Figure 2. By comparing these two regions and their relative turbidities, the performance of the chamber can be assessed.




Fig. 2 Depiction of Inlet and Outlet Regions

We implement a turbidity quantification algorithm in MATLAB. The steps of data processing are shown in Figure 3. At each frame of the video, the inlet and outlet regions are isolated and converted to grayscale. We then measure the average pixel value across the inlet frame and across the outlet frame. Repeating this process for each frame in a given experiment results in a vector of turbidity values (average pixel intensities) for both the inlet and outlet regions.




Fig. 3 Depiction of inlet and Output Regions

Due to difference in lighting and background between the inlet and outlet regions and differences in the initial silt dispersal between experiments, we calibrate the inlet and outlet turbidity values. A series of 14 separation experiments were used to determine the absolute minimum and maximum pixel value in both the inlet and outlet. Then for a given experiment the grayscale value of each region was normalized bounded by its respective maximum and minimum. This allowed for comparison of runs with different starting particle concentration, homogeneity, etc.

Additionally, to analyze the results of the experiment properly, the data first needed to be smoothed via a filtering process, so high frequency noise could be removed and clean time derivatives of the trends could be calculated. Data smoothing was

done via the `FILTFILT` function in MATLAB [18] so there would be no phase change associated with the smoothing; a third-order Butterworth filter was created from the `BUTTER` function with a normalized cutoff frequency of $0.33\overline{3}$ (from a cutoff frequency of 0.05 Hz and sampling frequency of 30 Hz). All difference plots utilize smoothed data, and other plots that utilize smoothed data are labeled as such. Calibration runs to characterize the silt settling without the ultrasonic transducer operating were taken. Three runs of each configuration were tested. The ranges tested were operating the pump at 2.5V and 3V, which corresponds to 6 to 7.4 mL/s. Additionally, the amplitude of the sine wave sent to the box was varied between 100, 150, and 200 mVpp which resulted in true power ranging from 1 to 10 W into the tank.

Trials were made comparable by subtracting the smoothed normalized outlet from the inlet box trend and starting at 0 initial difference. This approach enabled visualization of changes via the relative difference, as in the outlet box became clearer faster than the inlet box. One such plot is given in Figure 4. The average value of the difference was used as the metric to compare runs.




Fig. 4 Inlet Outlet Difference for 6 mL/s flow rate with 3W true power to acoustic chamber

Before each experiment, the same stirring technique was used. The output side of the tank was fully agitated first, then the input side. Once the input side was fully agitated, the output side of the tank was re-agitated for ten more seconds. Then, the whole system was left to settle for ten seconds to avoid high-rate settling of the particulate which would skew the results. The pump was left running during the stirring process to avoid any silt settling in the resonant chamber.

In addition, we retuned the UltraSep excitation frequency in each experiment to ensure that we excited a resonant frequency in the fluid—near the center frequency of the piezoelectric transducer (2 MHz)—to maximize the true power delivered to the UltraSep system. Because of the mismatch in the order of magnitude of the piezoelectric transducer impedance at resonances (single Ω), compared to the output impedance of the power amplifier (50 Ω), we include a 50:3 power transformer between the power amplifier and the transducer. Figure 5 shows an example impedance measurement, consisting of the UltraSep device impedance magnitude (blue) and phase (orange) as a function of frequency. Here, the impedance of the piezoelectric transducer and the fluid resonator are coupled - the gradual trend with minimum at 2 MHz and maximum at approximately 2.1 MHz correspond to the transducer, and the fine peaks spaced approximately every 44 kHz correspond to the fluid resonances. We select the excitation frequency to maximize the true power delivered to the transducer by selecting the fluid resonance near 2 MHz with impedance amplitude as close to 50 Ω $\angle 0^\circ$ as possible to match the power amplifier output impedance. In the example shown in 5, we identify 1.96 MHz to be the optimal operating frequency as it corresponds to an impedance of 57.07 Ω $\angle 0.03^\circ$, which was very close to the 50 Ω 0.00° phase impedance needed for perfect impedance matching.




Fig. 5 Impedance measurement of the UltraSep device

Numerical Simulations

Empirically, we observe fluctuations in the resonance frequency during the silt filtration process. This is problematic in practice because when the excitation and resonance frequency differ, the standing wave amplitude drops drastically, reducing the ability to filter the silt. We hypothesize that this is caused by the fluctuating and inhomogeneous density and sound speed in the UltraSep device as the particle distribution changes throughout the separation process. To confirm this hypothesis, we simulate the acoustic field in the UltraSep chamber and the acoustic force on the particles with different particle concentrations. COMSOL Multiphysics[®] was used to model the acoustic chamber in 2D from a horizontal cross section configuration.

Figure 6 shows the 2D simulation setup. The blue 17mm by 28mm rectangle represents the fluid. The glass walls are represented by two gray 1mm by 28mm rectangles, and the piezoelectric transducer is represented by the brown 1mm by 25mm rectangle. While this figure shows the whole acoustic chamber, in our work, we assumed symmetry across the $y = 14$ mm axis to decrease computation time.




Fig. 6 Model geometry of horizontal cross section view of acoustic chamber

We used COMSOL built-in material properties for the lead zirconate titanate (PZT-5H) transducer, silica glass walls, and water host fluid. For the silt particles, we used the density, $\rho_p = 2690 \text{ kg/m}^3$, and bulk modulus, $K_p = 3.6 \cdot 10^{10} \text{ Pa}$, values from Wang et al. 2018 [19] to calculate the speed of sound, $c_p = \sqrt{\frac{K_p}{\rho_p}} = 3658 \text{ m/s}$. Additionally, the silt particle sizes were measured experimentally using a Microtrac BLUEWAVE particle size analyzer. The distribution of particle sizes can be found in Figure 7. To simplify the ARP computation, we assume a single particle diameter of $5.5 \mu\text{m}$, the statistical mode of the particle distribution.




Fig. 7 Particle Size Distribution

The material properties of the bulk medium in the chamber is a function of particle concentration according to rule-of-mixtures formulae,

$$\begin{aligned} c_{eff}(x, y) &= \phi(x, y) \cdot c_p + (1 - \phi(x, y)) \cdot c_f, \text{ and} \\ \rho_{eff}(x, y) &= \phi(x, y) \cdot \rho_p + (1 - \phi(x, y)) \cdot \rho_f. \end{aligned} \quad (6)$$

$\phi(x, y)$ is the spatially varying particle concentration in the acoustic chamber. c_p and c_f are the speed of sound, and ρ_p and ρ_f are the density of the particle and fluid, respectively.

To model the coupled electro-mechanical-acoustic resonance behavior of the system, we used COMSOL's Solid mechanics, Pressure acoustics (frequency domain), and Electrostatics physics modules, as well as the Piezoelectric and Acoustic-structure interface multiphysics modules.

A free quadrilateral mesh was used for the transducer and glass walls, and a mapped mesh was used for the fluid domain. We performed a mesh convergence study, measuring the average von Mises stress in the structural domain, and the average absolute total acoustic pressure in the fluid domain at 2.075MHz.

We decreased the maximum element size of the structural domain from $\frac{1}{2}$ (transducer thickness) to $\frac{1}{10}$ (transducer thickness). The pressure and stress values computed at each mesh size demonstrate asymptotic convergence as shown in Figure 8. We chose the maximum element size to be $\frac{1}{10} \lambda = 0.1 \text{ mm}$.

We decreased the maximum element size of the structural domain from $\frac{1}{10} \lambda$ to $\frac{1}{19} \lambda$, for $\lambda = \frac{\text{fluid sound speed}}{\text{maximum frequency}}$. In a mesh refinement study, we fitted the values from the code to the model $y^* = y(\Delta x) + \beta(\Delta x^p)$, where β is a multiplication constant and p is the rate of convergence. Figure 9 shows the predicted value from the simulation $y(\Delta x)$, the extrapolated solution y^* from the model, truncation error $y - y^*$, and numerical uncertainty, which is related to the grid convergence index (GCI). In this paper, we used a mesh size of $\frac{1}{14} \lambda = 0.0506803 \text{ mm}$, which results in a GCI of 1.061% and 0.004% for the pressure and stress values respectively.

A frequency sweep study, from 1.9MHz to 2.1, was calculated with uniform particle concentrations of 0%, 1%, 5%, and 10%.

The COMSOL simulation described above can be also be used to solve Equation 7 using COMSOL's Coefficient Form PDE. This results in the particle concentration changing over time in the acoustic field. It is important to note that this method assumes that the resonance frequency and acoustic field are not changing as the particles move to the nodes. However, this serves as an approximation of how the particle distribution changes in a given acoustic field.

Dynamic Density Functional Theory (DDFT) is used to describe the change in particle concentration over time considering entropy $\mu_{ent}(\phi)$, external forces $\mu_{ex}(\phi)$, and interparticle crowding $\mu_{int}(\phi)$ [20]. Entropy describes the random motion of the particles in the fluid. The only external force we consider is the acoustic radiation force. We neglect viscous drag force, as this is a transient force that does not govern the location of the steady state solution. Interparticle crowding imposes a physical constraint on the particle behavior, accounting for particles colliding into each other. The change in particle concentration is described by the partial differential equation,

$$\frac{\partial \phi}{\partial t} = \nabla \cdot [\phi \nabla (\mu_{ent}(\phi) + \mu_{ex}(\phi) + \mu_{int}(\phi))], \quad (7)$$




Fig. 8 Mesh convergence of the transducer and glass geometries




Fig. 9 Mesh refinement study of the fluid geometry

where the forcing terms are

$$\begin{aligned}
 \mu_{ent}(\phi) &= k_B T \ln \phi \\
 \mu_{ex}(\phi) &= U \\
 \mu_{int}(\phi) &= \frac{1}{(1 - \phi)^2}.
 \end{aligned} \tag{8}$$

Future work in this area can include fine-tuning the time length and step that the solver solves for. Additionally, combining the frequency sweep and PDE studies in the iterative process shown in Figure 10 can give an approximation of how the resonance evolves with the changing particle concentration. In this workflow, given a particle concentration, the effective speed of sound and density in the UltraSep chamber is computed. A frequency sweep is calculated to find the resonance frequency at this particular particle concentration. The acoustic radiation potential (ARP), which governs the movement of the particles, is computed at the resonance frequency using Equation 2. The new distribution of particles is then found by solving Equation 7, as outlined previously. This new distribution of particles is used to find the new effective speed of sound and density. In this way, changes resonance frequency can be measured as the particles are forced to the nodes throughout the UltraSep process.




Fig. 10 Iterative process to track resonant behavior throughout UltraSep process

Results

Following the experimental procedure from Methods, we implement the UltraSep experimental method to separate silt from water.

Figure 11 shows the inlet (blue) and outlet (orange) turbidity measurements over time during an experimental and control trial. Both the inlet and outlet regions become less turbid over time, as a direct result of the agitated silt beginning to settle due to gravity. However, across their respective ranges of turbidity, the outlet box much more quickly becomes clearer than the inlet box, which can be attributed to the filtering effect of the acoustic chamber. The control trends mirror one another closely, and both become more clear at around the same rate as gravity settles the silt. Note that the starting conditions do not reflect the efficiency of the clarification and differing starting conditions are addressed by the difference method.




Fig. 11 Smoothed Normalized Inlet and Outlet Trends for 6 mL/s flow rate with 3W true power to acoustic chamber (experimental) and control conditions

In practice, operators will pass the UltraSep device through the water. As a result, it is likely that the device will be required to work for a range of fluid flow rates. At higher flow rates, the particles spend less time within the UltraSep device. Thus, to ensure that they are aggregated while in the device, a stronger acoustic field will be required. To understand how the required power changes with flow rate, we test different combinations of flow rate and true excitation power in the range of 6-7.4 mL/s and 2-10 W, respectively.

Overall, for each flow rate there was an optimal true excitation power, as an inadequate flow rate meant that the particles were trapped by the field, but an excessive flow rate would pull the particles through the chamber before they were trapped

by the acoustic chamber as demonstrated in Figures 12 and 13. For the slower flow rate of 6 mL/s or 2.5V pump there is a moderate negative correlation, which could be attributed to an increasing acoustic field strength preventing trapped particles from flocculating out. However, for the faster flow rate of 7.4 mL/s, better filtration occurs at a higher acoustic field strength. This is likely due to the increasing flow speed requiring a stronger field to force the particles together before they flow out of the chamber. Both of these trends were subjected to a T-test for correlation coefficient and were significant at the 5% significance level. The 6 mL/s flow trend had a $t=3.007$ and 7.4 mL/s flow trend had a $t=3.915$ both greater than the critical




Fig. 12 6 mL/s Flow Rate Trials




Fig. 13 7.4 mL/s Flow Rate Trials

t value of 2.365. We hypothesize that the variability in data is largely due to inhomogeneous mixing of the tanks as well as the fact that silt above the regions contributed to the amount of light in the inlet and outlet.

The trials were also compared with the excitation voltage to the piezoelectric. For each combination of parameters (flow rate and excitation voltage), a Student T test with $p = 0.05$ ($T_{crit} = 2.92$) was performed against control trials with the UltraSep device unpowered, using the same flow rate. The results are shown in Table 1. At the 5% significance level the 100 mVpp field strength at 6 mL/s and the 200 mVpp at 7.4 mL/s settings filter more effectively than their respective control groups of equal flow rate. This was conducted under the standard assumptions of the T-test.

Table 1 Average Maximum Pixel Difference Experimental Results

Pump Flow Rate (mL/s)	Wave Amplitude (mVpp)	Average Max Pixel Difference	T Value
6	100	0.0773	4.935
6	150	0.0247	1.445
6	200	0.0002	0.2536
7.4	100	-0.0078	-0.2077
7.4	150	0.0520	2.1035
7.4	200	0.0725	3.4200

Critical to the performance of the UltraSep device, however, was operating at a resonance peak. Although an initial frequency sweep was run in free flow conditions with the chamber turned off, once silt began to be captured by the chamber the density of the bulk fluid medium would become inhomogeneous, affecting the location of the resonance.

Following the simulation set up outlined in Methods, we obtained results of the COMSOL simulated frequency sweep with differing uniform particle concentrations. These results are shown in Figure 14. Each line represents a frequency sweep at different uniform particle concentrations of 0%, 1%, 5%, and 10%. Note each line is normalized to the maximum pressure amplitude. As can be seen in Figure 14, the resonance frequencies change with different uniform particle concentrations. This shows that the resonance frequency of the UltraSep device is related to the concentration of particles within the chamber. Further study is needed to investigate how changes in the particle distribution, including nonuniform distributions, impact the resonance frequency.




Fig. 14 Frequency sweeps with different uniform particle distributions

Figure 15 shows the initial uniform particle concentration of 5% as it changes in an acoustic field of 1.963MHz. Throughout the timeseries, the concentration of particles at the pressure nodes increases; this reflects the UltraSep process of particles aggregating at the nodes. The separation is most apparent at the midpoint of the chamber, where the standing wave is most stable. The open condition at the edge of the tank and termination of the piezoelectric make the reflection of the acoustic wave more imperfect, therefore the particle aggregation at the nodes is less pronounced.




Fig. 15 Particle concentration over time in an acoustic field at 1.963MHz, initial uniform particle concentration of 5%.

Conclusion

Experimentally, we were able to characterize the performance of UltraSep in varying conditions by adjusting flow rate and the strength of the signal delivered to the ultrasonic transducer. The results of the experiments thoroughly demonstrate that UltraSep is a viable technique for fluidic particle separation—relying on flocculation for the clarification of the fluid. The results also suggest that there are points of optimal performance for the process and further work should be done to characterize these optimum points.

Via simulations we demonstrated that the resonant frequency is highly dependent on particle concentrations. This suggests the need for an acoustic filtering device to have a feedback loop that constantly adjusts to where a resonant frequency occurs. Future work would involve an iterative simulation framework, outlined in Methods, where the results of a frequency sweep starting with a constant concentration are used to find where the particles would move, which would generate a new frequency sweep to continue the cycle. In this way patterns in the change in resonance could be uncovered as the distribution of particles changes. However, it has been demonstrated experimentally that even without this control system the acoustic filtration is successfully able to filter the particulate out at a faster rate than gravity alone.

Acknowledgments This research was funded by Los Alamos National Laboratory (LANL) through the Engineering Institute's Los Alamos Dynamics Summer School. The Engineering Institute is a research and education collaboration between LANL and the University of California San Diego's Jacobs School of Engineering. This collaboration seeks to promote multidisciplinary engineering research that develops and integrates advanced predictive modeling, novel sensing systems, and new developments in information technology to address LANL mission-relevant problems. The authors would like to thank Dr. François Hemez, Lawrence Livermore National Laboratory, for his technical advice about the mesh refinement study.

