

AI detection of damage and faults in concrete slabs using an LDV and projectile excitation.

Michiel Arnouts & Jasper Laforce
Supervisors: Nasser Ghaderi & Rhys Evans

2025-2026

Abstract

Internal defects in concrete structures, such as voids and delaminations, are difficult to detect reliably using conventional visual inspection or manual sounding, particularly when the damage is hidden below the surface. This study investigates an automated vibration-based approach for concrete defect detection by combining laser Doppler vibrometry, repeatable projectile excitation, and machine-learning-based signal interpretation. Three laboratory-scale concrete specimens were examined: one healthy reference specimen, one specimen containing artificial void-type defects, and one specimen containing artificial delamination-type defects. A pneumatic steel-ball launcher was used to apply localized impacts on a predefined measurement grid, while the resulting surface vibrations were measured without contact using an LDV. The measured time-domain responses were transformed into frequency-domain spectra and used as input for data-driven damage assessment.

Two neural-network strategies were evaluated. First, an autoencoder was trained on healthy reference data to detect deviations from the expected vibration response. However, the reconstruction errors of the healthy and damaged measurements were found to be comparable, indicating that the autoencoder did not provide reliable anomaly detection for the present dataset. Therefore, a supervised multilayer perceptron was developed to classify measurement points as healthy, void-related, or delamination-related. The classifier achieved an overall accuracy of 97.7%, with spatial prediction maps showing good agreement with known defect layouts. Additional validation indicated that the model did not rely solely on memorizing measurement positions, although performance decreased when there was less available spatial training data.

The results demonstrate that LDV-derived frequency spectra contain defect-sensitive information that can support the supervised classification of internal concrete defects under controlled laboratory conditions. Nevertheless, the limited number of specimens, the artificial defect configurations, and the laboratory environment mean that the proposed approach should be regarded as a proof of concept. More validation on larger datasets, more realistic damage scenarios, and field-like measurement conditions are required before practical deployment.

Contents

1	Introduction	3
2	Literature Study	5
2.1	State-of-the-Art NDT for concrete inspection	5
2.1.1	Laser Doppler Vibrometry	7
2.1.2	Limitations of the current NDT	7
2.2	Artificial Intelligence for NDT and SHM	7
2.3	AI models for vibration-based defect detection	8
2.3.1	Autoencoders	9
2.3.2	MLP Classifier	10
2.4	Research gap and positioning of the present work	10
3	Methodology	11
3.1	Materials	11
3.1.1	Obtaining the Young's modulus and the Poisson ratio	13
3.2	Siemens NX: Solver parameters and settings	14
3.3	Experimental setup	14
3.3.1	Final setup and working principle	17
3.3.2	LDV scanning measurement setup	18
3.3.3	Automated excitation system	19
3.3.4	Synchronization between LDV scanning and excitation	19
3.3.5	Safety considerations	20
3.4	Machine learning approaches	20
3.4.1	Working principle of AI-based models	20
3.4.2	1D-Autoencoder	22
3.4.3	MultiLayer Perceptron Classifier	25
4	Results	27
4.1	Numerical Simulations	27
4.1.1	Healthy concrete slabs	27
4.1.2	Delamination Concrete Slab	28
4.1.3	Voids defected concrete slab	29
4.2	Vibration Data	30
4.2.1	Post-Processing the Raw Vibration Data	30
4.2.2	Comparing	32
4.2.3	Visualizing the Eigenshapes	33
4.3	Machine learning models	35
4.3.1	1D-Autoencoder	35
4.3.2	MultiLayer Perceptron Classifier	36
5	Discussion	38
5.1	Relation to earlier numerical observations	38
5.2	Comparing the Data	40
5.2.1	Vibration data	40
5.2.2	Frequency domain as input for AI	41
5.3	Use of the types of AI models	41
5.4	Limitations and future work	42
5.4.1	Limitations	42
5.4.2	Future work	43
6	Conclusion	45
7	References	47

1 Introduction

Concrete structures form a major part of the built environment, but their condition can gradually deteriorate during their service life. Defects such as cracking, voids, segregation, and delamination may develop inside the material and remain hidden until visible damage appears on the surface. Reliable detection of these defects is therefore important for assessing structural condition and planning maintenance interventions. This is particularly relevant for safety-critical infrastructure such as bridges, tunnels, and large retaining structures, where early damage identification is a central objective of Structural Health Monitoring (SHM) and Non-Destructive Evaluation (NDE) [1, 2, 3, 4]. In the broader literature on SHM, damage assessment is commonly understood as a hierarchy of tasks that progresses from detection and localization to classification, quantification, and ultimately integrity assessment [1, 2]. This distinction is important for concrete inspection because the practical value of a sensing and analysis system depends not only on whether it can identify anomalies, but also on the type of engineering decision that its output can support.

Conventional inspection methods, including visual assessment and hammer sounding, remain widely used in practice because they are simple and inexpensive. However, they are inherently subjective, difficult to standardize and often ineffective for inaccessible or concealed defects [5, 6]. Contact-based NDE approaches such as ultrasonic testing and piezoelectric sensing provide more quantitative information, but their use may be labor-intensive and operationally restrictive when large surfaces, complex geometries, or hard-to-reach areas must be examined [6, 7]. In this context, a Laser Doppler Vibrometer (LDV) has emerged as a particularly attractive alternative because it enables non-contact vibration measurements with high spatial and temporal resolution, without sensor attachment, mass loading, or dense wiring [5, 8, 9]. The potential of laser-based concrete inspection has therefore been explored in a range of measurement configurations, including point-wise LDV, scanning LDV, remote acoustic excitation, and laser-ablation-based impact excitation [10, 11, 12, 13]. In our previous work, we systematically compared acoustic, mechanical and shaker excitation for LDV-based defect detection in concrete, showing that the excitation strategy strongly affects the signal-to-noise ratio, repeatability, and the visibility of defect-sensitive resonance features [14].

However, the main challenge is not the acquisition of vibration data itself, but the interpretation of that data in a reliable and physically meaningful manner. Defects do not appear directly in LDV measurements. Instead, they modify the local dynamic response through changes in stiffness, damping, resonance behavior, and wave propagation. In heterogeneous materials such as concrete, these effects are often subtle and can easily be masked by variability in excitation, surface reflectivity, measurement geometry, and environmental noise [5, 15, 16]. Earlier work already demonstrated that LDV can serve as a powerful receiver in ultrasonic and impact-echo-type concrete measurements, particularly when signal conditioning and spatial processing are used to improve detectability [8, 9]. More recent studies have extended this principle to a truly non-contact concrete inspection using acoustic-irradiation-induced vibration and laser-ablation excitation, showing that hidden defects can be detected, localized, and in some cases physically characterized from a distance [10, 17, 18, 11, 12, 13, 16].

Alongside these sensing advances, Artificial Intelligence (AI) and machine learning have become increasingly important for the interpretation of complex structural response data. Across SHM and NDE, recent reviews consistently report a shift from handcrafted feature engineering toward data-driven classification and representation learning, while also highlighting that greater diagnostic ambition requires stronger labels, more diverse datasets, and greater robustness to irrelevant sources of variability [6, 7, 3, 4]. Within the specific context of concrete inspection, AI-based methods have already been successfully applied to several NDE modalities. For example, convolutional neural networks have been used for defect detection in ultrasonic tomography images of concrete, while wavelet-packet-based neural and hybrid learning methods have been reported for the classification of defects in concrete blocks [19, 20, 21]. These studies demonstrate that automated defect recognition in concrete is feasible, but they also show that performance is closely related to data quality, measurement conditions, and the definition of the target classes.

The literature that is most directly relevant to the present work is the smaller but rapidly developing body of research that combines LDV with AI for concrete defect detection. Muramatsu et al. demonstrated that LDV frequency spectra can be used as input to a one-dimensional convolutional neural network for non-contact crack detection in concrete, reporting accuracies above 90% and up to 98.9% for the best pre-processing and network configuration [5]. Their results clearly show that AI can extract damage-sensitive information from vibrometry data even when spectral differences are too subtle for robust manual interpretation. Tian et al. extended this direction beyond binary screening by proposing a multi-class defect-classification framework in which LDV signals were used to distinguish between no defect, segregation, cavity, and foreign matter in concrete specimens [22]. In parallel, Tian et al. also proposed an LDV-based weighted spatial-spectral entropy method for internal defect imaging, thereby emphasizing that localization and signal denoising remain central issues for practical non-contact concrete inspection [15]. Additional LDV-based studies further showed that cavity size, burial depth, filling condition, and excitation frequency substantially affect the measured vibration response, indicating that the physical interpretation of LDV signals is strongly defect and context dependent, rather than reducible to defect name alone [17, 18, 16].

Taken together, the existing literature establishes the promise of LDV as a non-contact sensing modality for concrete diagnostics and confirms that AI can enhance the interpretation of the resulting signals. At the same time, the literature remains fragmented across different diagnostic objectives. Some studies are limited to binary damage screening, others focus on controlled laboratory-scale multi-class defect recognition, and others focus primarily on resonance-based localization or defect imaging [5, 22, 15, 12, 13, 16]. Consequently, a fundamental methodological question remains insufficiently resolved: for LDV-based AI in concrete inspection, should the primary goal be robust damage/no-damage screening, or is it already justified to pursue higher-level tasks such as defect-type classification and, eventually, severity-oriented characterization? This question is not merely algorithmic. It is directly related to the SHM damage-assessment hierarchy and to the level of physical interpretability, dataset quality, and generalization required for real inspection practice [1, 2, 6, 3, 4].

Building on our previous comparison of acoustic, mechanical, and shaker excitation for LDV-based defect detection in concrete [14], the present study investigates how LDV measurements can be combined with artificial intelligence for automated defect detection and characterization. Separately from the experimental excitation study on which the present work is directly based, an earlier numerical study indicated that internal concrete defects can produce shifts in eigenfrequencies and local changes in mode shapes, which supports the physical expectation that damaged samples can deviate from a healthy vibration-response baseline [23]. By positioning the problem within both the SHM/NDE literature and the emerging body of LDV-based concrete research, this work aims to clarify what level of diagnostic resolution is technically achievable, scientifically defensible and practically meaningful for non-contact inspection of concrete structures [5, 22, 15, 16, 1].

The relevance of this work lies in the transition from subjective, labor-intensive concrete inspection to automated, quantitative, and data-driven structural health monitoring. Current inspection practices, such as visual inspection and manual hammer sounding, remain valuable because they are simple and inexpensive, but their outcome strongly depends on the experience of the inspector, the accessibility of the structure (e.g. difficult due to asbestos [24, 25]) and the local testing conditions. For large concrete structures such as bridges, tunnels, and walls, this makes consistent large-area inspections difficult, time-consuming, and difficult to standardize. Moreover, early-stage or hidden defects may not produce clear visible or audible indications, even though they can already influence the local dynamic response of the structure.

The combination of laser Doppler vibrometry, automated projectile excitation, and artificial intelligence offers a promising route to overcome part of these limitations. LDV enables non-contact vibration measurements without attaching sensors to the concrete surface, while automated excitation improves the repeatability of the impact conditions compared with manual hammer impact. The resulting vibration signals can then be processed quantitatively and interpreted using AI-based methods.

In this way, inspection can evolve from a primarily experience-based judgment, toward a more reproducible measurement process in which each point on the structure is associated with a numerical damage-sensitive response.

This does not mean that such a system can already replace conventional inspection in all practical situations. Field conditions, environmental noise, surface variability, boundary conditions, and the diversity of real defects remain important challenges. However, the long-term potential is significant. If robustly validated, an automated LDV-based inspection system could reduce operator dependency, improve repeatability, enable faster scanning of large surfaces, and support earlier detection of internal defects before severe deterioration becomes visible. The present thesis should therefore be seen as a step toward future inspection methods in which human expertise is supported by objective vibration measurements and AI-based decision support, rather than replaced by it.

2 Literature Study

The purpose of this literature study is to position the proposed inspection method within the broader field of Non-Destructive Testing (NDT), Structural Health Monitoring (SHM), and data-driven damage assessment of concrete structures. Concrete defects such as voids, delaminations, cavities, segregation zones, and cracks are often difficult to detect because they may remain hidden below the surface until visible deterioration occurs. Reliable inspection methods must therefore be able to detect changes in material condition without damaging the structure itself. This requirement has led to the development of several NDT techniques, each with specific advantages, limitations, and suitable application domains.

For the present work, the most relevant literature can be grouped into three main categories. The first category concerns established NDT methods for concrete inspection, such as ultrasonic testing, infrared thermography, impact-echo testing, and vibration-based methods. The second category concerns laser Doppler vibrometry, which is particularly relevant because it enables non-contact vibration measurements. The third category concerns Artificial Intelligence (AI) and machine learning methods, which are increasingly used to interpret complex NDT signals and automate defect recognition. By combining these three research directions, the present study investigates whether LDV-derived vibration spectra, generated through repeatable projectile excitation, can be used for automated detection and classification of internal concrete defects.

2.1 State-of-the-Art NDT for concrete inspection

NDT methods are used to evaluate the internal condition of concrete without removing material or reducing structural capacity [26, 27]. In practice, visual inspection and hammer sounding are still widely used because they are simple, fast, and inexpensive [26, 6]. However, these methods strongly depend on the experience of the inspector and mainly provide qualitative information [26, 6]. Hidden defects, early-stage delamination, and internal voids may remain undetected when they do not yet produce visible surface symptoms or a clearly audible response [27, 26]. For this reason, more quantitative NDT techniques are required when the goal is to detect internal damage in a reproducible way. In the next sections, different techniques are discussed.

Ultrasonic Testing (UT)

Ultrasonic testing is one of the most established NDT techniques for concrete inspection [28, 27]. The method is based on the propagation of high-frequency stress waves through the material [28, 27]. In ultrasonic pulse velocity testing, a pulse is transmitted through the concrete and received at another location [28]. The measured travel time is used to calculate the wave velocity, which is influenced by material stiffness, density, cracking, voids, and other discontinuities [28, 26]. In general, sound and dense concrete produces higher wave velocities, whereas damaged or poorly compacted concrete tends to produce lower velocity or attenuated signals [28, 27].

Ultrasonic methods can provide quantitative information about the condition of concrete and are suitable for detecting internal defects, estimating material quality, and supporting tomography-based imaging [26, 27]. Recent work has also combined ultrasonic testing with machine learning methods. For example, ultrasonic tomography images and ultrasonic signals have been used as input for convolutional neural networks and wavelet-packet-based models for concrete flaw recognition [19, 20, 21, 29, 30]. These studies show that ultrasonic NDT can provide defect-sensitive information and that AI can improve the interpretation of complex ultrasonic data [19, 20, 21, 29, 30].

Nevertheless, ultrasonic testing also has practical limitations. Many ultrasonic configurations require physical contact between the transducers and the concrete surface [26, 27]. Coupling conditions, surface roughness, moisture, reinforcement, aggregate heterogeneity, and access to both sides of the structure can strongly influence the quality of the measurement [26, 28]. For large concrete surfaces, dense point-by-point ultrasonic measurements may become time-consuming [26]. Therefore, although ultrasonic testing is powerful, its practical implementation can be challenging for rapid large-area inspection.

Infrared Thermography

Infrared thermography is a non-contact NDT technique that evaluates temperature patterns on the surface of a structure [31, 32]. The principle is based on the fact that internal defects can disturb heat flow [31, 32]. When a concrete surface is heated or cooled, subsurface delaminations, voids, or debonded regions may produce local thermal contrasts because their thermal conductivity and heat capacity differ from those of sound concrete [31, 33]. These contrasts can be detected using an infrared camera [31].

A major advantage of infrared thermography is that it can inspect relatively large areas quickly and without physical contact [31, 33]. Therefore, it is attractive for bridge decks, tunnel linings, walls, and other concrete surfaces where fast screening is required [31, 32]. However, the technique is mainly sensitive to near-surface defects and strongly depends on environmental conditions [31, 32]. Solar radiation, wind, surface moisture, emissivity variations, shadowing, and the time of measurement can all influence the observed thermal pattern [32, 33]. As a result, thermography is effective as a rapid screening tool, but it may require careful interpretation or complementary methods when deeper internal defects must be characterized [33, 32].

Impact Echo

Impact-echo is another important NDT method for concrete structures [34, 35]. In this method, a short mechanical impact is applied to the surface of the concrete, generating stress waves that propagate through the structure [34, 35]. Reflections from boundaries, voids, delaminations, or other internal discontinuities affect the measured surface response [35, 36]. By analyzing the frequency content of the response, information can be obtained about plate thickness, internal defects, or changes in stiffness [36, 34].

Impact-echo is particularly relevant to the present study because both rely on the dynamic response of concrete after mechanical excitation [35, 34]. The main difference is that conventional impact-echo measurements often use contact sensors, whereas the present work uses an LDV to measure the surface vibration without contact [35, 37]. Recent studies have also combined impact-echo data with AI. For example, machine-learning and deep-learning models have been used to classify delaminations, honeycombing, and other defects from impact-echo spectra or time-frequency representations [38, 39, 40, 41]. These studies demonstrate that impact-generated vibration signals contain defect-sensitive information that can be exploited by data-driven models [38, 39, 40, 41].

However, impact-based methods are sensitive to the repeatability of the excitation. Manual hammer impacts can vary in force, contact duration, impact angle, and location [35, 42]. These variations can change the frequency content of the measured response and complicate automatic interpretation [35, 42]. This is one of the reasons why automated and repeatable excitation is important for AI-based vibration analysis.

2.1.1 Laser Doppler Vibrometry

Laser Doppler vibrometry is an optical measurement technique that determines surface vibration velocity from the Doppler shift of reflected laser light [37]. Because the measurement is contactless, LDV does not require sensor attachment, does not introduce mass loading, and can be used at locations where conventional sensors are difficult to install [37]. These properties make LDV highly attractive for concrete inspection, especially when large surfaces, inaccessible areas, or automated scanning procedures are considered [37, 43].

In concrete NDT, LDV has been used as a receiver in several vibration-based and acoustic inspection configurations. Earlier studies demonstrated that LDV can measure local vibration responses caused by impact excitation, acoustic irradiation, or laser-ablation-induced excitation [43, 9, 10, 12, 44, 13]. These studies show that defects such as cracks, cavities, voids, and delaminations can influence the measured vibration response through changes in local stiffness, damping, resonance behavior, and wave propagation [10, 12, 44, 13]. More recent work has also shown that defect size, burial depth, filling condition, and excitation frequency can affect the LDV response [45, 46, 47, 18]. This indicates that LDV signals contain physically meaningful information, but that the interpretation of these signals is not always straightforward [5, 22, 45].

2.1.2 Limitations of the current NDT

Although existing NDT methods are useful, none of them provides a complete solution for automated inspection of large concrete structures. Visual inspection and hammer sounding are inexpensive, but subjective and mainly qualitative [26, 6]. Ultrasonic testing can provide quantitative information, but often requires contact, coupling, and careful positioning of the sensors [28, 27]. Infrared thermography is fast and non-contact, but is strongly influenced by environmental conditions and is mainly sensitive to defects that produce measurable thermal contrasts near the surface [31, 32]. Impact-echo is sensitive to internal defects, but the results depend strongly on excitation repeatability, boundary conditions, signal processing, and operator experience [35, 34].

These limitations become more important when the inspection target is a large concrete structure such as a bridge, tunnel, wall, or retaining structure. In such cases, a practical inspection method should be repeatable, scalable, non-contact where possible, and capable of producing quantitative data over a spatial measurement grid. The method should also reduce operator dependency and provide results that can be interpreted consistently. LDV-based inspection addresses part of this need because it enables non-contact vibration measurement [37, 43]. However, LDV alone does not automatically provide a defect classification [5, 22]. The measured vibration signals must still be transformed into reliable diagnostic information. This is where AI-based interpretation becomes relevant [5, 22, 3].

2.2 Artificial Intelligence for NDT and SHM

Artificial Intelligence (AI) and machine learning have become increasingly important in Non-Destructive Testing (NDT) and Structural Health Monitoring (SHM), because modern inspection methods often generate large and complex datasets that are difficult to interpret using only manual judgment or simple threshold-based rules. In structural applications, damage is usually not measured directly. Instead, defects influence the measured response indirectly through changes in wave propagation, resonance behavior, damping, local stiffness, thermal response, or spectral content. This is especially relevant for concrete, where material heterogeneity, reinforcement, boundary conditions, surface quality, and environmental noise can mask defect-sensitive information [3, 4, 6].

In traditional NDT workflows, the interpretation of measurements often relies on expert experience, manually selected features, or predefined decision criteria. Although these approaches remain useful, they may be sensitive to operator dependency and measurement variability. Data-driven methods offer an alternative by learning relationships between measured signals and structural condition.

In concrete inspection, machine-learning and deep-learning methods have already been applied to ultrasonic measurements, impact-echo data, acoustic/hammering responses, thermographic images, and vibration based signals [7, 3, 6]. These studies show that AI can support automated defect detection and classification, provided that the measured data contain physically meaningful damage-sensitive information.

For vibration-based NDT, AI is particularly relevant because defects often produce distributed changes in the dynamic response rather than one isolated and easily interpretable signal feature. As mentioned above defects may lead to changes in resonance frequencies. Frequency-domain and time-frequency representations are therefore commonly used in learning-based approaches, since they can make defect-sensitive patterns more accessible than raw time-domain signals alone [3, 38, 39, 3].

Within LDV-based concrete inspection, the combination of non-contact vibration measurement and AI has shown promising results. Muramatsu et al. demonstrated that vibration signals obtained with an LDV can be used for non-contact concrete flaw detection with neural-network classification [5]. Tian et al. further showed that LDV-derived vibration signals can support multi-type concrete defect classification, distinguishing sound concrete from different internal defect categories [22]. These studies indicate that LDV measurements can contain defect-sensitive information suitable for data-driven interpretation. At the same time, they also show that model performance depends strongly on signal quality, pre-processing, excitation methods, and the representativeness of the dataset.

In vibration-based NDT, the way in which raw signals are represented before training is an important part of the diagnostic workflow. Time-domain signals are often transformed into the frequency domain because defects may influence resonance frequencies, spectral amplitudes, and damping-related characteristics rather than producing a clearly visible change in the raw signal. Frequency-domain analysis and spectral estimation are therefore commonly used in dynamic testing and vibration-based structural health monitoring [48, 49, 3]. At the same time, frequency ranges that are dominated by unrelated rigid-body motion, environmental noise, sensor artifacts, or low-energy content may be excluded when they do not contribute to the diagnostic objective. Such filtering or frequency-band selection should be justified by both physical reasoning and measured signal quality, since removing irrelevant components can improve the robustness of the features used for learning-based models [50, 48].

2.3 AI models for vibration-based defect detection

Artificial neural networks

Artificial neural networks are widely used for pattern recognition and signal interpretation because they can learn nonlinear relationships between input features and output predictions. A neural network consists of interconnected layers in which input data are transformed through trainable weights, biases, and activation functions. During training, the model updates these internal parameters to minimize a loss function, allowing it to learn patterns that are useful for classification, regression, or reconstruction tasks [51, 52].

The basic computational unit of a neural network is the artificial neuron. A neuron receives one or more input values, combines them using trainable weights and a bias term, and then applies an activation function to produce an output. In a fully connected layer, each neuron is connected to all outputs of the previous layer, allowing the network to combine information from the complete input feature vector. By stacking multiple fully connected layers, a multilayer perceptron can learn increasingly abstract internal representations of the input data [51, 52].

Activation functions are essential because they introduce nonlinearity into the network. Without nonlinear activation functions, a stack of fully connected layers would still behave like a single linear transformation and would therefore be unable to model complex relationships between vibration features and structural condition. Common activation functions include sigmoid, hyperbolic tangent, Rectified Linear Unit (ReLU), LeakyReLU, and Gaussian Error Linear Unit (GELU) [52, 51]. ReLU-type functions are widely used because they are computationally efficient and reduce some of the training difficulties associated with saturating activation functions.

LeakyReLU modifies the standard ReLU by allowing a small non-zero slope for negative inputs, which reduces the risk that neurons become inactive during training [53]. GELU is a smoother activation function that weights inputs according to their value and has been reported to perform well in several deep-learning tasks [54].

In vibration-based defect detection, the input to a neural network typically consists of signal-derived features, such as spectral amplitudes, frequency-domain descriptors, statistical features, or time-frequency representations. The advantage of neural networks is that they can learn complex decision boundaries that may be difficult to define manually. However, this flexibility also introduces risks. If the model is too simple, it may fail to capture relevant damage-sensitive information. If it is too complex relative to the available dataset, it may overfit the training data and generalize poorly to new measurements [51, 52]. This trade-off between model capacity and generalization is a central issue in AI-based NDT and SHM.

Besides the choice of activation function, neural-network training is strongly influenced by regularization and normalization strategies. Overfitting occurs when a model learns patterns that are specific to the training data instead of learning relationships that generalize to new measurements [52, 51]. Dropout is a commonly used regularization technique in which a fraction of neurons is randomly disabled during training. This reduces co-adaptation between neurons and can improve generalization [55]. Batch normalization is another commonly used technique that normalizes intermediate activations within the network, which can improve training stability and allow faster convergence [56]. These techniques are especially relevant for NDT datasets, where the number of available specimens is often limited and the risk of overfitting is therefore significant.

2.3.1 Autoencoders

Autoencoders are neural-network models that learn to reconstruct their input through a compressed internal representation. They consist of an encoder, which maps the input to a lower-dimensional latent space, and a decoder, which reconstructs the original input from this latent representation. This structure forces the model to retain the dominant patterns in the data while filtering out less relevant variation [57, 52].

The latent representation is the compressed internal description learned by the encoder. It is not simply a manually selected set of features, but a lower-dimensional representation that is optimized during training so that the decoder can reconstruct the input as accurately as possible. In this sense, the bottleneck forces the network to retain the most important structure in the input data while discarding less relevant variation [57, 52]. For vibration-based inspection, this is useful because a high-dimensional spectrum can be represented by a smaller set of learned latent variables that summarize the dominant spectral behavior of the reference condition.

When autoencoders are used for anomaly detection, the reconstruction error is interpreted as a measure of how well a measurement agrees with the learned reference behavior. If the model is trained mainly on healthy data, measurements from the same condition are expected to be reconstructed with relatively low error, whereas measurements that differ from the learned healthy pattern may produce higher error [58, 57]. However, a high reconstruction error is not automatically equivalent to structural damage. It can also be caused by noise, excitation variability, environmental effects, or changes in measurement conditions. Therefore, autoencoder-based anomaly detection should be interpreted together with physical insight and spatial consistency of the measured response [58, 52].

In NDT and SHM, autoencoders are commonly used for anomaly detection. When an autoencoder is trained on data representing a healthy or reference condition, it learns the expected behavior of the undamaged system. Measurements that differ from this learned reference response may produce a larger reconstruction error and can therefore be flagged as potentially anomalous [58, 57]. This is useful in structural inspection because healthy data are often easier to obtain than complete labeled datasets covering all possible defect types, depths, sizes and severity.

The main strength of an autoencoder-based approach is that it can detect deviations from a reference condition without requiring labeled examples of every possible damage scenario. This makes it suitable for screening applications and early-stage proof-of-concept studies. However, the method does not automatically identify the physical type of defect. A high reconstruction error may be caused by structural damage, but also by measurement noise, environmental variability, changes in excitation, or differences in surface condition. In contrast, if a defect produces only a subtle change in the measured response, the reconstruction error may remain close to that of healthy measurements. Therefore, autoencoder-based anomaly detection should be interpreted together with signal quality and physical understanding of the inspected structure [58, 52].

2.3.2 MLP Classifier

A Multilayer Perceptron (MLP) is a feed-forward neural network used for supervised learning. In classification problems, an MLP maps an input feature vector to a set of class probabilities, after which the class with the highest probability is selected as the model prediction. Unlike anomaly detection, supervised classification requires labeled examples for each class that the model is expected to recognize [51, 52].

In an MLP classifier, the final layer typically converts the learned internal representation into class scores or class probabilities. During supervised training, the predicted output is compared with the known label of each training example, and the network parameters are updated to reduce the classification loss [51, 52]. For vibration-based defect detection, this means that the classifier attempts to learn which spectral patterns are associated with each structural condition. Compared with an autoencoder, an MLP can provide more specific diagnostic information because it can distinguish between predefined classes, but this requires representative labeled data for every class included in the classification task.