References

1. Wang, Y., Song, W., Fortino, G., Qi, L.Z., Zhang, W., and Liotta, A. "An experimental-based review of image enhancement and image restoration methods for underwater imaging". *IEEE Access*, 7:140233–140251 (2019)ifnextchar.gobble.
2. Alajmi, E. and Alqenaei, A. "Underwater welding techniques". *International Journal of Engineering Research and Applications*, 7(2):14–17 (2017)ifnextchar.gobble.
3. Chemisky, B., Menna, F., Nocerino, E., and Drap, P. "Underwater survey for oil and gas industry: A review of close range optical methods". *Remote Sensing*, 13(14) (2021)ifnextchar.gobble.
4. Penrice, D. and Cooper, J.S. *Diving Casualties*. StatPearls Publishing, Treasure Island (FL) (2025)ifnextchar.gobble.
5. Tillmans, F. "Low-visibility diving - divers alert network" (2022)ifnextchar.gobble.
6. FOSTER, E.J. "A diver-operated underwater metal detector". *Archaeometry*, 12(2):161–166 (1970)ifnextchar.gobble.
7. Schultz, J.J., Healy, C.A., Parker, K., and Lowers, B. "Detecting submerged objects: The application of side scan sonar to forensic contexts". *Forensic Science International*, 231(1):306–316 (2013)ifnextchar.gobble.
8. Kapali, A., Neopane, H., Chitrakar, S., Kayastha, A., and Shrestha, O. "Experimental and cfd study of influence of sediment size on efficiency of hydrocyclone for use as sediment separation device". *Journal of physics. Conference series*, 1608(1):12014 (2020)ifnextchar.gobble.
9. Gao, Y., Wu, M., Lin, Y., and Xu, J. "Acoustic microfluidic separation techniques and bioapplications: A review". *Micromachines*, 11(10) (2020)ifnextchar.gobble.
10. Wu, M., Ozcelik, A., Rufo, J., Wang, Z., Fang, R., and Jun Huang, T. "Acoustofluidic separation of cells and particles". *Microsystems and Nanoengineering*, 5(1) (2019)ifnextchar.gobble.
11. Martinez-Marchese, A., Ansari, M., Marzo, A., Wang, M., Soo, S., and Toyserkani, E. "Ultrasound particle lensing: Powder stream control in directed energy deposition using acoustic radiation force fields". *Additive Manufacturing*, 73:103698 (2023)ifnextchar.gobble.
12. Greenhall, J. and Raeymaekers, B. "3d printing macroscale engineered materials using ultrasound directed self-assembly and stereolithography". *Advanced Materials Technologies*, 2(9):1700122 (2017)ifnextchar.gobble.
13. Hawkes, J.J., Coakley, W.T., Gröschl, M., Benes, E., Armstrong, S., Tasker, P.J., and Nowotny, H. "Single half-wavelength ultrasonic particle filter: Predictions of the transfer matrix multilayer resonator model and experimental filtration results". *The Journal of the Acoustical Society of America*, 111(3):1259–1266 (2002)ifnextchar.gobble.
14. Drinkwater, B.W. "Dynamic-field devices for the ultrasonic manipulation of microparticles". *Lab Chip*, 16:2360–2375 (2016)ifnextchar.gobble.
15. "Comsol multiphysics® v. 6.3. www.comsol.com. comsol ab, stockholm, sweden."
16. Zhang, Y. and Chen, X. "Particle separation in microfluidics using different modal ultrasonic standing waves". *Ultrasonics Sonochemistry*, 75:105603 (2021)ifnextchar.gobble.
17. Barmatz, M. and Collas, P. "Acoustic radiation potential on a sphere in plane, cylindrical, and spherical standing wave fields". *The Journal of the Acoustical Society of America*, 77(3):928–945 (1985)ifnextchar.gobble.
18. Inc., T.M. "Matlab version: 24.2.0.2806996 (r2024b update 3)" (2024)ifnextchar.gobble.
19. Wang, J., Liu, B., Kan, G., Li, G., Zheng, J., and Meng, X. "Frequency dependence of sound speed and attenuation in fine-grained sediments from 25 to 250 khz based on a probe method". *Ocean Engineering*, 160:45–53 (2018)ifnextchar.gobble.
20. Rauscher, M. *Dynamic Density Functional Theory (DDFT)*, pages 1–8. Springer US, Boston, MA (2013)ifnextchar.gobble.



Chapter 7

Extending Rattlesnake Vibration Controller to MIMO Sine Control and Other Significant Improvements

Cody M. Langston and Dan Rohe

Abstract With the field of multi-axis vibration exploding in recent years, new control strategies and algorithms are being simultaneously developed to provide more accurate test responses with less force, voltage, and current required. Since the test articles often undergo simultaneous environments while they are in service, significant research is being performed to combine these environments together for more realistic laboratory testing. The Rattlesnake Vibration Controller was developed as an open controller framework to allow researchers to perform multiple-input, multiple-output, combined environments research and development without being constrained to closed-source commercial software. Since its initial presentation four years ago, Rattlesnake has become the workhorse for MIMO vibration testing at Sandia National Laboratories. Recently, Rattlesnake has introduced a new control environment, the MIMO Swept Sine environment. The goal of this environment is to sweep through frequencies and excite the structure so that the response at specified control points matches a sine wave with a predetermined amplitude and phase. Rattlesnake uses a MIMO swept sine algorithm that combines both feedforward and feedback control to generate a drive signal that produces the desired response. This algorithm can adjust to nonlinearities in the structure and noise during measurement to produce consistent responses. This paper outlines this algorithm and provides an example workflow on how to use Rattlesnake to implement MIMO Swept Sine tests for vibration testing. Additionally, this paper gives an update on Rattlesnake's development with an emphasis on transitioning Rattlesnake into a production quality software and expanding Rattlesnake's open-source community.

Keywords Multi-Input Multi-Output · Swept Sine · Rattlesnake · Open-Source · Vibration Testing

Introduction

Multi-input, Multi-output (MIMO) control is becoming a widely adopted alternative to single-axis testing for environmental qualification and modal tests. For vibration tests, MIMO control allows engineers to test how structures respond to inputs at multiple degrees of freedom simultaneously which is a better representation of the environments the structure will experience [5, 4]. An added benefit of testing the structure in multiple axes at a time is that it reduces the number of tests required to fully quantify the structure resulting in both speed and cost benefits [2].

Rattlesnake is the in-house MIMO vibration controller software that Sandia uses to perform environmental testing. Rattlesnake is open-source and free to download at <https://github.com/sandialabs/rattlesnake-vibration-controller>. Rattlesnake communicates with existing data acquisition hardware to collect data from input devices (accelerometers, load cells, etc.), processes that data to calculate output signals, and sends those output signals to output devices (shakers, motors, etc.). The software currently supports National Instruments, Lan-XI, and Data Physics data acquisition hardware as well as Exodus, Sdynpy, and State Space virtual hardware. Rattlesnake is structured around multiple distinct testing environments which can be run individually or together to support tests such as random vibration, hammer, or sine sweep tests. Five different

Cody M. Langston
University of Georgia, College of Engineering
e-mail: cml34135@uga.edu

Dan Rohe
Sandia National Labs, Experimental Structural Dynamics
dprohe@sandia.gov

environments have been implemented in the current iteration of Rattlesnake: MIMO Random Vibration, MIMO Transient, MIMO Swept Sine, Modal, and Time History environments.

The most recent improvement to the Rattlesnake software is the addition of the MIMO Swept Sine environment. The goal of MIMO swept sine control is sweep through forcing frequencies to excite the structure in a way so that its response matches a sine wave with predetermined amplitude and phase. The sensors used to excite the structure are called drive signals and the sensors which respond with the desired sine waves are called control signals. One advantage of this testing technique is that the amplitudes of the measured responses are predetermined which allows engineers to design tests with high excitation levels and low signal to noise ratio resulting in higher quality FRFs. The aerospace industry regularly uses MIMO swept sine control when testing fully assembled structures as it allows them to guarantee that the structures will respond at safe levels [2]. This testing requires the usage of both open and closed loop control where the output signals are initially predicted by performing a system identification on the structure, and then adjusted during the test when the structure's response deviates from the desired amplitude and phase. The closed loop portion of the control can be used to classify non-linear behavior in the structure as it adjusts for differences between the response of the structure and the prediction made with linear assumptions. Some non-linear behavior is common during these tests due to their ability to safely excite the structure at higher levels than other tests.

Background

Rattlesnake's MIMO Swept Sine algorithm follows the structure outlined by [4]. This algorithm can be separated into four steps: specifying the desired response at the control channels, performing a system identification to predict the drive signals, choosing a method to track the amplitude and phase of the signal during the test, and performing the feedback control to adjust the drive signal when the control signal deviates from the desired signal.

The first step is to define the number of control channels, l , and number of drive channels, m , that will be used in this test. The desired signal is defined by specifying an amplitude and phase at each frequency for each control channel as well as the sweep rate between frequency breakpoints that determines the instantaneous frequency over time. Once the desired signal is defined, a system identification step must be performed where initial FRFs are computed with a low amplitude random vibration test. These initial FRFs are used to predict the drive signal at each frequency using the following relationship:

$$y = Hu, \quad (1)$$

where $y \in \mathbb{C}^{1 \times l}$ is the control response, $H \in \mathbb{C}^{l \times m}$ is the FRF matrix at that given frequency, and $u \in \mathbb{C}^{m \times 1}$ is the drive signal. Given a desired response, the required drive signal can be calculated with:

$$u_{opt} = \hat{Z}r, \quad (2)$$

where $u_{opt} \in \mathbb{C}^{1 \times m}$ is the predicted drive signal, $\hat{Z} \in \mathbb{C}^{l \times m}$ is the inverse of the FRF matrix, and $r \in \mathbb{C}^{l \times 1}$ is the desired control channel response. $\hat{Z}$ is usually computed with a Moore-Penrose pseudoinverse since the FRF is not always a square matrix. Various regularization and weighting parameters may be defined to tune the inversion process.

Due to non-linearities in the structure and error in the initial FRF calculation, the predicted drive signal will not generate the desired control signal. To account for this, the drive signal is adjusted in real time using a feedback control loop. The goal of this feedback control is to minimize the error:

$$e = r - y. \quad (3)$$

The desired response, r , is specified before hand, but the actual response, y , needs to be computed by tracking the amplitude and phase of the signal during data collection. A common way of doing this is by using an on-line frequency filter such as the harmonic estimator or digital tracking filter to estimate these values [4]. Another potential method of tracking the amplitude and phase is by using an Overlapped Vold-Kalman Filter outlined in [3, 6, 1].

The feedback control loop adjusts the drive signal by solving the following least squares optimization problem:

$$\min_u e^H e = u^H H^H H u - u^H H^H r - r^H H u + r^H r. \quad (4)$$

This optimization problem is typically solved using a gradient descent method which results in the step:

$$u_{n+1} = u_n + \alpha H^H e_n, \quad (5)$$

where α is the convergence coefficient which determines the step size. [4] showed that this feedback control is guaranteed to be Bounded-Input Bounded-Output stable when α is calculated as:

$$\alpha_{opt} = \frac{2}{\sigma_{max}^2}, \quad (6)$$

where σ_{max}^2 is the maximum eigenvalue resulting from the single value decomposition of the FRF matrix, H , which was calculated during the system identification of the structure.

Unlike many commercial sine control systems, Rattlesnake was designed from the start for "frame-based" data acquisition and output, similar to the ones that exist in standard Random Vibration Controllers. Instead of performing control on a sample-by-sample basis, Rattlesnake reads data in blocks (the size of which can be specified in the software). Additionally, when using the Vold-Kalman Filter, the measurement block may be constrained to a larger size due to processing and end-effects of the filter. Therefore, Rattlesnake will generally not be as responsive as other commercial control systems.

The current feedback implementation in Rattlesnake is a simple error correction routine. The complex error e at each frequency line in a measurement block is multiplied by $\hat{Z}$ to create a drive correction factor for that block. This is added to a running block drive correction factor after being scaled by the maximum drive correction Eq. 6. Future drive signals are modified by this correction factor.

Rattlesnake MIMO Sine Environment

Rattlesnake works as a wrapper around this control logic to allow MIMO Swept Sine tests to be set up and run seamlessly. To setup a MIMO Swept Sine test in Rattlesnake, choose the MIMO SINE VIBRATION option on the SELECT CONTROLLER TYPE dialog box shown in Fig. 1.




Fig. 1 Selecting MIMO Sine Vibration environment in Rattlesnake

Next, the user will setup their data acquisition hardware and its corresponding channels in the DATA ACQUISITION SETUP tab shown in Fig. 2. The Rattlesnake User's Manual found on GitHub has a comprehensive guide on how to set up a channel table. For this test, the channels that are defined with a FEEDBACK DEVICE will be used as the drive channels for the test.

The control channels are specified in the CONTROL CHANNELS table in the center of the ENVIRONMENTAL DEFINITION tab shown in Fig. 3. The relationship between the control and drive channels can be further defined by clicking on the TRANSFORMATION MATRICES button. The desired response from the control channels is defined with the BREAKPOINT table. This table consists of a series of breakpoints containing the desired amplitudes and phases of the control channels at the specified frequency. The desired response is generated by connecting these breakpoints through linear or logarithmic interpolation over time. Optionally, the warning and abort tables are used to define signals which will warn the user and abort the test if the system's response goes above those levels. During the test, the drive signals will start at the first breakpoint and stop at the last breakpoint. Additional sine tones are added by clicking the plus button on the tab group to the left of the BREAKPOINT table. These tones will run simultaneously and be superimposed on each other during the test. The SPECIFICATION PLOTS at the bottom of the window are used to track which tones will be playing at a given time.

The filter used to track the amplitude and phase of the sine tones during the test is chosen with the TRACKING FILTER PARAMETERS group. Rattlesnake currently supports a Digital Tracking Filter and an Overlapped Vold-Kalman Filter. After selecting a filter, the parameters required to define that filter will also be in this group. Clicking the EXPLORE... button will allow the user to vary the filter settings and apply the filters to the specification signal to visualize the results, including in the presence of noise. This can allow the user to dial in the filter settings without actually running the test.




Fig. 2 Setting up data acquisition hardware and channels on the DATA ACQUISITION SETUP tab of Rattlesnake




Fig. 3 Defining the desired sine signals and tracking filter on the ENVIRONMENT DEFINITION tab of Rattlesnake

The next step is to perform a system identification on the SYSTEM IDENTIFICATION tab shown in Fig. 4. Here, the user will set up the random vibration test used to generate the initial FRFs. Once the test is defined, hit the Start button to perform the random vibration test and compute the initial FRFs.

Similar to the MIMO Random Vibration and Transient environments in Rattlesnake, the Swept Sine environment allows users to load in custom control laws defined as Python classes with specific methods defined (`__init__`, `system_id_update`, `initialize_control`, `update_control`, `generate_signal`, and `finalize_control`), fundamentally changing how the Sine environment operates without having to modify the environment itself. These Python scripts can be assigned to the test by clicking the LOAD button within the CUSTOM PYTHON CONTROL SCRIPT section.




Fig. 4 Performing the system identification on the SYSTEM IDENTIFICATION tab of Rattlesnake

Performing a system identification will populate the TEST PREDICTIONS tab shown in Fig. 5. This tab gives the user an idea of the accuracy expected from the current test setup and acts as a sanity check to make sure that the test will not cause harm to the structure or devices.

If the predicted signals meet the required specifications, the software is ready to run the full test using the RUN TEST tab shown in Fig. 6. This is done by clicking on the ARM DATA ACQUISITION and START ENVIRONMENT buttons. As the test runs, the amplitude of the sine waves at each control channel will be displayed with the desired signal for that control channel. The real (x-axis) and imaginary (y-axis) adjustments made to the predicted drive signal by the feedback control loop are shown on the CONTROL UPDATES plot in the upper left corner. After a complete sine sweep, Rattlesnake provides the ability to update the feed-forward drive signal to make corrections for the next pass, allowing a user to run the entire sweep at a lower level to dial in the response prior to the full-level test. The drive and control signals from the test can be saved to a NetCDF4 file for further processing.

The MIMO Swept Sine environment has been shown to be very functional, especially in conjunction with the combined environments capabilities in Rattlesnake. Recently, the capability was utilized to run a mixed mode test controlling to 9 separate sine sweeps each controlling unique X,Y,Z table translations for a 6-DoF shaker (27 total sine signals), while simultaneously controlling to a random vibration profile in the background. Rattlesnake is able to synchronize these sine components within a single MIMO Sine Environment, allowing the instantaneous amplitudes and phases to be accurately computed for all sine tones. The implementation of the Vold-Kalman filter for order tracking allows for these sine tones to




Fig. 5 Plotted test predictions on the TEST PREDICTIONS tab of Rattlesnake




Fig. 6 Running the MIMO Swept Sine test on the TEST PREDICTIONS tab of Rattlesnake. The desired signal is orange and the measured signal is in blue.

cross in frequency without significantly impacting the ability to extract the amplitude and phases of the tones. Rattlesnake's flexibility allowed this challenging test to successfully be performed.

Rattlesnake Development Updates

Rattlesnake is an open-source codebase written in Python that allows the user to customize it to their specific needs. The current structure of Rattlesnake is that it collects data from a Hardware class which defines how to communicate with the data acquisition hardware. The collected data is sent out to Environment classes which process the data and determine the drive signals if necessary. The Environment classes send the drive signals back to the Hardware class where they are combined and output to the data acquisition hardware. This structure is key to Rattlesnake's usability for research applications as it allows the user to focus on programming their own control laws or environments, while using the Rattlesnake's existing implementation for aspects such as hardware communication, data streaming, etc.

In the past year, the Rattlesnake development team has made significant changes to transition Rattlesnake into a production quality product. These changes include a major overhaul to the development cycle and improvements to the ease of use, stability, and capabilities of Rattlesnake. Previously, Rattlesnake has been developed in house at Sandia with the codebase being pushed to the public GitHub whenever significant changes were made to the software. To avoid the complexity of maintaining two separate codebases, Rattlesnake will now use the public GitHub as its main repository. This change means that external users will get updates sooner and opens up the potential for collaboration without having to merge changes between two separate codebases.

With this transition, Rattlesnake is moving to a Continuous Integration/Continuous Delivery (CI/CD) development pipeline where changes to the software will be frequently committed to a central development branch. This CI/CD pipeline includes the introduction of unit tests which automatically run isolated portions of the software to detect possible bugs introduced when committing to the codebase. These tests must be maintained as they act as a safety net that allow Rattlesnake to be iterated on frequently with confidence that the changes will not introduce new issues into the software. Currently, these unit tests cover 45% of Rattlesnake's codebase.

Rattlesnake's transition to a production quality codebase also includes improvements to the software's ease of use. Rattlesnake has been added to the pip package manager to make it easier to install. Additionally, Rattlesnake is being designed to validate the user's inputs which improves error messaging and prevents the user from crashing the software. With this validation, the Rattlesnake UI is now capable of suggesting valid inputs for values such as data acquisition hardware, channel names, and voltage ranges. The Rattlesnake documentation now has its own website with improved search functionality. The documentation can be found at <https://sandialabs.github.io/rattlesnake-vibration-controller/book/>.

The development team hopes that with these changes, it will be easier to implement Rattlesnake into the workflow of other laboratories. The open-source nature of Rattlesnake allows others to make meaningful contributions to the software that will help others. The transition to the public GitHub and introduction of CI/CD development streamlines this process improving version control and bug tracking. It is now easier than ever to join the Rattlesnake open-source community.

Conclusion

This paper outlines the recent addition of MIMO Swept Sine control to the Rattlesnake vibration controller. MIMO Swept Sine control allows engineers to design the response of a structure to match specifications required for their testing purposes. This allows the test to reach higher excitation levels without the risk of damage to the structure or sensors. Rattlesnake streamlines the process of conducting a MIMO Swept Sine test by allowing the user to define their drive and control channels, shape the desired response at those control channels, and define parameters that are required for the test such as the tracking filter and system identification. This workflow has been used to conduct MIMO Swept Sine tests at Sandia National Laboratories. This paper also provides updates on Rattlesnake's development cycle with particular focus on improvements made to transition Rattlesnake into a production quality software.

Acknowledgments Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC (NTESS), a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration (DOE/NNSA) under contract DE-NA0003525. This written work is authored by an employee of NTESS. The employee, not NTESS, owns the right, title and interest in and to the written work and is responsible for its contents. Any subjective views or opinions that might be expressed in the written work do not necessarily represent the views of the U.S. Government. The publisher acknowledges

that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this written work or allow others to do so, for U.S. Government purposes. The DOE will provide public access to results of federally sponsored research in accordance with the DOE Public Access Plan.

References

1. J.R. Blough. Understanding the kalman/vold-kalman order tracking filters' formulation and behavior. 05 2007.
2. G. Gloth and M. Sinapius. Analysis of swept-sine runs during modal identification. *Mechanical Systems and Signal Processing*, 18(6):1421–1441, 2004.
3. H. Herlufsen, S. Gade, H. Konstantin-Hansen, and H. Vold. Characteristics of the vold-kalman order tracking filter. In *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.00CH37100)*, volume 6, pages 3895–3898 vol.6, 2000.
4. Umberto Musella, Bart Peeters, Francesco Marulo, and Patrick Guillaume. Multi-input multi-output swept sine control: A steepest descent solution for a challenging problem. In Nikolaos Dervilis, editor, *Special Topics in Structural Dynamics & Experimental Techniques, Volume 5*, pages 83–94. Springer International Publishing, 2020.
5. Daniel Rohe, Ryan Schultz, and Norman Hunter. Rattlesnake: An open-source multi-axis and combined environments vibration controller. Sandia National Laboratories (SNL-NM), Albuquerque, NM (United States), 10 2021.
6. Håvard Vold and Jan Leuridan. High resolution order tracking at extreme slew rates using kalman tracking filters. *Shock and Vibration*, 2:507–515, 01 1995.



Chapter 10

Experimental Evaluation of Three-Dimensional Mode Shapes and Frame-Consistent Comparison with Numerical Models

Umidjon Usmanov, Giuseppe Battiato, and Christian Maria Firrone

Abstract Experimental modal analysis with laser Doppler vibrometry is usually limited to the identification of mono-dimensional mode shapes, typically in the out-of-plane direction. This simplification restricts the understanding of the true three-dimensional deformation of a structure. In this work, we propose an experimental–computational framework to reconstruct full three-dimensional mode shapes from laser doppler vibrometry measurements. The method relies on three complementary test configurations, from which frequency response functions are extracted and combined through an algorithm to recover the complete spatial mode shape. A second contribution of this study addresses the challenge of reference frame inconsistency between experimental and numerical mode shapes. While numerical predictions are often defined in different coordinate systems, the relation between these frames is not trivial for components with complex geometry. We introduce a projection algorithm that transforms numerical mode shapes into the experimental reference frame, enabling a direct and reliable comparison. The combined approach improves the accuracy of experimental mode shape identification and provides a systematic procedure for validation against finite element models.

Keywords Experimental modal analysis · 3D mode shapes · Laser vibrometry · Reference frame projection · vibration analysis

Introduction

Laser Doppler vibrometry has become an essential tool in experimental vibration analysis because it enables non-contact and spatially dense measurements of structural response. In conventional step-scanning procedures, a laser spot is moved across a predefined measurement grid and remains at each point long enough to record the vibration signal. While this approach provides reliable frequency response functions (FRFs) and mode shapes, it is often limited to the out-of-plane direction and requires long acquisition times when the measurement grid is dense or the surface area is large.

To overcome these limitations, more advanced scanning strategies and multi-axis vibrometer systems have been developed to capture richer information about structural vibrations [1–4]. Nevertheless, in most experimental modal analysis campaigns, the evaluation of mode shapes remains essentially one-dimensional, neglecting in-plane components that are critical to fully understanding the dynamic behavior of complex structures. This restriction limits the accuracy of model validation and the reliability of predictive simulations.

In experimental practice, the extraction of complete three-dimensional vibration information remains challenging. A single scanning laser Doppler vibrometer (SLDV) is inherently limited, as it can only capture vibration along its line of sight and therefore cannot simultaneously measure both in-plane and out-of-plane components. To overcome this limitation, commercial three-dimensional SLDV systems have been developed, such as the Polytec PSV-400-3D, PSV-500-3D and PSV QTEC 3D [5]. These systems employ three scanning heads that synchronously direct laser beams to the same measurement point, allowing the full vibration vector to be reconstructed within a defined measurement coordinate system. Among the most relevant contributions to three-dimensional vibration measurement is the first-time work of Chen and Zhu, who investigated the use of a three-head SLDV system operating in a continuous and synchronous scanning mode [6]. In their study, three independent vibrometer heads (Top, Left, and Right) were arranged to simultaneously track the same point along a prescribed trajectory on the surface of the structure. By synchronizing the scanning paths of all three beams, the system was able to record vibration components along three distinct laser directions. Through appropriate geometrical calibration and

Umidjon Usmanov · Giuseppe Battiato · Christian Maria Firrone
Department of Mechanical and Aerospace Engineering
Polytechnic University of Turin, Turin, Italy 10129
e-mail: umidjon.usmanov@polito.it; giuseppe.battiato@polito.it; christian.firrone@polito.it

coordinate transformations, the complete three-dimensional response of the structure was reconstructed in a specified measurement coordinate system. However, this methodology depends on the availability of a specialized 3D-SLDV system, such as the Polytec PSV-500-3D, which integrates multiple scanning heads into a single measurement platform. While technically powerful, the cost of such equipment is extremely high, often exceeding the budget available in many academic or industrial laboratories. In addition, the physical space required to install and calibrate three independent laser heads can make the method impractical in certain test environments, particularly when the geometry of the component restricts the access.

The technique proposed in this work aims to overcome these limitations by relying on only a single scanning head. Instead of relying on simultaneous multi-axis measurements, the laser scanning head remains fixed in one position, while the component under test is arranged in three different configurations. The frequency response functions (FRFs) extracted from these separate tests are then utilized through a dedicated algorithm to reconstruct the complete three-dimensional mode shapes. This strategy significantly reduces instrumentation costs and expands the applicability of 3D experimental modal analysis to a broader range of laboratories and test setups. While it requires sequential rather than simultaneous measurements, the approach preserves the essential capability of identifying full 3D mode shapes, thus offering a practical and cost-effective alternative to multi-head vibrometer systems.

Methodology

The starting point of the proposed methodology is the definition of how a scanning laser Doppler vibrometer (SLDV) records the motion of a vibrating point. Figure 1 shows the schematic representation of the measurement principle. Let point C be located on the surface of the structure. Its motion is described with respect to a user-defined Cartesian reference frame (x, y, z) . This reference frame is body fixed and can be arbitrarily chosen by the operator and will be referred to as the experimental reference frame throughout the paper. The SLDV is positioned at point E , and its orientation relative to the experimental frame is described by two inclination angles: the horizontal angle β and the vertical angle α . These angles represent the deviation of the laser beam from the normal direction to the measured surface. Depending on the instrument type, the velocity measured by the vibrometer can be interpreted in different ways. In the most general case, a mono-dimensional vibrometer provides the projection of the velocity vector along the laser beam axis, i.e., along the line EC . This direction usually does not coincide with any axis of the experimental reference frame. For this reason, the raw measurement cannot be directly assigned to in-plane or out-of-plane vibration components. Some commercial vibrometers incorporate internal correction algorithms that account for the inclination of the laser head. In these cases, the system outputs the velocity component aligned with the out-of-plane direction of the measured surface, which in our notation corresponds to the z -axis. While this feature simplifies the interpretation of the measurements, it is not universally available, and its implementation varies between instruments.

The methodology proposed in this work follows a generic formulation that assumes the vibrometer does not perform such corrections. Accordingly, the measured signal is always considered as the projection of the velocity vector along the laser line. This general assumption allows the derivation of a set of equations that remain valid independently of the specific instrument. For users employing a vibrometer with built-in correction, the same formulation can be recovered by setting the inclination angles to zero, i.e., $\alpha = 0$ and $\beta = 0$. In this way, the proposed methodology provides a unified framework that can be applied both to corrected and uncorrected vibrometer outputs.

The experimental reference frame (x, y, z) is the basis in which all reconstructed 3D mode shapes are expressed. The SLDV records the projection of the motion of point C onto a laser line beam direction. By applying simple trigonometric relations to triangles ECB , EDC and ECA in Figure 1, the laser reading LR of the point C can be expressed as

$$LR = v_x \sin \gamma + v_y \sin \phi + v_z \cos \theta \quad (1)$$

where v_x , v_y and v_z are velocity components in the experimental reference frame. Since only the inclination angles α and β are typically available to the user from the laser equipment, it is necessary to express the beam projection angles γ , ϕ and θ in terms of α and β . The relations can be obtained by applying elementary trigonometric identities to the right-angled triangles EAB , EBC , EDC and EAC shown in Figure 1. The resulting expressions are

$$\sin \gamma = \frac{\tan \beta \cdot \cos \alpha}{\sqrt{1 + \cos^2 \alpha \cdot \tan^2 \beta}}; \quad \sin \phi = \frac{\sin \alpha}{\sqrt{1 + \cos^2 \alpha \cdot \tan^2 \beta}}; \quad \cos \theta = \frac{\cos \alpha}{\sqrt{1 + \cos^2 \alpha \cdot \tan^2 \beta}} \quad (2)$$

By substituting these trigonometric relations into the expression of laser reading (eq. 1), the projection can be compactly written as

$$LR = \frac{1}{\sqrt{1 + \cos^2 \alpha \cdot \tan^2 \beta}} \begin{bmatrix} \tan \beta \cdot \cos \alpha & \sin \alpha & \cos \alpha \end{bmatrix} \begin{Bmatrix} v_x \\ v_y \\ v_z \end{Bmatrix} \quad (3)$$




Fig. 1 Laser vibrometer measurement scheme in the neutral configuration

It is important to note that, in the special case where the vibrometer performs an internal correction of the inclination angles ($\alpha = 0$ and $\beta = 0$), the projection reduces to $LR = v_z$, which corresponds to the direct measurement of the pure out-of-plane motion of the component.




Fig. 2 Laser vibrometer measurement scheme in a rotated configuration

Since our goal is to recover the motion vector $\{v_x, v_y, v_z\}^T$, a single laser reading provides only one scalar projection, so the problem is underdetermined (three unknowns and one equation). To solve a system, we repeat measurement under at least three independent configurations. In practice, this is achieved by reorienting the component (e.g. rotating it about one or

more axis), while the laser head remains fixed. Figure 2 depicts a generic re-orientation of the specimen. Vibrometer has no information about how the component is rotated, the output of the vibrometer still refers to the original experimental frame defined at the beginning of the test. In the scheme, this is shown as the green frame (x', y', z') . For consistency, however, the vibration results must be expressed in the body-fixed frame that follows the orientation of the component $((x, y, z)$ in Figure 2). This requires introducing a transformation that accounts for the relative rotation between the experimental reference frame and the rotated body frame. Any generic rotation of the component can be represented as a sequence of three consecutive rotations about the principal axes of the experimental reference frame. This is commonly referred as the Euler rotation formalism. In this framework, a rotation about the z -axis by an angle δ_z , followed by a rotation about the y -axis by δ_y , and finally a rotation about the x -axis by δ_x , transforms the original reference frame (x, y, z) into the rotated frame (x', y', z') :

$$\begin{aligned} \mathbf{R}_{rot} &= \mathbf{R}_x \mathbf{R}_y \mathbf{R}_z = \begin{bmatrix} \cos\delta_z & \sin\delta_z & 0 \\ -\sin\delta_z & \cos\delta_z & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\delta_x & \sin\delta_x \\ 0 & -\sin\delta_x & \cos\delta_x \end{bmatrix} \begin{bmatrix} \cos\delta_y & 0 & -\sin\delta_y \\ 0 & 1 & 0 \\ \sin\delta_y & 0 & \cos\delta_y \end{bmatrix} \\ &= \begin{bmatrix} \cos\delta_y \cos\delta_z & \cos\delta_y \sin\delta_z & -\sin\delta_y \\ \sin\delta_x \sin\delta_y \cos\delta_z - \cos\delta_x \sin\delta_z & \sin\delta_x \sin\delta_y \sin\delta_z + \cos\delta_x \cos\delta_z & \sin\delta_x \cos\delta_y \\ \cos\delta_x \sin\delta_y \cos\delta_z + \sin\delta_x \sin\delta_z & \cos\delta_x \sin\delta_y \sin\delta_z - \sin\delta_x \cos\delta_z & \cos\delta_x \cos\delta_y \end{bmatrix} \end{aligned} \quad (4)$$

With the three independent configurations, the three laser readings provide three independent projections of the same motion at point C. Stacking these relations yield a 3 by 3 linear system for the unknown component $\{v_x, v_y, v_z\}^T$.

$$\begin{cases} LR_1 = \frac{1}{\sqrt{1+\cos^2\alpha_1 \cdot \tan^2\beta_1}} [\tan\beta_1 \cdot \cos\alpha_1 & \sin\alpha_1 & \cos\alpha_1] \cdot \mathbf{R}_{rot,1} \cdot \begin{Bmatrix} v_x \\ v_y \\ v_z \end{Bmatrix} \\ LR_2 = \frac{1}{\sqrt{1+\cos^2\alpha_2 \cdot \tan^2\beta_2}} [\tan\beta_2 \cdot \cos\alpha_2 & \sin\alpha_2 & \cos\alpha_2] \cdot \mathbf{R}_{rot,2} \cdot \begin{Bmatrix} v_x \\ v_y \\ v_z \end{Bmatrix} \\ LR_3 = \frac{1}{\sqrt{1+\cos^2\alpha_3 \cdot \tan^2\beta_3}} [\tan\beta_3 \cdot \cos\alpha_3 & \sin\alpha_3 & \cos\alpha_3] \cdot \mathbf{R}_{rot,3} \cdot \begin{Bmatrix} x \\ y \\ z \end{Bmatrix} \end{cases} \quad (5)$$

$$\begin{Bmatrix} LR_1 \\ LR_2 \\ LR_3 \end{Bmatrix} = \mathbf{T} \cdot \begin{Bmatrix} v_x \\ v_y \\ v_z \end{Bmatrix} \quad (6)$$

This system can be directly solved through matrix inversion, since the transformation matrix $\mathbf{T}$ is non-singular and well-conditioned as long as the three projection directions are not coplanar or nearly collinear. In practice, this condition is satisfied by choosing distinct rotations in the range of 20–45°. Under such conditions, the inversion is numerically stable and provides the three vibration components $\{v_x, v_y, v_z\}^T$. If more than three configurations are available, the system becomes overdetermined, and the unknowns can be estimated in a least-squares sense, which improves robustness against measurement noise and experimental uncertainties.

$$\begin{Bmatrix} v_x \\ v_y \\ v_z \end{Bmatrix} = \mathbf{T}^{-1} \begin{Bmatrix} LR_1 \\ LR_2 \\ LR_3 \end{Bmatrix} \quad (7)$$