For vibration-based concrete inspection, an MLP can be used to distinguish between different structural conditions if representative labeled data are available. This makes supervised classification more diagnostically specific than anomaly detection, because it can potentially separate healthy measurements from different defect types. However, this advantage comes with a stronger dependence on the quality and diversity of the training dataset. If the training data only represent a limited number of specimens or measurement configurations, the classifier may learn patterns that are specific to the laboratory setup rather than general defect characteristics [3, 4, 6].

This limitation is especially important for concrete structures, where variability in material composition, aggregate distribution, reinforcement, surface roughness, moisture, and boundary conditions can strongly influence measured vibration responses. High classification accuracy under controlled laboratory conditions should therefore be interpreted as evidence of feasibility rather than proof of field-ready generalization. For practical application, supervised classifiers must be validated on larger and more diverse datasets that include realistic defect scenarios and measurement conditions [3, 4].

2.4 Research gap and positioning of the present work

The literature shows that NDT methods for concrete inspection are moving from qualitative and operator-dependent assessment toward more automated, quantitative, and data-driven approaches. LDV is a promising sensing technique in this context because it enables non-contact vibration measurements without attaching sensors to the concrete surface. Previous LDV-based studies have demonstrated that defects can influence the measured vibration response and that AI can support the interpretation of these signals [5, 22]. However, the literature is still fragmented across different diagnostic objectives, including binary damage detection, defect localization, defect imaging, and multi-class defect classification.

A key open question is therefore what level of diagnostic resolution is realistic for LDV-based AI inspection of concrete. Detecting whether a measurement deviates from healthy behavior is a different task from classifying the physical defect type. Autoencoder-based anomaly detection is attractive when labeled defect data are limited, but its output is mainly a damage-sensitive indicator.

Supervised classification methods, such as MLP classifiers, can provide more detailed diagnostic information, but they require representative labeled data and careful validation. The choice between these approaches is therefore not only algorithmic, but also depends on data availability, physical interpretability, and the intended engineering use of the inspection result [3, 6, 4].

The present work is positioned within this gap by combining LDV-based vibration measurements with AI-based signal interpretation for internal concrete defect detection. Rather than treating AI as a standalone solution, the study evaluates whether vibration spectra obtained from concrete specimens contain sufficient defect-sensitive information to support automated interpretation. In this way, the work connects non-contact vibration-based NDT with data-driven SHM, while also addressing the practical distinction between anomaly detection and supervised defect-type classification.

3 Methodology

3.1 Materials

To evaluate the proposed LDV-based defect detection methodology, three concrete samples were manufactured with identical nominal dimensions of $60 \times 60 \times 10$ cm. All three concrete specimens were produced using the same concrete mixture to ensure that differences in the measured vibration response could primarily be attributed to the internal defect configuration rather than variations in material composition. For each sample, the mixture consisted of cement, water, sand, and two aggregates. The water-to-cement ratio was 0.50, based on 6.00 kg of water and 12.00 kg of cement. The total batch mass for one sample was approximately 95.38 kg. The composition of the mixture used for each concrete block is summarized in Table 1. The samples were designed to represent three structural conditions: a healthy reference sample, a sample containing artificial void-type defects, and a sample containing artificial delamination-type defects.

Constituent	Mass [kg]
Cement	12.00
Water	6.00
Sand	21.85
Aggregate K 216.3	27.46
Aggregate K 6.314	28.06
Total	~95.38

Table 1: Concrete mixture composition used for each specimen.

The first specimen was cast as a defect-free reference specimen and is referred to as the healthy specimen. This block was intended to represent intact concrete and to provide a baseline for comparison with the defective specimens.

The second specimen was fabricated with artificial void-type defects. These were introduced by embedding hollow plastic capsules and bubble wrap within the concrete prior to casting, thus creating localized internal discontinuities and air-filled regions intended to approximate void-like defects.

The void specimen contained eight artificially introduced closed void-type inclusions. These inclusions were intentionally designed with different characteristic dimensions in order to represent a heterogeneous internal void configuration rather than a single idealized defect geometry. Based on precast measurements, two main inclusion families were used: five capsule-type voids and three larger elongated voids.

The capsule-type voids were based on sealed plastic capsule shells and had nominal dimensions of approximately 40 mm in length and 30 mm in diameter. Therefore, these inclusions represented compact, near cylindrical voids. In contrast, the elongated voids were constructed as larger closed inclusions using bubble wrap and tape.

Their nominal overall length was approximately 110-120 mm with a width of 45-50 mm. As a result, these defects represented more extended void-like discontinuities with a larger enclosed air volume than the capsules.

The larger voids were not perfectly cylindrical since they were manually assembled using bubble wrap and tape. Small local deviations of a few millimeters may be present due to manual preparation and wrapping. The dimensions of these voids can be found in Table 2.

To preserve the intended positions during casting, the void inclusions were fixed within a support frame before concrete placement. This ensured that the artificial voids remained spatially distributed over the sample area and that their locations were known a priori for later comparison with the measured LDV response and anomaly maps. Figure 1a shows the frame projected onto the block.

Void family	Number	Geometry	Nominal length [mm]	Nominal transverse size [mm]
V1	5	Capsule-type closed inclusion	~ 40	~ 30
V2	3	Elongated closed inclusion	$\sim 100\text{--}110$	$\sim 40\text{--}45$

Table 2: Nominal dimensions of the artificial void inclusions used in the void specimen.

The third specimen was fabricated with artificial delamination-type defects. In this sample, polystyrene plates were embedded during casting to create planar internal interfaces within the concrete volume, mimicking delamination-like defects (Figure 1b). The dimensions of the polystyrene plates are shown in Table 3.

Insert type	Number	Nominal in-plane dimensions [mm]	Nominal thickness [mm]
Circular plate	1	$\varnothing \sim 100$	$\sim 12 - 15$
Square plate	1	$\sim 95 - 100 \times 95 - 100$	$\sim 12 - 15$
Medium rectangular plate	1	$\sim 150 - 160 \times 95 - 100$	$\sim 12 - 15$
Elongated rectangular plate	1	$\sim 280 - 290 \times 85 - 95$	$\sim 12 - 15$

Table 3: Nominal dimensions of the artificial delamination inserts used in the delamination specimen.

The delamination specimen contained four artificially introduced delamination-type inclusions made from polystyrene plates. In contrast to closed void inclusions, these defects were designed as thin planar discontinuities with relatively large in-plane dimensions and limited thickness, thereby approximating the geometric dimension of delaminations in concrete.

Based on pre-casting measurements, four delamination inserts with different geometries were used: one circular plate, one square plate, one medium rectangular plate, and one elongated rectangular plate. The circular insert had a diameter of approximately 100 mm. The square insert had dimensions of approximately 95-100 mm x 95-100 mm. The medium rectangular insert measured approximately 150-160 mm x 95-100 mm, whereas the elongated rectangular insert measured approximately 280-290 mm x 85-95 mm. All inserts had a nominal thickness of approximately 12-15 mm.

The delamination inserts were cut manually from polystyrene and therefore did not exhibit perfectly sharp edges. Minor local deviations are expected due to manual preparation and trimming.

To preserve their intended position during casting, the polystyrene inserts were attached to a reinforcement frame. This arrangement prevented significant displacement during concrete placement and ensured that approximate locations and geometries of the artificial delaminations were known a priori.

By keeping the external dimensions identical while varying the internal defect configuration, the experimental sample set provided a controlled basis for evaluating the sensitivity of the proposed LDV-based methodology to different defect scenarios.



(a) Void block. *This figure was generated using AI.*



(b) Delamination block. *This figure was generated using AI.*

Figure 1: Visualization of an approximate estimation where the defects are located inside the blocks. The steel frames kept the defects at approximately the same location, which could be useful later.

3.1.1 Obtaining the Young's modulus and the Poisson ratio

Two methods were used to obtain the Young's modulus and the Poisson ratio. Firstly, when comparing the results of our bachelor thesis from last year with a 3D model of the same block in Siemens NX [14], the Young's modulus and Poisson ratio were iterated to have the same mode shape and eigen frequency for the first eigenmode. This process took about 2 hours, but eventually a difference in frequency of 2 Hz was realized. This gave a Young's modulus equal to 47.1 GPa and a Poisson ratio equal to 0.25.

The second method was a real-life measurement of the Young's modulus with an ultrasonic tester from the SuPAR research group. This tester measures the time it takes for an ultrasonic wave to travel from one side of a concrete block to the other side. When the traveled distance is known (for example; thickness, height, or width of the block), a longitudinal speed can be calculated, which can be implemented in a formula to calculate the Young's modulus [59]. Between six and nine tests were performed on each side of the slab and on the blocks of the bachelor thesis to calculate an average.

$$E = V^2 \rho \frac{(1 + \nu)(1 - 2\nu)}{(1 - \nu)} \quad (1)$$

As seen in Equation 1, the calculation of the Young's modulus E is dependent on the Poisson ratio ν , which is unknown, the longitudinal speed of the ultrasonic waves V , and the density ρ . Normally a Poisson ratio of 0.2 is used in common concrete testing, which gave a modulus of 52.26 GPa for the blocks from [14] and 47.48 GPa for the slabs. When using a Poisson ratio of 0.25, as shown in Table 4, the results become 48.38 GPa for the blocks [14] and 43.97 GPa for the slabs. This gives a result which is much closer to the iteration when looking at the blocks from the bachelor thesis.

Because the slabs are made with the same concrete composition as the blocks [14], it can be assumed that the slabs will have a modulus of approximately 47.1 GPa and a Poisson ratio of 0.25 and these parameters will be used in Siemens NX to simulate concrete.

Variables	Healthy block[14]	Healthy slab (side)	Healthy slab (front)
Distance [m]	0.15	0.60	0.10
Time [10^{-6} s]	30.30	123.82	21.57
Velocity V [m/s]	4918.22	4829.78	4591.52
Poisson Ratio ν	0.25	0.25	0.25
Young's Modulus E [GPa]	48.38	46.65	42.17

Table 4: The values in this table are all averages taken from six to nine measurements. These averages were implemented into Equation 1, after which the Young's modulus is calculated (shown at the bottom). The Poisson ratio was iterated by using the simulations and comparing these to the measurements of the bachelor thesis[14]. A weighted average was calculated for the healthy slab to input into Siemens NX.

3.2 Siemens NX: Solver parameters and settings

To simulate the vibration response of the healthy concrete slab and damaged concrete slabs, it is important to use the correct solver and settings to obtain accurate results. Using a modal analysis, at different mode shapes under specified constraints, the corresponding resonant frequencies will be returned.

As seen in Figure 2, the following settings will be used to create the solution:

Solver: Simcenter Nastran

Analysis Type: Structural

Solution Type: SOL103: Real Eigenvalues




Figure 2: Solution parameters

As the eigenvalue method in this paper is chosen to work with Lanczos, the number of desired modes is set to 15, with a frequency range of 500 Hz to 4 kHz. Using SOL 103, the real natural frequencies and modeshapes are computed from an undamped free vibration of a structure. So, no loads or damping will be placed on the structure. This is a fast and stable method to get the eigenvalues. To link the meshes of the defects with the mesh of the concrete, mesh mating conditions are applied to the meshes. A free coincident constrain is used to ensure the correct working of the solver.

3.3 Experimental setup

The experimental setup was conducted under controlled laboratory conditions using the three concrete samples described in Section 3.1. During testing, each specimen was positioned vertically in a rigid support frame to ensure stable and repeatable boundary conditions throughout the measurement campaign. The vibration response of the concrete specimens was measured using a Polytec VibroFlex laser vibrometer system, consisting of a VibroFlex Connect unit acting as the configurable hub and signal interface, and a VibroFlex Neo non-contact vibration sensor used as the optical measurement head.

Laser Doppler vibrometry is an optical non-contact measurement technique used to determine surface vibration velocity through detection of the Doppler shift in reflected laser light. Because the measured frequency shift is proportional to the vibration velocity, LDV enables highly sensitive vibration measurements without introducing mass-loading effects. Scanning LDV systems further allow automated measurements over multiple spatial locations, enabling detailed modal analysis and high-resolution vibration characterization [14].

During the measurements, the VibroFlex Neo sensor was positioned parallel to the concrete specimen. A galvometric mirror system was used to redirect the laser beam toward the specimen surface, enabling spatially distributed vibration measurements across the surface of each sample. By using the same LDV configuration for all three specimens, differences in the recorded vibration response could be compared under identical measurement conditions.

A measurement grid was defined on the surface of each specimen to allow spatially distributed data acquisition, supporting the identification of regions whose behavior deviated from that of intact concrete. An overview of the complete automated measurement setup is shown in Figure 3. The approximate top-view layout is shown in Figure 4, illustrating the relative position of the concrete specimen, LDV, scanning mirror, projectile launcher, air compressor, control electronics and data-acquisition system.




Figure 3: Overview of the automated LDV measurement setup used for grid-based inspection of the concrete specimens. The system combines an air compressor, automated projectile launcher, control electronics, concrete specimen, Laser Doppler Vibrometer (LDV), and data-acquisition computer. The launcher moves over the predefined measurement grid and applies repeatable projectile impacts, while the LDV measures the resulting surface vibration response without contact. The acquired signals are stored and processed to generate spatial vibration-response maps for subsequent AI-based damage assessment. *This figure was generated using AI.*




Figure 4: Approximate top-view layout of the automated LDV inspection setup. The concrete specimen is mounted vertically, and the front surface contains a 40 cm $\times$ 40 cm scan area discretized into a 20 $\times$ 20 measurement grid. The LDV is positioned approximately parallel to the specimen and directs its laser beam toward a galvanometric scanning mirror, which redirects the beam approximately normal to the specimen surface. The pneumatic steel-ball launcher is positioned to the left of the LDV path and aimed at the specimen under an angle of approximately 20°. The LDV-specimen distance is approximately 1.6 m. The layout is schematic and not drawn to scale. *This figure was generated using AI.*

3.3.1 Final setup and working principle

The final setup was designed to perform automated point-by-point inspection of concrete specimens by combining projectile-based excitation, non-contact vibration measurement, and AI-based data interpretation. For each grid point, the pneumatic launcher is positioned and fires a steel sphere to generate a short and repeatable mechanical impact. The resulting surface vibration is measured without contact by the LDV, after which the time-domain response is stored and transformed into a frequency-domain spectrum used as input for AI-based analysis. After processing, the system proceeds to the next grid point and repeats the sequence until the entire grid has been completed. By assigning the processed response of each measurement point back to its original grid location, the setup enables the construction of spatial result maps over the concrete surface. The general working principle is illustrated in Figure 5.