The second part of this work addresses the challenge of comparing experimentally reconstructed mode shapes with numerical predictions. In many cases, finite element (numeric) mode shapes are defined with respect to a numerical reference frame that does not coincide with the experimental frame used during measurements. For simple geometries, the relation between these frames can often be established analytically. However, for components with complex geometry, the relation between the numerical and experimental reference frames is not straightforward and cannot be easily derived. The relation between the mode shape defined in numerical and experimental reference frame can be written in terms of the rotation matrix $\mathbf{R}_{num}$:

$$\begin{Bmatrix} v_x \\ v_y \\ v_z \end{Bmatrix}_{exp} = \mathbf{R}_{num} \begin{Bmatrix} v_x \\ v_y \\ v_z \end{Bmatrix}_{num} \quad (8)$$

the transformation between the numerical and experimental reference frames can be described by a single rotation matrix, denoted $\mathbf{R}_{num}$. This matrix is of size 3×3 and defines how the numerical coordinates $(v_x, v_y, v_z)_{num}^T$ are rotated in order

to align with the experimental coordinates $(v_x, v_y, v_z)_{exp}^T$. For each measurement point of the component, the same rotation applies, meaning that the transformation is consistent across the entire structure. When the mode shape vector is written for all measurement points, the global transformation matrix takes a block-diagonal form. Each block is identical and equal to $\mathbf{R}_{num}$, while the size of the full matrix scales with the number of measurement points. In this way, the entire numerical mode shape can be projected into the experimental frame simply by applying this block-diagonal transformation:

$$\begin{pmatrix} v_{x1} \\ v_{y1} \\ v_{z1} \\ v_{x2} \\ v_{y2} \\ v_{z2} \\ \vdots \\ v_{xn} \\ v_{yn} \\ v_{zn} \end{pmatrix}_{exp} = \begin{bmatrix} R_{11} & R_{12} & R_{13} & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ R_{21} & R_{22} & R_{23} & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ R_{31} & R_{32} & R_{33} & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & R_{11} & R_{12} & R_{31} & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & R_{21} & R_{22} & R_{23} & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & R_{31} & R_{32} & R_{33} & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & R_{11} & R_{12} & R_{13} \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & R_{21} & R_{22} & R_{32} \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & R_{31} & R_{32} & R_{33} \end{bmatrix} \cdot \begin{pmatrix} v_{x1} \\ v_{y1} \\ v_{z1} \\ v_{x2} \\ v_{y2} \\ v_{z2} \\ \vdots \\ v_{xn} \\ v_{yn} \\ v_{zn} \end{pmatrix}_{num} \quad (9)$$

Since the transformation matrix $\mathbf{R}_{num}$ is not known, the objective is to determine it. Both numerical and experimental mode shape vectors are available, expressed in their respective reference frames. By rewriting the relation between them, the numerical mode shape vector can be arranged in matrix form, while the entries of $\mathbf{R}_{num}$ are treated as unknowns to be solved. This formulation leads to a linear system where the coefficients come from the numerical mode shapes and the unknown vector contains the nine components of the transformation matrix:

$$\begin{pmatrix} v_{x1} \\ v_{y1} \\ v_{z1} \\ v_{x2} \\ v_{y2} \\ v_{z2} \\ \vdots \\ v_{xn} \\ v_{yn} \\ v_{zn} \end{pmatrix}_{exp} = \begin{bmatrix} v_{x1} & v_{y1} & v_{z1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & v_{x1} & v_{y1} & v_{z1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & v_{x1} & v_{y1} & v_{z1} \\ v_{x2} & v_{y2} & v_{z2} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & v_{x2} & v_{y2} & v_{z2} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & v_{x2} & v_{y2} & v_{z2} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ v_{xn} & v_{yn} & v_{zn,n} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & v_{xn} & v_{yn} & v_{zn} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & v_{xn} & v_{yn} & v_{zn} \end{bmatrix}_{num} \begin{pmatrix} R_{11} \\ R_{12} \\ R_{13} \\ R_{21} \\ R_{22} \\ R_{23} \\ R_{31} \\ R_{32} \\ R_{33} \end{pmatrix} = \mathbf{V} \begin{pmatrix} R_{11} \\ R_{12} \\ R_{13} \\ R_{21} \\ R_{22} \\ R_{23} \\ R_{31} \\ R_{32} \\ R_{33} \end{pmatrix} \quad (10)$$

This system is generally overdetermined, since multiple measurement points contribute independent equations for the same unknown transformation matrix. Therefore, the problem cannot be solved exactly but can be approached using the least squares method.

$$\begin{pmatrix} R_{11} \\ R_{12} \\ R_{13} \\ R_{21} \\ R_{22} \\ R_{23} \\ R_{31} \\ R_{32} \\ R_{33} \end{pmatrix} = \mathbf{V}^+ \begin{pmatrix} v_{x1} \\ v_{y1} \\ v_{z1} \\ v_{x2} \\ v_{y2} \\ v_{z2} \\ \vdots \\ v_{xn} \\ v_{yn} \\ v_{zn} \end{pmatrix}_{exp} \quad (11)$$

Applications

The experiments were performed on a mechanical component designed to exhibit clear three-dimensional vibration behavior, in order to assess the effectiveness of the proposed algorithm. The measurements were carried out using a PolyTec QTEC PSV-500 scanning laser Doppler vibrometer. The selected component is a collector from an air-compressor. A total of 18 measurement points were defined on its surface, as shown in Figure 3a.




Fig. 3 Experimental setup and orientations of the tested component: (a) neutral position with no rotation; (b) rotation of 30° about the x -axis (YZ-plane rotation, clockwise); (c) rotation of 30° about the y -axis (XZ-plane rotation, counterclockwise)

To enable the reconstruction of the full 3D mode shapes, the component was tested in three distinct orientations with respect to the laser beam:

- Neutral position – the main plane of the component was kept orthogonal to the laser beam (Figure 3a).
- Rotation about the x -axis – the component was rotated by 30° (Figure 3b).
- Rotation about the y -axis – the component was rotated by 30° (Figure 3c).

The chosen orientations ensured that the laser projections were sufficiently distinct, thereby providing the independent information necessary for solving the system of equations and extracting the 3D motion. The experimental reference frame used for all measurements is illustrated in Figure 3a.

An example of a measured frequency response function (FRF) for a single measurement point in the neutral configuration is shown in Figure 4. Peaks in the response correspond to the resonance frequencies of the component, which are later used for mode shape extraction.




Fig. 4 An example of a measured FRF in the neutral configuration (no rotation)

It is important to note that the laser vibrometer directly measures velocity, whereas mode shapes are commonly defined in terms of displacement. However, for a harmonic response at a given frequency, velocity and displacement are simply

related by the angular frequency. This implies that the measured velocity is equal to the corresponding displacement scaled by angular frequency. Since scaling a mode shape vector does not alter its spatial distribution, the mode shapes can be evaluated directly from the velocity measurements without additional transformation.

Regarding the identification of modal parameters, the measured FRFs exhibit distinct, well-separated resonance peaks with sharp profiles, indicating low modal damping and the absence of modal coupling. Under these conditions, a simple peak-picking method was adopted for the extraction of the modal frequencies and associated mode shapes. For the present study, the analysis was limited to the three modes. The determination of the rotation matrix in eq. (10-11) was carried out by employing the mode shape measurements obtained from all three identified modes. Using multiple modes in the formulation ensures consistency and improves the robustness of the estimation, as the transformation matrix must remain valid across different mode shapes of the same structure.




Fig. 5 Comparison of experimental and numerical 3D mode shapes for the first three modes: (a) Mode 1, (b) Mode 2, (c) Mode 3

The comparison between experimental and numerical 3D mode shapes is presented in Figure 5 for the first three modes. In all cases, the undeformed geometry is shown in gray, while the experimental and numerical deformed shapes are overlaid. The results indicate a very good agreement in the overall deformation patterns, with the experimental shapes closely following the numerical predictions. Small local deviations are visible at certain measurement points, which can be attributed to




Fig. 6 Modal Assurance Criterion (MAC) matrix comparing experimental and numerical 3D mode shapes for the first three modes

experimental noise and limitations in spatial resolution. Nevertheless, the global characteristics of the vibration modes are consistently reproduced.

The quality of the correlation was further assessed using the Modal Assurance Criterion (MAC), shown in Figure 6. The diagonal values are very close to unity for all three modes (0.93, 0.94, and 0.85, respectively), confirming a strong correspondence between experimental and numerical shapes. The off-diagonal terms remain close to zero, demonstrating a clear separation between the identified modes and the absence of significant cross-correlation. These results validate the effectiveness of the proposed methodology in reconstructing experimental 3D mode shapes and provide confidence in their use for comparison against numerical predictions.

Conclusion

This work has presented a novel methodology for the experimental evaluation of three-dimensional mode shapes using a single scanning laser Doppler vibrometer. Unlike conventional approaches that require multiple laser heads for simultaneous multi-directional measurements, the proposed method relies on repositioning the test component in different orientations.

By combining three independent measurement configurations and applying a dedicated reconstruction algorithm, the full 3D mode shapes can be extracted. An additional algorithm was introduced to transform numerical mode shapes into the experimental reference frame, enabling a consistent and meaningful comparison between numerical predictions and experimental data even when the reference frames are not directly aligned.

The main advantages of the method are its simplicity, cost-effectiveness, and flexibility. The requirement of only one laser scanning head makes the approach accessible in laboratory environments where multi-head systems are impractical due to high cost or setup complexity. Furthermore, the framework is general and does not rely on prior knowledge of the component geometry, making it suitable for both simple and complex structures.

The methodology was validated through experiments performed on a mechanical component exhibiting clear 3D vibration behavior. Three orientations of the component were used to generate independent measurements, from which the 3D mode shapes were reconstructed. Comparison with numerical mode shapes showed very good agreement. The overlaid mode shape plots confirmed the accuracy of the reconstructed deformation patterns, while the MAC analysis yielded high diagonal values (0.93, 0.94, and 0.85) and near-zero off-diagonal values, demonstrating strong correlation and proper mode separation.

References

1. Sriram, P., and Craig, J., "Mode Shape Measurement Using a Scanning Laser Doppler Vibrometer," *International Journal of Analytical and Experimental Modal Analysis*, Vol. 7, No. 3, pp. 169–178, 1992.
2. Stanbridge, A., and Ewins, D., "Modal Testing Using a Scanning Laser Doppler Vibrometer," *Mechanical Systems and Signal Processing*, Vol. 13, No. 2, pp. 255–270, 1999.
3. Allen, M. S., and Sracic, M. W., "A New Method for Processing Impact Excited Continuous-Scan Laser Doppler Vibrometer Measurements," *Mechanical Systems and Signal Processing*, Vol. 24, No. 3, pp. 721–735, 2010.
4. Halkon, B. J., Frizzel, S., and Rothberg, S., "Vibration Measurements Using Continuous Scanning Laser Vibrometry: Velocity Sensitivity Model Experimental Validation," *Measurement Science and Technology*, Vol. 14, No. 6, pp. 773, 2003.
5. Polytec GmbH, "PSV-500-3D Scanning Vibrometer Product Page," Polytec Vibrometry Products – Full-Field Vibrometers, accessed August 2025. <https://www.polytec.com/int/vibrometry/products/full-field-vibrometers/psv-500-3d-scanning-vibrometer>.
6. Chen, D.-M., and Zhu, W. D., "Investigation of Three-Dimensional Vibration Measurement by Three Scanning Laser Doppler Vibrometers in a Continuously and Synchronously Scanning Mode," *Journal of Sound and Vibration*, Vol. 498, pp. 115950, 2021.



Chapter 11

Collapsing Distributed Fiber Sensors to Point Measurements to Boost Signal-to-noise using Custom-designed Fiber Coils

Angel H. Chafino, Serena A. Vu, Reganne M. Watts, Steven R. Anton, Carly M. Donahue, Andrew K. Harter, Nathan T. Maier, and William A. Redford

Abstract After structural failures like those caused from natural events, the ability for search and rescue teams to locate people is essential to saving lives yet one of the most difficult to complete. Distributed Acoustic Sensing (DAS) has emerged as a versatile and innovative technology that uses a pulsed laser to turn a regular fiber optic cable into an array of vibrational sensors that can prove useful for these scenarios. We propose a novel application of DAS to detect ground propagated waves using surface deployed fiber sensors which could prove to be pivotal for search and rescue efforts. However, two main challenges still exist with this technique; coupling the fiber to the medium of interest, and the reduction of unwanted noise which dictates signal-to-noise ratio (SNR). Here we present experimental results testing different fiber sensor designs which can be integrated into fiber and used to improve SNR for surface deployed scenarios. The sensors are based on fiber coils, which both amplify the signal and allow for noise mitigation techniques. We characterize coil geometry effects on SNR values and the effectiveness of various coupling techniques for detecting ground propagated waves in a field environment. A strain gauge coil geometry greatly improved the SNR and implementation of sensor compression allowed for a compact point sensor; this was successfully implemented into a weighted roll-able noise dampening fabric, significantly improving the amplitude of impulse signals.

Keywords Distributed acoustic sensing (DAS) · Fiber optics · Earthquake monitoring · Emergency response · Point sensors

Angel H. Chafino
New Mexico State University, Las Cruces, New Mexico, USA;
e-mail: chafino@nmsu.edu

Serena A. Vu
Brown University, Providence, Rhode Island, USA;
e-mail: serena.vu@brown.edu

Reganne M. Watts
University of Michigan, Ann Arbor, Michigan, USA;
e-mail: regannew@umich.edu

Steven R. Anton
Tennessee Technological University, Cookeville, Tennessee, USA;
e-mail: santon@tntech.edu

Carly M. Donahue · Andrew K. Harter · Nathan T. Maier
Los Alamos National Laboratory, Los Alamos, New Mexico, USA;
e-mail: cmd@lanl.gov; aharter@lanl.gov; nmaier@lanl.gov

William A. Redford
Georgia Tech to Los Alamos National Laboratory, Atlanta, Georgia, USA
e-mail: wredford@lanl.gov

Introduction

Distributed acoustic sensing (DAS) has rapidly gained popularity in various fields including seismology, structural health monitoring, and disaster response due to the ability to collect large amounts of data from one fiber [1, 2]. While many researchers have used existing underground telecommunication optical fibers for DAS, these fibers may not exist in the desired sensing areas and are limited to their existing installation configurations. Due to its versatility and live imaging capabilities, we believe DAS could play a pivotal role in rapid response rescue scenarios as it would allow for coverage of a large area with the potential of being rapidly deployed with only one single optic fiber cable. Surface deployment of optical fiber is an emerging sensing method as it allows for temporary deployment of an array of sensors and can be done in areas where trenching is infeasible, such as in emergency response scenarios or environmentally protected land. Rapid installation of surface deployed optical fiber has its own challenges due to fiber sensitivity to wind noise and poor coupling of the fiber to the ground. This paper focuses on evaluating various methods of coiling, coupling, and wind reduction for surface deployed fiber optic cables.

Background

The DAS system consists of an interrogator, which is an electrical unit connected to one end of an optical fiber. The interrogator emits a laser pulse down the fiber and measures the phase shift of Rayleigh backscattering to determine the strain rate of the fiber. Strain rate is averaged over gauge lengths and binned into channels, which are equally spaced in the fiber. The gauge length for the interrogator used in this work is 10 meters. Small perturbations in the fiber location due to acoustic signals alter the interaction of the pulse with the fiber and thus the signal that is received [1].

DAS has found wide applicability in the field of geology as a tool for seismology, subsurface structure imaging, fluid pressure sensing, and studying various other underground phenomena [2, 3]. It has also been used to detect hypersonic vehicle entry and bird calls which shows its versatility and desirability among a wide range of disciplines [4, 5].

While it has shown its versatility, there are still challenges to fully apply this system for rapid deployment using surface fiber. DAS is directionally sensitive, or only able to detect disturbances that cause longitudinal strain on the fiber. It is also highly impacted by the quality of coupling to the desired sensing medium and the configuration of the fibers. Various researchers have found that coiled fibers increase the signal-to-noise ratio, making it more advantageous over straight fiber. This coiling allows a standard gauge length of fiber to be collapsed into a single point sensor, increasing its amplitude and decreasing directional sensitivity as it allows for a portion of the fiber to always have longitudinal strain [5, 6, 7, 8].

Traditional surface deployment methods of fiber optic cables for DAS involve unspooling the fiber on the ground and pressing it into the surface by walking over it [9]. While quick to execute, this results in poor coupling with the desired sensing medium. Various methods have been proposed to improve surface deployment methods for distributed sensing arrays, including using soft robotics, towing the cable behind a vehicle, and a roll-able blanket of sensors [10, 11]. However, these methods can be bulky, time-consuming, and often unfeasible for rapid response.

Our goal is to design an adaptable and easily deployable array of fiber optic point sensors. While there has been research in coiling of fiber and its applications, there is a lack of research into rapid deployment techniques and the viability of only using surface deployed fibers. The goal of this study is to evaluate various methods of coiling, coupling, and wind reduction for coiled optical fiber sensors to guide future work on a rapid deployment mechanism.

This paper begins with a discussion of our instrumentation and coil design. Then we discuss data collection and field layouts. Afterwards we discuss the results of the various tests and conclude with a discussion of future improvements to the rapid deployment system.

Methods

Fabrication

All the coil geometries were fabricated using ABS filament printed with an infill of 15% using a grid pattern in order to reach a higher first natural frequency based on previous research [12]. This shifted the first resonance frequency outside of excitation range of our shaker and any frequency of interest to ensure that the printed components did not impart any external behaviors into the coil system.

We coiled a single-mode optical fiber from Thorlabs (SMF-28-100), which has an operating wavelength of 1260 to 1625nm, around the geometries presented in Figure 1. This fiber has a 900 μ m diameter jacket and a minimum bend radius of 3.75cm for long-term storage. The radius for all curves used in the coil designs is 4cm to satisfy this requirement and

prevent the fiber from experiencing unnecessary strain. Additionally, the wall thickness of these designs are 1cm thick. We spliced all coiled sensors as a continuous fiber to enable the evaluation of all sensors simultaneously, if desired.




Fig. 1 Top row left to right: plain coil, standard disk, two gauge, three gauge, rounded square, layering disk. Bottom row left to right: infinity, 1 to 2 aspect ratio, 1 to 2.5 aspect ratio, 2x diameter, 1x diameter, low frequency.