Figure 5: Working principle of the automated LDV inspection workflow. A scan grid is first defined on the concrete surface. For each grid point, the pneumatic launcher is positioned and fires a steel ball to generate a localized mechanical impact. The resulting surface vibration is measured without contact by the LDV, after which the time-domain signal is stored and transformed into a frequency-domain spectrum. These spectra are subsequently used for AI-based anomaly detection or defect classification, and the point-wise results are projected back onto the scan grid to obtain a spatial result map. *This figure was generated using AI.*

3.3.2 LDV scanning measurement setup

A Laser Doppler Vibrometer (LDV) was used to measure the vibrations of the concrete block. The model of the LDV is the Polytec VibroFlex Xtra, which was connected to the Polytec VibroFlex Connect for the processing of the data acquisition. The settings of the VibroFlex Connect are shown in Table 5.

Bandwidth	Frequency resolution	Max Velocity	Sensitivity
10 kHz	2.5 Hz	10 mm/s	5 mm/s/V

Table 5: The settings used for the VibroFlex Connect

The scan area covered approximately $40 \text{ cm} \times 40 \text{ cm}$ on the front surface of each concrete specimen, discretized into a 20×20 measurement grid resulting in 400 measurement points per specimen. At each grid point, two repeated samples were recorded to assess repeatability and reduce the influence of incidental measurement variability. The nominal spacing between adjacent measurement points was approximately 2 cm. The acquired time-domain signals were subsequently transformed into frequency-domain representations, since defect-sensitive information is often more clearly reflected in resonance shifts and local amplitude variations than in the raw time histories. Prior to further analysis, the signals were preprocessed to obtain comparable spectral representations at all measurement points, including the removal of irrelevant frequency content and normalization of the spectra to reduce the influence of absolute excitation strength.

In order to monitor and control the grid-based measurement procedure, a custom Python interface was developed, which is shown in Figure 6. This interface displayed the predefined scan grid in real time and indicated the current measurement position during operation, with each grid point visually updated after completion. The scan sequence followed a snake pattern across the grid, reducing unnecessary repositioning of the excitation system. The interface also allowed the measurement campaign to be interrupted and resumed: if the experiment was stopped, the last completed grid coordinates could be entered manually when the system was restarted, after which the scan continued along the predefined path. Since two repeated measurements were recorded at each point, a completed point was only accepted after both acquisitions had been performed.




Figure 6: Python-based measurement-control interface used during the automated LDV scanning. The interface displays the current scan grid, the completed measurement points, and a live FFT of the acquired vibration signal. The scan can be paused and resumed, and the last completed coordinates can be entered manually to continue the measurement sequence from a previous stopping point.

In addition to grid tracking, the interface displayed a live FFT of the measured vibration response for each acquisition, providing an immediate quality check during the measurement campaign. If a measurement showed excessive noise, dropouts, or an unrealistic frequency response, the acquisition could be repeated before continuing with the rest of the grid. All the parameters used for the scanning grid can be found in Table 6.

Parameter	Value
Scan area	$40 \times 40 \text{ cm}^2$
Measurement grid	20×20 points
Total measurement points per specimen	400
Repeated samples per point	2
Nominal point spacing	$\sim 2 \text{ cm}$
LDV-specimen distance	$\sim 160 \text{ cm}$
LDV measurement direction	Approximately normal to specimen surface
Projectile launcher angle	$\sim 20^\circ$

Table 6: Overview of the LDV measurement parameters.

3.3.3 Automated excitation system

To generate repeatable mechanical excitation during LDV measurements, a custom pneumatic impact excitation device was used. The device was originally developed as part of a bachelor thesis project and was further integrated into the present experimental setup for automated scanning measurements. In contrast to manual impact excitation, the proposed system enables controlled and repeatable triggering at predefined measurement locations, thereby reducing operator dependence and improving measurement efficiency.

The excitation device consists of an automated pneumatic launcher mounted on a two-axis positioning system, allowing movement in both horizontal and vertical directions. The launcher is electronically connected to a Raspberry Pi, which serves as the communication interface between the excitation system and the LDV control environment, allowing the excitation sequence to be synchronized with the generated scan grid in Python. For each measurement point, the launcher is moved to the corresponding position and triggered to apply mechanical impact locally near the LDV measurement location, after which it automatically advances to the next position in the scan sequence.

The launcher operates pneumatically at an approximate working pressure of 2 bar and uses steel spheres with a diameter of 10 mm as projectiles, stored in an integrated magazine with a capacity of approximately 20 spheres. A solenoid-based feeding and firing mechanism allows the projectiles to be released sequentially, enabling automated repeated excitation without manual reloading after each individual measurement point. The impacts of the projectile on the concrete provided sufficient excitation energy to induce a measurable local vibration response in the concrete specimens. The main components of the automated projectile excitation system are shown in Figure 7.

3.3.4 Synchronization between LDV scanning and excitation

A key feature of the experimental setup is the synchronization between the LDV scan routine and the automated excitation device. The measurement grid was defined in Python and communicated to the control system, after which the excitation device followed the same sequence of spatial positions. For each predefined point, the launcher was aligned with the target location and triggered in coordination with the LDV acquisition routine. This synchronized workflow ensured that excitation and measurement were spatially coupled throughout the experiment, reducing manual intervention and improving the repeatability of the experimental procedure. All the programs to make the measurement grid, run the connection to the LDV and the Raspberry Pi, perform data acquisition, process the data, and training the AI were all done in PyCharm 2025.3.2.1. with Python 3.11 .




Figure 7: Component overview of the automated projectile excitation system. The launcher is mounted on a ballscrew-driven linear positioning system actuated by NEMA 23 stepper motors, allowing controlled movement over the measurement grid. Limit switches are used for reference positioning and travel protection, while the pneumatic splitter distributes compressed air to the projectile-launching mechanism.

3.3.5 Safety considerations

Because the excitation system uses pneumatically launched 10 mm steel spheres and operates in a fully automated manner, safety constituted a major aspect of the design of the project. Without appropriate containment, the projectiles could rebound or travel beyond the intended impacted zone, creating a significant hazard for personnel and equipment in the laboratory environment. For this reason, the experimental setup was designed and operated with safety as a primary constraint, with particular attention paid to shielding of the impact zone, confinement of the projectile trajectory, and controlled access to the test area during operation. Since the launcher can automatically move across the scan grid and fire sequentially without direct manual intervention, the interaction between LDV control, launcher positioning, and firing logic had to be carefully managed to prevent unintended triggering or unsafe motion. Consequently, the safe use of the system depended not only on mechanical protection, but also on reliable synchronization and control, and safety was therefore treated as an integral part of the experimental methodology.

3.4 Machine learning approaches

3.4.1 Working principle of AI-based models

The AI-based models used in this work are based on artificial neural networks. A neural network is a data-driven model that learns the relationships between input data and output predictions by adjusting internal parameters during training. In this study, the input consists of features extracted from LDV vibration signals, mainly frequency-domain information obtained after pre-processing. The purpose of the network is to learn patterns in these spectra that are related to the structural condition of the concrete.

A neural network consists of layers of interconnected neurons. The input layer receives the feature vector of one measurement point, while the output layer provides the model result, such as a reconstructed signal, an anomaly, or a class prediction. Between these layers, hidden layers transform the input data into increasingly useful internal representations.

Each neuron combines the values from the previous layer using trainable weights and a bias term, after which an activation function is applied. These activation functions introduce nonlinearity, allowing the network to model complex relationships in the vibration response.

During training, the model compares its output with the desired output using a loss function. The weights and biases are then updated to minimize this loss. In this way, the network gradually learns which spectral features are important to represent healthy concrete or distinguish between different defect conditions. However, the complexity of the model must be chosen carefully. A model with too few neurons or layers may not capture the relevant patterns, while an overly complex model may overfit the training data and generalize poorly to new measurements.

In this paper, two neural-network approaches are investigated. The autoencoder is used for anomaly detection and consists of an encoder and a decoder. The encoder compresses the input into a lower-dimensional latent representation, also called the bottleneck, while the decoder reconstructs the original input from this compressed representation. When trained data from the healthy slab, the autoencoder learns the typical spectral behavior of intact concrete. Measurements that cannot be accurately reconstructed are interpreted as deviations from the healthy baseline.

The second approach is a Multilayer Perceptron (MLP) Classifier, which is used for supervised classification. In this case, the model is trained with labeled examples and learns to assign each input spectrum to a predefined class, such as healthy concrete, void-type defect or delamination-type defect.

In the following sections, several signal-processing and machine-learning terms are used. A frequency bin refers to one discrete frequency component obtained after applying the Fast Fourier Transform (FFT). Since the FFT converts a time-domain vibration signal into a frequency-domain spectrum, the continuous frequency range is represented by a finite number of frequency bins. Each bin therefore contains information about the vibration amplitude at a specific frequency or within a very small frequency interval.

A feature is a numerical quantity extracted from the measured signal and used as input for the AI model. When multiple features are combined, they form a feature vector. In this work, such features include logarithmic spectral magnitudes, spectral flux values, and band energy ratios. These features summarize the vibration response of one measurement point in a form that can be processed by the neural network.

Standardization means that each feature is rescaled using its mean and standard deviation. This is done to make the different input features comparable and to reduce the influence of variations in impact force or signal amplitude. The reconstruction error is the difference between the original input and the reconstructed output of the autoencoder. A high reconstruction error indicates that the measured signal is difficult for the model to reproduce and may therefore deviate from the learned healthy response.

For the neural-network models, additional regularization and normalization techniques are used. Dropout randomly disables a fraction of neurons during training, which reduces the risk that the model becomes too dependent on specific neurons. Batch normalization rescales intermediate activations inside the network, improving training stability. Finally, the Fisher score is used as a feature-selection criterion. It evaluates how strongly a frequency bin differs between classes, and therefore helps retain the frequency components that are most relevant for classification.

3.4.2 1D-Autoencoder

The first machine-learning approach investigated in this work was an autoencoder-based anomaly detection model. This model was selected because the available dataset is strongly unbalanced from a practical inspection perspective: measurements from an intact reference condition are easier to define than measurements covering all possible defect types, sizes, depths and geometries. Instead of explicitly learning each defect class, the based approach learns a compact representation of the vibration response of healthy concrete. Measurements that cannot be represented or accurately reconstructed from this learned healthy baseline are then interpreted as potentially anomalous.

In this framework, the frequency-domain data obtained from the healthy sample was used for model training. Firstly, the raw signal was cut into heavily overlapping windows, and a Hanning window was applied to smooth the edges. The Hanning window is a widely used mathematical function to smooth the data and reduce spectral leakage when an FFT is calculated[60, 61, 62]. Then, an FFT was done on each window. Only the magnitude was kept, and only within the 200 - 10,000 Hz range. This range was established after examining the simulations and measurements. The simulations indicate that the first eigenfrequency appears at approximately 730 Hz. And after examining the measurements, it was clear that there were some peaks around 30 Hz, but these originate from the block wobbling rather than the vibrational response of the block itself. This is why it was chosen to only look from 200 Hz to higher frequencies to ensure that the wobbling would not be implemented in the training data.

When the FFT is calculated, three features are extracted from the FFT: logarithmic magnitude, spectral flux, and band energy ratios. The logarithmic magnitude captures the spectral shape, the spectral flux is an indicator on how steep the peaks are, and the band energy ratios shows how the energy of the vibration is distributed along some predetermined frequency bands. The logarithmic magnitude and the spectral flux are both calculated per frequency bin. This means that the output is not a single number, but a vector. The band energy ratios is divided in five frequency bands. Five bands were chosen because the more bands, the higher the risk of including noise or duplication in the training data. For example, if a high number of frequency bands is chosen, then that data will provide similar information as the log magnitude feature, therefore duplicating the data and rendering the extra frequency bands useless. The reasoning behind five bands comes from the simulation results. The first modes cluster around 700 - 1000 Hz, the next modes around 1200 - 1700 Hz and the last modes from 2800 - 3000 Hz. The first three frequency bands were chosen around those ranges, with the last two ranges being from 4500 Hz to 10,000 Hz, to ensure that the entire frequency range is captured.

These features are then concentrated in one feature vector. Lastly, the values are standardized using the mean and standard deviation of each feature. This is done to counteract any variations in the input force of the projectile. The entire pre-processing workflow is visualized in Figure 8.

After the pre-processing, the encoder maps the high-dimensional spectral input to a lower-dimensional latent representation, also known as the bottleneck, forcing the model to retain only the most relevant characteristics of the healthy vibration response. A decoder then reconstructs the original spectrum from this latent representation. After training, the reconstruction error between the input spectrum and the reconstructed spectrum is used as an anomaly score. The data is split into training data and validation data. 15% of the data is used for validation, which is performed each epoch. The other 85% is used for training the model. In this way, the model is validated on data it has not seen yet, eliminating the risk of the model memorizing measurements and not actually learning anything about the defect signatures.

The reconstruction error was computed for each measurement point individually. Low reconstruction errors indicate that the measured spectrum is consistent with the learned healthy response, whereas higher reconstruction errors indicate that the local vibration response deviates from the healthy baseline. In this way, the based model provides a point-wise damage-sensitive indicator without requiring the model to be trained on every possible defect configuration. The anomaly scores can be subsequently projected back onto the original measurement grid to obtain a spatial anomaly map of the concrete specimen.

Pre-processing pipeline of the training data




Figure 8: Pre-processing pipeline of the LDV vibration measurements. The raw time-domain response is segmented into overlapping windows, windowed using a Hanning window, transformed to the frequency domain using the FFT, and converted into standardized frequency-domain features for AI-based anomaly detection and defect classification. *This figure was generated using AI.*

This approach is particularly suitable for the present proof-of-concept study because only one healthy reference specimen and two artificially damaged specimens were available. A fully supervised multi-class classifier would require a larger and more diverse dataset, than the autoencoder, to distinguish reliably between different defect types and severity. The autoencoder-based anomaly detection strategy therefore represents a more conservative first step: it focuses on detecting deviations from healthy behavior rather than directly assigning a specific physical defect label. The autoencoder should detect damages, but will not be able to classify them. The encoder-decoder anomaly detection workflow is illustrated in Figure 9.