Additionally, we used the Finite Element Analysis (FEA) package in SolidWorks to design a coil geometry that combined the aspect ratio and layering disk designs and could be excited by the frequencies supplied by the shaker. We hoped that the vibrations and mode shapes would increase the strain in the fiber and therefore the signal we were detecting. The resulting design and first natural frequency is shown in Figure 2. The accuracy between the FEA model and the environment was not of high concern, and so the analysis was done using a free-free boundary condition when designing for the lowest frequency. This design was printed with 100% infill to best simulate the finite element model. It is understandable that once printed and placed on the ground the fidelity of the model to the actual component will vary. However, we believed that a small exploration of this concept would still prove valuable.




Fig. 2 (A) The fabricated coil design and (B) the finite element analysis showing the 7th mode within those being excited by the shaker.

Taking inspiration from geophone covers and other wind dampening mechanism, we 3D printed a 4 prong basket shape that could easily be placed on top of the coil without disturbing its boundary conditions to the ground. The material is easily hot glued onto the basket which has a minimal area to ensure that the basket itself is not inducing any dampening. As shown in Figure 3, we evaluated four different materials: a quilted fabric, felt, rubber, and an acoustic foam.

We also looked at a high density foam (0.048 g/cm^3) and a low density foam (0.029 g/cm^3) cover with weighted edges and a divot to completely encase the coil, shown in Figure 4.

For coupling methods, we tested a 25lb weight, a 50lb weight, a stake, and covering the coil with a small shovelful of sand. The stake was designed to be modular with the layering coil disk. This was done so we could use the same sensor for each test and only modify the coupling mechanism. These methods are shown in Figure 5.

Our rapid deployment design was inspired by seismic land streamers which are arrays of geophones towed behind a vehicle. We also took elements from our coil geometry studies and coupling/wind tests to build our land streamer shown in




Fig. 3 (A) The rubber basket, (B) the quilt basket, (C), the acoustic foam basket, (D) the felt basket, and (E) an underside view showing the 3D printed base for each of the baskets.




Fig. 4 (A) Top view of the two foams with low density on the left and high density on the right. (B) Bottom view of the two foams with a disk inside low density on the left and high density on the right.

Figure 6. The body of the land streamer was made from an acoustic deadening fabric cut into 7in wide strips and then pieced together into a 21ft long strip. We glued six layering disk coils to the bottom and taped four 3oz weights around each coil to help couple it to the ground. To deploy the streamer, one end was rolled up so a bundle was created that could quickly be rolled out on the field.

Indoor Testing Environment

To minimize external noise sources, we conducted individual coil SNR tests in a 10m by 0.53m wooden box filled with basalt sand shown in Figure 7. Previous work had found that the box has a resonance at 30Hz which greatly impacted the response of the fibers within. In order to ensure that comparisons can be made between different coil geometries, a particular interest is set on frequencies lower than 30Hz and all coil geometries are placed in the same location to mitigate areas of nodes and anti-nodes within the box.

We used a shaker located at one end of the box to generate a step sine wave in the frequency range of 15 to 40 Hz. The shaker was controlled as an open loop using TiePie to supply 10-second steps in 1Hz increments to achieve steady state. Smaller step increments were not considered due to time constraints.




Fig. 5 (A) The modular stake design, (B) adding medium to simulate buried conditions, (C) the 25lb shot bag, and (D) the 50lb shot bag.




Fig. 6 One of the coils attached to the bottom of the land streamer



Fig. 7 The 10m long wooden box used as an indoor testing environment

Outdoor Testing Environment

For outdoor testing, we used a large field located on laboratory campus, shown in Figure 8. We dug a trench 6in deep and buried 100m of tactical optic fiber to serve as a reference for our rapid deployment measurements. This is largely informed from previous research showing that the best coupled fibers are those that are buried underground, as it allows for a more direct measurement of the ground strain and reduces the wind noise. We wanted to compare the best known coupling and wind reduction method to our rapid deployment designs. Existing fibers that are underground and used for DAS are generally in conduit reducing its coupling with the ground. We then deployed a single coil on the fiber to perform our coupling methods on to make the assumption that manufacturing variations had no impact on the response of the coil.

We used a PCB 086D20 modal sledgehammer to generate our impulses. The medium tip was selected to excite frequencies between 0 to 300Hz. To avoid damaging the modal hammer and creating a hole from repeated impacts, an aluminum block was placed to take the impact at the 2m, 5m, and 10m distances away from the surface deployed coil. Three impacts were captured at each location, using a proprietary Matlab script and a National Instruments USB-4431 to make sure analyzed impacts were between 9,000 and 11,000 Newtons. For the coupling methods, we only tested impacts that were perpendicular to the fiber while the land streamer had impacts both perpendicular to and parallel to the fiber. The interrogator recorded the entire testing session of 9 impacts, and the data from the impulse collection was used to obtain a 4 second range of where the impulse from the hammer occurred.

Results and Discussion

Indoor Testing

Coil SNR

To calculate the SNR for the various coil geometries, we took a 5 second sample of the signal at each excited frequency and a 5 second snippet of noise at the same frequency. We then used Equation 1:




Fig. 8 A ground view of the outdoor testing environment. Triangles represent the location of perpendicular to the fiber impulse locations (2, 5 and 10 meters) while the circle represents the sensor used to test the coupling methods and wind mitigation.

$$SNR = 10 \log_{10} \frac{A_S}{A_N} \quad (1)$$

where A_s is the mean amplitude of the signal at the given frequency and A_N is the mean amplitude of the noise at the same frequency.

We repeated this process for each frequency that was excited. We then calculated the mean SNR below 30 Hz, which was the first resonant frequency of the box, and the mean SNR above 30 Hz, which included the box resonant frequency. As shown in Figure 9, we compared the SNRs for each of our coil geometries. All of the geometries performed better than a plain coil looped without any base geometry both above and below the box resonance. The coils generally followed the same trends even above the box resonance, however particular emphasis will be taken for those below this region where the box is well characterized.

When comparing the usage of having multiple gauges, the two and three gauge coils provide a marginal SNR increase when compared to the standard and plain coils. We did not find the increase in SNR to be significant enough to justify taking the time and manufacturing effort to spool the larger amount of fiber onto one coil. Interestingly, it was also found that layering the fibers on top of each other did not negatively impact the SNR of the coils which allows us to design more compact sensors which are easier to rapidly deploy without greatly affecting the response, when compared to the standard single layer coils.

For the aspect ratio geometries inspired by strain gauges, we observed an increase in SNR both above and below the box resonance when compared to the plain coil. However, there was not a substantial difference between the 1:2 and 1:2.5 designs likely due to only an inch being added to the longitudinal sides. These findings held up with previous research that used fiber as a strain gauge and supports the principle of the fiber having higher response to longitudinal strain. The low frequency design did not perform better than the regular aspect ratio however, it did produce a smaller standard deviation in the mean SNR both above and below the first box resonance frequency.

When comparing the standard and plain coils to the 1.5x diameter and 2x diameter (where 8cm is 1x diameter), we see significant variations above and below the box resonance. The 1.5x diameter has a smaller standard deviation both above and below the box resonance. When only looking at the SNR below the box resonance, the 1.5x diameter has a greater mean




Fig. 9 A Plot of all calculated SNRs for the various coil designs. The boxes are separated into two data sets corresponding to responses above or below the box resonance of 30Hz. The bars represent one standard deviation of the mean amplitude.

SNR than the 2x, while the converse is true for SNR above the box resonance. We expected to observe an increase in the SNR as the diameter got larger. However, only a small sample of diameters were tested, so trends, plateau, or correlations are indiscernible and require more testing to characterize.

We create novel coil geometries to explore the idea of having coils with longitudinal sections at different angles in an attempt to increase the signal by having more fiber in regions that will be responsive to a perpendicular wave. The infinity has the highest mean SNR both above and below the box resonance frequency. The rounded square has a slightly larger SNR than the double infinity. This suggests that utilizing a variety of angles in the sensor can positive impact the sensor's ability to detect impacts.

Outdoor Testing

Coupling Methods

From the output of the modal sledgehammer, we found that the magnitude of the force input and the amplitude of the signal recorded by the coil have a highly nonlinear relationship. We averaged all three impulse time series recorded by the coil at each impact distance (2, 5, and 10m) from the coil. The results of this averaging are shown in Figure 10. Due to fluctuations in the noise floor and an uncontrolled environment, we couldn't calculate the signal to noise ratio for the different coupling methods.

Both the 25lb and 50lb weights substantially increased the SNR when compared to the plain coil, especially at the 10m impact distance where there are still prominent peaks in the data. Both the stake and scoop of sand did not increase the amplitude and the impacts were indiscernible from the noise at the further distances. This suggests that weighing down the coils helps them couple to the ground and allows them to detect signals from farther away.

Wind Reduction Methods

To control the wind noise source, we placed an industrial fan approximately 1.5m away from the uncovered coil. We then compared the mean amplitude of 5s of noise taken after the last hit was produced. In Figure 11, the values were then normalized by the uncovered coil so values <1 indicate a reduction in noise while the values >1 indicate an increase in noise floor.

The basket materials that show promise for wind reduction were the acoustic foam and the felt as they all had smaller mean noise values and error bars than the uncovered coil. The quilt material performed similarly to the plain coil while the rubber had a slightly higher mean noise than the plain coil.




Fig. 10 A series of plots showing the average response of the coil and the standard deviation for 0.5 seconds after the impulse at 2m, 5m, and 10m.




Fig. 11 A comparison of the normalized noise values for the various basket coverings. The error bars representing one standard deviation of the mean.

When comparing the two weighted foam squares, low density foam had a higher mean noise and a higher standard deviation than the high density foam as seen in Figure 12. One possible reason for this increase is that this foam block was approximately an inch taller than the high density foam block which would make it less aerodynamic, possibly introducing more wind noise into the system.

Land Streamer Deployment

The land streamer was able to clearly detect impacts at all three distances, as shown in Figure 13. While there is more noise than the ground fiber, the impacts are still prominent among the spectra. Due to the prominence of the peaks at 10m, we believe the streamer could clearly detect impacts even further away.




Fig. 12 A comparison of the normalized noise values for low and high density foam blocks as wind reduction methods. The error bars represent one standard deviation from the mean.




Fig. 13 A comparison of the waterfall plots for the land streamer and buried fiber. The first group of impacts are at 2m, the second group are at 5m, and the last group are at 10m.

Figure 14 shows the mean amplitude for the impacts tangential (along the fiber) and normal (perpendicular to the fiber) for the buried fiber and the land streamer. The land streamer has a similar amplitude for impacts in both directions, showing that we were able to reduce the directionality of the fiber and make a more omnidirectional sensor. We also noted that, despite rolling and unrolling the land streamer several times, there were no fractures in the fiber showing promise for future rapid deployment designs using this method.




Fig. 14 A comparison between the average amplitude of the land streamer and the buried fiber at the given distances.

Conclusion

Design Recommendations

From the various experimental results, we have several design recommendations regarding future fiber optic point sensor coiling geometries. One of the most appealing is that layering fibers does not have a significant impact on SNR. This implies that smaller compact coils can be created and implemented into rapid deployment methods. We believe this is applicable to any coiling geometry and thus could be combined with the most responsive geometry, if layering can be integrated into the design. Although fiber is inherently directional, this can be used advantageously to create more sensitive coils by forming long longitudinal sections. This is particularly supported by the aspect ratio and infinity geometries, as these contained long sections of fiber that could still be compressed to a point sensor and sense ground propagated waves. This can be used both for directionality if the source direction is known or omnidirectional sensing by ensuring symmetry of the longitudinal sections like that of the infinity and double infinity. Additionally, it is clear that weighing down the coils leads to an observable improvement in the amplitude of the signal as both the 25lb weight and the 50lb weight improved the signal at all distances.

Rapid Deployment Feasibility

We were able to easily manufacture the land streamer and meet all of the design requirements. Our design showed significant reduction in the noise floor of the fiber, especially when compared to a plain coil, while also improving the response at larger distances. We were able to easily roll and unroll the land streamer without damaging the fiber, showing promise for future designs that involve unrolling long strings of fiber point sensors.

Future Work and Experimental Recommendations

Future work should be done to better optimize the coiling geometry, looking into more complex simulations and optimization algorithms based on the findings presented here. Additionally, more thorough laboratory studies should be done on the wind reduction methods in order to confidently quantify the response of various materials.

Acknowledgments This research was funded by Los Alamos National Laboratory (LANL) through the Engineering Institute's Los Alamos Dynamics Summer School. The Engineering Institute is a research and education collaboration between LANL and the University of California San Diego's Jacobs School of Engineering. This collaboration seeks to promote multidisciplinary engineering research that develops and integrates

advanced predictive modeling, novel sensing systems, and new developments in information technology to address LANL mission relevant problems. The authors also thank Jason Martinez for providing 3D printers and guidance throughout the project.

References

1. Zhan, Z. “Distributed acoustic sensing turns fiber-optic cables into sensitive seismic antennas”. *Seismological Research Letters*, 91(1):1–15 (2019).
2. Lindsey, N.J. and Martin, E.R. “Fiber-optic seismology”. *Annual Review of Earth and Planetary Sciences*, 49(1):309–336 (2021).
3. Becker, M., Coleman, T., Ciervo, C., Cole, M., and Mondanos, M. “Fluid pressure sensing with fiber-optic distributed acoustic sensing”. *The Leading Edge*, 36(12):1018–1023 (2017).
4. Carr, C.G., Donahue, C.M., Viens, L., Beardslee, L.B., McGhee, E.A., and Danielson, L.R. “Detection of a space capsule entering Earth’s atmosphere with distributed acoustic sensing (DAS)”. *Seismological Research Letters* (2025).
5. Jensen, A.L., Redford, W.A., Shergill, N.P., Beardslee, L.B., and Donahue, C.M. “Identification of bird species in large multi-channel data streams using distributed acoustic sensing”. In Matarazzo, T., Hemez, F., Tronci, E.M., and Downey, A., editors, *Data Science in Engineering Vol. 10*, pages 97–107. Springer Nature Switzerland (2024).
6. Stepanov, K.V., Zhirnov, A.A., Koshelev, K.I., Chernutsky, A.O., Khan, R.I., and Pnev, A.B. “Sensitivity improvement of phi-OTDR by fiber cable coils”. *Sensors*, 21(21) (2021).
7. Harmon, N., Rychert, C.A., Davis, J., Brambilla, G., Buffet, W., Chichester, B., Dai, Y., Bogiatzis, P., Snook, J., van Putten, L., and Masoudi, A. “Surface deployment of DAS systems: Coupling strategies and comparisons to geophone data”. *Near Surface Geophysics*, 20(5):465–477 (2022).
8. Innanen, K., Lawton, D., Hall, K.W., Bertram, K.L., Bertram, M., and Bland, H. “Design and deployment of a prototype multicomponent DAS sensor”. *CREWES* (2018).
9. Mjehovich, J., Jin, G., Martin, E.R., and Shragge, J. “Rapid surface deployment of a DAS system for earthquake hazard assessment”. *Journal of Geotechnical and Geoenvironmental Engineering*, 149(5):04023027 (2023).
10. Gruebele, A.M., Zerbe, A.C., Coad, M.M., Okamura, A.M., and Cutkosky, M.R. “Distributed sensor networks deployed using soft growing robots”. In *2021 IEEE 4th International Conference on Soft Robotics (RoboSoft)*, pages 66–73. IEEE (2021).
11. Pandey, A., Shragge, J., Chambers, D., and Girard, A.J. “Developing a distributed acoustic sensing seismic land streamer: Concept and validation”. *Geophysics*, 88(6):WC59–WC67 (2023).
12. Parpala, R.C., Popescu, D., and Pupaza, C. “Infill parameters influence over the natural frequencies of ABS specimens obtained by extrusion-based 3D printing”. *Rapid Prototyping Journal*, 27(6):1273–1285 (2021).



Chapter 12

Reduced Order Structural Modeling for Control System Simulation

Max Dommers, Matthew S. Pearson, and Adam C. Johnson

Abstract Reduced order structural modeling is a highly effective tool for efficient system level simulation of mechanical systems. It has become heavily utilized in the aerospace industry for simulations that involve integration of a large number of subcomponent models or iteration across a large number of load cases. One common application is to incorporate flexible effects in an inertial body active control simulation. Industry examples of this include optical payload pointing simulation, space and aircraft GNC simulation, and fine robotic arm positioning. In these applications, flexible body dynamics are often considered secondary effects, but in many cases can have a substantial impact on performance. This study identifies a technical workflow for preparing a system level reduced-order structural model for efficient use in a closed-loop control simulation. The proposed procedure addresses common workflow challenges such as integrating multiple component models and converting the reduced model into a format usable in common coding languages. Further, the approach leverages previously established methods that retain model accuracy while reducing the degrees of freedom (DOFs) for efficient simulation. As an example problem, this study demonstrates the development of a flexible spacecraft structural plant model for optical pointing simulation. A broadly applicable approach is outlined for incorporating flexible body structural dynamics into control systems simulations.

Keywords State Space, Substructuring, Hurty/Craig-Bampton, Closed-loop control, Disturbance Rejection

Introduction

The baseline design of a controller for an inertial system focuses on guiding rigid body dynamics. High precision systems often require that the flexible body effects be considered in the finalized controller design. In the case of optical pointing or robotic arm control, seemingly miniscule vibrations can render the hardware ineffective at performing its designated task. Table 1 shows common control simulations that require assessment of flexible effects.

Table 1 Dynamic Control Simulations that May Require Flexible Body Representation

Engineering Problem	Rigid Representation	Impact of Flexible Body
Aerodynamic Steering	Aero forces acting on system inertia	IMU Vibration Feedback
Optical Pointing	Disturbance Forces acting on system inertia	Unwanted disturbance of chief ray (aka Jitter)
Robotic Arm Control	Joint motor forces acting on system inertia	Tool center point instability / vibration
Rocket / LV Steering	Thrust and aero forces acting on system inertia	IMU Feedback & Slosh instability

Though modern Finite Element Models (FEM) can offer accurate results, they are too computationally expensive to directly incorporate into a transient control simulation that may need to provide results quickly for tuning of controller

Max Dommers · Matthew S. Pearson · Adam C. Johnson
Quartus Engineering Incorporated, San Diego, CA 92121
e-mail: mdommers@quartus.com; matt.pearson@quartus.com; adam.johnson@quartus.com

parameters through iteration or checking a large number of load cases. This problem can become even more prohibitive for system level simulations that involve combining models from multiple vendors amounting to large total DOF counts. Often times, simplifying dynamics models are incorporated to approximate flexible effects. Simplified flexible representations are sometimes used to alleviate this issue (e.g. mass spring models, pendulum models, simple flexible beam models, etc.). This report summarizes an approach for creating a reduced-order structural model that only retains a small number of DOFs without sacrificing model accuracy. The resulting model perfectly matches the response of an explicit coupled FEA system model. While the methods in this report are well established and based on common practice, the purpose of this work is to provide a single transparent guide for applying these methods specifically to the problem of controls modeling of dynamic systems.

Background

An example problem was developed based on the well-known ShockSat test article developed by Dr. Sam Yunis [13] for the purposes of sharing analysis and test results within the structural dynamics community. This article was chosen for this study because of its open source nature allowing transparent sharing of results. Further, the structure consists of two primary structural elements that resemble an optical payload and spacecraft bus structure, lending itself to demonstration of the line of sight (LoS) control problem. Figure 1 shows the FEA model that was previously developed of the ShockSat test article. In large spacecraft system level analysis, it is common for payload and bus vendors to be separate organizations, thus deliver models separately on independent timelines. To reflect this common logistical challenge, the Optical model and Spacecraft model were separated into two individual component models and reduced to a superelement in Hurty/Craig-Bampton (HCB) format before performing the study. The HCB format was selected for this study due to its common use in model sharing. However the described procedures will work on any mechanical system generally defined by a mass and stiffness matrix. A previous publication by Wender and Griebel [3] describes a general procedure for reducing a FEM in various software formats into this common HCB format from various independent FEA software suites.




Fig. 1 FEA Model of ShockSat Test Article

Once the component models are assembled into a system level model and checked, the model can be used for rapid and efficient simulation. In this example, Figure 3 shows the key components of the coupled model for performing pointing error control simulation. The model allows motion of the dish structure to be recovered as an output due to 6 DOF disturbance forces applied to the frame of the structure. This input-output framework is analogous to a common jitter problem where operational noise from reaction wheels causes excessive pointing error in an imaging system. Figures 2 and 3 illustrate the assembly of the components that are used as inputs and outputs for this problem.

Finally, the Control Simulation section of this report shows how a reduced structural model like this can be installed into a control loop for the purposes of identifying issues in controller design and parameters during realistic operational load cases. In this report, only a single time domain control simulation is performed by incorporating a fast steering mirror (FSM) as a jitter mitigation approach.




Fig. 2 Illustration of Coupling SC and Optic HCB Models at Interface DOFs




Fig. 3 Reduced ShockSat Model for Pointing Error Control Simulation

Component Mode Synthesis

In order to capture system-level dynamics properly, all component substructures must be connected at their respective interfaces. As a prerequisite for this work, it is assumed that the user has received all component models in Hurty/Craig-Bampton (HCB) format. This is a very common format for passing models efficiently among structural dynamics analysts particularly in the aerospace industry. The advantage to the HCB format is that it leverages the efficiency of a modal model while retaining physical degrees of freedom at the interface with other components. This allows for use of very small mass and stiffness matrices without substantial sacrifice in solution accuracy when compared to the explicitly modeled coupled system. A previous publication by Wender and Griebel [3] describes a general procedure for reducing a FEM into the common HCB format from various independent FEA software suites.

The initial component HCB models are treated as ordinary mass matrices (M_{HCB}), stiffness matrices (K_{HCB}), and damping matrices (Z_{HCB}) used in second order differential equations of motion as shown in (Equation 1).

$$[M_{HCB}] \{\ddot{P}\} + [Z_{HCB}] \{\dot{P}\} + [K_{HCB}] \{P\} = \{F\} \quad (1)$$

The mass, damping, and stiffness matrices have been transformed to take the form shown in (Equations 2-4), where ω is a diagonal matrix of constrained modal natural frequencies and various K and M terms define physical stiffness and mass at the boundaries and interfaces (B-Set) as well as cross terms between the boundaries and modal DOFs (Q-Set).

A few special considerations are noted. First, it is important to note that the B-Set of each component model should contain both the interface DOFs between the components as well as DOFs that are going to be used as inputs to the simulation. Second, this problem uses a simple modal damping formulation in which damping terms are only assigned to modal DOFs as shown in (Equation 4). This is common practice for simple structures, but other methods may be more appropriate for systems with non-homogenous or frequency dependent damping. In these cases, more tailored forms of component mode synthesis such as that of Hruda-Benfield [12] may be required.

$$[M_{HCB}] = \begin{bmatrix} M_{bb} & M_{bq} \\ M_{qb} & I \end{bmatrix} \quad (2)$$

$$[K_{HCB}] = \begin{bmatrix} K_{bb} & 0 \\ 0 & \omega^2 \end{bmatrix} \quad (3)$$

$$[Z_{HCB}] = \begin{bmatrix} 0 & 0 \\ 0 & 2\xi\omega \end{bmatrix} \quad (4)$$

Two HCB models are provided in this example, and they must be coupled together to form a single connected structure. The two structures will be called structure *a* and structure *b* respectively. Various approaches can be used to impose coupling constraints on components systems [1]. For this example, the rigid dual assembly method is used [2]. The first step is adding the HCB stiffness matrices of structure *a* and structure *b* to a single uncoupled stiffness matrix, $[K_{uncoupled}]$, as shown in (Equation 5).

$$[K_{uncoupled}] = \begin{bmatrix} K_{HCB,a} & 0 \\ 0 & K_{HCB,b} \end{bmatrix} \quad (5)$$

The coupling terms are added to the uncoupled stiffness matrix in the form of constraint equations for each coupled pair of DOFs. To create the constraint equations, corresponding DOFs to be coupled are identified. One from structure *a* and one from structure *b*. These two degrees of freedom are set equal as in (Equation 6).

$$p_a = p_b \quad (6)$$

The equation is rearranged such that it can be written into the stiffness matrix as show in (Equation 7). The augmented stiffness matrix shown in (Equation 8) now includes the constraint equation from (Equation 7) (now notated as “CE”) and effectively enforces equivalence between the selected DOFs. This process is repeated for every DOF that needs to be linked between the two HCBs.

$$p_a - p_b = 0 \quad (7)$$

$$[K_{coupled}] = \begin{bmatrix} K_{HCB,a} & 0 & CE \\ 0 & K_{HCB,b} & -CE \\ CE & -CE & 0 \end{bmatrix} \quad (8)$$

The mass matrix is padded with zeros as shown in (Equation 9) to maintain the same size as the coupled stiffness matrix. This causes the mass matrix to be singular.

$$[M_{coupled}] = \begin{bmatrix} M_{HCB,a} & 0 & 0 \\ 0 & M_{HCB,b} & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (9)$$

The coupled damping matrix is created similarly to the mass matrix as shown in (Equation 10). This study assumes a 0.5% modal damping is applied to all modes.

$$[Z_{coupled}] = \begin{bmatrix} Z_{HCB,a} & 0 & 0 \\ 0 & Z_{HCB,b} & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (10)$$

Data recovery matrices (DRMs) are output transformation matrices describing the required transformation between HCB output and desired physical quantities. Any DRMs of the two HCB models are appended as shown in (Equation 11). Extra

zeros are added to the end of the DRM to ensure it has the same number of rows as the $[K_{coupled}]$, $[Z_{coupled}]$, and $[M_{coupled}]$ matrices.

$$[DRM_{coupled}] = [\begin{matrix} DRM_a & DRM_b & 0 \end{matrix}] \quad (11)$$

Once a model is assembled, the component mode synthesis is verified by checking free-free modes of the coupled system. In general, there should be six rigid body modes of the system near zero Hz. If there are more, some components may not have been correctly connected. If there are less, the model may be grounded. In either case, the individual models and coupling equations should be checked for errors. Further, it is recommended that the DRM matrices contain sufficient output to visualize and verify rigid body motions and basic flexible modes of the coupled structure. Table 2 shows a free-free modes check for the ShockSat example problem. Though it is often not possible in practice, for this example the modes of the dual assembly component mode synthesis (right column) was compared against the explicitly assembled FEA model in Nastran (left column). The flexible modes agreed between the two methods within 0.001% and the rigid body modes were more than five orders of magnitude lower than the flexible modes in both cases.

Table 2 Free-free Modes Check

	Mode Frequency [Hz]		% Error
	FEA Coupled	Superelement After Dual Assembly	
Rigid Body Rx	3.221E-04	9.357E-05	N/A
Rigid body Tx	1.464E-04	2.180E-04	N/A
Rigid Body Rz	2.115E-04	3.462E-04	N/A
Rigid Body Ty	4.799E-04	3.617E-04	N/A
Rigid Body Ry	5.273E-04	3.917E-04	N/A
Rigid Body Tz	5.600E-04	4.528E-04	N/A
Axial 1	1.070E + 02	1.070E + 02	0.000%
X Bending 1	1.128E + 02	1.128E + 02	0.001%
Axial 2	1.149E+02	1.149E+02	0.000%
Y Bending 1	1.393E+02	1.393E+02	0.001%

State Space Conversion

For closed loop control simulations, a first order state space is typically desired for easy integration with other elements of the mechanical system. To achieve this form, the assembled second order system model is converted into the descriptor state space format shown in (Equation 12-13). It should be noted that the descriptor state space E matrix is only required due to the presence of a singular mass matrix (from the previous step). To avoid using the E matrix, the second order system can first be diagonalized and transformed into modal coordinates with mass normalized mode shapes. This will result in an identity E matrix allowing it to be neglected. Alternatively, different assembly methods may be used that avoid creating a singular mass matrix [1].

$$[E] \{\dot{x}\} = [A] \{x\} + [B] \{u\} \quad (12)$$

$$y = [C] \{x\} + [D] \{u\} \quad (13)$$

The second order equations of motion include the displacement column vector $\{P\}$ and it's second derivative $\{\ddot{P}\}$. The second order system is rewritten using a new state vector $\{x\}$. The top half of the state vector contains the displacement column vector $\{P\}$ and the lower half of the state vector contains the velocity column vector $\{\dot{P}\}$ as shown in (Equation 14).

$$\{x\} = \begin{Bmatrix} P \\ \dot{P} \end{Bmatrix} \quad (14)$$

The equations of motion from (Equation 1), with a damping matrix $[Z]$, is rewritten and reorganized using the state variable $\{x\}$ and its first derivative $\{\dot{x}\}$.

$$[M] \{\dot{x}\} = -[K] \{x\} - [Z] \{\dot{x}\} + \{F\} \quad (15)$$

Once expanded, the previous equation takes this form shown in (Equation 16).

$$\begin{bmatrix} 0 & 0 \\ 0 & M \end{bmatrix} \begin{Bmatrix} \dot{P} \\ \ddot{P} \end{Bmatrix} = \begin{bmatrix} 0 & 0 \\ -K & -Z \end{bmatrix} \begin{Bmatrix} P \\ \dot{P} \end{Bmatrix} + \begin{Bmatrix} 0 \\ F \end{Bmatrix} \quad (16)$$

Note that the bottom half of (Equation 16) is a restatement of the equations of motion from (Equation 1). And the top half is empty. New equations that relate the derivative of the state vector to the state vector can fill the empty top half as shown in (Equation 17).

$$\begin{bmatrix} I & 0 \\ 0 & M \end{bmatrix} \begin{Bmatrix} \dot{P} \\ \ddot{P} \end{Bmatrix} = \begin{bmatrix} 0 & I \\ -K & -Z \end{bmatrix} \begin{Bmatrix} P \\ \dot{P} \end{Bmatrix} + \begin{Bmatrix} 0 \\ F \end{Bmatrix} \quad (17)$$

This is the first order descriptor state space form of the system where:

$$[E] = \begin{bmatrix} I & 0 \\ 0 & M \end{bmatrix} \quad (18)$$

$$[A] = \begin{bmatrix} 0 & I \\ -K & -Z \end{bmatrix} \quad (19)$$

$$[B] = \begin{Bmatrix} 0 \\ F \end{Bmatrix} \quad (20)$$

$$[C] = [\text{DRM}_a \quad \text{DRM}_b \quad 0] \quad (21)$$

$$[D] = 0 \quad (22)$$

And

$$[E] \{\dot{x}\} = [A] \{x\} + [B] \{u\} \quad (23)$$

$$y = [C] \{x\} + [D] \{u\} \quad (24)$$

Transfer Function Calculation

Now that the model is in state space format, various analyses can be performed. Frequency domain transfer functions are often desired to understand potential input frequencies that will cause excessive motion or feedback into a system. Transfer functions are the ratio of input to output in the frequency domain. In state space form, the state variable $\{x(t)\}$ and the input $\{u(t)\}$ are both functions of time. These equations can be converted into the complex frequency domain using the Laplace transformation.

$$[E] s \{X(s)\} = [A] \{X(s)\} + [B] \{U(s)\} \quad (25)$$

$$Y(s) = [C] \{X(s)\} + [D] \{U(s)\} \quad (26)$$

Rearranging and substituting the equations above will yield the equation to calculate transfer functions. If the model was transformed into modal coordinates before conversion into state space, then the descriptor matrix $[E]$ will be replaced by the identity matrix.

$$G(s) = C(Es - A)^{-1}B + D \quad (27)$$

Figure 4 shows an example of transfer function inspection for X axis reaction wheel disturbance to X axis optic displacement. This plot shows potentially problematic resonances in the structural system. Care must also be taken to recognize any modal truncation that has been applied to the model. In this case, modal response starts to roll off after 300 Hz. Though it is often not possible in practice, for this example the response of the reduced state space model (blue line) was compared against explicitly FEA simulation in Nastran. These two solutions overlay with no observable disagreement emphasizing the accuracy of the reduced model.




Fig. 4 Example Transfer Function Inspection Optic Tx / Reaction Wheel Fx

Control Simulation

A time domain simulation was created using Simulink to illustrate how the resulting state space model can be integrated into a control simulation. The ShockSat dish characterized previously is treated as the primary element of an optical system, in this case, a refractive lens. A simple optical payload consisting of a fast-steering mirror (FSM) and a generic sensing element are paired with the optic to represent a remote-sensing style optical system (see Figure 5). Typically, these types of systems utilize the FSM to steer the beam and reject optical jitter disturbances that may negatively impact the system's line of sight (LoS). The following section will demonstrate the rejection of disturbances from the onboard reaction wheel systems that impact system LoS using closed loop control.




Fig. 5 Schematic of Simple Optical Payload for LoS Simulation

It is well known that for a thin lens, radial (or decenter) motions of the lens cause an angular steering of the output beam with respect to its nominal direction. To a first order linear approximation, this is the only impact to LoS [7]. Angular motions of the optic are insignificant to LoS and axial motion only causes a change in axial beam focus. As such, only the components of the state space model associated with the radial degrees of freedom, the disk x and y directions, need to be included in the simulation of this element. A small radial motion of the optic, δx , produces a beam tilt, $\delta\theta$, proportional to the inverse of the lens focal length [7].

$$\delta\theta = \frac{\delta x}{f} \quad (28)$$

For small motions of the flat FSM mirror, the change in the reflected light angle θ is related to the change in the mirror angle, α , by the following relationship [7]:

$$\theta = 2\alpha \quad (29)$$

At the sensing element, the deviation of the focused spot moves by the relationship:

$$r \approx f\theta = f(\delta\theta + 2\alpha) \quad (30)$$




Fig. 6 LoS Pointing Simulink Model

The previous equations give the geometric relationship between the perturbed optics and the beam motion on the sensor. Now consider the components that interact dynamically with this geometric relationship. First, the FSM mechanism is considered. For the FSM model, idealized dynamics are assumed for simplicity, a simple 2nd order harmonic oscillator representation is used for each axis. While this is a simplification, other works have shown that this is a good approximation [10]. For both axes, the transfer function between the input voltage and the output FSM angle is:

$$T_x(s) = T_y(s) = \frac{\alpha(s)}{V(s)} = \frac{\omega_{FSM}^2}{s^2 + 2\xi_{FSM}\omega_{FSM}s + \omega_{FSM}^2} \quad (31)$$

This can be assembled into the 2x2 transfer function format:

$$\alpha_{FSM}(s) = \begin{bmatrix} \frac{\omega_{FSM}^2}{s^2 + 2\xi_{FSM}\omega_{FSM}s + \omega_{FSM}^2} & 0 \\ 0 & \frac{\omega_{FSM}^2}{s^2 + 2\xi_{FSM}\omega_{FSM}s + \omega_{FSM}^2} \end{bmatrix} V(s) \quad (32)$$

This can be represented in Simulink using either a transfer function or State Space form. The system was converted to State Space form from the transfer function using standard realization methods [6].