The model used in this paper consisted of a total of 10 layers. Figure 10 shows the structure of the autoencoder. This number of layers ensured that less information was lost and that enough features were captured. The first layer consists of the same number of neurons as there are input features. This is done to let the model transform the features so that it is easier to compress in the upcoming layers. LeakyRELU was used to ensure that when negative input is given, the neuron does not go to 0 and stays 0 for the whole training process. Instead, LeakyRELU will change the value to 0.01 times the negative value. The first two and last two layers use the dropout function as well. This function converts 10% of the neurons to zero for each training cycle. This makes the neurons less dependent on each other and makes the overall model better.




Figure 9: Schematic overview of the encoder-decoder anomaly detection workflow. The model is trained only on the frequency spectrum from the healthy reference specimen, where the encoder maps the frequency-domain input to a lower-dimensional latent representation and the decoder reconstructs the original spectrum. During inference, spectra with low reconstruction error are interpreted as consistent with the healthy baseline, whereas spectra with high reconstruction error are flagged as potentially anomalous. The point-wise anomaly scores are projected back onto the measurement grid to obtain a spatial anomaly map. *This figure was generated using AI.*




Figure 10: The autoencoder model consists of 10 layers in total. The number of neurons for each layer is displayed on the top and the input functions on the bottom.

3.4.3 MultiLayer Perceptron Classifier

As a second machine-learning approach, an Multilayer Perceptron (MLP) was investigated for supervised classification of the vibration measurements. In contrast to the autoencoder, which is trained only on healthy data and detects deviations from a learned reference condition, the MLP uses labeled data to assign each measurement point directly to a predefined class. In the present study, these classes correspond to the investigated structural conditions: healthy concrete, void-type defects, and delamination-type defects.

The input to the MLP is the pre-processed frequency-domain of a single measurement point. First, a Hanning window is applied on the measurements, and then an FFT is done, just like the pre-processing of the autoencoder. The difference between the autoencoder and the MLP is that only one feature is extracted from the FFTs, not three. For the MLP, only the logarithmic magnitude is kept for each frequency bin.

To ensure that only the relevant features remain, the bins will be selected according to their Fisher score. The Fisher score compares the mean value of a feature throughout the dataset to the value of that feature in the frequency bin selected. If the Fisher score is low, the features match quite well and that feature will probably not give any information on whether there is a defect or not. If the Fisher score is high, then that feature is not present in all the measurements and therefore could be an indicator of defects in the concrete sample. The Fisher score is calculated for three comparisons: void versus no defect, delamination versus no defect, and void versus delamination. All the bins with at least a Fisher score of 5% of the maximum score of all the bins are kept. If this results in more than 120 bins, then the 120 top scoring bins are kept. This entire process will ensure that only the frequency bins with useful information get inputted into the MLP.

The last step of the pre-processing is to standardize the feature values to counteract the varying input force, just like it was done for the autoencoder. Now that pre-processing is done, the MLP then maps this input vector through one or more hidden layers with nonlinear activation functions to an output layer that provides a class probability for each considered condition. The predicted class is selected as the class with the highest output probability. The input data are split up just as it's done for the autoencoder. 20% of the data is used for validation, and the other 80% for training.

The multilayer perceptron consists of four layers. The first layer contains 120 neurons because 120 bins are kept after Fisher scoring. The last layer has 3 neurons, one for each class: void, delamination, and healthy. The layout of the model can be seen in Figure 11.




Figure 11: The multilayer perceptron model consist of four layers in total. The number of neurons for each layer is displayed on the top and the input functions on the bottom.

Instead of using the LeakyRELU function like the autoencoder, the GELU was used. GELU is the most popular choice when it comes to classification models, because it is more aggressive in shaping activations, and therefore it is better in making decisions, which is important for classification. However, before the GELU, the activations were normalized by applying batch normalization. This stabilizes and speeds up training.

The MLP is useful in the present work because it addresses a different diagnostic objective than the autoencoder. Whereas the autoencoder mainly performs binary damage screening by identifying spectra that deviate from the healthy baseline, the MLP investigates whether the spectral response contains sufficient information to distinguish between different types of defect. This is relevant because voids and delaminations can influence the local dynamic response in different ways through changes in stiffness, damping, resonance behavior, and wave propagation.

After classification, the predicted class labels can be projected back onto the original measurement grid. This results in a spatial classification map that indicates the most likely structural condition at each measurement point. Such a representation is useful for evaluating whether the model predictions correspond to the known defect locations in the samples and for assessing the potential of vibration-based NDT for defect-type classification. The supervised MLP classification workflow is illustrated in Figure 12.




Figure 12: Schematic overview of the Multilayer Perceptron (MLP) workflow for supervised classification of vibration measurements. Each measurement point is represented by a pre-processed frequency-domain spectrum, which is used as the input feature vector for the MLP. During training, the network learns to map spectra to predefined class labels, such as healthy concrete, void-type defects, and delamination-type defects. During inference, the MLP outputs a probability for each class, and the class with the highest probability is assigned to the measurement point. The point-wise predictions are then projected back onto the measurement grid to obtain a spatial classification map. *This figure was generated using AI.*

4 Results

4.1 Numerical Simulations

In this section, the results of the different simulations performed in Siemens NX will be explained [23]. To create the slabs in Siemens NX, the measurements of the real-life slabs will be used, so that the results can be compared in this paper. A convergence indicator of 2% is used to determine if the result is within convergence when changing the element size of the model.

4.1.1 Healthy concrete slabs

For this simulation, a 3D Tetrahedral mesh, CTETRA(10), is used with an element size of 50 mm. Using a mesh collector, the slab is made of concrete with the following specifications listed in Table 7.

Table 7: Material properties used for the healthy concrete slab.

Property	Value
Young's modulus	47.1 GPa
Poisson's ratio	0.25
Density	2400 kg/m ³

Without constraints or placing a load on the concrete slab, a solution can be found as shown in Figure 13.




Figure 13: Mode 1: 728.50 Hz

After performing a convergence study, by refining the element size of the mesh, the results in Table 8 were found and provide an overview of the different modes with their accompanying natural frequency.

Mode	Natural frequency [Hz]
1	728
2	1076
3	1272
4	1738
5	1738
6	2865
7	2865
8	2905
9	2911
10	3051

Table 8: Solution of SOL103 simulation in Siemens NX on a healthy concrete slab.

4.1.2 Delamination Concrete Slab

To simulate this in Siemens NX, two different types of materials are used. The slab is made out of concrete with the specific parameters listed in *Materials*, and the defects in the slab are assigned to the material "Air". For both the concrete slab and the defects, a 3D Tetrahedral mesh, CTETRA(10), is used with an element size of 55 mm for the concrete and 24.1 mm for the air-filled defects. After simulation and conducting a convergence study, the results found are shown in Table 9.

Mode	Natural frequency [Hz]
1	679
2	1001
3	1195
4	1615
5	1618
6	2615
7	2621
8	2653
9	2708
10	2854

Table 9: Solution of SOL103 simulation in Siemens NX on a delamination-defected concrete slab.

As shown in Figure 14a, the first mode shape looks the same as the first mode shape of the healthy slab. However, in Figure 14b, it can be seen that the maximum displacement is lower than the healthy slab.




Figure 14: Mode 1 and Mode 7 of the concrete slab with delamination defects, simulated in Siemens NX.

4.1.3 Voids defected concrete slab

The process for simulating the slab with voids is the same as the one with delaminations. The only difference is shape and placement of the artificial defects. Like previous, a 3D Tetrahedral mesh, CTETRA(10), was used with an element size of 50 mm for the concrete and 20.7 mm for the defects. A convergence study was done and there were no changes bigger than 2%, which means that there is convergence. The results are shown in Table 10.

The welded frame of the voids is not exactly parallel to the front surface of the slab. This has been accounted for, and therefore the defects were placed on a plane of 3° from the front plane, just like the real slab. All defects are filled with air and will have the material "Air" selected in Siemens NX.

Mode	Natural frequency [Hz]
1	719
2	1061
3	1268
4	1723
5	1730
6	2709
7	2834
8	2850
9	2884
10	3003

Table 10: Solution of SOL103 simulation in Siemens NX on a voids-defected concrete slab.

As shown in Figure 15a, the first mode shape does not differ from the healthy slab and the delamination slab. However, looking at Figure 15b, it can be observed that there are some irregularities in the modeshape.




Figure 15: Mode 1 and Mode 15 of the concrete slab with voids as defects, simulated in Siemens NX.

4.2 Vibration Data

4.2.1 Post-Processing the Raw Vibration Data


After measuring, it is important to post-process the raw vibration data. It was very important to check the quality of the signals before they were used to train the AI models. A very noisy measurement could slow down or even disrupt model training. The filtering of the data was done manually by using a Python script, which represented the data in both the time and frequency domain. This way, it was very clear which measurements had a lot of noise. Two buttons were used to classify the measurement as good or bad. The noisy measurements were removed from the database automatically to ensure it not slipping into the training dataset. After the quality check, another script was used to identify if there were any missing measurement points after removing the noisy and low-quality measurements. If a few data points next to each other were missing, then the measurements of some of those points were re-done. This way, there was enough data to derive conclusions or to interpolate the missing points.




Figure 16: Manual checking of the raw vibration data was done with a Python script to make it as efficient as possible. The buttons on the bottom were used to classify them very easily. Bad measurements were deleted to ensure that they did not slip into the training data.

Averaging


After filtering, most of the measurements look like Figure 17a. But some of the 'good' measurements still contain some noise, as in Figure 17b. This noise can be reduced by taking the average of the measurements done at that point. Each point had two measurements taken. By averaging the two measurements, the result becomes slightly less noisy, as seen in Figure 17c. When the amount of averages is increased, for example 10 averages, then the noise decreases significantly, as shown in Figure 17d.




(a)



(b)



(c)



(d)

Figure 17: Most of the measurements have a similar quality like (a). (b) shows a measurement with more noise which can still be used. (c) was achieved by doing two averages on (b) and (d) by doing 10 averages. More averages gives better quality, but increases the measuring time a lot.

4.2.2 Comparing

Time Domain

When a point is measured, the output is a graph in the time domain with the measured velocity as the Y-axis, as shown in Figure 18a. When comparing Figure 18a & 18b, it is clear that 18a will have more noise than 18b. This can be deduced by looking at the amount of speckle noise, the large spikes when the vibration has already been dampened, also known as dropouts. The difference in noise can be seen when an FFT is applied to both time domain measurements, seen in Figure 18c & 18d.




Figure 18: (a) will contain more noise when an FFT is performed than (b). This can be expected by noticing the speckle noise, also known as dropouts. The difference in the amount of noise can be clearly seen when comparing (c) & (d)

Frequency Domain

After post-processing, an FFT was applied on the measurements in MATLAB R2024b to compare the data. This makes it easier to plot multiple measurements on top of each other and to make the graphs look nicer. A qualitative comparison between representative frequency-domain responses is shown in Figure 19. The healthy reference response can be interpreted as the dynamic baseline behavior of intact concrete. In contrast, spectra obtained near void-type or delamination-type regions deviate from this baseline, with differences visible in peak position, peak amplitude, bandwidth, and overall spectral shape.





Figure 19: An FFT comparison on the same grid point between the healthy, void-defected, and delamination-defected block. Every block has a different frequency response because of the effects due to the defects.


4.2.3 Visualizing the Eigenshapes

A Python script was written to visualize the eigenshapes of the measurement grids to get a clearer understanding of the vibration response of the specimens and to compare them with the simulations as well. As mentioned in *Post-Processing the Raw Vibration Data*, some points were missing because the quality of both the measurements was poor. That is why interpolation was used to ensure that a full grid could be projected, which is easier to interpret.


There is one problem with visualizing the eigenshapes: there was no phase reference, which means that it is uncertain to which direction the displacement is moving when the visualization is created. This can be seen by looking at the difference between Figure 13 and 20. Then there is one disadvantage to the measurement grid in regard to the visualization: the shapes show up over the entire block, but the grid only covers $40\text{ cm} \times 40\text{ cm}$ which means that some part of the shapes cannot be seen. This is not really a problem because it can still be deduced which eigenshape it is with the measurement grid.



(a) Healthy slab: Mode 1



(b) Delamination slab: Mode 1



(c) Void slab: Mode 1

Figure 20: The first modeshape from each block is shown. When a measurement point was missing, it gets interpolated to establish a full grid visualization.

4.3 Machine learning models

4.3.1 1D-Autoencoder

As seen in the *Methodology* section, the autoencoder should be able to replicate the healthy vibration signals and have more difficulty with reconstructing the damaged signal, which then would be detected by a higher reconstruction error. In theory, this would work very well to find damages. In practice, this was not the case.

The first training test on the autoencoder was performed by utilizing the time domain as input. The same pre-processing is done like explained in the *Methodology* section, with the windowing, the standardization, and then inputting the data into the model. While training, the value of the cost function was saved together with the weight applied to the neurons. This way, the model could be trained until a minima of the cost function was found, and then the weights would be restored to the values that gave the low cost function value. After training on all the healthy data, and optimizing some of the parameters through trial and error, the following output was produced (Figure 21).




Figure 21: The attempted reconstruction of one healthy and one damaged window, done by the autoencoder. The reconstruction of the damaged window is slightly better than the healthy window, which makes it a problem to detect damaged samples when using the cost function to detect them.

When examining Figure 21, it is clear that there is a problem: both healthy and damaged windows get reconstructed quite well. Actually, the damaged window gets reconstructed a little better. So, if the damaged signal has to be detected by using the loss value, it is more likely that healthy samples will be flagged instead of damaged ones. The RMS values of all the measurements are shown in Table 11.

Healthy On top of void	Average RMS [mm/s]	Std RMS [mm/s]	Min RMS [mm/s]	Max RMS [mm/s]
	0.605	0.106	0.138	1.457
	0.652	0.105	0.278	1.017

Table 11: The average RMS values for the healthy measurements and the measurements on top of the void. The distribution of both are very close to each other, which explains why the autoencoder is able to reconstruct both so well.

After concluding that an input in the time domain is not working, a switch was made to the frequency domain. This gave some new hope because the shape of the frequency domain changes more drastically than the time domain, which was proven in [14], meaning that the difference between healthy and damaged windows may be more noticeable, which would hopefully make it more difficult to reconstruct damaged windows so that damages can be detected. Unfortunately, this was not the case. The damaged window was replicated almost as good as the healthy one, as seen in Figure 22. The autoencoder only flagged 0.3% of damaged windows as damaged, but it also flagged 0.3% of the healthy windows as damaged.