A simple proportional controller was used to stabilize the FSM. More sophisticated servo control schemes are possible but are unnecessarily complicated for this stable LoS pointing example and are left to future work.

$$C(s) = \begin{bmatrix} K_{px} & 0 \\ 0 & K_{py} \end{bmatrix} V(s) \quad (33)$$

To provide the disturbance input to the system, a single reaction wheel was modeled at one of the input positions of dynamic structural model. A common tonal disturbance model with harmonics proposed by Gutierrez [8][9] was used to capture realistic reaction wheel dynamics. This model computes the forces $\{m_{ij} | j = 1, 2, 3\}$ and moments $\{m_{ij} | j = 4, 5, 6\}$ of the i th wheel with the following equation:

$$m_{ij}(t) = \sum_{k=1}^{n_j} C_{ijk} \Omega_i^2 \sin(2\pi h_{jk} \Omega_i t + \phi_{ijk}) \quad (34)$$

where C_{ijk} is the amplitude of the k th harmonic, Ω_i is the frequency in Hz of the i th wheel, h_{jk} is the k th harmonic of the j th component of the i th wheel, and ϕ_{ijk} is the k th phase of the j th component of the i th wheel. To make sense mechanically,

it was assumed that the x and y components of the forces and moments have a phase difference of 90° . This is because the imbalance in the wheel is expected to be acting at different times in each direction as the wheel completes a full rotation.

Finally, the disturbance inputs are fed into the structural dynamic model and the outputs in the directions of the ShockSat dish are extracted:

$$\{\dot{x}\} = [A] \{x\} + [B_1 \quad B_{2,3,4}] \begin{Bmatrix} u_1 \\ u_{2,3,4} \end{Bmatrix} \quad (35)$$

$$y = \begin{bmatrix} C_{other} \\ C_{dish \ x.y} \\ C_{dish \ z,rx,ry,rz} \end{bmatrix} \{x\} + \begin{bmatrix} D_{1 \ other} & D_{2,3,4 \ other} \\ D_{1 \ dish \ x.y} & D_{2,3,4 \ dish \ x.y} \\ D_{1 \ dish \ z,rx,ry,rz} & D_{2,3,4 \ dish \ z,rx,ry,rz} \end{bmatrix} \begin{Bmatrix} u_1 \\ u_{2,3,4} \end{Bmatrix} \quad (36)$$




Fig. 7 Time and Spatial Domain Results for $\Omega_1 = 20\text{Hz}$

Selecting the necessary components for reaction wheel 1, and the dish x,y outputs, the following 6 input, 2 output MIMO state space model remains:

$$\{\dot{x}\} = [A] \{x\} + [B_1] \{u_1\} \quad (37)$$

$$y = [C_{dish \ x.y}] \{x\} + [D_{1 \ dish \ x.y}] \{u_1\} \quad (38)$$

Note that the full A matrix is still used though further inspection of the states could allow for additional model reduction of A . Notably, the 1st mode of the full state space model at 107Hz is no longer present because it appears in the z axis direction only. The resulting output of this state space model, δ_{xy} , is divided by the focal length and added into the LoS equation as an output disturbance.

With all these sub-components, a complete LoS pointing model with structurally amplified output disturbances has been implemented. The disturbances are model output disturbances (rather than input disturbances) and as such, the controller can be tuned to attenuate the disturbance but only within the bandwidth of the controller.

The following parameters were used in the simulation:

Table 3 Simulation Parameters

Parameter Name	Value [Units]	Description
f	1.5 [m]	Optical System focal length
K_{px}, K_{py}	1800 [V/m]	Proportional Gain Factors
ω_{FSM}	100 [rad/s]	FSM axis natural frequency
ξ_{FSM}	0.3	FSM axis damping ratio
C_{ijk}	5×10^{-6} [N/Hz ²] [Nm/Hz ²] for $\{i = 1, 2, 3\}$	Tonal amplitude scale factor
h_{jk}	1, 2, 3	First 3 harmonics included
ϕ_{ijk}	$\{\phi_{i1k}, \phi_{i2k}, \phi_{i3k}, \phi_{i4k}, \phi_{i5k}, \phi_{i6k}\} = \{0, 90, 0, 0, 90, 0\}$ [°]	Phases of each force/torque component

Results

Now that the model is constructed, there are a few different ways it can be used to analyze the system's performance. By running the simulation in open loop and closed loop mode, it is possible to understand how effective the control system is at attenuating the disturbances. The results can be analyzed in a time domain, spatial domain, and frequency domain to assess closed loop performance.




Fig. 8 Time and Spatial Domain Results for $\Omega_1 = 56.4\text{Hz}$ (2nd harmonic at 1st x/y mode)

In the time domain, it is possible to see that sensor signals x and y are significantly less noisy with the presence of closed loop control. Pixel sizes on sensing element are often in the micron range. Submicron attenuation is achieved for this system using the simple closed loop controller proposed in the previous section. These results can also be analyzed in the spatial domain of x vs. y which is equivalent to the motion of the spot on the sensor. The open loop spot motion shows a limit cycle of large magnitude while the closed loop spot motion is significantly attenuated. As expected, there is a frequency dependent relationship between the wheel velocity and the rejected disturbances. It is useful to view this in the time domain to understand the absolute magnitude of the response, but the frequency dependent nature of the results is best understood in the frequency domain. This can be seen by computing the H_∞ norm of a frequency sweep across the structural model from the State space inputs to sensor measurement outputs and computing the singular values of the resulting frequency responses.[11]

The following plots show the frequency characteristics of the FSM control subsystem. The singular value plot shows a comparison of the worst case frequency response of the open loop and closed loop systems. Using the high bandwidth controller, it is possible to achieve meaningful attenuation at these high frequencies to reduce the LoS motion at the sensor.




Fig. 9 Frequency Domain Results (top) Singular Value plot from force/torque input at reaction wheel to sensor measurement output, (bottom left) closed loop Bode plot of FSM control system, (bottom right) closed loop output sensitivity FSM control system

The closed loop Bode plot shows the closed loop response with a -3dB closed loop bandwidth around 1.05kHz. There is significant peaking at the 675Hz resonance frequency of the closed loop system. This peak also appears in the singular value plot on the far right side. The presence of this peak is a tradeoff of the closed loop approach. Further attenuation of this peak may be possible with additional control design or mechanical optimization of the FSM mechanism. Finally, the output sensitivity Bode plot shows the relationship between the output disturbance input to the control system and the measurement output as a function of frequency. Disturbances are significantly attenuated up to around 430Hz where the curve passes -3dB.

Conclusion

The outlined procedure successfully develops a reduced order system level structural model that can be efficiently utilized in closed loop time domain simulation of a mechanical system with a flexible structural element while retaining model accuracy. As an example, the outlined procedure is demonstrated on a spacecraft structure presented with the challenge of reducing optical pointing error due to structural vibration during operation. A system level spacecraft model is assembled from two component superelements. The coupled model passes basic checks and is shown to match the result when explicit FEA is used. The second order model is then converted to a first order state space representation for use in control system simulation. Frequency domain response inspection reveals potentially problematic resonances. The structural model is then integrated into a full system level control model and a controller is utilized to reduce the impact of the identified structural resonances. The functionality of the controller design is demonstrated through a closed loop time domain simulation of pointing error given a reaction-wheel-like disturbance input.

Acknowledgments

Dr. Ashkan Arianpour

Dr. Sam Yunis

Chris Flanigan

References

1. Kabe, Kabe, A. M. (2020). *Structural dynamics fundamentals and advanced applications: Volume II*. Academic Press.
2. Klerk, D. d. (2008). *General framework for dynamic substructuring: History, review, and classification of techniques*. AIAA.
3. Griebel, M., & Wender, T. (2025). *Software agnostic Hurty/Craig-Bampton dynamic model reduction*. In IMAC XLIV.
4. Scholz, H. (2015). *Control theory of descriptor systems lecture notes*. TU Berlin.
5. Li, Y., Green, T. C., & Gu, Y. (2024). *Descriptor state space modeling of power systems*. IEEE Transactions on Power Systems, 39(4), 5495-5508.
6. Hespanha, J. P. (2018). *Linear systems theory* (2nd ed.). Princeton University Press.
7. Burge, J. H. (2006). *An easy way to relate optical element motion to system pointing stability*. *Current Developments In Lens Design And Optical Engineering VII, Proc. SPIE*, 6288, 62880I-1. DOI: 10.1117/OL.12.683809.
8. Gutierrez, H. L., Jr. (1999). *Performance assessment and enhancement of precision controlled structures during conceptual design* (Doctoral dissertation, Massachusetts Institute of Technology). Massachusetts Institute of Technology.
9. Liu, K.-C., Maghami, P., & Blaurock, C. (2008, August 18). Reaction wheel disturbance modeling, jitter analysis, and validation tests for Solar Dynamics Observatory (AIAA Paper No. 2008-7232). In *AIAA Guidance, Navigation, and Control Conference and Exhibit*.
10. Hawe, L. E. II. (2006). *Control of a fast steering mirror for laser-based satellite communication* (Master's thesis, Massachusetts Institute of Technology). Massachusetts Institute of Technology.
11. Zhou, K., & Doyle, N. (1998). *Essentials of robust control*. Prentice Hall, Inc.
12. Blelloch, P., & Tenger, J. C. (1998). *Application of Hruda-Benfield method to component damping for Space Shuttle coupled loads*. AIAA.
13. Yunis, Sam. "ShockSat." NESCAcademy Online, NASA Engineering and Safety Center, n.d., <https://nescacademy.nasa.gov/workshops/ShockSat>, Accessed 13 May 2026.



Chapter 13

Assessing Floquet Theory as an Alternative to Multi-Blade Coordinate Transformation in Wind Turbine Stability Analysis

Jose J. Mora Amaro, Emmanuel Branlard, and Riccardo Riva

Abstract We investigate the use of Floquet theory as an alternative to the Multi-Blade Coordinate transformation for wind turbine stability analysis. A simplified 5-DOF rotor model demonstrates the method and supports the development of an open-source Floquet analysis tool for future integration with multi-physics platforms such as OpenFAST.

Keywords Stability analysis · Coleman transformation · Floquet theory

Introduction

In the analysis of rotormachinery, the Multi-Blade Coordinate (MBC) transformation [1–3] is widely used to decouple periodic systems into fixed-frame representations, facilitating modal and stability analyses. However, MBC has limitations—particularly for anisotropic rotors and for systems with strong coupling or time-periodic behavior not well represented in the fixed frame [3, 4]. Current state-of-the-art wind turbine analysis tools typically perform stability assessments under isotropic assumptions, neglecting gravity effects and azimuthal wind variations [5]. However, as field observations reveal emerging instability issues, the importance of accounting for anisotropy and exploring alternatives to MBC is increasingly recognized.

In this work, we focus on exploring the application of Floquet theory [6] to perform stability analyses, and setting the stage for future application to multi-physics wind turbine softwares such as OpenFAST [5]. Our objective is to investigate the strengths and weaknesses of Floquet-based approaches relative to MBC and develop an open-source Floquet analysis tool tailored for wind turbine applications. Beyond establishing a stability criterion, our goal is also to extract modal characteristics (frequencies, damping ratios, and mode shapes) to guide design iterations when instability regions are identified. While related studies have addressed periodic stability analysis [7, 8], key questions remain regarding the generalization of the method, and, to the authors' knowledge, no open-source implementation is currently available.

In this preliminary work, we use a simplified 5-degree-of-freedom system, consisting of two degrees of freedom for the support structure and one edgewise degree of freedom per blade, following the structure of the original Coleman model. Using this model, we investigate the differences between Coleman-based and Multi-Blade Coordinate (MBC) transformations. Our goal is to identify the conditions under which Floquet theory provides clear advantages and to lower the barrier for its adoption in the wind energy modeling community.

Jose J. Mora Amaro · Emmanuel Branlard

Wind Energy Center, University of Massachusetts - Amherst, Amherst, MA 01002, USA

e-mail: jmoraamaro@umass.edu; ebanlard@umass.edu

Riccardo Riva

Technical University of Denmark, Roskilde, Denmark

e-mail: ricriv@dtu.dk

Summary of Floquet Theory

We consider a periodic coefficient system of period T defined as:

$$\dot{\underline{x}} = \underline{A}(t)\underline{x} \quad (1)$$

where the periodicity implies $\underline{A}(t + T) = \underline{A}(t)$. The Floquet transformation [6], performs a change of coordinates from $\underline{x}$ to a new set of coordinates $\underline{z}$, assuming that in this system, the state-space equation is a linear time invariant (LTI) system. The transformation writes:

$$\underline{x}(t) = \underline{Q}(t)\underline{z}(t) \quad (2)$$

where $\underline{Q}(t) = \underline{X}(t)e^{-\underline{B}t}$ is the Floquet transformation matrix, $\underline{B} = \frac{1}{T} \log(\underline{X}(T))$ is the matrix of Floquet factors [9], $\underline{X}$ is the fundamental matrix, and $\underline{X}(T)$ is the monodromy matrix (the fundamental matrix evaluated at T). After applying the transformation given in Equation 2 to Equation 1, we obtain the LTI system [9]:

$$\dot{\underline{z}} = \underline{B}\underline{z} \quad (3)$$

Simple stability analysis can be performed by looking at the eigenvalues of the monodromy matrix, called the characteristic multipliers. The eigenvectors of the monodromy matrix are written $\underline{S}$, and its eigenvalues $\underline{\theta}$. The eigenvalues of $\underline{B}$, denoted by $\underline{\eta}$, are called the Floquet exponents, and its eigenvectors are denoted by $\underline{V}$. It can be shown that $\underline{\eta}$ and $\underline{\theta}$ are related as [9]:

$$\theta_j = e^{\eta_j T}, \quad \Rightarrow \quad \eta_j = \frac{1}{T} \ln \theta_j \quad (4)$$

The complex logarithmic is multivalued and not invertible. Therefore, there is an indeterminacy when finding the Floquet exponents. Resolving this indeterminacy is challenging: stability can be determined directly, but obtaining mode shapes, frequencies, and Campbell diagrams requires additional effort.

Methods

Implementation details

The indeterminacy in Equation 4 can be lifted by considering the physical meaning of harmonic participation of each mode. A complete discussion of the significance of harmonic indeterminacy is given in [4]. The most accepted resolution is done by quantifying the modal participation across each possible frequency [7, 8, 10]. Riva et al. [7, 11] consider the relationship between the Floquet multipliers and exponents, directly constructing the eigenvalues of the LTI space by inference using Equation 4. To evaluate the modes in the frequency domain, the author infers a state space constructed on an equivalent eigenvector basis as the initial system, considering that $\underline{X}(0) = \underline{I}$. In this space, $\underline{V} = \underline{S}$. Based on that, the eigenvector basis for the LTP is:

$$\underline{\Xi}(t) = \underline{Q}(t)\underline{V} \quad (5)$$

Using a Fourier decomposition of $\underline{\Xi}$ and normalizing the coefficients, they quantify the modal participation of each frequency, thereby identifying the dominant oscillation frequency of each mode.

In this work, we solve the stability problem directly on the LTI system (Equation 3) by finding the eigenvalues of $\underline{B}$ (i.e., $\underline{\eta}$) in that space. Similar to the work of Riva, we project the eigenvectors on the rotating frame to quantify the modal participation. However, in our framework, we lift the inference made on the prior implementations by projecting the eigenvectors of the floquet exponent matrix rather than the eigenvectors of the monodromy matrix, generalizing the stability analysis to the LTI system. We will publish our framework on an open-source GitHub repository. Floquet analysis tool tailored for wind turbine applications.

5 DOFs system

We use the 5 degrees of freedom (DOF) system from Coleman [1] as a test case for our library. The model consists of two ground frame DOFs for the base motion for the rotor, and one edgewise degree of freedom for each blade.

Results

In Figure 1, we illustrate the application of our framework to the problem given in subsection 3.2. On this figure, we compare the MBC transformation and the Floquet theory results, focusing on the dominant frequencies.




Fig. 1 Comparison between symmetric modes obtained by MBC vs Floquet stability analysis

In the figure, the regressive and progressive modes of the blades are obtained by inference, adding ± 1 frequency period to the dominant frequency. And overall agreement with the MBC method is confirmed.

Conclusion and future work

We verified our methodology and will publish our framework as an open-source repository before the conference. We intend to apply our library to study wind turbine anisotropy further. In particular, we plan to apply this framework to OpenFAST, an open-source multi-physics tool to simulate wind turbines. This will enable use to conduct stability analyses of full wind turbine models capturing aero-hydro-servo-elastic interactions in a high-fidelity setting, and investigating where current MBC methodology fails.

Acknowledgments We would like to thank Justin Porter from National Renewable Energy Laboratory for his continuous support and fruitful discussions.

References

1. R. P. Coleman and A. M. Feingold. *Theory of Self-Excited Mechanical Oscillations of Helicopter Rotors with Hinged Blades*. Tech. rep. TR 1351. National Advisory Committee for Aeronautics (NACA), 1958.
2. W. Johnson. *Helicopter Theory*. Princeton, NJ: Princeton University Press, 1980.
3. G. S. Bir. “Multi-Blade Coordinate Transformation and Its Application to Wind Turbine Analysis”. In: *Proceedings of the AIAA/ASME Wind Energy Symposium*. Reno, NV, Jan. 2008.
4. D. A. Peters, S. M. Lieb, and L. A. Ahaus. “Interpretation of Floquet eigenvalues and eigenvectors for periodic systems”. In: *Journal of the American Helicopter Society* 56.3 (2011), pp. 1–11. DOI: [10.4050/JAHS.56.032002](https://doi.org/10.4050/JAHS.56.032002).
5. J. Jonkman and el. al. *OpenFAST v3.5.1. Open-source wind turbine simulation tool*, available at <http://github.com/OpenFAST/OpenFAST/>. 2023.
6. G. Floquet. “Sur les équations différentielles linéaires à coefficients périodiques”. In: *Annales scientifiques de l’École Normale Supérieure* 12.2 (1883), pp. 47–88. DOI: [10.24033/asens.220](https://doi.org/10.24033/asens.220). URL: <https://doi.org/10.24033/asens.220>.
7. R. Riva, S. Cacciola, and C. L. Bottasso. “Periodic stability analysis of wind turbines operating in turbulent wind conditions”. In: *Wind Energy Science* 1.2 (2016), pp. 177–203. DOI: [10.5194/wes-1-177-2016](https://doi.org/10.5194/wes-1-177-2016).

8. B. Pamfil, H. Bredmose, and T. Kim. “Coleman-free aero-elastic stability methods for three- and two-bladed floating wind turbines”. In: *Wind Energy Science* 10.4 (2025), pp. 827–856. DOI: [10.5194/wes-10-827-2025](https://doi.org/10.5194/wes-10-827-2025).
9. E. Folkers. *Floquet’s Theorem*. Bachelor’s Thesis. 2018.
10. C. L. Bottasso and S. Cacciola. “Model-independent periodic stability analysis of wind turbines”. In: *Wind Energy* 18.5 (2015), pp. 865–887. DOI: [10.1002/we.1735](https://doi.org/10.1002/we.1735).
11. R. Riva. *System ID library*, last accessed September 2025. URL: <https://gitlab.windenergy.dtu.dk/wtstab/sysid>.



Chapter 14

Using Minimum Entropy Decomposition for Rolling Element Bearing Fault Diagnosis in a Permanent Magnet Motor

Surendra (Suri) Ganeriwala

Abstract The rolling element bearing is one of the most vulnerable components in a permanent magnet motor because it is most often under high load and high-speed running conditions. Prompt diagnostics of rolling element bearing faults is critical not only for the safe operation of machines, but also for the reduction of maintenance cost. This work studies the diagnostics of bearing inner race and outer race faults. Rolling element bearing fault diagnosis using vibration signals relies profoundly on using the well-developed methodology of envelope analysis. Envelope analysis is traditionally performed by amplitude demodulating an excited frequency band using Hilbert transformation to determine the fault frequencies. The enveloping technique suffers due to a low signal-to-noise ratio and not knowing the frequency band where to apply it. In this paper we present the fault analysis results of bearing with inner race and outer race defects in a permanent magnet motor using traditional enveloping combined with advanced signal processing techniques. Minimum Entropy Decomposition (MED) is used to enhance bearing signal in cases with a low signal-to-noise ratio. Spectral Kurtosis (SK) is used to identify the demodulation band. Results show that the efficiency of the envelope analysis in providing a correct diagnosis is increased when it is augmented with bearing signal enhancing and frequency band determination techniques.

Keywords Outer race bearing fault · Inner race bearing fault · Minimum entropy deconvolution · Spectral kurtosis · Envelop analysis

Introduction

The rolling element bearing is widely used in today's rotating machinery. Unlike the long-life fluid bearing, the rolling element bearing often encounters metal-to-metal contacts under high load and high-speed running conditions, which makes it one of the most vulnerable components in a machine. Diagnostics of rolling element bearing faults can guarantee safe operation of the machine and reduce the maintenance cost. Faults usually occur on the four main components of a rolling element bearing, i.e., the outer race, the inner race, the ball, and the cage. Hence, it is crucial to detect the corresponding fault frequencies, i.e., ball pass frequency outer race (BPFO), ball pass frequency inner race (BPFI), ball spin frequency (BSF), and fundamental train frequency (FTF) [1, 2]. The values of them are determined by the geometry of the bearing and the rotating speed of the shaft.

When a local fault occurs in a bearing, a series of impacts are generated when the rollers or balls rolling over the fault. Like a hammer test in modal analysis, these impacts will excite the high frequency resonances of the machine. Depending on the location of the fault, different features of the vibration signal are generated by the fault, as shown in Figure 1. The outer race fault usually occurs in the load zone and are not modulated by any rotating speed. In contrast, the inner race fault is typically modulated by the shaft speed and the impact has the largest amplitude when the fault passes the load zone. Like the inner race fault, the rolling element ball fault is modulated by the FTF.

Surendra (Suri) Ganeriwala
Spectra Quest, Inc., Richmond, VA 23228
e-mail: suri@spectraquest.com




Fig. 1 Features of vibration signals generated by different locations of a fault in a rolling element bearing (stationary outer race) [3]

Envelope Analysis

In order to detect the bearing fault frequencies from a vibration signal measured by an accelerometer, envelope analysis is the main method to be applied. For a measured vibration signal $y(t)$, the first step of envelope analysis is to perform complex demodulation of the signal, which can be obtained by

$$Y(t) = y(t) \exp(-jp(f_l + f_u)t) \quad (1)$$

where f_l and f_u are the lower and upper limits of a selected frequency band. The second step is to lowpass filter the demodulated signal $Y(t)$ to compute the analytic signal $z(t)$. The cutoff frequency of the lowpass filter is $(f_u - f_l)/2$.

Minimum Entropy Deconvolution (MED)

One can see that the selection of a proper frequency band is the key of envelope analysis. For a rotating machinery, besides the bearing fault frequencies, there are other frequency components in the vibration signal, especially the rotating frequency and its harmonics. Those discrete frequency components contain the major energy in the signal and could mask the bearing fault frequencies, especially in the early stage of the bearing fault, which causes some difficulty to choose the proper frequency band. Hence, it is important to separate those discrete frequency components first. A lot of signal processing methods have been investigated for this purpose, such as time synchronization averaging, cepstrum, and adaptive filtering [4]. They require extra efforts to process the signal and are not in the scope of this paper. Since the vibration signal is measured by an accelerometer on the housing of a bearing instead of directly from the bearing, the transfer path from bearing to the accelerometer also has an effect on the acquired signal, such as reducing the amplitude of the impact generated by the bearing fault. In order to minimize the effect from the transfer path, the MED technique is used here [5, 6]. The output of a linear time-invariant system can be expressed as the convolution integral of the system input and system's impulse response

$$y(t) = x(t) * h(t) = \int_{-8}^8 h(t-t)x(t)dt \quad (2)$$

where $y(t)$ is the measured output signal, $x(t)$ is the original signal from the bearing, and $h(t)$ is the impulse response of the system. The deconvolution is to find an inverse filter to remove the effect of $h(t)$ and recover the original signal $x(t)$. In order to obtain the coefficients of the optimal filter, an optimization problem based on kurtosis is formulated. Kurtosis is a measure of “peakedness” of a signal. Since a series of impacts are generated by the local fault of the bearing, the optimization problem is to maximize the kurtosis of the output signal $y(t)$. More details on the MED technique can be found in Ref. [5].

Spectral Kurtosis (SK)

After the kurtosis of the vibration signal is maximized by the MED technique, a proper frequency band can be selected through the SK using a kurtogram [7, 8]. The SK can indicate the nonstationary or non-Gaussian frequency components of the signal in the frequency domain. In the kurtogram, small values of SK can be seen at frequencies where only stationary Gaussian noise exists, while large values of SK can be seen at frequencies where transient occurs. The SK of a signal $y(t)$ can be computed based on the short-time Fourier transform (STFT) of the signal

$$S(t, f) = \int_{-8}^8 y(t)w(t-t)e^{-2\pi f t} dt \quad (3)$$

where $w(t)$ is the window function used in STFT. The SK is calculated by

$$SK(f) = \frac{\left(\sum_{i=1}^n |S(t_i, f)|^4\right)/n}{\left(\left(\sum_{i=1}^n |S(t_i, f)|^2\right)/n\right)^2} - 2, f \neq 0 \quad (4)$$

STFT-based kurtogram is associated with constant time and frequency resolution. Sawalhi [3] used complex Morlet wavelets to construct different filter banks to calculate kurtogram. The wavelets have the property of uniform resolution on a logarithmic frequency scale. Antoni [9] proposed a fast kurtogram algorithm using bandpass filtering along with a simplified computation to approximate the SK for each window size and frequency rather than compute the STFT, which makes it more suitable for on-line industrial applications. This paper will use the fast algorithm to calculate kurtogram.

Experimental Investigation

Experiments were conducted on a drive-train diagnostics/prognostics test rig, as shown in Figure 2. This test rig can be used to study the transmission diagnostics/prognostics and gear dynamics. One bearing on the drive motor was seeded with outer race fault and inner race fault separately. The bearing information is shown in Table 1. Two single-axis accelerometers that were used to measure horizontal and vertical vibration were mounted on the housing of the motor, as shown in Figure 3. One can see that there is a transfer path between the bearing and actual measurement location. The bearing fault analysis were investigated at three rotating speeds, including 1800 RPM, 3600 RPM and 5000 RPM. Corresponding bearing fault frequencies are shown in Table 2. The data was acquired by SpectraQuest DAQ and in-house software VibraQuest. The sampling frequency was 102.4 kHz and the acquisition time was 12.8 s. Vibration data in the vertical direction is used for analysis in the following sections.




Fig. 2 Drive-train diagnostics/prognostics test rig

Table 1 Bearing information

Bearing	# of rolling elements	Pitch diameter	Rolling element diameter	Contact angle
6307	8	2.264 in	0.531 in	0 degree




Fig. 3 Faulted rolling element bearing and measurement location

Table 2 Bearing fault frequencies at different rotating speeds

Rotating Speed	BPFO	BPFI	BSF	FTF
1800 RPM	91.855 Hz	148.145 Hz	60.437 Hz	11.482 Hz
3600 RPM	183.710 Hz	296.290 Hz	120.873 Hz	22.964 Hz
5000 RPM	255.153 Hz	411.514 Hz	167.880 Hz	31.894 Hz

Outer Race Fault

The raw vertical vibration data between 0 s and 0.1 s at the three rotating speeds are shown in Figure 4. A series of impacts corresponding to the BPFO can be seen in Figure 4(a) at the rotating speed of 1800 RPM, though not very prominent. As the rotating speed increases, the impacts are less prominent and are mixed with general vibration. The raw data was then processed by the MED technique, and the results are shown in Figure 4. Compared to Figure 4(a), the impacts in Figure 4(a) are more prominent, which makes it easier to identify the corresponding BPFO. Appearance of the impacts in Figure 4(b) and (c) is also improved. From Figures 4 and 5, one can see that the transfer path from the bearing to the measurement




Fig. 4 Raw vertical vibration data corresponding to the outer race fault at (a) 1800 RPM, (b) 3600 RPM and (c) 5000 RPM




Fig. 5 Data processed by the MED technique at (a) 1800 RPM, (b) 3600 RPM and (c) 5000 RPM

location has a negative effect on the impacts generated by the bearing fault and the MED technique can reduce this negative effect. The kurtosis was calculated before and after applying the MED technique and is shown in Table 3. One can see that the kurtosis is greatly increased after applying the MED technique at all the three rotating speeds. Since the kurtosis is a measure of “peakedness”, larger kurtosis indicates more obvious “peakedness” in the data, which agrees well with the results in Figure 5.

Kurtogram of the processed data is shown in Figure 6 and Figure 7, respectively. The proper frequency band for envelope analysis can be obtained from the kurtogram. The amplitude spectrum of the envelop signal was calculated using the frequency band obtained from the kurtogram by each method, as shown in Figures 6 and 7. Though the proper frequency bands obtained by the two methods are different, the BPFO and its harmonics can be clearly identified at each rotating speed in Figures 6 and 7. Since this is an outer race fault, there is no sideband around the BPFO.




Fig. 6 Kurtogram corresponding to the outer race fault at (a) 1800 RPM, (b) 3600 RPM and (c) 5000 RPM calculated by the fast algorithm




Fig. 7 Amplitude spectrum of the square envelope signal from the fast algorithm corresponding to the outer race fault at (a) 1800 RPM, (b) 3600 RPM and (c) 5000 RPM

Table 3 Kurtosis before and after applying the MED technique

Rotating Speed	Before	After
1800 RPM	3.0254	42.5602
3600 RPM	3.8301	16.5063
5000 RPM	2.3856	8.7642

Inner Race Fault

The raw vertical vibration data corresponding to the inner race fault is shown in Figure 8. It is difficult to tell the location of an impact in Figure 8(a). Some impacts can be seen in Figure 8(b) and (c). The effect of amplitude modulation by the rotating shaft is not obvious. The raw data was then processed by the MED technique and is shown in Figure 9. One can see that the impacts are more obvious in the processed data compared to those in the raw data. The kurtosis was calculated before and after applying the MED technique and is shown in Table 4. One can see that the kurtosis is slightly increased, but




Fig. 8 Raw vertical vibration data corresponding to the inner race fault at (a) 1800 RPM, (b) 3600 RPM and (c) 5000 RPM




Fig. 9 Data processed by the MED technique at (a) 1800 RPM, (b) 3600 RPM and (c) 5000 RPM

Table 4 Kurtosis before and after applying the MED technique

Rotating Speed	Before	After
1800 RPM	0.5715	4.1604
3600 RPM	1.6744	4.5812
5000 RPM	3.2212	6.8436

not as much as the case in Sec. 2.1. One possible reason is that the transfer path in this case has more components involved since the fault is from inner race. Impacts generated by the inner race fault is weaker at the measurement location.

Kurtogram of the processed data is shown in Figures 8 and 9. One can see that the bandwidths are narrower than those in Sec. 2.1, which indicates that the energy of the impacts generated by the inner race fault is smaller. The frequency bands obtained from the two methods are similar at 3600 RPM and 5000 RPM. The proper frequency band in Figure 8(a) is low and at the range of rotating frequency. The amplitude spectrum of the envelope signal was calculated using the frequency band obtained from the kurtogram, as shown in Figures 10 and 11. One can see that they are similar at 3600 RPM and 5000 RPM. However, the amplitude spectrum in Figure 10(a) has a low signal-to-noise ratio at 1800 RPM due to the low center frequency in Figure 10(a). The BPFI can be seen from the amplitude spectrum. However, the rotating frequency and its harmonics also play a major role in the amplitude spectrum, which may hinder identification of the BPFI. Compared Figure 10 with Figure 8 and Figure 11 with Figure 9, one can see that the identification of BPFI is not as obvious as that of BPFO.




Fig. 10 Kurtogram corresponding to the inner race fault at (a) 1800 RPM, (b) 3600 RPM and (c) 5000 RPM calculated by the fast algorithm




Fig. 11 Amplitude spectrum of the square envelope signal from the fast algorithm corresponding to the inner race fault at (a) 1800 RPM, (b) 3600 RPM and (c) 5000 RPM

Conclusion

Fault analysis of rolling element bearings using MED and SK is presented in this paper. A brief introduction of the methodology is given. Experiments on a drive-train test rig were conducted to investigate the methodology. Some observations are summarized as follows:

- (1) By maximizing the SK of a measured signal, an inverse filter can be designed by the MED technique to reduce the effect of the transfer path on the measured signal.
- (2) The kurtogram can help select a proper frequency band for envelope analysis. Two methods to calculate the kurtogram are investigated in this paper. Though the frequency bands obtained from the wavelet method and the fast algorithm are different, the amplitude spectra of the square envelope signals from the two methods are similar and both of them can be used to identify bearing fault frequencies.
- (3) Compared to the BPFO, the rotating frequency and its harmonics have more negative effect on the identification of the BPFI.

The combination of MED and SK shows a great potential for an automated diagnosis procedure to identify rolling element bearing faults.

References

1. Randall, R. B., & Antoni, J. (2011). Rolling element bearing diagnostics—A tutorial. *Mechanical systems and signal processing*, 25(2), 485-520.
2. Smith, W. A., & Randall, R. B. (2015). Rolling element bearing diagnostics using the Case Western Reserve University data: A benchmark study. *Mechanical Systems and Signal Processing*, 64, 100-131.
3. Sawalhi, N. (2007). Diagnostics, prognostics and fault simulation for rolling element bearings. *The University of New South Wales, School of Mechanical and Manufacturing Engineering, Australia, Doctor of Philosophy thesis*.
4. Randall, R. B., Sawalhi, N., & Coats, M. (2011). A comparison of methods for separation of deterministic and random signals. *International Journal of Condition Monitoring*, 1(1), 11-19.
5. Endo, H., & Randall, R. B. (2007). Enhancement of autoregressive model based gear tooth fault detection technique by the use of minimum entropy deconvolution filter. *Mechanical Systems and Signal Processing*, 21(2), 906-919.
6. Sawalhi, N., Randall, R. B., & Endo, H. (2007). The enhancement of fault detection and diagnosis in rolling element bearings using minimum entropy deconvolution combined with spectral kurtosis. *Mechanical Systems and Signal Processing*, 21(6), 2616-2633.
7. Antoni, J. (2006). The spectral kurtosis: a useful tool for characterising non-stationary signals. *Mechanical Systems and Signal Processing*, 20(2), 282-307.
8. Antoni, J., & Randall, R. B. (2006). The spectral kurtosis: application to the vibratory surveillance and diagnostics of rotating machines. *Mechanical systems and signal processing*, 20(2), 308-331.
9. Antoni, J. (2007). Fast computation of the kurtogram for the detection of transient faults. *Mechanical Systems and Signal Processing*, 21(1), 108-124.

Special Topics & Experimental Techniques in Structural Dynamics, Volume 5

Dario Di Maio ■ Matthew Allen ■ Matthew Stefanski ■ Chad Walber

Special Topics & Experimental Techniques in Structural Dynamics, Volume 5: Proceedings of the 44th IMAC, A Conference and Exposition on Structural Dynamics, 2026, the fifth volume of ten from the Conference brings together contributions to this important area of research and engineering. The collection presents early findings and case studies on fundamental and applied aspects of Structural Dynamics, including papers on:

- Experimental Techniques
- Dynamics of Rotating Machinery
- Metamaterials, Metastructures and Additive Manufacturing
- Dynamics of Biomedical and Bioinspired Systems

ISBN 978-87-438-1444-3



9 788743 814443



River Publishers