Waar de fuck komt dees vandaan?




Figure 22: Even when the autoencoder is trained on a frequency domain as input, the reconstruction is still equally as good for both healthy and damaged windows. This makes damage detection unfeasible.

4.3.2 MultiLayer Perceptron Classifier

The trained MultiLayer Perceptron (MLP) classifier was trained on the 3 samples. For each sample, the predicted labels were projected back onto the original measurement positions, resulting in a spatial classification map. This representation allows the model performance to be evaluated not only through the numerical accuracy but also through the spatial agreement between the predicted classes and the known defect configuration of each specimen. For the healthy reference block, all evaluated

measurement points were classified as healthy/on-slab, resulting in 684/684 correct predictions and an accuracy of 100%. The classification map therefore contains no false void or delamination regions, as shown in Figure 23.





Figure 23: MLP-based spatial classification result for the healthy reference specimen. All evaluated measurement points were classified as healthy/on-slab, resulting in 684/684 correct predictions and an accuracy of 100.0%. This indicates that the model did not introduce false defect regions in the intact specimen under the controlled laboratory conditions of this study.

For the void specimen, the classifier achieved 1102/1134 correct predictions, corresponding to an accuracy of 97.2%. The predicted over-void regions show strong spatial agreement with the known locations of the embedded void inclusions. Most regions outside the artificial defects were classified as healthy/on-slab, while the red over-void labels are concentrated in the expected defective zones, as shown in Figure 24b. The few misclassified points appear mainly as isolated points or near transitions between healthy and defective regions.



(a) Known artificial void layout before casting. *This figure was generated using AI.* (b) MLP classification map obtained from LDV measurements.


Figure 24: Comparison between the known artificial void layout and the MLP-based spatial classification result for the void specimen. The predicted over-void regions show strong spatial agreement with the embedded void locations, indicating that the frequency spectra contain defect-sensitive information related to the internal void configuration.

For the delamination sample, the classifier achieved 1132/1169 correct predictions, which corresponds to an accuracy of 96.8%. The yellow over-delamination regions form spatially coherent zones that correspond well with the known positions of the embedded polystyrene delamination inserts, as seen in Figure 25b. Similarly to the void specimen, most misclassified points occur near the boundaries between on-slab and over-defect regions or as isolated local deviations.


In general, the MLP achieved 2918/2987 correct predictions in all three specimens, corresponding to an overall accuracy of 97.7%, as shown in Table 12. The spatial classification maps show that the healthy block is classified entirely as intact, while the predicted void and delamination regions correspond well with the known embedded defect layouts.

Specimen	Correct predictions	Total samples	Accuracy
Healthy block	684	684	100.0%
Void block	1102	1134	97.2%
Delamination block	1132	1169	96.8%
Overall	2918	2987	97.7%

Table 12: Classification performance of the MLP model for the three concrete specimens.



(a) Known artificial delamination layout before casting. *This figure was generated using AI.*



(b) MLP classification map obtained from LDV measurements.

Figure 25: Comparison between the known artificial delamination layout and the MLP-based spatial classification result for the delamination specimen. The predicted over-delamination regions show strong spatial agreement with the embedded delamination inserts, indicating that the frequency spectra contain defect-sensitive information related to the internal delamination configuration.

To evaluate whether the MLP detects defect signatures in the frequency domain rather than memorizing measurement locations, an additional validation was carried out. The model was trained on measurements from only one half of the blocks and then tested on the other half, which was unknown to the model during training. When trained on the left half and tested on the right half, it achieved a correctness of 81.8%. When trained on the right half and tested on the left half, it achieved a correctness of 80.1%.

5 Discussion

5.1 Relation to earlier numerical observations

In the bachelor thesis[14] it was apparent that when there are defects present in a concrete block, the eigenfrequencies shift compared to a healthy block. To achieve a better understanding of the shifting of the eigenfrequencies on the slabs that were used for this thesis, it was clear that simulations could provide this understanding before real testing started. The results that were extracted out of Siemens NX could then be compared to the real life measurements, as shown in Table 13. Although some frequencies are not detected in real life while present in the Siemens NX simulations, it is clear that those that are present come real close to each other. The higher the frequency, the further apart the simulations and real life values get. This could be explained by the nonlinear character of vibration propagation through the concrete block, which becomes more difficult to predict in higher frequencies.

Healthy Slab		Delamination Slab	
Real Life [Hz]	Siemens NX [Hz]	Real Life [Hz]	Siemens NX [Hz]
728	728	738	679
1058	1076	1085	1001
1220	1272	1258	1195
1718	1738	1702	1615
2705	2865	2665	1618
2803	2905	2735	2615
3322	2911	2820	2621
	3051	3268	2653
			2708
			2854
Void Slab			
Real Life [Hz]	Siemens NX [Hz]		
725	719		
1062	1061		
1225	1268		
1688	1723		
2723	1730		
2797	2709		
3243	2834		
	2850		
	2884		
	3003		

Table 13: In the left columns, it shows the eigenfrequencies of the real life measurements that are below 4 kHz. The right columns show the predicted eigenfrequencies that were simulated in Siemens NX.

When comparing the eigenshapes of the real life measurements and the simulations, in Figure 26, it is clear that they are very similar. This is the case for most of the eigenshapes when compared with the eigenshapes from roughly the same eigenfrequency of the measurements. Like mentioned previously in *Results*, the measurement grid does not span over the entire block, which means that the eigen-shape cannot be visualized entirely. Another comment mentioned was that there was no reference for the phase when measuring. This means that the visualization of the real slabs will not look like the simulated modes, because it is uncertain in which direction the displacement moves for each eigenshape.

The decrease in natural frequency observed in the delamination and void slabs compared to the healthy slab is consistent with a reduction in local stiffness caused by the presence of internal discontinuities. This can indicate that the material is less stiff and that defects are present. When examining the simulated void slab, the irregularities visible in higher mode shapes, such as Figure 15b, correspond spatially to the known defect locations, which could be an indication of the defects within the slab.




Figure 26: The first eigenmode of the simulated Healthy block compared to eigenmode of the slabs in real life. They look very alike, which means that the simulation predicted the eigenshapes correctly.

5.2 Comparing the Data

5.2.1 Vibration data

The LDV measurements were mainly compared in the frequency domain because defect-sensitive information is often more clearly visible in spectral features than in the raw time-domain response. Internal defects do not appear directly in the measured vibration signal, but can influence the local dynamic response by changing the stiffness, damping, resonance behavior, and wave propagation. Void-type defects may cause localized changes in resonance peaks and amplitude levels, while delamination-type defects may lead to broader resonance features or a modified spectral energy distribution. These changes are not necessarily visible in the raw time signal, but they can become more apparent after transformation to the frequency domain, as seen in Figure 19. For example, at ~ 5800 Hz, there is a clear shift between the eigenfrequency of the healthy sample compared to the two samples with defects.

The reason why only two measurements were taken per grid point was to decrease the measurement time. As seen in Figure 17c, the effect of averaging two measurements is negligible. The purpose of the two measurements was rather a redundancy, to ensure that there was at least one good measurement for each point, rather than to decrease the noise. If both measurements were of reasonable quality, then they are both used for the training or validation when training the model.

5.2.2 Frequency domain as input for AI

The comparison of frequency-domain spectra provides the physical motivation for using them as input for AI-based models. If the spectral response above defective regions differs systematically from the response of intact concrete, then these differences can be used as defect-sensitive features. For the autoencoder, such deviations were expected to result in higher reconstruction errors. For the MLP classifier, the same spectral differences were used to distinguish between healthy regions, void-type defects, and delamination-type defects.

5.3 Use of the types of AI models

In this paper, two types of models were used: the 1D-autoencoder and MultiLayer Perceptron (MLP) classifier. Both models have a different output and working principle, which have their own advantages and disadvantages. Both of these models are used in recent machine learning NDT approaches[63, 64, 65].

The autoencoder’s main advantage is that it does not need a lot of data to get good results. The disadvantage of the autoencoder is that being a binary anomaly detector, it can only say if a measurement is healthy or damaged. It cannot say which type of defect is detected, which could be beneficial when repairing the structure. However, the autoencoder proved inefficient after numerous tests. The complexity of the measurements performed on a damaged sample is almost less complex than or equally complex as the measurements performed on a healthy sample. This makes the reconstruction error for damaged and healthy samples approximately the same, as seen in Figure 21 & 22, which means that damaged samples will not be detected with the use of the autoencoder. The distributions of the healthy and void RMS values shown in Table 11 are very close to each other, which confirms this: the damaged signal has less or equal complexity compared to the healthy signal, meaning that when the autoencoder is able to reconstruct the healthy measurements, it will automatically be able to reconstruct the damaged ones as well. Even in the frequency domain, the complexity of the damaged signal is less than that of the healthy signal, which makes it easier to reconstruct for the autoencoder, therefore not detecting any damages.

However, even if the autoencoder worked correctly, it is necessary to keep in mind that the interpretation of the encoder output requires caution. A high reconstruction error does not necessarily prove the presence of a defect, since it may also be caused by measurement noise, local surface variability, imperfect excitation repeatability, laser alignment effects, or boundary-condition differences. For this reason, the anomaly maps should be interpreted together with the known defect layout, repeated measurements, and physical understanding of the experimental setup. The autoencoder-based model should therefore be regarded as a screening tool that highlights suspicious regions for further inspection, rather than as a standalone diagnostic method.

After concluding that the autoencoder would not work for this experiment, the switch was made to the MLP model. The advantage of the MLP is that it can classify the types of defects. The biggest disadvantage is that it needs more data than the autoencoder, and the data has to be labeled. If the location of the defects inside the block were unknown, a different approach should be considered. The MLP proved to be a good choice, with a total accuracy of 97.7%.

However, the MLP approach requires careful interpretation. Since it is a supervised model, its performance depends strongly on the quality and representativeness of the labeled training data. With only a limited number of specimens and artificially introduced defects, the model can learn specimen-specific patterns rather than general defect characteristics. Therefore, the MLP should be interpreted as an exploratory classification model in this proof-of-concept study, rather than as a fully validated defect-identification method.

The validation experiment in which the model was trained on one half of the blocks and tested on the other showed correctness values of 81.8% and 80.1% respectively. This indicates that the model is not memorizing measurement positions but is instead detecting spectral signatures of the defects.

It is interesting to notice that there is a gain of $\sim 16\%$ achieved when the model is trained on the full block. This means that the model needs enough measurements to learn the difference between defect signatures and position-dependent spectral variations.

The autoencoder and MLP classifier are two models that were used for this experiment, but in the literature research, it is clear that these are not the only feasible solutions. The autoencoder [63] and MLP [64, 65] are used in other studies, but most of the studies found use 1D-Convolutional Neural Networks (CNN)[66, 67], 2D-CNN[67, 41], Support Vector Machine (SVM)[67], Extreme Learning Machine (ELM)[68], or other models[65].

The MLP of this model obtained an accuracy of 97.7%. This result should be compared with other studies to ensure that it is an expected result. Table 14 shows the comparison between the accuracy of other studies compared to the model used in this experiment. A total accuracy of 97.7% seems like a good result, especially compared to Dethof et al.[64] who also used an MLP for defect detection.

Paper	Model used	Accuracy
This paper	MLP	97.7% (total)
Dethof et al. 2023 [64]	MLP	90% - 97%
Dorafshan et al. 2020 [66]	1D-CNN	88%
Zhang et al. 2016 [68]	Extreme Learning Machine (ELM)	$\sim 99\%$
Qu et al. 2025 [41]	2D-CNN	98.4%
Jafari et al. 2022 [67]	Support Vector Machine (SVM)	97%

Table 14: Comparison between the accuracy of the MLP of this paper and the accuracy of other studies done.

5.4 Limitations and future work

5.4.1 Limitations

Several limitations of the present study should be acknowledged. First, the experimental dataset was obtained from a limited number of concrete specimens representing one healthy condition and two artificial damage configurations. Although this design is appropriate for a controlled proof-of-concept investigation, it does not capture the full variability encountered in real concrete structures, where boundary conditions, environmental influences, and defect morphology may vary substantially.

Second, the investigated defects were artificially introduced and therefore simplified approximations of void and delamination-type damage. While such controlled defect configurations are useful for establishing a measurable contrast between healthy and damaged states, they cannot fully reproduce the geometric complexity, and uncertain formation mechanisms of naturally occurring defects in practice.

Third, all measurements were performed under controlled laboratory conditions using a fixed experimental setup, automated excitation, and predefined scan grids. The resulting data are therefore well suited for methodological development, but the robustness of the proposed approach under field conditions remains to be demonstrated. In particular, practical deployment may be affected by environmental noise, surface condition, accessibility, alignment constraint, and variability in excitation coupling.

A further limitation concerns the excitation process itself. Although the automated pneumatic system was designed to improve repeatability, impact-based measurements may still exhibit variability in excitation amplitude and local contact conditions. Such variability can influence the measured spectral response and may contribute to increased errors. For this reason, further refinement of excitation remains important for future work.

Another limitation is related to the AI-based interpretation itself. The supervised MLP classifier achieved high accuracy under the controlled conditions of this study, but was trained and evaluated on a limited dataset with known artificial defect locations. As a result, the model may still be influenced by sample-specific patterns, measurement-grid effects, or laboratory-specific conditions rather than only by general defect characteristics. This is particularly relevant because the investigated defects were artificially introduced and the number of concrete specimens was limited. Therefore, the reported classification performance should be interpreted as evidence of feasibility under controlled conditions, rather than as proof of general applicability to real concrete structures. Future work should validate the model on larger and more diverse datasets, including different concrete mixtures, defect depths, defects sizes, boundary conditions, and naturally occurring damage.

Finally, the present results should be interpreted as an initial step toward automated LDV-based screening of concrete defects rather than as a fully generalized inspection solution. Further work should therefore include larger datasets, additional specimens, more realistic defect conditions, improved normalization strategies, and validation under less controlled measurement environments.

5.4.2 Future work

A promising future direction is the integration of the proposed LDV-based inspection setup on a remotely operated or semi-autonomous robotic platform for field applications, especially in tunnels and other hazardous or difficult-to-access environments. In the present study, non-contact LDV measurements, automated projectile excitation, grid-based scanning, and AI-based anomaly detection were combined in a controlled laboratory setup. This workflow could be extended to a mobile platform carrying the LDV sensor, projectile launcher, control electronics, power supply, and wireless communication unit.

During operation, the robot could position itself in front of a tunnel-lining segment, align the LDV with a predefined measurement grid, trigger the projectile launcher to generate repeatable mechanical impact, and transmit the measured vibration response to a remote operator station for AI-based analysis. Such a system would reduce the need for personnel to enter hazardous zones, for example, in tunnels with restricted access or possible asbestos contamination, while also improving measurement repeatability and enabling more standardized inspection of large concrete surfaces.

This concept is consistent with previous developments in robotic tunnel inspection, where autonomous platforms have been investigated for navigation, visual defect detection, 3D reconstruction, and structural assessment of tunnel linings [69, 70]. In contrast to purely vision-based systems, the proposed LDV-based approach would provide information on the local dynamic response of the concrete, which may be sensitive to internal defects such as voids, cavities, or delaminations before they become visible on the surface. Recent studies on remote non-contact LDV measurements for tunnel-lining cavities further support the feasibility of this direction [16].

Nevertheless, several challenges must be addressed before practical field deployment is possible. These include robotic platform stability, LDV alignment on curved or rough tunnel surfaces, safe containment of the projectile excitation system, wireless synchronization, environmental noise, and validation under realistic field conditions. Therefore, such a robotic LDV inspection system should be regarded as a promising long-term extension of the present laboratory setup, rather than as an already validated inspection solution. An overview of the proposed future robotic inspection concept is shown in Figure 27.

A further limitation is that the present excitation mechanism is not fully non-destructive. Although the vibration response is measured without contact using the LDV, the pneumatic launcher uses steel spheres as projectiles, which produced local surface damage to the concrete specimens. Consequently, the current system should be interpreted as a non-contact vibration measurement method with mechanical projectile excitation, rather than as a fully non-destructive testing method.

Future work should therefore investigate less invasive excitation strategies. Possible options include lower-density or more compliant projectiles, such as polymer, rubber-coated, or ice-based spheres.




Figure 27: Conceptual future-work extension of the LDV-based inspection setup toward a remotely operated robotic tunnel-inspection platform. The mobile platform carries the LDV sensor, automated projectile launcher, control electronics, and wireless communication unit. At each inspection location, the robot aligns the LDV with a predefined measurement grid, triggers a controlled projectile impact, records the vibration response, and transmits the data to a remote operator station for AI-based anomaly mapping and condition assessment. *This figure was generated using AI.*

In impact testing, the stiffness and material of the impact interface influence the duration of contact, the force amplitude, and the frequency content of the excitation pulse [71]. Similar optimization could be applied to the projectile launcher by varying projectile material, diameter, velocity, and impact angle, with the aim of minimizing surface indentation while preserving sufficient signal-to-noise ratio and repeatability. Alternatively, fully contactless excitation methods, such as laser ablation or laser-induced plasma shock waves, could be explored, since these have already been investigated for concrete defect detection in combination with vibration-based measurements [13, 72].



Figure 28: Surface damage observed after repeated projectile excitation with steel spheres. Although the LDV measurement itself is non-contact, the current excitation mechanism produced local impact marks and small surface indentations on the concrete specimen. This illustrates that the present configuration should be regarded as a non-contact vibration measurement setup with mechanical impact excitation, rather than as a fully non-destructive inspection method.

6 Conclusion

This work presented a proof-of-concept framework for automated damage screening in concrete structures using a Laser Doppler Vibrometer (LDV), projectile-based excitation, and AI-based interpretation of vibration data. The proposed setup combines a non-contact LDV measurement system with an automated pneumatic steel-ball launcher, allowing point-by-point excitation and vibration measurement over a predefined scan grid. By transforming the acquired time-domain signals into frequency-domain spectra, each measurement point can be associated with a local vibration response that is suitable for further data-driven analysis.

A controlled experimental setup was designed using three concrete samples: a healthy reference block, a block containing artificial void-type defects, and a block containing artificial delamination-type defects. This specimen configuration provided a structured basis for investigating whether changes in the local dynamic response can be used to distinguish between intact and defective concrete. The use of a predefined measurement grid and an automated excitation system improved the repeatability and consistency of the measurement procedure compared to fully manual impact testing.

These specimens were replicated in Siemens NX to simulate the vibrational response. These simulations gave information on the expected eigenfrequencies and their corresponding eigenshapes. These were later compared to the real life measurements. The simulated frequencies lay very close to the measured eigenfrequencies and the simulated eigenshapes were predicted correctly.

Two AI-based strategies were investigated and showed clearly different performances. First, an autoencoder-based anomaly detection approach was evaluated to learn the spectral response of healthy concrete and identify deviations from this learned baseline. Although this strategy is attractive when only limited defective data are available, it did not provide reliable anomaly detection for the present dataset. Healthy and damaged measurements were reconstructed with comparable regions from intact concrete. As a result, the autoencoder could not be used as a standalone damage-detection method in this study.

In contrast, the supervised Multilayer Perceptron (MLP) classifier provided substantially better performance for defect-type classification. Using selected frequency-domain LDV features, the MLP was able to distinguish between healthy regions, void-type defects, and delamination-type defects. The model achieved an overall classification accuracy of 97.7%, with an classification accuracy of 100.0% for the healthy sample, 97.2% for the void sample, and 96.8% for the delamination sample. Moreover, the resulting spatial classification maps showed good agreement with the known artificial defect layouts, indicating that the measured frequency spectra contained defect-sensitive information relevant for supervised classification under controlled laboratory conditions.

The proposed methodology demonstrates the potential to combine automated excitation, non-contact vibration measurement, and AI-based signal interpretation for concrete inspection. Rather than relying solely on visual inspection or subjective manual hammer sounding, the developed workflow assigns a quantitative vibration-based response to each measurement point. This enables the construction of spatial maps that can support the localization and interpretation of potentially damaged regions. In particular, the spatial MLP predictions suggest that the proposed workflow is not limited to global damage screening, but can also provide local defect-type information when representative labeled training data are available.

Several limitations should be acknowledged. The experimental dataset was limited to a small number of laboratory specimens with artificial defect configurations. Therefore, the present work should be interpreted as an initial proof of concept rather than as a fully generalized inspection method. In addition, although the LDV measurement itself is non-contact, the current excitation mechanism uses steel projectiles and can produce local surface damage. As a result, the present setup should not yet be considered fully non-destructive in the strict practical sense. Future work should therefore investigate less invasive excitation strategies, such as polymer, rubber-coated, or ice-based projectiles, as well as fully contactless excitation methods. Furthermore, the high classification accuracy of the MLP should be interpreted with caution, since the model was trained and evaluated on a limited dataset with known artificial defect locations. The results therefore demonstrate feasibility under controlled conditions, but do not yet prove generalization to real structures with unknown defect geometries, variable material properties, changing boundary conditions, or environmental disturbances.

Further development should also focus on larger datasets, additional defect types, more realistic boundary conditions, improved normalization strategies, and validation under field-like conditions. A promising long-term direction is the integration of the LDV sensor, excitation system, and AI-based analysis pipeline on a remotely operated robotic platform for inspection of tunnels and infrastructure. Such a system could reduce operator exposure in hazardous or inaccessible environments while improving measurement repeatability and inspection standardization.

Overall, this study demonstrates that automated projectile excitation combined with non-contact LDV measurements can generate spatially resolved vibration data suitable for AI-based interpretation. Although the autoencoder did not provide reliable anomaly detection for the present dataset, the MLP results show that the LDV features in the frequency-domain can support defect-type classification in controlled concrete specimens. Therefore, the work represents a step toward automated, quantitative and spatially resolved inspection of large concrete structures. Future research should focus on larger and more diverse datasets, naturally occurring defects, less invasive excitation strategies, and validation under field-like conditions before the approach can be considered for practical structural health monitoring applications.

7 References

- [1] Christoph Kralovec and Martin Schagerl. “Review of Structural Health Monitoring Methods Regarding a Multi-Sensor Approach for Damage Assessment of Metal and Composite Structures”. In: *Sensors* 20.3 (2020), p. 826. DOI: [10.3390/s20030826](https://doi.org/10.3390/s20030826).
- [2] Diego A. Tibaduiza Burgos et al. “Damage Identification in Structural Health Monitoring: A Brief Review from its Implementation to the Use of Data-Driven Applications”. In: *Sensors* 20.3 (2020), p. 733. DOI: [10.3390/s20030733](https://doi.org/10.3390/s20030733).
- [3] Onur Avci et al. “A review of vibration-based damage detection in civil structures: From traditional methods to Machine Learning and Deep Learning applications”. In: *Mechanical Systems and Signal Processing* 147 (2021), p. 107077. DOI: [10.1016/j.ymssp.2020.107077](https://doi.org/10.1016/j.ymssp.2020.107077).
- [4] Jing Jia and Ying Li. “Deep Learning for Structural Health Monitoring: Data, Algorithms, Applications, Challenges, and Trends”. In: *Sensors* 23.21 (2023), p. 8824. DOI: [10.3390/s23218824](https://doi.org/10.3390/s23218824).
- [5] Masayasu Muramatsu, Shigeaki Uchida, and Yasuo Takahashi. “Noncontact Detection of Concrete Flaws by Neural Network Classification of Laser Doppler Vibrometer Signals”. In: *Engineering Research Express* 2.2 (2020), p. 025017. DOI: [10.1088/2631-8695/ab8ba4](https://doi.org/10.1088/2631-8695/ab8ba4).
- [6] Hossein Taheri, Maria Gonzalez Bocanegra, and Mohammad Taheri. “Artificial Intelligence, Machine Learning and Smart Technologies for Nondestructive Evaluation”. In: *Sensors* 22.11 (2022), p. 4055. DOI: [10.3390/s22114055](https://doi.org/10.3390/s22114055).
- [7] Mohsen Azimi, Armin Dadras Eslamlou, and Gokhan Pekcan. “Data-Driven Structural Health Monitoring and Damage Detection through Deep Learning: State-of-the-Art Review”. In: *Sensors* 20.10 (2020), p. 2778. DOI: [10.3390/s20102778](https://doi.org/10.3390/s20102778).
- [8] B. Koehler, G. Hentges, and W. Mueller. “Improvement of ultrasonic testing of concrete by combining signal conditioning methods, scanning laser vibrometer and space averaging techniques”. In: *NDT & E International* 31.4 (1998), pp. 281–287. DOI: [10.1016/S0963-8695\(98\)00005-X](https://doi.org/10.1016/S0963-8695(98)00005-X).
- [9] D. Algernon et al. “Imaging of the elastic wave propagation in concrete using scanning techniques: Application for impact-echo and ultrasonic echo methods”. In: *Journal of Nondestructive Evaluation* 27.1–3 (2008), pp. 83–97. DOI: [10.1007/s10921-008-0034-4](https://doi.org/10.1007/s10921-008-0034-4).
- [10] Ryo Akamatsu et al. “Proposal of Non Contact Inspection Method for Concrete Structures Using High-Power Directional Sound Source and Scanning Laser Doppler Vibrometer”. In: *Japanese Journal of Applied Physics* 52.7S (2013), 07HC12. DOI: [10.7567/JJAP.52.07HC12](https://doi.org/10.7567/JJAP.52.07HC12).
- [11] Tsuneyoshi Sugimoto. “Non-Destructive Inspection of Concrete Structures by Noncontact Acoustic Inspection Technology Using Acoustic Irradiation Induced Vibration”. In: *IEICE ESS Fundamentals Review* 16.3 (2023), pp. 156–166. DOI: [10.1587/essfr.16.3_156](https://doi.org/10.1587/essfr.16.3_156).
- [12] Naotoshi Yasuda et al. “Detection and characteristics estimation of defects in concrete structures using laser ablation-induced vibration”. In: *Tunnelling and Underground Space Technology* 103 (2020), p. 103460. DOI: [10.1016/j.tust.2020.103460](https://doi.org/10.1016/j.tust.2020.103460).
- [13] Johannes F. Scherr et al. “Detection of Delaminations in Concrete Plates Using a Laser Ablation Impact Echo Technique”. In: *Journal of Nondestructive Evaluation* 42 (2023), p. 11. DOI: [10.1007/s10921-022-00921-x](https://doi.org/10.1007/s10921-022-00921-x).
- [14] Michiel Arnouts et al. *Non-Contact Damage Detection in Concrete Using Laser Doppler Vibrometry and Various Excitation Methods*. Preprints.org, Preprint, 2026. DOI: [10.3390/metrology6020035](https://doi.org/10.3390/metrology6020035).
- [15] Xu Tian et al. “An Internal Defect Detection Algorithm for Concrete Blocks Based on Local Mean Decomposition-Singular Value Decomposition and Weighted Spatial-Spectral Entropy”. In: *Entropy* 25.7 (2023), p. 1034. DOI: [10.3390/e25071034](https://doi.org/10.3390/e25071034).
- [16] Hongyun Yang et al. “Experimental research on remote non-contact laser vibration measurement for tunnel lining cavities”. In: *Scientific Reports* 15 (2025), p. 105. DOI: [10.1038/s41598-024-83819-0](https://doi.org/10.1038/s41598-024-83819-0).
- [17] Kazuko Sugimoto et al. “Detection of resonance frequency of both the internal defects of concrete and the laser head of a laser Doppler vibrometer by spatial spectral entropy for noncontact acoustic inspection”. In: *Japanese Journal of Applied Physics* 58 (2019), SGGB15. DOI: [10.7567/1347-4065/ab1a2f](https://doi.org/10.7567/1347-4065/ab1a2f).

- [18] Kazuko Sugimoto and Tsuneyoshi Sugimoto. “Detection of Internal Defects in Concrete and Evaluation of a Healthy Part of Concrete by Noncontact Acoustic Inspection Using Normalized Spectral Entropy and Normalized SSE”. In: *Entropy* 24.2 (2022), p. 142. DOI: [10.3390/e24020142](https://doi.org/10.3390/e24020142).
- [19] Marek Słoiński, Krzysztof Schabowicz, and Ewa Krawczyk. “Detection of Flaws in Concrete Using Ultrasonic Tomography and Convolutional Neural Networks”. In: *Materials* 13.7 (2020), p. 1557. DOI: [10.3390/ma13071557](https://doi.org/10.3390/ma13071557).
- [20] Tianyu Hu et al. “Ultrasonic based concrete defects identification via wavelet packet transform and GA-BP neural network”. In: *PeerJ Computer Science* 7 (2021), e635. DOI: [10.7717/peerj-cs.635](https://doi.org/10.7717/peerj-cs.635).
- [21] Jinhui Zhao, Tianyu Hu, and Qichun Zhang. “A Wavelet Packet Transform and Convolutional Neural Network Method Based Ultrasonic Detection Signals Recognition of Concrete”. In: *Sensors* 22.10 (2022), p. 3863. DOI: [10.3390/s22103863](https://doi.org/10.3390/s22103863).
- [22] Xu Tian et al. “Concrete Multi-Type Defect Classification Algorithm Based on MSSMA-SVM”. In: *Sensors* 22.23 (2022), p. 9145. DOI: [10.3390/s22239145](https://doi.org/10.3390/s22239145).
- [23] Jasper Laforce. “Simulating the frequency response of decaying concrete slabs vs healthy concrete slabs in Siemens NX”. Unpublished manuscript, Mechanical Systems, Universiteit Antwerpen. 2025.
- [24] OVAM. *Wat is verboden?* 2026.
- [25] OVAM. *Asbest en uw gezondheid.* 2026.
- [26] ACI Committee 228. *Report on Nondestructive Test Methods for Evaluation of Concrete in Structures*. Tech. rep. ACI 228.2R-13. Farmington Hills, MI: American Concrete Institute, 2013.
- [27] V. M. Malhotra and Nicholas J. Carino, eds. *Handbook on Nondestructive Testing of Concrete*. 2nd ed. Boca Raton, FL: CRC Press, 2004. ISBN: 9780849314858.
- [28] ASTM International. *ASTM C597-22: Standard Test Method for Ultrasonic Pulse Velocity Through Concrete*. West Conshohocken, PA: ASTM International, 2022. URL: <https://www.astm.org/C0597-22.html>.
- [29] Jinhui Zhao et al. “Defect Recognition in Concrete Ultrasonic Detection Based on Wavelet Packet Transform and Stochastic Configuration Networks”. In: *IEEE Access* 9 (2021), pp. 9284–9295. DOI: [10.1109/ACCESS.2021.3050715](https://doi.org/10.1109/ACCESS.2021.3050715).
- [30] Lara Guizi Anoni, Vladimir Guilherme Haach, and Lev Khazanovich. “Image Reconstruction in Concrete Ultrasound Tomography: A Systematic Review”. In: *Construction and Building Materials* 441 (2024), p. 137472. DOI: [10.1016/j.conbuildmat.2024.137472](https://doi.org/10.1016/j.conbuildmat.2024.137472).
- [31] M. R. Clark, D. M. McCann, and M. C. Forde. “Application of Infrared Thermography to the Non-Destructive Testing of Concrete and Masonry Bridges”. In: *NDT & E International* 36.4 (2003), pp. 265–275. DOI: [10.1016/S0963-8695\(02\)00060-9](https://doi.org/10.1016/S0963-8695(02)00060-9).
- [32] J. H. A. Rocha and Y. V. Póvoas. “Infrared Thermography as a Non-Destructive Test for the Inspection of Reinforced Concrete Bridges: A Review of the State of the Art”. In: *Revista AL-CONPAT* 7.3 (2017), pp. 200–214. DOI: [10.21041/ra.v7i3.223](https://doi.org/10.21041/ra.v7i3.223).
- [33] Seong-Hoon Kee et al. “Nondestructive Bridge Deck Testing with Air-Coupled Impact-Echo and Infrared Thermography”. In: *Journal of Bridge Engineering* 17.6 (2012), pp. 928–939. DOI: [10.1061/\(ASCE\)BE.1943-5592.0000350](https://doi.org/10.1061/(ASCE)BE.1943-5592.0000350).
- [34] Mary J. Sansalone and William B. Streett. *Impact-Echo: Nondestructive Evaluation of Concrete and Masonry*. Jersey Shore, PA: Bullbrier Press, 1997. ISBN: 9780961261061.
- [35] Nicholas J. Carino. “The Impact-Echo Method: An Overview”. In: *Proceedings of the 2001 Structures Congress and Exposition*. Washington, DC, 2001.
- [36] ASTM International. *ASTM C1383-15(2022): Standard Test Method for Measuring the P-Wave Speed and the Thickness of Concrete Plates Using the Impact-Echo Method*. West Conshohocken, PA: ASTM International, 2022. URL: <https://www.astm.org/c1383-15r22.html>.
- [37] Paolo Castellini, Milena Martarelli, and Enrico P. Tomasini. “Laser Doppler Vibrometry: Development of Advanced Solutions Answering to Technology’s Needs”. In: *Mechanical Systems and Signal Processing* 20.6 (2006), pp. 1265–1285. DOI: [10.1016/j.ymssp.2005.11.015](https://doi.org/10.1016/j.ymssp.2005.11.015).

- [38] Jing-Kui Zhang, Weizhong Yan, and De-Mi Cui. “Concrete Condition Assessment Using Impact-Echo Method and Extreme Learning Machines”. In: *Sensors* 16.4 (2016), p. 447. DOI: [10.3390/s16040447](https://doi.org/10.3390/s16040447).
- [39] Faezeh Jafari and Sattar Dorafshan. “Comparison between Supervised and Unsupervised Learning for Autonomous Delamination Detection Using Impact Echo”. In: *Remote Sensing* 14.24 (2022), p. 6307. DOI: [10.3390/rs14246307](https://doi.org/10.3390/rs14246307).
- [40] Sattar Dorafshan and Hoda Azari. “Deep Learning Models for Bridge Deck Evaluation Using Impact Echo”. In: *Construction and Building Materials* 263 (2020), p. 120109. DOI: [10.1016/j.conbuildmat.2020.120109](https://doi.org/10.1016/j.conbuildmat.2020.120109).
- [41] Feng Qu et al. “Structural damage identification and classification from time-frequency impact echo data using deep learning”. In: *Nondestructive Testing and Evaluation* 40.11 (2025), pp. 5341–5362. DOI: [10.1080/10589759.2024.2440816](https://doi.org/10.1080/10589759.2024.2440816).
- [42] Michael Kobusch, Ludwig Klaus, and Luis Muñiz Mendoza. “Investigation of Impact Hammer Calibrations”. In: *Proceedings of the 23rd IMEKO TC3, 13th TC5 and 4th TC22 International Conference*. Helsinki, Finland, 2017.
- [43] Bernd Koehler, G. Hentges, and W. Mueller. “Improvement of Ultrasonic Testing of Concrete by Combining Signal Conditioning Methods, Scanning Laser Vibrometer and Space Averaging Techniques”. In: *NDT & E International* 31.4 (1998), pp. 281–287. DOI: [10.1016/S0963-8695\(98\)00005-X](https://doi.org/10.1016/S0963-8695(98)00005-X).
- [44] Yutaka Nakagawa et al. “Study on Concrete Internal Defect Detection from a Moving Cart Equipped with LDVs and a Sound Source”. In: *2023 IEEE International Ultrasonics Symposium (IUS)*. IEEE, 2023, pp. 1–4. DOI: [10.1109/IUS51837.2023.10306719](https://doi.org/10.1109/IUS51837.2023.10306719).
- [45] Xu Tian et al. “An Internal Defect Detection Algorithm for Concrete Blocks Based on Local Mean Decomposition-Singular Value Decomposition and Weighted Spatial-Spectral Entropy”. In: *Entropy* 25.7 (2023), p. 1034. DOI: [10.3390/e25071034](https://doi.org/10.3390/e25071034).
- [46] Hongyun Yang et al. “Experimental Research on Remote Non-Contact Laser Vibration Measurement for Tunnel Lining Cavities”. In: *Scientific Reports* 15 (2025), p. 105. DOI: [10.1038/s41598-024-83819-0](https://doi.org/10.1038/s41598-024-83819-0).
- [47] Kazuko Sugimoto et al. “Detection of Resonance Frequency of Both the Internal Defects of Concrete and the Laser Head of a Laser Doppler Vibrometer by Spatial Spectral Entropy for Noncontact Acoustic Inspection”. In: *Japanese Journal of Applied Physics* 58 (2019), SGGB15. DOI: [10.7567/1347-4065/ab1a2f](https://doi.org/10.7567/1347-4065/ab1a2f).
- [48] Julius S. Bendat and Allan G. Piersol. *Random Data: Analysis and Measurement Procedures*. 4th ed. Hoboken, NJ: Wiley, 2010. ISBN: 9780470248775.
- [49] David J. Ewins. *Modal Testing: Theory, Practice and Application*. 2nd ed. Baldock, England: Research Studies Press, 2000. ISBN: 9780863802188.
- [50] Alan V. Oppenheim and Ronald W. Schaffer. *Discrete-Time Signal Processing*. 3rd ed. Upper Saddle River, NJ: Pearson, 2010. ISBN: 9780131988422.
- [51] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. New York: Springer, 2006. ISBN: 978-0-387-31073-2.
- [52] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Cambridge, MA: MIT Press, 2016. ISBN: 978-0-262-03561-3. URL: <https://www.deeplearningbook.org/>.
- [53] Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. “Rectifier Nonlinearities Improve Neural Network Acoustic Models”. In: *Proceedings of the 30th International Conference on Machine Learning*. 2013. URL: https://ai.stanford.edu/~amaas/papers/relu_hybrid_icml2013_final.pdf.
- [54] Dan Hendrycks and Kevin Gimpel. “Gaussian Error Linear Units (GELUs)”. In: *arXiv preprint arXiv:1606.08415* (2016). URL: <https://arxiv.org/abs/1606.08415>.
- [55] Nitish Srivastava et al. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958. URL: <https://jmlr.org/papers/v15/srivastava14a.html>.

- [56] Sergey Ioffe and Christian Szegedy. “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift”. In: *Proceedings of the 32nd International Conference on Machine Learning*. Vol. 37. Proceedings of Machine Learning Research. 2015, pp. 448–456. URL: <https://proceedings.mlr.press/v37/ioffe15.html>.
- [57] Geoffrey E. Hinton and Ruslan R. Salakhutdinov. “Reducing the Dimensionality of Data with Neural Networks”. In: *Science* 313.5786 (2006), pp. 504–507. DOI: [10.1126/science.1127647](https://doi.org/10.1126/science.1127647).
- [58] Yang Deng et al. “Abnormal data detection for structural health monitoring: State-of-the-art review”. In: *Developments in the Built Environment* 17 (2024), p. 100337. DOI: [10.1016/j.dibe.2024.100337](https://doi.org/10.1016/j.dibe.2024.100337).
- [59] Jan Stoop. *Labo beproeven van bouwmaterialen: Practica B1+B2*. Last accessed 5 December 2025. 2024.
- [60] Wang Gong-bao et al. “A precise algorithm for non-integer harmonics analysis based on FFT and neural network”. en. In: *Wuhan Univ. J. Nat. Sci.* 9.4 (July 2004), pp. 454–458. DOI: [10.1007/BF02830442](https://doi.org/10.1007/BF02830442).
- [61] Kui Fu Chen and Yan Feng Li. “Combining the Hanning windowed interpolated FFT in both directions”. en. In: *Comput. Phys. Commun.* 178.12 (June 2008), pp. 924–928. DOI: [10.1016/j.cpc.2008.02.008](https://doi.org/10.1016/j.cpc.2008.02.008).
- [62] Gamze Cabadag, Ali Degirmenci, and Omer Karal. “Comparative analysis of interpolation techniques for FFT-based frequency estimation”. In: *J. Univers. Comput. Sci.* 32.2 (Feb. 2026), pp. 267–285.
- [63] Chi Lu and Yoshimi Sonoda. “Application of light-weighted CNN for diagnosis of internal concrete defects using hammering sound”. In: *Nondestructive Testing and Evaluation* 39.8 (2024), pp. 2426–2449. DOI: [10.1080/10589759.2023.2301032](https://doi.org/10.1080/10589759.2023.2301032).
- [64] Florian Dethof and Sylvia Kessler. “Automated honeycomb detection during Impact Echo inspections in concrete using AI trained by simulation data”. In: *ReJNDT* 1.1 (2023).
- [65] Addison Salazar and Arturo Serrano. “Automatic classification of impact-echo spectra I”. In: *Encyclopedia of Artificial Intelligence*. IGI Global, 2009, pp. 192–198.
- [66] Sattar Dorafshan and Hoda Azari. “Deep learning models for bridge deck evaluation using impact echo”. In: *Construction and Building Materials* 263 (2020), p. 120109. DOI: [10.1016/j.conbuildmat.2020.120109](https://doi.org/10.1016/j.conbuildmat.2020.120109).
- [67] Faezeh Jafari and Sattar Dorafshan. “Comparison between Supervised and Unsupervised Learning for Autonomous Delamination Detection Using Impact Echo”. In: *Remote Sensing* 14.24 (2022), p. 6307. DOI: [10.3390/rs14246307](https://doi.org/10.3390/rs14246307).
- [68] Jing-Kui Zhang, Weizhong Yan, and De-Mi Cui. “Concrete Condition Assessment Using Impact-Echo Method and Extreme Learning Machines”. In: *Sensors* 16.4 (2016), p. 447. DOI: [10.3390/s16040447](https://doi.org/10.3390/s16040447).
- [69] E. Protopapadakis et al. “Autonomous robotic inspection in tunnels”. In: *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* III-5 (2016), pp. 167–174. DOI: [10.5194/isprsannals-III-5-167-2016](https://doi.org/10.5194/isprsannals-III-5-167-2016).
- [70] Juan G. Victores et al. “Robot-aided tunnel inspection and maintenance system by vision and proximity sensor integration”. In: *Automation in Construction* 20.5 (2011), pp. 629–636. DOI: [10.1016/j.autcon.2010.12.005](https://doi.org/10.1016/j.autcon.2010.12.005).
- [71] Michael Kobusch et al. “Investigation of Impact Hammer Calibrations”. In: *Proceedings of the 23rd IMEKO TC3 Conference*. 2017, pp. 1–6.
- [72] Sho Wakata et al. “Defect detection of concrete in infrastructure based on Rayleigh wave propagation generated by laser-induced plasma shock waves”. In: *International Journal of Mechanical Sciences* 218 (2022), p. 107039. DOI: [10.1016/j.ijmecsci.2021.107039](https://doi.org/10.1016/j.ijmecsci.2021.107039).